

ROLL Working Group
Internet-Draft

Intended status: Informational
Expires: July 13, 2014

A. Junior
R. Sofia
COPELABS, University Lusofona
January 10, 2014

Energy-awareness metrics global applicability guidelines
draft-ajunior-roll-energy-awareness-01

Abstract

This document describes a new set of energy-awareness metrics which have been devised to be applicable to any multihop routing protocol having in mind LLNs, including the Routing for Low Power and Lossy Networks (RPL) protocol.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 13, 2014.

Copyright and License Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Energy-awareness Routing Metrics	4
2.1. Energy Node Ranking: ENR	4
2.2. Father-Son Association Ranking: Energy-awareness Father-Son (EFS)	5
2.3. Design Aspects	5
3. Applicability of the Proposed Metrics	5
3.1. RPL Applicability Guidelines	5
3.1.1. Impact in Node Energy Object	6
3.2. OLSR Applicability Guidelines	6
3.2.1. Impact in HELLO Messages	7
3.2.2. Impact in TC Messages	7
3.2.3. OLSR Link Tuple	8
3.2.4. Routing Table	8
3.2.5. MPR Selection	9
3.3. AODV Applicability Guidelines	9
3.3.1. Route Request (RREQ) Message Format	10
3.3.2. Route Reply (RREP) Message Format	10
3.3.3. HELLO Message Format	11
3.3.4. Route Selection	11
3.3.5. Routing Table	12
4. Acknowledgments	12
5. Security Considerations	12
6. IANA Considerations	12
7. References	12
7.1. Normative References	12
7.2. Informative References	12
Authors' Addresses	14

1. Introduction

Low Power and Lossy Networks (LLNs) routing requirements have been specified in [RFC5548], [RFC5673], [RFC5826], [RFC5867], and [RFC6719]. Additional aspects concerning routing metrics and also constraints in design are available in [RFC6551]. Path computation algorithms for single metrics have already been proposed and used in [RFC6552], and [RFC6719].

Within the context of LLNs, we consider the specific case of User-centric Networks (UCNs) [ULoop], i.e., networks partially or completely based on equipment that is owned and carried by regular Internet end-users. Concrete examples of UCNs can be Wi-Fi networks established on-the-fly after a disaster of some nature (e.g., disaster networks); a municipality network where networking nodes are provided by the Internet end-user, who is willing to share network resources (e.g. Internet access; radio spectrum) at the exchange of specific incentives.

The intention of this document is two-fold. Firstly, we describe energy-awareness metrics that can be applied to any multihop protocol currently being considered in LLNs. Secondly, we provide design guidelines concerning the applicability of such metrics for the specific case of RPL.

The effectiveness and performance validation of the metrics described in this document is out of the scope of the document, but can be found in detail in [AJUNIOR1], [AJUNIOR2] and [AJUNIOR3].

1.1. 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 [RFC2119].

This document makes use of the terminology defined in [I-D.ietf-roll-terminology]. Moreover, this document defines the following terms, in accordance with [RFC5835] terminology:

Optimal path: is defined as a path in the DAG that minimizes (or maximizes, respectively) the Rank value between any given pair of source-destination nodes, as well as its sub-paths.

Path weight: a value representing link or/and node characteristics of a path. This definition coincides with "path cost" defined in [RFC6719]. Path weight is used by RPL to compare different paths.

Idle mode: When a node is not receiving or transmitting, the node is still listening to the shared medium (overhearing) and is said to be in Idle mode.

Transmission mode: When a node is transmitting information.

Reception mode: When a node is receiving information.

Node lifetime: Corresponds to the period of time since a node becomes active until the node is said to be dead, i.e., from a network perspective, the node ceases to exist.

Network lifetime: Associated to the time period since a topology becomes active, until the topology becomes disconnected, from a destination reachability perspective.

Energy cost: The cost associated to the node or to the association between two nodes which consider the energy parameters.

2. Energy-awareness Routing Metrics

This section describes the routing metrics proposed, from an operational perspective. Conceptual aspects and validation of the metrics, as well as concrete performance indicators can be found in [AJUNIOR1], [AJUNIOR2] and [AJUNIOR3].

2.1. Energy Node Ranking: ENR

The Energy Node Ranking (ENR) metric is a node weight which ranks a node in terms of its energy consumption stability. We explore the fact that nodes may be in idle mode for a long time. Nodes that have been in idle mode for a long period of time in the past and that still have a reasonable large estimated lifetime are better candidates to be elements in an optimal path. In other words, over time we estimate how much of its lifetime has node i been in idle mode, to then provide an estimate towards the node's future energy expenditure, as this will for sure impact the node's lifetime.

Hence, we consider the total period in idle time T_Idle , over the full lifetime expected for a specific node, which is given by the sum of the elapsed time period T with the estimated lifetime of the node, as provided in equation 1. The estimated lifetime $C(i)$ consider the ratio between residual energy and drain rate which can capture the heterogeneous energy capability of nodes [J.J.GARCIA-LUNA-ACEVES].

$$ENR(i) = (T - T_Idle)/(T * C(i)) \quad (1)$$

2.2. Father-Son Association Ranking: Energy-awareness Father-Son (EFS)

Based on ENR, we consider a composition of the ENRs of both a father and successor nodes (association between two nodes), as specified in equation 2.

$$\text{EFS}(i,j) = \text{ENR}(i) * \text{ENR}(j) \quad (2)$$

EFS provides a ranking which we believe is useful to assist the routing algorithm to converge quickly in multipath environments, as the selection on which successor to consider shall be made up from, by the father node. The goal is, similarly to ENR, to improve the network lifetime without disrupting the overall network operation. Hence, the smaller EFS(i,j) is, the more likelihood a link has to become part of a path.

2.3. Design Aspects

This section describes aspects concerning the applicability of the metrics, e.g. messaging aspects.

The energy cost ranking (ENR or EFS) are recorded in reserved field of control messages of any routing protocol occupying 16 bits.

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|   Energy Cost (ENR or EFS)   |
+---+---+---+---+---+---+---+---+

```

3. Applicability of the Proposed Metrics

This section describes how the proposed metrics can be applied to the most popular multihop routing protocols in LLNs. We start by RPL applicability guidelines, to then consider OLSR (actually work in progress OLSRv2 [OLSRv2]) as a concrete example of Link-state approaches, and AODV (actually work in progress AODVv2 [AODVv2]) as a concrete example of distance-vector approaches.

3.1. RPL Applicability Guidelines

In order to use the metrics described in this document on the Routing Protocol for Low-Power and Lossy Networks (RPL), no changes or adaptation to the protocol are needed. By separating the packet processing and forwarding processes from the routing path selection, RPL provides a very flexible way of using and incorporating different metrics.

RPL operates upon the concept of Destination-Oriented Directed Acyclic Graph (DODAG), where routes are calculated from all nodes to a single destination in the topology (root node). Each node in the topology has a Rank, that is basically a value that represents its distance to the topology root.

According to specific LLNs applications, such routes are calculated in order to achieve different objectives that may be desired (e.g. minimize delay, maximize throughput, minimize energy usage), so different Objective Functions (OF) may be defined. An OF defines how routing metrics, constraints and related functions are used, in order to define the route between the nodes towards a single destination in the topology. That is, an OF, in conjunction with routing metrics and constraints, allows for the selection of a DODAG to join (if there is more than one), and a number of peers in that DODAG as parents (that is, an ordered list of parents). The OF is also responsible to compute the Rank of the node.

The [RFC6551] defines a very flexible mechanism for the advertisement of routing metrics and constraints used by RPL, even though no OF is presented. A high degree of flexibility is offered by that mechanism, and a set of routing metrics and constraints are also described in the document.

3.1.1. Impact on <object>

In order to use the metrics described in this document, the Node Energy object (NE), as defined in [RFC6551], can be used without the need for any changes or adaptation. The NE structure is composed by a set of flags (8 bits), and an 8-bits field (E_E) used for carrying the value of the estimated energy.

```

0                                     1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+-----+-----+-----+-----+
| Flags | I | T | E | Energy Cost |
+-----+-----+-----+-----+
```

To use the NE object with the metrics described in this document, the value of ENR or EFS metrics should be placed in the E_E field, and the flag 'E' (Estimation) should be set, indicating that a value for the estimated energy is provided in the E_E field. The other flags of the NE should be filled as defined in the standard.

3.2. OLSR Applicability Guidelines

The applicability of the proposed metrics does not imply significant operation changes to OLSR standard as defined in [RFC3626]. The only

change required is the creation of a special field or considering the Reserved field to hold the energy cost information of the nodes. This information will be used as basis to calculate the nodes routing tables, and must be stored in the neighbors information and in the routing table. This section describes a few design guidelines to apply the proposed metrics to OLSR.

3.2.1. Impact in <scope/context>

In OLSR, the HELLO messages are used mainly to conduct link sensing, neighbor detection and MPR selection. Therefore, to inform the other nodes about its energy-aware cost, a node sends ENR or EFS via HELLO messages. The metrics can be sent in the Reserved field in the beginning of the HELLO message body defined in the standard.

```

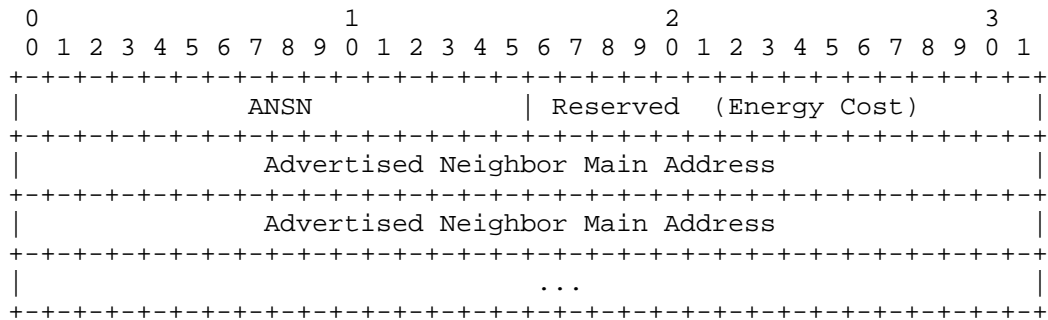
      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|   Energy Cost (ENR or EFS)   |   Htime   |   Willingness   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               ...                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

If the node is configured to use a node-based metric - ENR, then the energy cost received via HELLO messages is enough to represent the cost of the links towards those neighbors. If the node considers a Father-Son composition such as EFS then the information received is used to compute the final energy cost associated to the link based on the neighbor's energy cost and its own.

An OLSR node uses TC messages to disseminate links between itself and its neighbors. This information is spread throughout all the network, and based on this information, each node can build its own network topology. Furthermore, the topology information of each node is used to calculate its routing table.

TCs messages are used to spread the energy cost of nodes in order to compute the routing table using the Reserved field.



For each advertised link in the TC message, the Energy Cost can again be carried in the Reserved field.

3.2.3. OLSR Link Tuple

An OLSR node stores a set of information about its neighbors. This set of information, named "link tuple", is defined in [RFC3626] as (L_local_iface_addr, L_neighbor_iface_addr, L_SYM_time, L_ASYM_time, L_time), where L_local_iface_addr is the interface address of the local node, L_neighbor_iface_addr is the interface address of the neighbor node, L_SYM_time is the time until which the link is considered symmetric, L_ASYM_time is the time until which the neighbor interface is considered heard, and L_time specifies the time at which the record expires.

In order to use the energy-aware metrics defined in this document, a new field should be added to the link tuple. This extra field, named "L_energy", stores the energy cost sent by the neighbor node in the HELLO messages (in case of node-based metrics) or the calculated energy cost related to the link towards that node (in case of successor-based metrics).

When a node receives a HELLO message, the link set (set of link tuples) is updated. If the node receives a HELLO message from a neighbor node that does not exist in the link set, a new link tuple is created. In both cases, the information carried in the Energy Cost field of the HELLO message body must be considered. In case a link tuple exists, the L_energy value should be updated; if the tuple is created, the value of the L_energy field should be based on the Energy Cost field of the HELLO message received.

3.2.4 Routing Table

Each OLSR node maintains a routing table with information which allows it to route packets destined to other nodes in the network. As defined in the OLSR standard, the routing table is composed by

entries with the following information: `R_dest_addr`, `R_next_addr`, `R_dist`, `R_iface_addr`, where `R_dest_addr` is the final destination, `R_next_addr` is the next hop towards the destination, `R_dist` is the distance in number of hops, and `R_iface_addr` is the address of the local interface through which the node is reachable.

Using energy-aware metrics, the field `R_dist` no longer holds the distance in terms of hops, but in terms of energy cost. Therefore, the `R_dist` field holds the energy cost of the total path to reach that specific destination. All the other fields remain without any changes.

3.2.5. MPR Selection

The MPR selection criteria is also relevant in the contest of path computation based on the proposed Energy Cost metrics. Therefore, one simple approach (of many that can be designed) for selecting the MPRs based on the energy cost of the neighboring nodes.

Basically, when choosing the MPR, a node should take into consideration not only the number of 2-hop neighbors each of its 1-hop neighbors has; it should also take into consideration the energy cost of the neighbor nodes. Therefore, when there are more than one 1-hop neighbors covering the same number of uncovered 2-hop neighbors, the one with the lowest energy cost weight to the current node is selected as MPR.

3.3. AODV Applicability Guidelines

In contrast to link-state routing, distance-vector routing protocols work by having each node sharing its routing table with its neighbors. Routers using distance-vector protocol do not have knowledge of the entire path to a destination. Instead, distance-vector uses two key information: i) the direction in which or interface to which a packet should be forwarded; and ii) the distance from it to the final destination, where distance means number of hops.

To use energy-aware metrics, the concept of distance based on number of hops must be adapted to be based on a per-hop calculated energy cost. Therefore, the routing table of distance-vector routing protocols using energy-aware metrics does not hold the distance in number of hops to the destination; it holds the energy cost calculated for all the route from it to the destination node instead.

Energy-aware metrics can be applied to AODV without major changes. As the optimal path is chosen reactively based on the hop-count of request/reply messages, in order to use the energy cost to make a

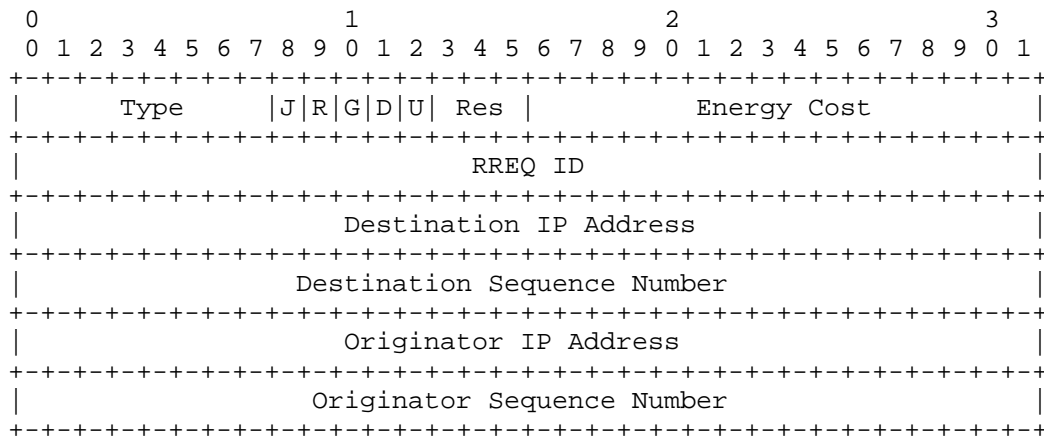
decision on the more energy efficient route from source node to destination node, the calculated energy cost value must be transmitted from one node to another, during the route discovery procedure.

The calculated energy cost, transmitted from node to node when searching for a route, is a cumulative value that represents the energy cost calculated for the path from the source until the current node.

3.3.1. Route Request (RREQ) Message Format

When a route to a new destination node is required, the source node broadcasts RREQ messages to its neighbors. Those messages are broadcasted to other nodes throughout the network until one of them eventually reaches the destination node. For energy-aware metrics, the energy cost of the route is calculated as the RREQ message is re-broadcasted; this information is carried in the RREQ messages, and when those messages reach the destination, they carry the energy cost of the entire route, from source to destination.

In order to carry the energy cost value, a slight change needs to be applied to the RREQ message format. The space originally used for the field Hop Count will be used for carrying the cumulative energy cost calculated throughout the path. The Energy Cost field will take place using the 8 bits previously used for the Hop Count value, and using 8 bits of the Reserved field. This change does not increase the packet size, not increasing the routing control overhead in the network.



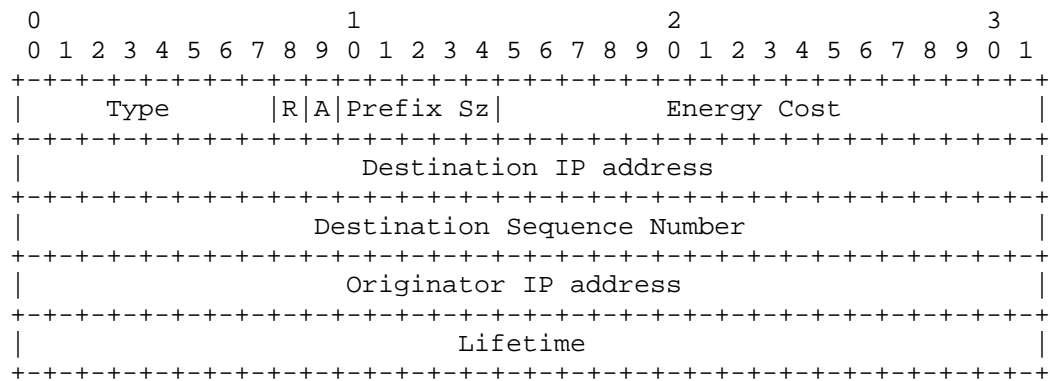
As the RREQ is broadcasted by the nodes in the network, the Energy Cost field is updated with the energy of the local node. When the RREQ message reaches the destination node, the Energy Cost field will

have the energy cost for the entire path, from source node to destination node.

3.3.2. Route Reply (RREP) Message Format

RREP messages are used when an intermediate node receives an RREQ message, but it already knows a route to the destination node specified in that message, or when an RREP message reaches the destination node. In both cases, an RREP message is created and sent back to the source node.

In order to transmit the information about the route energy cost back to the source node, the RREP message must carry a cumulative energy cost value, calculated throughout the path back to the source. This information is carried in the field Energy Cost, added to the RREP message structure, taking place of the Hop Count field of 8 bits, and taking 8 bits from the Reserved field.



As the RREP message is sent back to the node which originated the RREQ message, the Energy Cost field accumulates the energy cost calculated throughout the path. Thus, when the RREP message reaches the originator node, the Energy Cost represents the total energy cost of the path from destination back to the originator.

3.3.3. HELLO Message Format

In AODV, HELLO messages are used to offer connectivity information and also for exchange the energy cost to the case of successor based metric. HELLO messages are broadcasted locally having the same format as RREP messages, with TTL = 1, the Hop Count field set to 0, and the Destination IP Address set to its own IP address. For energy-aware metrics, the HELLO message would have the format of the RREP message described in subsection 3.3.2, and the Energy Cost field would carry

the energy cost of the node originating the message.

3.3.4. Route Selection

When a route to a new destination node is required, the source node broadcasts RREQ messages to its neighbors. Those messages are usually broadcasted by the neighbors to other nodes throughout the network until one of them eventually reaches the destination node. When an RREQ message reaches the destination (or an intermediate node that has a route to the destination), the RREQ message is not broadcasted anymore. Each intermediate node caches the information about the source of the RREQ message, in order to route back to the originator.

Through this process, the originator node selects the shortest-path based on energy cost field of the routing table to the desired destination node.

3.3.5. Routing Table

According to [RFC3561], AODV uses the following fields with each route table entry: Destination IP Address; Destination Sequence Number; Valid Destination Sequence Number flag; Other state and routing flags (e.g., valid, invalid, repairable, being repaired); Network Interface; Hop Count (number of hops needed to reach destination); Next Hop; List of Precursors; Lifetime (expiration or deletion time of the route).

For the usage of energy-aware metrics, the field Hop Count is replaced by a new field, named Energy Cost. This field holds the energy cost calculated to reach the destination, through the Next Hop specified.

4. Acknowledgments

This draft is supported by national fundings via Fundacao para Ciencia e Tecnologia (FCT), in the context of the UCR project PTDC/EEA-TEL/103637/2008. Thanks to Universidade Federal de Goias (UFG/CAC) and Fundacao de Amparo a Pesquisa do Estado de Goias (FAPEG).

5. Security Considerations

There are no new security implications related to this draft.

6. IANA Considerations

None.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP14, RFC2119, March 1997.

7.2. Informative References

[RFC3626] Clausen, T., Jacquet, P. (Ed.), "Optimized Link State Routing Protocol (OLSR)", RFC3626, October 2003.

[RFC3561] Perkins, C., Belding-Royer, E., Das, S., "Ad hoc On-Demand Distance Vector (AODV) Routing", RFC3561, July 2003.

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

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

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

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

[I-D.ietf-roll-applicability-ami] D. Popa, M. Gillmore, L. Toutain, J. Hui, R. Ruben, K. Monden, "Applicability Statement for the Routing Protocol for Low Power and Lossy Networks (RPL) in AMI Networks", draft-ietf-roll-applicability-ami-07, July, 2013.

[RFC6550] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. Vasseur, and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.

[RFC6551] JP. Vasseur, M. Kim, K. Pister, N. Dejean, D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power

and Lossy Networks", RFC 6551, March 2012.

- [RFC6551] JP. Vasseur, M. Kim, K. Pister, N. Dejean, D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, March 2012.
- [RFC6719] O. Gnawali, P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, September 2012.
- [ULOOP] "ULOOP: User-centric Wireless Local-Loop," EU IST FP7 Project (Grant 257418).
- [AJUNIOR1] A. Junior, R. Sofia, and A. Costa, "Energy-awareness metrics for multihop wireless user-centric routing" in The 2012 International Conference on Wireless Networks (ICWN'12), July 2012.
- [AJUNIOR2] A. Junior, R. Sofia, and A. Costa, "Energy-efficient heuristics for multihop routing in user-centric environments" in 12th International Conference on Next Generation Wired/Wireless Networking (NEW2AN), August 2012.
- [AJUNIOR3] A. Junior, R. Sofia, and A. Costa, "Energy-awareness in Multihop Routing" in 2012 IFIP Wireless Days conference (WD'12), November 2012.
- [I-D.ietf-roll-terminology] JP. Vasseur, "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-13.txt, September 30, 2013.
- [RFC5835] A. Morton, S. Van den Berghe, "Framework for Metric Composition", RFC 5835, April 2010.
- [J.J.GARCIA-LUNA-ACEVES] D. Kim, J. J. Garcia-Luna-Aceves, K. Obraczka, J.-C. Cano, and P. Manzoni, "Routing mechanisms for mobile ad hoc networks based on the energy drain rate," IEEE Transactions on Mobile Computing, vol. 2, no. 2, pp. 161-173, 2003.
- [OLSRv2] T. Clausen, C. Dearlove, P. Jacquet, U. Herberg, "The Optimized Link State Routing Protocol version 2", draft-ietf-manet-olsrv2-19, March 23, 2013.
- [AODVv2] C. Perkins, S. Ratliff, J. Dowdell, "Dynamic MANET On-demand (AODVv2) Routing", draft-ietf-manet-aodvv2-02, July 15, 2013.

Authors' Addresses

Antonio Junior
COPELABS, University Lusofona
Building U, 1st Floor
Campo Grande, 376
1749-024 Lisboa - Portugal
Email: antoniocojr@gmail.com

Rute Sofia
COPELABS, University Lusofona
Building U, 1st Floor
Campo Grande, 376
1749-024 Lisboa - Portugal
Email: rute.sofia@ulusofona.pt

roll
Internet-Draft
Intended status: Standards Track
Expires: September 6, 2014

Y. Doi
TOSHIBA Corporation
M. Gillmore
Itron, Inc
March 5, 2014

MPL Parameter Configuration Option for DHCPv6
draft-doi-roll-mpl-parameter-configuration-05

Abstract

This draft defines a way to configure MPL parameter set via DHCPv6 option. MPL has a set of parameters to control its behavior, and the parameter set is often configured as a network-wide parameter because the parameter set should be identical for each MPL forwarder in an MPL domain. Using the MPL Parameter Configuration Option defined in this document, a network can be configured with a single set of MPL parameter easily.

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 6, 2014.

Copyright Notice

Copyright (c) 2014 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. MPL Parameter Configuration Option	3
2.1. Unsigned Short Floating Point	4
2.2. MPL Parameter Configuration Option Format	5
2.3. DHCPv6 Client Behavior	6
2.4. MPL Forwarder Behavior	6
2.5. DHCPv6 Server Behavior	7
2.6. DHCPv6 Relay Behavior	7
3. IANA Considerations	7
4. Security Considerations	8
5. References	8
5.1. Normative References	8
5.2. Non-Normative References	8
Appendix A. Update History	8
Appendix B. Acknowledgements	9
Authors' Addresses	9

1. Introduction

Multicast Protocol for Low power and Lossy Networks (MPL) [I-D.ietf-roll-trickle-mcast] defines a protocol to make a multicast network among low power and lossy network i.e. wireless mesh networks. MPL has a set of parameters to control its behavior and tradeoff between end-to-end delay and network utilization. In most environments, the default parameters are acceptable. However, in some environments, the parameter set must be configured carefully in order to meet the requirements of each environment. According to the MPL draft section 5.4, each parameter in the set should be same for all nodes within an MPL domain. And the MPL draft does not define a method to configure the MPL parameter set.

Some managed wireless mesh networks may have a DHCP server to configure network parameters. MPL parameter set shall be considered as a part of network parameters (nodes in an MPL domain should use an identical parameter set). This document is to define the way to distribute parameter sets for MPL forwarders as a simple DHCPv6 [RFC3315] option.

2. MPL Parameter Configuration Option

Per MPL domain, there are following 10 parameters. An MPL domain is defined by an MPL domain address.

- o PROACTIVE_FORWARDING
- o SEED_SET_ENTRY_LIFETIME
- o DATA_MESSAGE_IMIN
- o DATA_MESSAGE_IMAX
- o DATA_MESSAGE_K
- o DATA_MESSAGE_TIMER_EXPIRATIONS
- o CONTROL_MESSAGE_IMIN
- o CONTROL_MESSAGE_IMAX
- o CONTROL_MESSAGE_K
- o CONTROL_MESSAGE_TIMER_EXPIRATIONS

One network may have multiple MPL domains with different

configurations. To configure more than one MPL domain via DHCP, there may be more than one MPL Parameter Configuration Option given to DHCP clients from a DHCP server.

2.1. Unsigned Short Floating Point

MPL has many timer parameters. Expected range of the timers depends on the network topology or MAC/PHY nature. To accommodate wide range of timer values efficiently, the MPL Parameter Configuration Option uses base-10 unsigned short floating point number with 3-bit exponent and 13-bit significand defined as follows (exp. stands for exponent).

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
| exp. |           significand         |
+---+---+---+---+---+---+---+---+

```

The represented value is (significand) * 10^(exp.). The minimum exponent is 0 (binary 000) and the maximum is 6 (binary 110). exp=7 (binary 111) is reserved for future use. The minimum significand is 0 (all 0) and the maximum is 8191 (all 1).

Unlike IEEE754 half precision floating point (binary16), there is no sign bit (no negative value for a timer), exponent is not biased (no fractional value for a timer), no implicit leading 1 in significand, and base is 10. Therefore, there could be more than one representation for a value.

Followings are examples of common timer values represented by unit of millisecond.

One second (1,000 milliseconds): exp = 3, significand = 1, 0x6001.

One minute (60,000 milliseconds): exp = 4, significand = 6, 0x8006.

One hour (3,600,000 milliseconds): exp = 5, significand = 36, 0xa024.

One day (86,400,000 milliseconds): exp = 5, significand = 864, 0xa360

Maximum timer length represented by an unsigned short floating point with millisecond precision is 8191 * 10⁶ milliseconds (13 weeks 3 days 19 hours 16 minutes 40 seconds).

With exponent and significand, an unsigned short floating point (usfp) can be encoded as follows.

```
usfp = (exponent << 13)|(0x1fff & significand);
```

2.2. MPL Parameter Configuration Option Format

To distribute a configuration of an MPL domain or a default value for all MPL domains (wildcard) under the network managed by the DHCP server, this document defines a DHCPv6 option format as follows. Short floating point format is used to describe wide range of timer values.

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										
OPTION_MPL_PARAMETERS																				option_len																					
P	Z									C_K					Z2					DM_K					SE_LIFETIME																
										DM_IMIN																				DM_IMAX											
										DM_T_EXP																				C_IMIN											
										C_IMAX																				C_T_EXP											

(if option_len = 32)

MPL Domain Address (128bits)																															
(cont'ed)																															
(cont'ed)																															
(cont'ed)																															

OPTION_MPL_PARAMETERS: DHCPv6 option identifier (not yet assigned).

option_len: Length of the option. It SHOULD be 16 (without MPL domain address) or 32 (with MPL domain address)

P (1 bit): A flag to indicate PROACTIVE_FORWARDING

Z (1 bit) Reserved. Should be 0.

C_K (5 bits): CONTROL_MESSAGE_K.

Z2 (3 bits) Reserved. Should be all 0.

DM_K (5 bits): DATA_MESSAGE_K.

SE_LIFETIME: SEED_SET_ENTRY_LIFETIME. The unit is millisecond and the type is unsigned short floating point.

DM_IMIN: DATA_MESSAGE_IMIN. The unit is millisecond and the type is unsigned short floating point.

DM_IMAX: DATA_MESSAGE_IMAX. The unit is millisecond and the type is unsigned short floating point.

DM_T_EXP: DATA_MESSAGE_TIMER_EXPIRATIONS. The unit is millisecond and the type is unsigned short floating point.

C_IMIN: CONTROL_MESSAGE_IMIN. The unit is millisecond and the type is unsigned short floating point.

C_IMAX: CONTROL_MESSAGE_IMAX. The unit is millisecond and the type is unsigned short floating point.

C_T_EXP: CONTROL_MESSAGE_TIMER_EXPIRATIONS. The unit is millisecond and the type is unsigned short floating point.

2.3. DHCPv6 Client Behavior

Clients MAY request MPL Parameter Configuration Option, as described in RFC3315 [RFC3315], sections 17.1.1, 18.1.1, 18.1.3, 18.1.4, 18.1.5 and 22.7. As a convenience to the reader, we mention here that the client includes requested option codes in Option Request Option.

Clients MUST discard MPL Parameter Configuration Option if it is invalid (i.e. it sets reserved bits or it has timers with reserved exp=7 in Unsigned Short Floating Point).

2.4. MPL Forwarder Behavior

If a DHCPv6 client requests and receives MPL Parameter Configuration Option, the node SHOULD join the MPL domain given by the option and act as an MPL forwarder. Each node SHOULD configure its MPL forwarder with the given parameter set for the MPL domain.

The priority of MPL Parameter Configuration applied for an MPL Domain is as follows (high to low).

- o Specific MPL Parameter Configuration to the MPL Domain (optlen=32)
- o Wildcard MPL Parameter Configuration (optlen=16)
- o Default configuration given in the MPL specification.

There SHALL be no more than one MPL Parameter Configuration Option for a MPL domain or the wildcard. Thus, the order of DHCPv6 options in the packet has no effect on precedence.

A node MAY leave from an MPL domain if the following two conditions are satisfied. 1) The MPL domain is configured by a DHCPv6 option from a DHCPv6 server previously. 2) The node has received an updated MPL Parameter Configuration Option without a configuration for the MPL domain.

MPL parameter may be updated occasionally. With stateful DHCPv6, updates can be done when the renewal timer expires. Information Refresh Time Option [RFC4242] shall be used to keep each forwarders updated.

To reduce periodical update traffic a node may try to use very long interval between updates. In the case, reconfigure shall be used to keep forwarder parameter sets synchronized. For stateless DHCPv6, [I-D.jiang-dhc-stateless-reconfiguration] may be used (if approved).

2.5. DHCPv6 Server Behavior

Sections 17.2.2 and 18.2 of RFC3315 [RFC3315] govern server operation in regards to option assignment. As a convenience to the reader, we mention here that the server will send MPL Parameter Configuration Option only if configured with specific value for MPL Parameter Configuration Option and the client requested it.

Servers MUST ignore incoming MPL Parameter Configuration Option.

2.6. DHCPv6 Relay Behavior

It's never appropriate for a relay agent to add options to a message heading toward the client, and relay agents don't actually construct Relay-Reply messages anyway. There are no additional requirements for relays.

3. IANA Considerations

A DHCPv6 option code for MPL Parameter Configuration Option needs to be assigned from IANA.

4. Security Considerations

A forged option may cause excessive layer-2 broadcasting. Implementations should set reasonable bounds for each parameter. For example, not too high K, not too low IMIN, etc. These may be implementation dependent or may be derived from MAC/PHY specifications. DHCP server or the network itself shall be trusted by some means including network access control or DHCP authentications.

5. References

5.1. Normative References

- [I-D.ietf-roll-trickle-mcast]
Hui, J. and R. Kelsey, "Multicast Forwarding Using Trickle", draft-ietf-roll-trickle-mcast-07 (work in progress), February 2014.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.

5.2. Non-Normative References

- [I-D.ietf-dhc-option-guidelines]
Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", draft-ietf-dhc-option-guidelines-17 (work in progress), January 2014.
- [I-D.jiang-dhc-stateless-reconfiguration]
Jiang, S. and B. Liu, "Stateless Reconfiguration in Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", draft-jiang-dhc-stateless-reconfiguration-01 (work in progress), February 2014.

Appendix A. Update History

Updates on 04 to 05:

- o Editorial fix and some clarification (thanks to Ted Lemon and folks in IESG/RFC-Editor Language Editing Session)
- o Added reference to RFC4242 and corrected renewal descriptions on MPL forwarder behavior (thanks to Tatsuya Jinmei for pointing it out)

Updates on 03 to 04:

- o Added more sections according to dhc-options-guidelines
- o Removed 'no update' requirement on MPL forwarder behavior
- o Added reference to I-D.jiang-dhc-stateless-reconfiguration

Updates on 02 to 03:

- o C flag is removed and wildcard shall be identified by optlen
- o Added some description on update of MPL parameters
- o Clearly stated there shall not be two or more configuration for an MPL domain and option order is not significant.

Updates on 01 to 02:

- o Added co-author

Updates on 00 to 01:

- o Corrected target area, track, etc.

Appendix B. Acknowledgements

The authors thank Richard Kelsey and Yoshi Ohba for technical advices to improve this draft.

Authors' Addresses

Yusuke Doi
TOSHIBA Corporation
Komukai Toshiba Cho 1
Saiwai-Ku
Kawasaki, Kanagawa 2128582
JAPAN

Phone: +81-45-342-7230
Email: yusuke.doi@toshiba.co.jp
URI:

Matthew Gillmore
Itron, Inc
2111 N Molter Rd.
Liberty Lake, WA 99019
USA

Email: matthew.gillmore@itron.com

Roll
Internet-Draft
Intended status: Standards Track
Expires: April 9, 2017

N. Cam-Winget, Ed.
Cisco Systems
J. Hui
Nest
D. Popa
Itron, Inc
October 6, 2016

Applicability Statement for the Routing Protocol for Low Power and Lossy
Networks (RPL) in AMI Networks
draft-ietf-roll-applicability-ami-15

Abstract

This document discusses the applicability of RPL in Advanced Metering Infrastructure (AMI) networks.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 9, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
1.2. Required Reading	3
1.3. Out of scope requirements	3
2. Routing Protocol for LLNs (RPL)	4
3. Description of AMI networks for electric meters	4
3.1. Deployment Scenarios	5
4. Smart Grid Traffic Description	7
4.1. Smart Grid Traffic Characteristics	7
4.2. Smart Grid Traffic QoS Requirements	8
4.3. RPL applicability per Smart Grid Traffic Characteristics	9
5. Layer 2 applicability	9
5.1. IEEE Wireless Technology	9
5.2. IEEE PowerLine Communication (PLC) technology	9
6. Using RPL to Meet Functional Requirements	10
7. RPL Profile	11
7.1. RPL Features	11
7.1.1. RPL Instances	11
7.1.2. DAO Policy	11
7.1.3. Path Metrics	11
7.1.4. Objective Function	11
7.1.5. DODAG Repair	12
7.1.6. Multicast	12
7.1.7. Security	12
7.2. Description of Layer-two features	13
7.2.1. IEEE 1901.2 PHY and MAC sub-layer features	13
7.2.2. IEEE 802.15.4 (g + e) PHY and MAC features	14
7.2.3. IEEE MAC sub-layer Security Features	15
7.3. 6LowPAN Options	16
7.4. Recommended Configuration Defaults and Ranges	17
7.4.1. Trickle Parameters	17
7.4.2. Other Parameters	18
8. Manageability Considerations	18
9. Security Considerations	19
9.1. Security Considerations during initial deployment	19
9.2. Security Considerations during incremental deployment	19
9.3. Security Considerations based on RPL's Threat Analysis	20
10. Privacy Considerations	20
11. IANA Considerations	20
12. Acknowledgements	20
13. References	21
13.1. Normative References	21
13.2. Informative references	22

Authors' Addresses	23
--------------------	----

1. Introduction

Advanced Metering Infrastructure (AMI) systems enable the measurement, configuration, and control of energy, gas and water consumption and distribution, through two-way scheduled, on exception, and on-demand communication.

AMI networks are composed of millions of endpoints, including meters, distribution automation elements, and eventually home area network devices. They are typically inter-connected using some combination of wireless and power-line communications, forming the so-called Neighbor Area Network (NAN) along with a backhaul network providing connectivity to "command-and-control" management software applications at the utility company back office.

1.1. Requirements Language

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

1.2. Required Reading

[surveySG] gives an overview of Smart Grid architecture and related applications.

NAN can use wireless communication technology in which case is using, from the IEEE 802.15.4 standard family, the [IEEE802.15.4g] PHY Layer amendment and [IEEE802.15.4e] MAC sub-layer amendment, specifically adapted to smart grid networks.

NAN can also use PLC (Power Line Communication) technology as an alternative to wireless communications. Several standards for PLC technology have emerged, such as [IEEE1901.2].

NAN can further use a mix of wireless and PLC technologies to increase the network coverage ratio, a critical requirement for AMI networks.

1.3. Out of scope requirements

The following are outside the scope of this document:

- o Applicability statement for RPL (Routing Protocol for Low Power and Lossy Networks) [RFC6550] in AMI networks composed of battery-powered devices (i.e., gas/water meters).
- o Applicability statement for RPL in AMI networks composed of a mix of AC powered devices (i.e., electric meters) and battery-powered meters (i.e., gas/water meters).
- o Applicability statement for RPL storing mode of operation in AMI networks.

2. Routing Protocol for LLNs (RPL)

RPL provides routing functionality for mesh networks that can scale up to thousands of resource-constrained devices, interconnected by low power and lossy links, and communicating with the external network infrastructure through a common aggregation point(s) (e.g., a LLN Border Router or LBR).

RPL builds a Directed Acyclic Graph (DAG) routing structure rooted at a LBR (LLN Border Router), ensures loop-free routing, and provides support for alternate routes, as well as, for a wide range of routing metrics and policies.

RPL was designed to operate in energy-constrained environments and includes energy-saving mechanisms (e.g., Trickle timers) and energy-aware metrics. RPL's ability to support multiple different metrics and constraints at the same time enables it to run efficiently in heterogeneous networks composed of nodes and links with vastly different characteristics [RFC6551].

This document describes the applicability of RPL non-storing mode (as defined in [RFC6550]) to AMI deployments. The Routing Requirements for Urban Low-Power and Lossy Networks are applicable to AMI networks as well. The terminology used in this document is defined in [RFC7102].

3. Description of AMI networks for electric meters

In many deployments, in addition to measuring energy consumption, the electric meter network plays a central role in the Smart Grid since the device enables the utility company to control and query the electric meters themselves and can serve as a backhaul for all other devices in the Smart Grid, e.g., water and gas meters, distribution automation and home area network devices. Electric meters may also be used as sensors to monitor electric grid quality and to support applications such as Electric Vehicle charging.

Electric meter networks can be composed of millions of smart meters (or nodes), each of which is resource-constrained in terms of processing power, storage capabilities, and communication bandwidth, due to a combination of factors including regulations on spectrum use, and on meter behavior and performance, on heat emissions within the meter, form factor and cost considerations. These constraints result in a compromise between range and throughput, with effective link throughput of tens to a few hundred kilobits per second per link, a potentially significant portion of which is taken up by protocol and encryption overhead when strong security measures are in place.

Electric meters are often interconnected into multi-hop mesh networks, each of which is connected to a backhaul network leading to the utility company network through a network aggregation point, e.g., an LBR.

3.1. Deployment Scenarios

AMI networks are composed of millions of endpoints distributed across both urban and rural environments. Such endpoints can include electric, gas, and water meters, distribution automation elements, and home area network devices.

Devices in the network communicate directly with other devices in close proximity using a variety of low-power and/or lossy link technologies that are both wireless and wired (e.g., IEEE 802.15.4g, IEEE 802.15.4e, IEEE 1901.2, IEEE 802.11). In addition to serving as sources and destinations of packets, many network elements typically also forward packets and thus form a mesh topology.

In a typical AMI deployment, groups of meters within physical proximity form routing domains, each in the order of a 1,000 to 10,000 meters. Thus, each electric meter mesh typically has several thousand wireless endpoints, with densities varying based on the area and the terrain.

where nodes belonging to the same DODAG (Destination Oriented Directed Acyclic Graph) can be connected to the grid through different substations. If narrowband PLC technology is used, it will follow more or less the physical tree structure since diaphony may allow one phase to communicate with the other. This is particularly true near the LBR. Some mixed topology can also be observed, since some LBRs may be strategically installed in the field to avoid all the communications going through a single LBR. Nevertheless, the short propagation range forces meters to relay the information.

4. Smart Grid Traffic Description

4.1. Smart Grid Traffic Characteristics

In current AMI deployments, metering applications typically require all smart meters to communicate with a few head-end servers, deployed in the utility company data center. Head-end servers generate data traffic to configure smart data reading or initiate queries, and use unicast and multicast to efficiently communicate with a single device (i.e. Point-to-Point (P2P) communications) or groups of devices respectively (i.e., Point-to-Multipoint (P2MP) communication). The head-end server may send a single small packet at a time to the meters (e.g., a meter read request, a small configuration change, service switch command) or a series of large packets (e.g., a firmware download across one or even thousands of devices). The frequency of large file transfers (e.g., firmware download of all metering devices) is typically much lower than the frequency of sending configuration messages or queries. Each smart meter generates Smart Metering Data (SMD) traffic according to a schedule (e.g., periodic meter reads), in response to on-demand queries (e.g., on-demand meter reads), or in response to some local event (e.g., power outage, leak detection). Such traffic is typically destined to a single head-end server. The SMD traffic is thus highly asymmetric, where the majority of the traffic volume generated by the smart meters typically goes through the LBRs, and is directed from the smart meter devices to the head-end servers, in a MP2P (Mesh Peer to Peer)fashion. Current SMD traffic patterns are fairly uniform and well-understood. The traffic generated by the head-end server and destined to metering devices is dominated by periodic meter reads, while traffic generated by the metering devices is typically uniformly spread over some periodic read time-window.

Smart metering applications typically do not have hard real-time constraints, but they are often subject to bounded latency and stringent reliability service level agreements.

Distribution Automation (DA) applications typically involve a small number of devices that communicate with each other in a Point-to-

Point (P2P) fashion, and may or may not be in close physical proximity. DA applications typically have more stringent latency requirements than SMD applications.

There are also a number of emerging applications such as electric vehicle charging. These applications may require P2P communication and may eventually have more stringent latency requirements than SMD applications.

4.2. Smart Grid Traffic QoS Requirements

As described previously, the two main traffic families in a NAN are:

- A) Meter-initiated traffic (Meter-to-head-end - M2HE)
 - B1) request is sent in point-to-point to a specific meter
 - B2) request is sent in multicast to a subset of meters
 - B3) request is sent in multicast to all meters

The M2HE are event-based, while the HE2M are mostly command-response. In most cases, M2HE traffic is more critical than HE2M one, but there can be exceptions.

Regarding priority, traffic may also be decomposed into several classes :

- C1) Highly Priority Critical traffic for Power System Outage, Pricing Events and Emergency Messages require a 98%+ packet delivery under 5 s. Payload size < 100 bytes
- C2) Critical Priority traffic Power Quality Events, Meter Service Connection and Disconnection require 98%+ packet delivery under 10s. Payload size < 150 bytes
- C3) Normal Priority traffic for System Events including Faults, Configuration and Security require 98%+ packet delivery under 30s. Payload size < 200 bytes
- C4) Low Priority traffic for Recurrent Meter Reading require 98%+ packet 2 hour delivery window 6 times per day. Payload size < 400 bytes

- C5) Background Priority traffic for firmware/software updates processed to 98%+ of devices within 7 days. Average firmware update is 1 MB.

4.3. RPL applicability per Smart Grid Traffic Characteristics

RPL non-storing mode of operation naturally support upstream and downstream forwarding of unicast traffic between the DODAG root and each DODAG node, and between DODAG nodes and DODAG root, respectively.

Group communication model used in smart grid requires RPL non-storing mode of operation to support downstream forwarding of multicast traffic with a scope larger than link-local. The DODAG root is the single device that injects multicast traffic, with a scope larger than link-local, into the DODAG.

5. Layer 2 applicability

5.1. IEEE Wireless Technology

IEEE Std. 802.15.4g and IEEE 802.15.4e amendments to 802.15.4 standard have been specifically developed for smart grid networks. They are the most common PHY and MAC layers used for wireless AMI network. 802.15.4g specifies multiple modes of operation (FSK, OQPSK and OFDM modulations) with speeds from 50kbps to 600kbps, and allows for transport of a full IPv6 packet (i.e., 1280 octets) without the need for upper layer segmentation and re-assembly.

IEEE Std. 802.15.4e is an amendment to IEEE Std 802.15.4 that specifies additional media access control (MAC) behaviors and frame formats that allow IEEE 802.15.4 devices to support a wide range of industrial and commercial applications that were not adequately supported prior to the release of this amendment. It is important to notice that 802.15.4e does not change the link-layer security scheme defined in the last two updates to 802.15.4 (e.g. 2006 and 2011 amendments).

5.2. IEEE PowerLine Communication (PLC) technology

The IEEE Std. 1901.2 standard specifies communications for low frequency (less than 500 kHz) narrowband power line devices via alternating current and direct current electric power lines. IEEE Std P1901.2 standard supports indoor and outdoor communications over low voltage line (line between transformer and meter, less than 1000 V), through transformer low-voltage to medium-voltage (1000 V up to 72 kV) and through transformer medium-voltage to low-voltage power

lines in both urban and in long distance (multi- kilometer) rural communications.

IEEE Std. 1901.2 defines the PHY layer and the MAC sub-layer of the data link layer. The MAC sub-layer endorses a sub-set of 802.15.4 and 802.15.4e MAC sub-layer features.

IEEE Std. 1901.2 PHY Layer bit rates are scalable up to 500 kbps depending on the application requirements and type of encoding used.

IEEE Std. 1901.2 MAC layer allows for transport of a full IPv6 packet (i.e., 1280 octets) without the need for upper layer segmentation and re-assembly.

IEEE Std. 1901.2 specifies the necessary link-layer security features that fully endorse 802.15.4 MAC sub-layer security scheme.

6. Using RPL to Meet Functional Requirements

The functional requirements for most AMI deployments are similar to those listed in [RFC5548]. This section informallly highlights some of the similarities:

- o The routing protocol MUST be capable of supporting the organization of a large number of nodes into regions containing on the order of 10^2 to 10^4 nodes each.
- o The routing protocol MUST provide mechanisms to support configuration of the routing protocol itself.
- o The routing protocol SHOULD support and utilize the large number of highly directed flows to a few head-end servers to handle scalability.
- o The routing protocol MUST dynamically compute and select effective routes composed of low-power and lossy links. Local network dynamics SHOULD NOT impact the entire network. The routing protocol MUST compute multiple paths when possible.
- o The routing protocol MUST support multicast and unicast addressing. The routing protocol SHOULD support formation and identification of groups of field devices in the network.

RPL supports the following features:

- o Scalability: Large-scale networks characterized by highly directed traffic flows between each smart meter and the head-end servers in

the utility network. To this end, RPL builds a Directed Acyclic Graph (DAG) rooted at each LBR.

- o Zero-touch configuration: This is done through in-band methods for configuring RPL variables using DIO (DODAG Information Object) messages, and DIO message options [RFC6550].
- o The use of links with time-varying quality characteristics: This is accomplished by allowing the use of metrics that effectively capture the quality of a path (e.g., Expected Transmission Count (ETX)) and by limiting the impact of changing local conditions by discovering and maintaining multiple DAG parents, and by using local repair mechanisms when DAG links break.

7. RPL Profile

7.1. RPL Features

7.1.1. RPL Instances

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

7.1.2. DAO Policy

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

7.1.3. Path Metrics

Smart metering deployments utilize link technologies that may exhibit significant packet loss and thus require routing metrics that take packet loss into account. To characterize a path over such link technologies, AMI deployments can use the Expected Transmission Count (ETX) metric as defined in [RFC6551].

Additional metrics may be defined in companion RFCs.

7.1.4. Objective Function

RPL relies on an Objective Function for selecting parents and computing path costs and rank. This objective function is decoupled from the core RPL mechanisms and also from the metrics in use in the network. Two objective functions for RPL have been defined at the

time of this writing, OF0 [RFC6552] and MRHOF [RFC6719], both of which define the selection of a preferred parent and backup parents, and are suitable for AMI deployments.

Additional objective functions may be defined in companion RFCs.

7.1.5. DODAG Repair

To effectively handle time-varying link characteristics and availability, AMI deployments SHOULD utilize the local repair mechanisms in RPL. Local repair is triggered by broken link detection. The first local repair mechanism consists of a node detaching from a DODAG and then re-attaching to the same or to a different DODAG at a later time. While detached, a node advertises an infinite rank value so that its children can select a different parent. This process is known as poisoning and is described in Section 8.2.2.5 of [RFC6550]. While RPL provides an option to form a local DODAG, doing so in AMI for electric meters is of little benefit since AMI applications typically communicate through a LBR. After the detached node has made sufficient effort to send notification to its children that it is detached, the node can rejoin the same DODAG with a higher rank value. The configured duration of the poisoning mechanism needs to take into account the disconnection time applications running over the network can tolerate. Note that when joining a different DODAG, the node need not perform poisoning. The second local repair mechanism controls how much a node can increase its rank within a given DODAG Version (e.g., after detaching from the DODAG as a result of broken link or loop detection). Setting the DAGMaxRankIncrease to a non-zero value enables this mechanism, and setting it to a value of less than infinity limits the cost of count-to-infinity scenarios when they occur, thus controlling the duration of disconnection applications may experience.

7.1.6. Multicast

Multicast support for RPL in non-storing mode are being developed in companion RFCs (see [RFC7731]).

7.1.7. Security

AMI deployments operate in areas that do not provide any physical security. For this reason, the link layer, transport layer and application layer technologies utilized within AMI networks typically provide security mechanisms to ensure authentication, confidentiality, integrity, and freshness. As a result, AMI deployments may not need to implement RPL's security mechanisms; they MUST include, at a minimum, link layer security such as that defined by IEEE 1901.2 and IEEE 802.15.4.

7.2. Description of Layer-two features

7.2.1. IEEE 1901.2 PHY and MAC sub-layer features

The IEEE Std. 1901.2 PHY layer is based on OFDM modulation and defines a time frequency interleaver over the entire PHY frame coupled with a Reed Solomon and Viterbi Forward Error Correction for maximum robustness. Since the noise level in each OFDM sub-carrier can vary significantly, IEEE 1901.2 specifies two complementary mechanisms allowing to fine-tune the robustness/performance tradeoff implicit in such systems. More specifically, the first (coarse-grained) mechanism, defines the modulation from several possible choices (robust (super-ROBO, ROBO), BPSK, QPSK,...). The second (fine-grained) maps the sub-carriers which are too noisy and deactivates them.

The existence of multiple modulations and dynamic frequency exclusion renders the problem of selecting a path between two nodes non-trivial, as the possible number of combinations increases significantly, e.g. use a direct link with slow robust modulation, or use a relay meter with fast modulation and 12 disabled sub-carriers. In addition, IEEE 1901.2 technology offers a mechanism (adaptive tone map) for periodic exchanges on the link quality between nodes to constantly react to channel fluctuations. Every meter keeps a state of the quality of the link to each of its neighbors by either piggybacking the tone mapping on the data traffic, or by sending explicit tone map requests.

IEEE 1901.2 MAC frame format shares most in common with the IEEE 802.15.4 MAC frame format [IEEE802.15.4], with a few exceptions described below.

- o IEEE 1901.2 MAC frame is obtained by prepending a Segment Control Field to the IEEE 802.15.4 MAC header. One function of the Segment Control Field is to signal the use of the MAC sub-layer segmentation and reassembly.
- o IEEE 1901.2 MAC frames uses only the 802.15.4 MAC addresses with a length of 16 and 64 bits.
- o IEEE 1901.2 MAC sub-layer endorses the concept of Information Elements, as defined in [IEEE802.15.4e]. The format and use of Information Elements are not relevant to RPL applicability statement.

The IEEE 1901.2 PHY frame payload size varies as a function of the modulation used to transmit the frame and the strength of the Forward Error Correction scheme.

The IEEE 1901.2 PHY MTU size is variable and dependent on the PHY settings in use (e.g. bandwidth, modulation, tones, etc). As quoted from the IEEE 1901.2 specification: For CENELEC A/B, if MSDU size is more than 247 octets for robust OFDM (ROBO) and Super-ROBO modulations or more than 239 octets for all other modulations, the MAC layer shall divide the MSDU into multiple segments as described in 5.3.7. For FCC and ARIB, if the MSDU size meets one of the following conditions: a) For ROBO and Super-ROBO modulations, the MSDU size is more than 247 octets but less than 494 octets, b) For all other modulations, the MSDU size is more than 239 octets but less than 478 octets.

7.2.2. IEEE 802.15.4 (g + e) PHY and MAC features

IEEE Std 802.15.4g defines multiple modes of operation, where each mode uses different modulation and has multiple data rates. Additionally, 802.15.4g PHY layer includes mechanisms to improve the robustness of the radio communications, such as data whitening and Forward Error Correction coding. The 802.15.4g PHY frame payload can carry up to 2048 octets.

The IEEE Std 802.15.4g defines the following modulations: MR-FSK (Multi-Rate FSK), MR-OFDM (Multi-Rate OFDM) and MR-O-QPSK (Multi-Rate O-QPSK). The (over-the-air) bit rates for these modulations range from 4.8 to 600kbps for MR-FSK, from 50 to 600kbps for MR-OFDM and from 6.25 to 500kbps for MR-O-QPSK.

The MAC sub-layer running on top of a 4g radio link is based on IEEE 802.15.4e. The 802.15.4e MAC allows for a variety of modes for operation. These include:

- o Timetimeslotslotted channel hopping (TSCH): specifically designed for application domains such as process automation
- o Low latency deterministic networks (LLDN): for application domains such as factory automation
- o Deterministic and synchronous multi-channel extension (DSME): for general industrial and commercial application domains that includes Channel diversity to increase network robustness
- o Asynchronous multi-channel adaptation (AMCA): for large infrastructure application domains.

The MAC addressing scheme supports short (16-bits) addresses along with extended (64-bits) addresses. These addresses are assigned in different ways and is specified by specific standards organizations.

Information Elements, Enhanced Beacons and frame version 2, as defined in 802.15.4e, MUST be supported.

Since the MAC frame payload size limitation is given by the 4g PHY frame payload size limitation (i.e., 2048 bytes) and MAC layer overhead (headers, trailers, Information Elements and security overhead), the MAC frame payload MUST be able to carry a full IPv6 packet of 1280 octets without upper layer fragmentation and re-assembly.

7.2.3. IEEE MAC sub-layer Security Features

Since IEEE 1901.2 standard is based on the 802.15.4 MAC sub-layer and fully endorses the security scheme defined in 802.15.4, we only focus on description of IEEE 802.15.4 security scheme.

The IEEE 802.15.4 specification was designed to support a variety of applications, many of which are security sensitive. The IEEE 802.15.4 provides four basic security services: message authentication, message integrity, message confidentiality, and freshness checks to avoid replay attacks.

The 802.15.4 security layer is handled at the media access control layer, below 6LoWPAN layer. The application specifies its security requirements by setting the appropriate control parameters into the radio/PLC stack. The 802.15.4 defines four packet types: beacon frames, data frames, acknowledgments frame, and command frames for the media access control layer. The 802.15.4 specification does not support security for acknowledgement frames; data frames, beacon frames and command frames can support integrity protection and confidentiality protection for the frames's data field. An application has a choice of security suites that control the type of security protection that is provided for the transmitted MAC frame. Each security suite offers a different set of security properties and guarantees, and ultimately different MAC frame formats. The 802.15.4 specification defines eight different security suites, outlined below. We can broadly classify the suites by the properties that they offer: no security, encryption only (AES-CTR), authentication only (AES-CBC-MAC), and encryption and authentication (AES-CCM). Each category that supports authentication comes in three variants depending on the size of the MAC (Message Authentication Control) that it offers. The MAC can be either 4, 8, or 16 bytes long. Additionally, for each suite that offers encryption, the recipient can optionally enable replay protection.

- o Null = No security.
- o AES-CTR = Encryption only, CTR mode.

- o AES-CBC-MAC-128 = No encryption, 128-bit MAC.
- o AES-CBC-MAC-64 = No encryption, 64-bit MAC.
- o AES-CCM-128 = Encryption and 128-bit MAC.
- o AES-CCM-64 = Encryption and 64-bit MAC.
- o AES-CCM-32 = Encryption and 32-bit MAC.

Note that AES-CCM-32 is the most commonly used cipher in these deployments today.

To achieve authentication, any device can maintain an Access Control List (ACL) which is a list of trusted nodes from which the device wishes to receive data. Data encryption is done by encryption of Message Authentication Control (MAC) frame payload using the key shared between two devices, or among a group of peers. If the key is to be shared between two peers, it is stored with each entry in the ACL list; otherwise, the key is stored as the default key. Thus, the device can make sure that its data can not be read by devices that do not possess the corresponding key. However, device addresses are always transmitted unencrypted, which makes attacks that rely on device identity somewhat easier to launch. Integrity service is applied by appending a Message Integrity Code (MIC) generated from blocks of encrypted message text. This ensures that a frame can not be modified by a receiver device that does not share a key with the sender. Finally, sequential freshness uses a frame counter and key sequence counter to ensure the freshness of the incoming frame and guard against replay attacks.

A cryptographic MAC is used to authenticate messages. While longer MACs lead to improved resiliency of the code, they also make packet size larger and thus take up bandwidth in the network. In constrained environments such as metering infrastructures, an optimum balance between security requirements and network throughput must be found.

7.3. 6LowPAN Options

AMI implementations based on 1901.2 and 802.15.4(g+e) can utilize all of the IPv6 Header Compression schemes specified in [RFC6282] Section 3 and all of the IPv6 Next Header compression schemes specified in [RFC6282] Section 4, if reducing over the air/wire overhead is a requirement.

7.4. Recommended Configuration Defaults and Ranges

7.4.1. Trickle Parameters

Trickle [RFC6206] was designed to be density-aware and perform well in networks characterized by a wide range of node densities. The combination of DIO packet suppression and adaptive timers for sending updates allows Trickle to perform well in both sparse and dense environments. Node densities in AMI deployments can vary greatly, from nodes having only one or a handful of neighbors to nodes having several hundred neighbors. In high density environments, relatively low values for Imin may cause a short period of congestion when an inconsistency is detected and DIO updates are sent by a large number of neighboring nodes nearly simultaneously. While the Trickle timer will exponentially backoff, some time may elapse before the congestion subsides. While some link layers employ contention mechanisms that attempt to avoid congestion, relying solely on the link layer to avoid congestion caused by a large number of DIO updates can result in increased communication latency for other control and data traffic in the network. To mitigate this kind of short-term congestion, this document recommends a more conservative set of values for the Trickle parameters than those specified in [RFC6206]. In particular, DIOIntervalMin is set to a larger value to avoid periods of congestion in dense environments, and DIORedundancyConstant is parameterized accordingly as described below. These values are appropriate for the timely distribution of DIO updates in both sparse and dense scenarios while avoiding the short-term congestion that might arise in dense scenarios. Because the actual link capacity depends on the particular link technology used within an AMI deployment, the Trickle parameters are specified in terms of the link's maximum capacity for transmitting link-local multicast messages. If the link can transmit m link-local multicast packets per second on average, the expected time it takes to transmit a link-local multicast packet is $1/m$ seconds.

DIOIntervalMin: AMI deployments SHOULD set DIOIntervalMin such that the Trickle Imin is at least 50 times as long as it takes to transmit a link-local multicast packet. This value is larger than that recommended in [RFC6206] to avoid congestion in dense urban deployments as described above.

DIOIntervalDoublings: AMI deployments SHOULD set DIOIntervalDoublings such that the Trickle Imax is at least 2 hours or more.

DIORedundancyConstant: AMI deployments SHOULD set DIORedundancyConstant to a value of at least 10. This is due to the larger chosen value for DIOIntervalMin and the proportional

relationship between `Imin` and `k` suggested in [RFC6206]. This increase is intended to compensate for the increased communication latency of DIO updates caused by the increase in the `DIOIntervalMin` value, though the proportional relationship between `Imin` and `k` suggested in [RFC6206] is not preserved. Instead, `DIORedundancyConstant` is set to a lower value in order to reduce the number of packet transmissions in dense environments.

7.4.2. Other Parameters

- o AMI deployments SHOULD set `MinHopRankIncrease` to 256, resulting in 8 bits of resolution (e.g., for the ETX metric).
- o To enable local repair, AMI deployments SHOULD set `MaxRankIncrease` to a value that allows a device to move a small number of hops away from the root. With a `MinHopRankIncrease` of 256, a `MaxRankIncrease` of 1024 would allow a device to move up to 4 hops away.

8. Manageability Considerations

Network manageability is a critical aspect of smart grid network deployment and operation. With millions of devices participating in the smart grid network, many requiring real-time reachability, automatic configuration, and lightweight network health monitoring and management are crucial for achieving network availability and efficient operation. RPL enables automatic and consistent configuration of RPL routers through parameters specified by the DODAG root and disseminated through DIO packets. The use of Trickle for scheduling DIO transmissions ensures lightweight yet timely propagation of important network and parameter updates and allows network operators to choose the trade-off point they are comfortable with respect to overhead vs. reliability and timeliness of network updates. The metrics in use in the network along with the Trickle Timer parameters used to control the frequency and redundancy of network updates can be dynamically varied by the root during the lifetime of the network. To that end, all DIO messages SHOULD contain a Metric Container option for disseminating the metrics and metric values used for DODAG setup. In addition, DIO messages SHOULD contain a DODAG Configuration option for disseminating the Trickle Timer parameters throughout the network. The possibility of dynamically updating the metrics in use in the network as well as the frequency of network updates allows deployment characteristics (e.g., network density) to be discovered during network bring-up and to be used to tailor network parameters once the network is operational rather than having to rely on precise pre-configuration. This also allows the network parameters and the overall routing protocol behavior to evolve during the lifetime of the network. RPL specifies

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

9. Security Considerations

Smart grid networks are subject to stringent security requirements as they are considered a critical infrastructure component. At the same time, they are composed of large numbers of resource- constrained devices inter-connected with limited-throughput links. As a result, the choice of security mechanisms is highly dependent on the device and network capabilities characterizing a particular deployment.

In contrast to other types of LLNs, in smart grid networks centralized administrative control and access to a permanent secure infrastructure is available. As a result, smart grid networks are deployed with security mechanisms such as link-layer, transport layer and/or application-layer security mechanisms; while it is best practice to secure all layers, using RPL's secure mode may not be necessary. Failure to protect any of these layers can result in various attacks; without strong authentication of devices in the infrastructure can lead to uncontrolled and unauthorized access. Similarly, failure to protect the communication layers can enable passive (in wireless mediums) attacks as well as man-in-the-middle and active attacks.

As this document describes the applicability of RPL non-storing mode, the security considerations as defined in [RFC6550] also applies to this document and to AMI deployments.

9.1. Security Considerations during initial deployment

During the manufacturing process, the meters are loaded with the appropriate security credentials (keys, certificates). The configured security credentials during manufacturing are used by the devices to authenticate with the system and to further negotiate operational security credentials, for both network and application layers.

9.2. Security Considerations during incremental deployment

If during the system operation a device fails or is known to be compromised, it is replaced with a new device. The new device does not take over the security identity of the replaced device. The security credentials associated with the failed/compromised device are removed from the security appliances.

9.3. Security Considerations based on RPL's Threat Analysis

[RFC7416] defines a set of security considerations for RPL security. This document defines how it leverages the device's link layer and application layer security mechanisms to address the threats as defined in Section 6 of [RFC7416].

Like any secure network infrastructure, an AMI deployment's ability to address node impersonation, active man-in-the-middle attacks relies on mutual authentication and authorization process. The credential management from the manufacturing imprint of security credentials of smart meters to all credentials of nodes in the infrastructure, all credentials must be appropriately managed and classified through the authorization process to ensure beyond the identity of the nodes, that the nodes are behaving or 'acting' in their assigned roles.

Similarly, to ensure that data has not been modified, confidentiality and integrity at the suitable layers (e.g. link layer, application layer or both) should be used.

To provide the security mechanisms to address these threats, an AMI deployment MUST include the use of the security schemes as defined by IEEE 1901.2 (and IEEE 802.15.4). With IEEE 802.15.4 defining the security mechanisms to afford mutual authentication, access control (e.g. authorization) and transport confidentiality and integrity.

10. Privacy Considerations

Privacy of information flowing through smart grid networks are subject to consideration. An evolving set of recommendations and requirements are being defined by different groups and consortiums; for example, the U.S. Department of Energy issued a document [DOEVCC] defining a process and set of recommendations to address privacy issues. As this document describes the applicability of RPL, the privacy considerations as defined in [I-D.ietf-6lo-privacy-considerations] and [EUPR] apply to this document and to AMI deployments.

11. IANA Considerations

This memo includes no request to IANA.

12. Acknowledgements

Matthew Gillmore, Laurent Toutain, Ruben Salazar, and Kazuya Monden were contributors and noted as authors in earlier drafts. The authors would also like to acknowledge the review, feedback, and

comments of Jari Arkko, Dominique Barthel, Cedric Chauvenet, Yuichi Igarashi, Philip Levis, Jeorjeta Jetcheva, Nicolas Dejean, and JP Vasseur.

13. References

13.1. 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, <<http://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, <<http://www.rfc-editor.org/info/rfc6550>>.
- [IEEE802.15.4]
"IEEE Standard for Information technology-- Local and metropolitan area networks-- Specific requirements-- Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPANs)", IEEE Standard 802.15.4, September 2006.
- [IEEE802.15.4e]
"IEEE Standard for Local and metropolitan area networks-- Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer", IEEE Standard 802.15.4e, April 2012.
- [IEEE802.15.4g]
"IEEE Standard for Local and metropolitan area networks-- Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 3: Physical Layer (PHY) Specifications for Low-Data-Rate, Wireless, Smart Metering Utility Networks", IEEE Standard 802.15.4g, November 2012.
- [IEEE1901.2]
"IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", IEEE Standard 1901.2, December 2013.

[surveySG]

Gungor, V., Sahin, D., Kocak, T., Ergut, S., Buccella, C., Cecati, C., and G. Hancke, "A Survey on Smart Grid Potential Applications and Communication Requirements", Feb 2013.

13.2. Informative references

[RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <<http://www.rfc-editor.org/info/rfc7731>>.

[RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.

[RFC5548] Dohler, M., Ed., Watteyne, T., Ed., Winter, T., Ed., and D. Barthel, Ed., "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, DOI 10.17487/RFC5548, May 2009, <<http://www.rfc-editor.org/info/rfc5548>>.

[RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<http://www.rfc-editor.org/info/rfc6206>>.

[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, <<http://www.rfc-editor.org/info/rfc6551>>.

[RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, DOI 10.17487/RFC6719, September 2012, <<http://www.rfc-editor.org/info/rfc6719>>.

[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, <<http://www.rfc-editor.org/info/rfc6552>>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.

[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, <<http://www.rfc-editor.org/info/rfc7416>>.

[I-D.ietf-6lo-privacy-considerations] Thaler, D., "Privacy Considerations for IPv6 over Networks of Resource-Constrained Nodes", draft-ietf-6lo-privacy-considerations-03 (work in progress), September 2016.

[DOEVCC] "Voluntary Code of Conduct (VCC) Final Concepts and Principles", Jan 2015, <http://energy.gov/sites/prod/files/2015/01/f19/VCC%20Concepts%20and%20Principles%202015_01_08%20FINAL.pdf>.

[EUPR] "Information for investors and data controllers", Jun 2016, <<https://ec.europa.eu/energy/node/1748>>.

Authors' Addresses

Nancy Cam-Winget (editor)
Cisco Systems
3550 Cisco Way
San Jose, CA 95134
US

Email: ncamwing@cisco.com

Jonathan Hui
Nest
3400 Hillview Ave
Palo Alto, CA 94304
USA

Email: jonhui@nestlabs.com

Daniel Popa
Itron, Inc
52, rue Camille Desmoulins
Issy les Moulineaux 92130
FR

Email: daniel.popa@itron.com

Roll
Internet-Draft
Intended status: Standards Track
Expires: January 22, 2016

A. Brandt
Sigma Designs
E. Baccelli
INRIA
R. Cragie
ARM Ltd.
P. van der Stok
Consultant
July 21, 2015

Applicability Statement: The use of the RPL protocol suite in Home
Automation and Building Control
draft-ietf-roll-applicability-home-building-12

Abstract

The purpose of this document is to provide guidance in the selection and use of protocols from the RPL protocol suite to implement the features required for control in building and home environments.

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 January 22, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Relationship to other documents	4
1.2.	Terminology	4
1.3.	Required Reading	5
1.4.	Out of scope requirements	5
2.	Deployment Scenario	5
2.1.	Network Topologies	6
2.2.	Traffic Characteristics	7
2.2.1.	General	8
2.2.2.	Source-sink (SS) communication paradigm	8
2.2.3.	Publish-subscribe (PS, or pub/sub)) communication paradigm	9
2.2.4.	Peer-to-peer (P2P) communication paradigm	9
2.2.5.	Peer-to-multipeer (P2MP) communication paradigm	10
2.2.6.	Additional considerations: Duocast and N-cast	10
2.2.7.	RPL applicability per communication paradigm	10
2.3.	Layer-2 applicability	11
3.	Using RPL to meet Functional Requirements	12
4.	RPL Profile	13
4.1.	RPL Features	13
4.1.1.	RPL Instances	13
4.1.2.	Storing vs. Non-Storing Mode	14
4.1.3.	DAO Policy	14
4.1.4.	Path Metrics	14
4.1.5.	Objective Function	14
4.1.6.	DODAG Repair	14
4.1.7.	Multicast	15
4.1.8.	Security	16
4.1.9.	P2P communications	19
4.1.10.	IPv6 address configuration	19
4.2.	Layer 2 features	19
4.2.1.	Specifics about layer-2	19
4.2.2.	Services provided at layer-2	19
4.2.3.	6LowPAN options assumed	20
4.2.4.	Mesh Link Establishment (MLE) and other things	20
4.3.	Recommended Configuration Defaults and Ranges	20
4.3.1.	Trickle parameters	20
4.3.2.	Other Parameters	20
5.	MPL Profile	21
5.1.	Recommended configuration Defaults and Ranges	21
5.1.1.	Real-Time optimizations	21

5.1.2. Trickle parameters	21
5.1.3. Other parameters	22
6. Manageability Considerations	23
7. Security Considerations	23
7.1. Security considerations during initial deployment	23
7.2. Security Considerations during incremental deployment	24
7.3. Security Considerations for P2P uses	25
7.4. MPL routing	25
7.5. RPL Security features	25
8. Other related protocols	25
9. IANA Considerations	26
10. Acknowledgements	26
11. Changelog	26
12. References	28
12.1. Normative References	28
12.2. Informative References	32
Appendix A. RPL shortcomings in home and building deployments	33
A.1. Risk of undesired long P2P routes	33
A.1.1. Traffic concentration at the root	34
A.1.2. Excessive battery consumption in source nodes	34
A.2. Risk of delayed route repair	34
A.2.1. Broken service	34
Appendix B. Communication failures	35
Authors' Addresses	36

1. Introduction

The primary purpose of this document is to give guidance in the use of the Routing Protocol for Low power and lossy networks (RPL) protocol suite in two application domains:

- o Home automation
- o Building automation

The guidance is based on the features required by the requirements documents "Home Automation Routing Requirements in Low-Power and Lossy Networks" [RFC5826] and "Building Automation Routing Requirements in Low-Power and Lossy Networks" [RFC5867] respectively. The Advanced Metering Infrastructure is also considered where appropriate. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most likely network structures. An abstract set of distinct communication paradigms is then used to frame the applicability domains.

Home automation and building automation application domains share a substantial number of properties:

- o In both domains, the network can be disconnected from the ISP and must still continue to provide control to the occupants of the home/building. Routing needs to be possible independent of the existence of a border router
- o Both domains are subject to unreliable links but require instant and very reliable reactions. This has impact on routing because of timeliness and multipath routing.

The differences between the two application domains mostly appear in commissioning, maintenance and the user interface, which do not typically affect routing. Therefore, the focus of this applicability document is on reliability, timeliness, and local routing.

It should be noted that adherence to the guidance does not necessarily guarantee fully interoperable solutions in home automation networks and building control networks and that additional rigorous and managed programs will be needed to ensure interoperability.

1.1. Relationship to other documents

The Routing Over Low power and Lossy networks (ROLL) working group has specified a set of routing protocols for Low-Power and Lossy Networks (LLN) [RFC6550]. This applicability text describes a subset of those protocols and the conditions under which the subset is appropriate and provides recommendations and requirements for the accompanying parameter value ranges.

In addition, an extension document has been produced specifically to provide a solution for reactive discovery of point-to-point routes in LLNs [RFC6997]. The present applicability document provides recommendations and requirements for the accompanying parameter value ranges.

A common set of security threats are described in [RFC7416]. The applicability statements complement the security threats document by describing preferred security settings and solutions within the applicability statement conditions. This applicability statement recommends lighter weight security solutions appropriate for home and building environments and indicates why these solutions are appropriate.

1.2. 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].

Additionally, this document uses terminology from [RFC6997], [I-D.ietf-roll-trickle-mcast], [RFC7102], [IEEE802.15.4], and [RFC6550].

1.3. Required Reading

Applicable requirements are described in [RFC5826] and [RFC5867]. A survey of the application field is described in [BCsurvey].

1.4. Out of scope requirements

The considered network diameter is limited to a maximum diameter of 10 hops and a typical diameter of 5 hops, which captures the most common cases in home automation and building control networks.

This document does not consider the applicability of Routing Protocol for Low-Power and Lossy Networks (RPL)-related specifications for urban and industrial applications [RFC5548], [RFC5673], which may exhibit significantly larger network diameters.

2. Deployment Scenario

The use of communications networks in buildings is essential to satisfy energy saving regulations. Environmental conditions of buildings can be adapted to suit the comfort of the individuals present inside. Consequently when no one is present, energy consumption can be reduced. Cost is the main driving factor behind deployment of wireless networking in buildings, especially in the case of retrofitting, where wireless connectivity saves costs incurred due to cabling and building modifications.

A typical home automation network is comprised of less than 100 nodes. Large building deployments may span 10,000 nodes but to ensure uninterrupted service of light and air conditioning systems in individual zones of the building, nodes are typically organized in sub-networks. Each sub-network in a building automation deployment typically contains tens to hundreds of nodes, and for critical operations may operate independently from the other sub-networks.

The main purpose of the home or building automation network is to provide control over light and heating/cooling resources. User intervention via wall controllers is combined with movement, light and temperature sensors to enable automatic adjustment of window blinds, reduction of room temperature, etc. In general, the sensors and actuators in a home or building typically have fixed physical locations and will remain in the same home or building automation network.

People expect an immediate and reliable response to their presence or actions. For example, a light not switching on after entry into a room may lead to confusion and a profound dissatisfaction with the lighting product.

Monitoring of functional correctness is at least as important as timely responses. Devices typically communicate their status regularly and send alarm messages notifying a malfunction of controlled equipment or network.

In building control, the infrastructure of the building management network can be shared with the security/access, the Internet Protocol (IP) telephony, and the fire/alarm networks. This approach has a positive impact on the operation and cost of the network; however, care should be taken to ensure that the availability of the building management network does not become compromised beyond the ability for critical functions to perform adequately.

In homes, the entertainment network for audio/video streaming and gaming has different requirements, where the most important requirement is the need for high bandwidth not typically needed for home or building control. It is therefore expected that the entertainment network in the home will mostly be separate from the control network, which also lessens the impact on availability of the control network

2.1. Network Topologies

In general, the home automation network or building control network consists of wired and wireless sub-networks. In large buildings especially, the wireless sub-networks can be connected to an IP backbone network where all infrastructure services are located, such as Domain Name System (DNS), automation servers, etc.

The wireless sub-network can be configured according to any of the following topologies:

- o A stand-alone network of 10-100 nodes without border router. This typically occurs in the home with a stand-alone control network, in low cost buildings, and during installation of high end control systems in buildings.
- o A connected network with one border router. This configuration will happen in homes where home appliances are controlled from outside the home, possibly via a smart phone, and in many building control scenarios.

- o A connected network with multiple border routers. This will typically happen in installations of large buildings.

Many of the nodes are battery-powered and may be sleeping nodes which wake up according to clock signals or external events.

In a building control network, for a large installation with multiple border routers, sub-networks often overlap both geographically and from a wireless coverage perspective. Due to two purposes of the network, (i) direct control and (ii) monitoring, there may exist two types of routing topologies in a given sub-network: (i) a tree-shaped collection of routes spanning from a central building controller via the border router, on to destination nodes in the sub-network; and/or (ii) a flat, un-directed collection of intra-network routes between functionally related nodes in the sub-network.

The majority of nodes in home and building automation networks are typically class 0 devices [RFC7228], such as individual wall switches. Only a few nodes (such as multi-purpose remote controls) are more expensive Class 1 devices, which can afford more memory capacity.

2.2. Traffic Characteristics

Traffic may enter the network originating from a central controller or it may originate from an intra-network node. The majority of traffic is light-weight point-to-point control style; e.g. Put-Ack or Get-Response. There are however exceptions. Bulk data transfer is used for firmware update and logging, where firmware updates enter the network and logs leave the network. Group communication is used for service discovery or to control groups of nodes, such as light fixtures.

Often, there is a direct physical relation between a controlling sensor and the controlled equipment. For example the temperature sensor and room controller are located in the same room sharing the same climate conditions. Consequently, the bulk of senders and receivers are separated by a distance that allows one-hop direct path communication. A graph of the communication will show several fully connected subsets of nodes. However, due to interference, multipath fading, reflection and other transmission mechanisms, the one-hop direct path may be temporally disconnected. For reliability purposes, it is therefore essential that alternative n-hop communication routes exist for quick error recovery. (See Appendix B for motivation.)

Looking over time periods of a day, the networks are very lightly loaded. However, bursts of traffic can be generated by e.g.

incessant pushing of the button of a remote control, the occurrence of a defect, and other unforeseen events. Under those conditions, the timeliness must nevertheless be maintained. Therefore, measures are necessary to remove any unnecessary traffic. Short routes are preferred. Long multi-hop routes via the border router, should be avoided whenever possible.

Group communication is essential for lighting control. For example, once the presence of a person is detected in a given room, lighting control applies to that room only and no other lights should be dimmed, or switched on/off. In many cases, this means that a multicast message with a 1-hop and 2-hop radius would suffice to control the required lights. The same argument holds for Heating, Ventilating, and Air Conditioning (HVAC) and other climate control devices. To reduce network load, it is advisable that messages to the lights in a room are not distributed any further in the mesh than necessary based on intended receivers.

An example of an office surface is shown in [office-light], and the current use of wireless lighting control products is shown in [occuswitch].

2.2.1. General

Whilst air conditioning and other environmental-control applications may accept response delays of tens of seconds or longer, alarm and light control applications may be regarded as soft real-time systems. A slight delay is acceptable, but the perceived quality of service degrades significantly if response times exceed 250 ms. If the light does not turn on at short notice, a user may activate the controls again, thus causing a sequence of commands such as `Light{on,off,on,off,...}` or `Volume{up,up,up,up,up,...}`. In addition the repetitive sending of commands creates an unnecessary loading of the network, which in turn increases the bad responsiveness of the network.

2.2.2. Source-sink (SS) communication paradigm

This paradigm translates to many sources sending messages to the same sink, sometimes reachable via the border router. As such, source-sink (SS) traffic can be present in home and building networks. The traffic may be generated by environmental sensors (often present in a wireless sub-network) which push periodic readings to a central server. The readings may be used for pure logging, or more often, processed to adjust light, heating and ventilation. Alarm sensors may also generate SS style traffic. The central server in a home automation network will be connected mostly to a wired network segment of the home network, although it is likely that cloud

services will also be used. The central server in a building automation network may be connected to a backbone or be placed outside the building.

With regards to message latency, most SS transmissions can tolerate worst-case delays measured in tens of seconds. Fire detectors, however, represent an exception; For example, special provisions with respect to the location of the Fire detectors and the smoke dampers need to be put in place to meet the stringent delay requirements measured in seconds.

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

This paradigm translates to a number of devices expressing their interest for a service provided by a server device. For example, a server device can be a sensor delivering temperature readings on the basis of delivery criteria, like changes in acquisition value or age of the latest acquisition. In building automation networks, this paradigm may be closely related to the SS paradigm given that servers, which are connected to the backbone or outside the building, can subscribe to data collectors that are present at strategic places in the building automation network. The use of PS will probably differ significantly from installation to installation.

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

This paradigm translates to a device transferring data to another device often connected to the same sub-network. Peer-to-peer (P2P) traffic is a common traffic type in home automation networks. Most building automation networks rely on P2P traffic, described in the next paragraph. Other building automation networks rely on P2P control traffic between controls and a local controller box for advanced group control. A local controller box can be further connected to service control boxes, thus generating more SS or PS traffic.

P2P traffic is typically generated by remote controls and wall controllers which push control messages directly to light or heat sources. P2P traffic has a stringent requirement for low latency since P2P traffic often carries application messages that are invoked by humans. As mentioned in Section 2.2.1, application messages should be delivered within a few hundred milliseconds - even when connections fail momentarily.

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

This paradigm translates to a device sending a message as many times as there are destination devices. Peer-to-multipeer (P2MP) traffic is common in home and building automation networks. Often, a thermostat in a living room responds to temperature changes by sending temperature acquisitions to several fans and valves consecutively. This paradigm is also closely related to the PS paradigm in the case where a single server device has multiple subscribers.

2.2.6. Additional considerations: Duocast and N-cast

This paradigm translates to a device sending a message to many destinations in one network transfer invocation. Multicast is well-suited for lighting where a presence sensor sends a presence message to a set of lighting devices. Multicast increases the probability that the message is delivered within the strict time constraints. The recommended multicast algorithm (e.g. [I-D.ietf-roll-trickle-mcast]) provides a mechanism for delivering messages to all intended destinations.

2.2.7. RPL applicability per communication paradigm

In the case of the SS paradigm applied to a wireless sub-network to a server reachable via a border router, the use of RPL [RFC6550] in non-storing mode is appropriate. Given the low resources of the devices, source routing will be used from the border router to the destination in the wireless sub-network for messages generated outside the mesh network. No specific timing constraints are associated with the SS type messages so network repair does not violate the operational constraints. When no SS traffic takes place, it is good practice to load only RPL code enabling P2P mode of operation [RFC6997] to reduce the code size and satisfy memory requirements.

P2P-RPL [RFC6997] is required for all P2P and P2MP traffic taking place between nodes within a wireless sub-network (excluding the border router) to assure responsiveness. Source and destination devices are typically physically close based on room layout. Consequently, most P2P and P2MP traffic is 1-hop or 2-hop traffic. Appendix A explains why P2P-RPL is preferable to RPL for this type of communication. Appendix B explains why reliability measures such as multi-path routing are necessary even when 1-hop communication dominates.

Additional advantages of P2P-RPL for home and building automation networks are, for example:

- o Individual wall switches are typically inexpensive class 0 devices [RFC7228] with extremely low memory capacities. Multi-purpose remote controls for use in a home environment typically have more memory but such devices are asleep when there is no user activity. P2P-RPL reactive discovery allows a node to wake up and find new routes within a few seconds while memory constrained nodes only have to keep routes to relevant targets.
- o The reactive discovery features of P2P-RPL ensure that commands are normally delivered within the 250 ms time window. When connectivity needs to be restored, discovery is typically completed within seconds. In most cases, an alternative (earlier discovered) route will work and route rediscovery is not necessary.
- o Broadcast storms typically associated with route discovery for Ad hoc On-Demand Distance Vector (AODV) [RFC3561] are less disruptive for P2P-RPL. P2P-RPL has a "STOP" bit which is set by the target of a route discovery to notify all other nodes that no more Directed Acyclic Graph (DAG) Information Option (DIO) messages should be forwarded for this temporary DAG. Something looking like a broadcast storm may happen when no target is responding however, in this case, the Trickle suppression mechanism kicks in, limiting the number of DIO forwards in dense networks.

Due to the limited memory of the majority of devices, P2P-RPL SHOULD be deployed with source routing in non-storing mode as explained in Section 4.1.2.

Multicast with Multicast Protocol for Low power and Lossy Networks (MPL) [I-D.ietf-roll-trickle-mcast] is preferably deployed for N-cast over the wireless network. Configuration constraints that are necessary to meet reliability and timeliness with MPL are discussed in Section 4.1.7.

2.3. Layer-2 applicability

This document applies to [IEEE802.15.4] and [G.9959] which are adapted to IPv6 by the adaption layers [RFC4944] and [RFC7428]. Other layer-2 technologies, accompanied by an "IP over Foo" specification, are also relevant provided there is no frame size issue, and there are link layer acknowledgements.

The above mentioned adaptation layers leverage on the compression capabilities of [RFC6554] and [RFC6282]. Header compression allows small IP packets to fit into a single layer 2 frame even when source routing is used. A network diameter limited to 5 hops helps to achieve this even while using source routing.

Dropped packets are often experienced in the targeted environments. Internet Control Message Protocol (ICMP), User Datagram Protocol (UDP) and even Transmission Control Protocol (TCP) flows may benefit from link layer unicast acknowledgments and retransmissions. Link layer unicast acknowledgments SHOULD be enabled when [IEEE802.15.4] or [G.9959] is used with RPL and P2P-RPL.

3. Using RPL to meet Functional Requirements

Several features required by [RFC5826], [RFC5867] challenge the P2P paths provided by RPL. Appendix A reviews these challenges. In some cases, a node may need to spontaneously initiate the discovery of a path towards a desired destination that is neither the root of a DAG, nor a destination originating Destination Advertisement Object (DAO) signalling. Furthermore, P2P paths provided by RPL are not satisfactory in all cases because they involve too many intermediate nodes before reaching the destination.

P2P-RPL [RFC6997] SHOULD be used in home automation and building control networks, as point-to-point style traffic is substantial and route repair needs to be completed within seconds. P2P-RPL provides a reactive mechanism for quick, efficient and root-independent route discovery/repair. The use of P2P-RPL furthermore allows data traffic to avoid having to go through a central region around the root of the tree, and drastically reduces path length [SOFT11] [INTEROP12]. These characteristics are desirable in home and building automation networks because they substantially decrease unnecessary network congestion around the root of the tree.

When more reliability is required, P2P-RPL enables the establishment of multiple independent paths. For 1-hop destinations this means that one 1-hop communication and a second 2-hop communication take place via a neighbouring node. Such a pair of redundant communication paths can be achieved by using MPL where the source is a MPL forwarder, while a second MPL forwarder is 1 hop away from both the source and the destination node. When the source multicasts the message, it may be received by both the destination and the 2nd forwarder. The 2nd forwarder forwards the message to the destination, thus providing two routes from sender to destination.

To provide more reliability with multiple paths, P2P-RPL can maintain two independent P2P source routes per destination, at the source. Good practice is to use the paths alternately to assess their existence. When one P2P path has failed (possibly only temporarily), as described in Appendix B, the alternative P2P path can be used without discarding the failed path. The failed P2P path, unless proven to work again, can be safely discarded after a timeout

(typically 15 minutes). A new route discovery is done when the number of P2P paths is exhausted due to persistent link failures.

4. RPL Profile

P2P-RPL SHOULD be used in home automation and building control networks. Its reactive discovery allows for low application response times even when on-the-fly route repair is needed. Non-storing mode SHOULD be used to reduce memory consumption in repeaters with constrained memory when source routing is used.

4.1. RPL Features

An important constraint on the application of RPL is the presence of sleeping nodes.

For example, in a stand-alone network, the master node (or coordinator) providing the logical layer-2 identifier and unique node identifiers to connected nodes may be a remote control which returns to sleep once new nodes have been added. Due to the absence of the border router, there may be no global routable prefixes at all. Likewise, there may be no authoritative always-on root node since there is no border router to host this function.

In a network with a border router and many sleeping nodes, there may be battery powered sensors and wall controllers configured to contact other nodes in response to events and then return to sleep. Such nodes may never detect the announcement of new prefixes via multicast.

In each of the above mentioned constrained deployments, a link layer node (e.g. coordinator or master) SHOULD assume the role of authoritative root node, transmitting unicast Router Advertisement (RA) messages with a Unique Local Address (ULA) prefix information option to nodes during the joining process to prepare the nodes for a later operational phase, where a border router is added.

A border router SHOULD be designed to be aware of sleeping nodes in order to support the distribution of updated global prefixes to such sleeping nodes.

4.1.1. RPL Instances

When operating P2P-RPL on a stand-alone basis, there is no authoritative root node maintaining a permanent RPL Direction-Oriented Directed Acyclic Graph (DODAG). A node MUST be able to join at least one RPL instance, as a new, temporary instance is created

during each P2P-RPL route discovery operation. A node MAY be designed to join multiple RPL instances.

4.1.2. Storing vs. Non-Storing Mode

Non-storing mode MUST be used to cope with the extremely constrained memory of a majority of nodes in the network (such as individual light switches).

4.1.3. DAO Policy

Nodes send DAO messages to establish downward paths from the root to themselves. DAO messages are not acknowledged in networks composed of battery operated field devices in order to minimize the power consumption overhead associated with path discovery. The DAO messages build up a source route because the nodes MUST be in non-storing mode.

If devices in LLNs participate in multiple RPL instances and DODAGs, both the RPLInstance ID and the DODAGID SHOULD be included in the DAO.

4.1.4. Path Metrics

Expected Transmission Count (ETX) is the RECOMMENDED metric. [RFC6551] provides other options.

Packets from asymmetric and/or unstable links SHOULD be deleted at layer 2.

4.1.5. Objective Function

Objective Function 0 (OF0) MUST be the Objective Function. Other Objective Functions MAY be used when dictated by circumstances.

4.1.6. DODAG Repair

Since P2P-RPL only creates DODAGs on a temporary basis during route repair or route discovery, there is no need to repair DODAGs.

For SS traffic, local repair is sufficient. The accompanying process is known as poisoning and is described in Section 8.2.2.5 of [RFC6550]. Given that the majority of nodes in the building do not physically move around, creating new DODAGs should not happen frequently.

4.1.7. Multicast

Commercial lighting deployments may have a need for multicast to distribute commands to a group of lights in a timely fashion. Several mechanisms exist for achieving such functionality; [I-D.ietf-roll-trickle-mcast] is the RECOMMENDED protocol for home and building deployments. This section relies heavily on the conclusions of [RT-MPL].

At reception of a packet, the MPL forwarder starts a series of consecutive trickle timer intervals, where the first interval has a minimum size of I_{min} . Each consecutive interval is twice as long as the former with a maximum value of I_{max} . There is a maximum number of intervals given by $max_expiration$. For each interval of length I , a time t is randomly chosen in the period $[I/2, I]$. For a given packet, p , MPL counts the number of times it receives p during the period $[0, t]$ in a counter c . At time t , MPL re-broadcasts p when $c < k$, where k is a predefined constant with a value $k > 0$.

The density of forwarders and the frequency of message generation are important aspects to obtain timeliness during control operations. A high frequency of message generation can be expected when a remote control button is incessantly pressed, or when alarm situations arise.

Guaranteeing timeliness is intimately related to the density of the MPL routers. In ideal circumstances the message is propagated as a single wave through the network, such that the maximum delay is related to the number of hops times the smallest repetition interval of MPL. Each forwarder that receives the message passes the message on to the next hop by repeating the message. When several copies of a message reach the forwarder, it is specified that the copy need not be repeated. Repetition of the message can be inhibited by a small value of k . To assure timeliness, the value of k should be chosen high enough to make sure that messages are repeated at the first arrival of the message in the forwarder. However, a network that is too dense leads to a saturation of the medium that can only be prevented by selecting a low value of k . Consequently, timeliness is assured by choosing a relatively high value of k but assuring at the same time a low enough density of forwarders to reduce the risk of medium saturation. Depending on the reliability of the network links, it is advisable to choose the network such that at least 2 forwarders per hop repeat messages to the same set of destinations.

There are no rules about selecting forwarders for MPL. In buildings with central management tools, the forwarders can be selected, but in the home is not possible to automatically configure the forwarder topology at the time of writing this document.

4.1.8. Security

RPL MAY use unsecured RPL messages to reduce message size. If there is a single node that uses unsecured RPL messages, link-layer security MUST be used on all nodes. Therefore all RPL messages MUST be secured using either:

- o RPL message security, or
- o Link-layer security, or
- o Both RPL message security and link-layer security

A symmetric key is used to secure a RPL message using either RPL message security or link-layer security. The symmetric key MUST be distributed or established in a secure fashion. There may be more than one symmetric key in use by any node at any one time. The same symmetric key MUST NOT be used for both RPL message security and link-layer security between two peer nodes.

4.1.8.1. Symmetric key distribution

The scope of symmetric key distribution MUST be no greater than the network itself, i.e. a group key. This document describes what needs to be implemented to meet this requirement. The scope of symmetric key distribution MAY be smaller than the network, for example:

- o A pairwise symmetric key between two peers.
- o A group key shared between a subset of nodes in the network.

4.1.8.2. Symmetric key distribution mechanism

The authentication mechanism as described in Section 6.9 of [ZigBeeIP] SHALL be used to securely distribute a network-wide symmetric key.

The purpose of the authentication procedure is to provide mutual authentication resulting in:

- o Preventing untrusted nodes without appropriate credentials from joining the trusted network.
- o Preventing trusted nodes with appropriate credentials from joining an untrusted network.

There is an Authentication Server, which is responsible for authenticating the nodes on the network. If the authentication is

successful, the Authentication Server sends the network security material to the joining node through the PANA protocol ([RFC5191], [RFC6345]). The joining node becomes a full participating node in the network and is able to apply layer 2 security to RPL messages using the distributed network key.

The joining node does not initially have access to the network security material. Therefore, it is not able to apply layer 2 security for the packets exchanged during the authentication process. The enforcement point rules at the edge of the network ensure that the packets involved in the PANA authentication are processed even though they are unsecured at MAC layer. The rules also ensure that any other incoming traffic that is not secured at the MAC layer is discarded and is not forwarded.

4.1.8.2.1. Authentication Stack

Authentication can be viewed as a protocol stack as a layer encapsulates the layers above it.

- o TLS [RFC5246] MUST be used at the highest layer of the authentication stack and carries the authentication exchange. There is one cipher suite based on pre-shared key [RFC6655] and one cipher suite based on ECC [RFC7251].
- o EAP-TLS [RFC5216] MUST be used at the next layer to carry the TLS records for the authentication protocol.
- o The Extensible Authentication Protocol [RFC3748] MUST be used to provide the mechanisms for mutual authentication. EAP requires a way to transport EAP packets between the joining node and the node on which the Authentication Server resides. These nodes are not necessarily in radio range of each other, so it is necessary to have multi-hop support in the EAP transport method. The PANA protocol [RFC5191], [RFC6345], which operates over UDP, MUST be used for this purpose. [RFC3748] specifies the derivation of a session key using the EAP key hierarchy; only the EAP Master Session Key shall be derived, as [RFC5191] specifies that it is used to set up keys for PANA authentication and encryption.
- o PANA [RFC5191] and PANA relay [RFC6345] MUST be used at the next layer:
 - * The joining node MUST act as the PANA Client (PaC)
 - * The parent edge router node MUST act as a PANA relay (PRE) according to [RFC6345], unless it is also the Authentication

Server. All routers at the edge of the network MUST be capable of functioning in the PRE role.

- * The Authentication Server node MUST act as the PANA Authentication Agent (PAA). The Authentication Server MUST be able to handle packets relayed according to [RFC6345].

This network authentication process uses link-local IPv6 addresses for transport between the new node and its parent. If the parent is not the Authentication Server, it MUST then relay packets from the joining node to the Authentication Server and vice-versa using PANA relay mechanism [RFC6345]. The joining node MUST use a link-local address based on its EUI-64 as the source address for initial PANA authentication message exchanges.

4.1.8.2.2. Applicability Statements

The applicability statements describe the relationship between the various specifications.

4.1.8.2.2.1. Applicability Statement for PSK TLS

[RFC6655] contains AEAD TLS cipher suites that are very similar to [RFC5487] whose AEAD part is detailed in [RFC5116]. [RFC5487] references both [RFC5288] and the original PSK cipher suite document [RFC4279], which references [RFC5246], which defines the TLS 1.2 messages.

4.1.8.2.2.2. Applicability Statement for ECC TLS

[RFC7251] contains AEAD TLS cipher suites that are very similar to [RFC5289] whose AEAD part is detailed in [RFC5116]. [RFC5289] references the original ECC cipher suite document [RFC4492], which references [RFC5246], which defines the TLS 1.2 messages.

4.1.8.2.2.3. Applicability Statement for EAP-TLS and PANA

[RFC5216] specifies how [RFC3748] is used to package [RFC5246] TLS records into EAP packets. [RFC5191] provides transportation for the EAP packets and the network-wide key carried in an encrypted AVP specified in [RFC6786]. The proposed PRF and AUTH hashes based on SHA-256 are represented as in [RFC5996] and detailed in [RFC4868].

4.1.8.2.3. Security using RPL message security

If RPL is used with secured messages [RFC6550], the following RPL security parameter values SHOULD be used:

- o Counter Time Flag (T) = 0: Do not use timestamp in the Counter Field. Counters based on timestamps are typically more applicable to industrial networks where strict timing synchronization between nodes is often implemented. Home and building networks typically do not implement such strict timing synchronization therefore a monotonically increasing counter is more appropriate.
- o Algorithm = 0: Use Counter with Cipher Block Chaining Message Authentication Code (CBC-MAC Mode) (CCM) with Advanced Encryption Standard (AES)-128. This is the only assigned mode at present.
- o Key Identifier Mode (KIM) = 10: Use group key, Key Source present, Key Index present. Given the relatively confined perimeter of a home or building network, a group key is usually sufficient to protect RPL messages sent between nodes. The use of the Key Source field allows multiple group keys to be used within the network.
- o Security Level (LVL) = 0: Use MAC-32. This is recommended as integrity protection for RPL messages is the basic requirement. Encryption is unlikely to be necessary given the relatively non-confidential nature of RPL message payloads.

4.1.9. P2P communications

[RFC6997] MUST be used to accommodate P2P traffic, which is typically substantial in home and building automation networks.

4.1.10. IPv6 address configuration

Assigned IP addresses MUST be routable and unique within the routing domain [RFC5889].

4.2. Layer 2 features

No particular requirements exist for layer 2 but for the ones cited in the IP over Foo RFCs (see Section 2.3).

4.2.1. Specifics about layer-2

Not applicable

4.2.2. Services provided at layer-2

Not applicable

4.2.3. 6LowPAN options assumed

Not applicable

4.2.4. Mesh Link Establishment (MLE) and other things

Not applicable

4.3. Recommended Configuration Defaults and Ranges

The following sections describe the recommended parameter values for P2P-RPL and Trickle.

4.3.1. Trickle parameters

Trickle is used to distribute network parameter values to all nodes without stringent time restrictions. The recommended Trickle parameter values are:

- o DIOIntervalMin 4 = 16 ms
- o DIOIntervalDoublings 14
- o DIORedundancyConstant 1

When a node sends a changed DIO, this is an inconsistency and forces the receiving node to respond within I_{min} . So when something happens which affects the DIO, the change is ideally communicated to a node, n hops away, within n times I_{min} . Often, dependent on the node density, packets are lost, or not sent, leading to larger delays.

In general we can expect DIO changes to propagate within 1 to 3 seconds within the envisaged networks.

When nothing happens, the DIO sending interval increases to 4.37 minutes, thus drastically reducing the network load. When a node does not receive DIO messages during more than 10 minutes it can safely conclude the connection with other nodes has been lost.

4.3.2. Other Parameters

This section discusses the P2P-RPL parameters.

P2P-RPL [RFC6997] provides the features requested by [RFC5826] and [RFC5867]. P2P-RPL uses a subset of the frame formats and features defined for RPL [RFC6550] but may be combined with RPL frame flows in advanced deployments.

The recommended parameter values for P2P-RPL are:

- o MinHopRankIncrease 1
- o MaxRankIncrease 0
- o MaxRank 6
- o Objective function: OF0

5. MPL Profile

MPL is used to distribute values to groups of devices. Using MPL, based on the Trickle algorithm, timeliness should also be guaranteed. A deadline of 200 ms needs to be met when human action is followed by an immediately observable action such as switching on lights. The deadline needs to be met in a building where the number of hops from seed to destination varies between 1 and 10.

5.1. Recommended configuration Defaults and Ranges

5.1.1. Real-Time optimizations

When the network is heavily loaded, MAC delays contribute significantly to the end to end delays when MPL intervals between 10 to 100 ms are used to meet the 200 ms deadline. It is possible to set the number of buffers in the MAC to 1 and set the number of Back-off repetitions to 1. The number of MPL repetitions compensates for the reduced probability of transmission per MAC invocation [RT-MPL].

In addition, end to end delays and message losses are reduced, by adding a real-time layer between MPL and MAC to throw away the earliest messages (exploiting the MPL message numbering) and favour the most recent ones.

5.1.2. Trickle parameters

This section proposes values for the Trickle parameters used by MPL for the distribution of packets that need to meet a 200 ms deadline. The probability of meeting the deadline is increased by (1) choosing a small Imin value, (2) reducing the number of MPL intervals thus reducing the load, and (3) reducing the number of MPL forwarders to also reduce the load.

The consequence of this approach is that the value of k can be larger than 1 because network load reduction is already guaranteed by the network configuration.

Under the condition that the density of MPL repeaters can be limited, it is possible to choose low MPL repeat intervals (I_{min}) connected to k values such that $k > 1$. The minimum value of k is related to:

- o Value of I_{min} . The length of I_{min} determines the number of packets that can be received within the listening period of I_{min} .
- o Number of repeaters receiving the broadcast message from the same forwarder or seed. These repeaters repeat within the same I_{min} interval, thus increasing the c counter.

Within the first MPL interval a limited number, q , of messages can be transmitted. Assuming a 3 ms transmission interval, q is given by $q = I_{min}/3$. Assuming that at most q message copies can reach a given forwarder within the first repeat interval of length I_{min} , the related MPL parameter values are suggested in the following sections.

5.1.2.1. I_{min}

The recommended value is $I_{min} = 10$ to 50 ms.

When I_{min} is chosen much smaller, the interference between the copies leads to significant losses given that q is much smaller than the number of repeated packets. With much larger intervals the probability that the deadline will be met decreases with increasing hop count.

5.1.2.2. I_{max}

The recommended value is $I_{max} = 100$ to 400 ms.

The value of I_{max} is less important than the value of $max_expiration$. Given an I_{min} value of 10 ms, the 3rd MPL interval has a value of $10 * 2 * 2 = 40$ ms. When I_{min} has a value of 40 ms, the 3rd interval has a value of 160 ms. Given that more than 3 intervals are unnecessary, the I_{max} does not contribute much to the performance.

5.1.3. Other parameters

Other parameters are the k parameter and the $max_expiration$ parameter.

$k > q$ (see condition above). Under this condition and for small I_{min} , a value of $k=2$ or $k=3$ is usually sufficient to minimize the losses of packets in the first repeat interval.

$max_expiration = 2 - 4$. Higher values lead to more network load while generating copies which will probably not meet their deadline.

6. Manageability Considerations

At this moment it is not clear how homenets will be managed. Consequently it is not clear which tools will be used and which parameters must be exposed for management.

In building control, management is mandatory. It is expected that installations will be managed using the set of currently available tools (including IETF tools like Management Information Base (MIB) modules, NETCONF modules, Dynamic Host Configuration Protocol (DHCP) and others) with large differences between the ways an installation is managed.

7. Security Considerations

This section refers to the security considerations of [RFC6997], [RFC6550], [I-D.ietf-roll-trickle-mcast], and the counter measures discussed in sections 6 and 7 of [RFC7416].

Communications network security is based on providing integrity protection and encryption to messages. This can be applied at various layers in the network protocol stack based on using various credentials and a network identity.

The credentials which are relevant in the case of RPL are: (i) the credential used at the link layer in the case where link layer security is applied (see Section 7.1) or (ii) the credential used for securing RPL messages. In both cases, the assumption is that the credential is a shared key. Therefore, there MUST be a mechanism in place which allows secure distribution of a shared key and configuration of network identity. Both MAY be done using: (i) pre-installation using an out-of-band method, (ii) delivered securely when a device is introduced into the network or (iii) delivered securely by a trusted neighbouring device as described in Section 4.1.8.1. The shared key MUST be stored in a secure fashion which makes it difficult to be read by an unauthorized party.

This document mandates that a layer-2 mechanism be used during initial and incremental deployment. Please see the following sections.

7.1. Security considerations during initial deployment

Wireless mesh networks are typically secured at the link layer in order to prevent unauthorized parties from accessing the information exchanged over the links. It is a basic practice to create a network of nodes which share the same keys for link layer security and exclude nodes sending unsecured messages. With per-message data

origin authentication, it is possible to prevent unauthorized nodes joining the mesh.

At initial deployment the network is secured by consecutively securing nodes at the link layer, thus building a network of secured nodes. Section 4.1.8.2 describes a mechanism for building a network of secured nodes.

This document does not specify a multicast security solution. Networks deployed with this specification will depend upon layer-2 security to prevent outsiders from sending multicast traffic. It is recognized that this does not protect this control traffic from impersonation by already trusted devices. This is an area for a future specification.

For building control an installer will use an installation tool that establishes a secure communication path with the joining node. It is recognized that the recommendations for initial deployment of Section 7 and Section 7.1 do not cover all building requirements such as selecting the node-to-secure independent of network topology.

It is expected that a set of protocol combinations will evolve within currently existing alliances of building control manufacturers. Each set satisfies the installation requirements of installers, operators, and manufacturers of building control networks in a given installation context, e.g. lighting deployment in offices, HVAC installation, incremental addition of equipment in homes, and others.

In the home, nodes can be visually inspected by the home owner and a simple procedure, e.g. pushing buttons simultaneously on an already secured device and an unsecured joining device is usually sufficient to ensure that the unsecured joining device is authenticated and configured securely, and paired appropriately.

This recommendation is in line with the countermeasures described in section 6.1.1 of [RFC7416].

7.2. Security Considerations during incremental deployment

Once a network is operational, new nodes need to be added, or nodes fail and need to be replaced. When a new node needs to be added to the network, the new node is joined to the network via an assisting node in the manner described in Section 7.1.

On detection of a compromised node, all trusted nodes need to have their symmetric keys known to be shared with the compromised node re-keyed, and the trusted network is built up as described in Section 7.1.

7.3. Security Considerations for P2P uses

Refer to the security considerations of [RFC6997].

7.4. MPL routing

The routing of MPL is determined by the enabling of the interfaces for specified Multicast addresses. The specification of these addresses can be done via a Constrained Application Protocol (CoAP) application as specified in [RFC7390]. An alternative is the creation of a MPL MIB and use of Simple Network Management Protocol (SNMP)v3 [RFC3411] or equivalent techniques to specify the Multicast addresses in the MIB. For secure dissemination of MPL packets, layer 2 security SHOULD be used and the configuration of multicast addresses as described in this section MUST be secure.

7.5. RPL Security features

This section follows the structure of section 8, "RPL security features" of [RFC7416]. [RFC7416] provides a thorough analysis of security threats and proposed counter measures relevant to RPL and MPL.

In accordance with section 8.1 of [RFC7416], "Confidentiality features", RPL message security implements payload protection, as explained in Section 7 of this document. The attributes key-length and life-time of the keys depend on operational conditions, maintenance and installation procedures.

Section 7.1 and Section 7.2 of this document recommend link-layer security to assure integrity in accordance with section 8.2 of [RFC7416], "Integrity features".

The provision of multiple paths recommended in section 8.3 "Availability features" of [RFC7416] is also recommended from a reliability point of view. Randomly choosing paths MAY be supported.

A mechanism for key management, discussed in section 8.4, "Key Management" of [RFC7416], is provided in Section 4.1.8.2.

Section 7.5, "Considerations on Matching Application Domain Needs" of [RFC7416] applies as such.

8. Other related protocols

Application and transport protocols used in home and building automation domains are expected to mostly consist in CoAP over UDP, or equivalents. Typically, UDP is used for IP transport to keep down

the application response time and bandwidth overhead. CoAP is used at the application layer to reduce memory footprint and bandwidth requirements.

9. IANA Considerations

No considerations for IANA pertain to this document.

10. Acknowledgements

This document reflects discussions and remarks from several individuals including (in alphabetical order): Stephen Farrell, Mukul Goyal, Sandeep Kumar, Jerry Martocci, Catherine Meadows, Yoshihira Ohba, Charles Perkins, Yvonne-Anne Pignolet, Michael Richardson, Ines Robles, Zach Shelby, and Meral Sherazipour.

11. Changelog

RFC editor, please delete this section before publication.

Changes from version 0 to version 1.

- o Adapted section structure to template.
- o Standardized the reference syntax.
- o Section 2.2, moved everything concerning algorithms to section 2.2.7, and adapted text in 2.2.1-2.2.6.
- o Added MPL parameter text to section 4.1.7 and section 4.3.1.
- o Replaced all TODO sections with text.
- o Consistent use of border router, monitoring, home- and building network.
- o Reformulated security aspects with references to other publications.
- o MPL and RPL parameter values introduced.

Changes from version 1 to version 2.

- o Clarified common characteristics of control in home and building.
- o Clarified failure behaviour of point to point communication in appendix.

- o Changed examples, more hvac and less lighting.
- o Clarified network topologies.
- o replaced reference to smart_object paper by reference to I-D.roll-security-threats
- o Added a concise definition of secure delivery and secure storage
- o Text about securing network with PANA

Changes from version 2 to version 3.

- o Changed security section to follow the structure of security threats draft.
- o Added text to DODAG repair sub-section

Changes from version 3 to version 4.

- o Renumbered sections and moved text to conform to applicability template
- o Extended MPL parameter value text
- o Added references to building control products

Changes from version 4 to version 5.

- o Large editing effort to streamline text
- o Rearranged Normative and Informative references
- o Replaced RFC2119 terminology by non-normative terminology
- o Rearranged text of section 7, 7.1, and 7.2 to agree with the intention of section 7.2

Changes from version 5 to version 6.

- o Issues #162 - #166 addressed

Changes from version 6 to version 7.

- o Text of section 7.1 edited for better security coverage.

Changes from version 7 to version 8.

- o Requirements language paragraph removed
- o Acronyms clarified
- o MPL parameters clarified

Changes from version 8 to version 9.

- o More acronyms clarified
- o References updated

Changes from version 9 to version 10.

- o Changes due to IESG and security review
- o Requirements language reinstated
- o RPL security parameter selection clarified
- o Removed multicast security reference

Changes from version 10 to 11.

- o Further changes due to IESG and security review
- o ZigBee IP authentication and key establishment specified

Changes from version 11 to 12.

- o Further clarifications added

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004, <<http://www.rfc-editor.org/info/rfc3748>>.
- [RFC4279] Eronen, P., Ed. and H. Tschofenig, Ed., "Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)", RFC 4279, DOI 10.17487/RFC4279, December 2005, <<http://www.rfc-editor.org/info/rfc4279>>.

- [RFC4492] Blake-Wilson, S., Bolyard, N., Gupta, V., Hawk, C., and B. Moeller, "Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS)", RFC 4492, DOI 10.17487/RFC4492, May 2006, <<http://www.rfc-editor.org/info/rfc4492>>.
- [RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", RFC 4868, DOI 10.17487/RFC4868, May 2007, <<http://www.rfc-editor.org/info/rfc4868>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC5116] McGrew, D., "An Interface and Algorithms for Authenticated Encryption", RFC 5116, DOI 10.17487/RFC5116, January 2008, <<http://www.rfc-editor.org/info/rfc5116>>.
- [RFC5191] Forsberg, D., Ohba, Y., Ed., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, DOI 10.17487/RFC5191, May 2008, <<http://www.rfc-editor.org/info/rfc5191>>.
- [RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, DOI 10.17487/RFC5216, March 2008, <<http://www.rfc-editor.org/info/rfc5216>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5288] Salowey, J., Choudhury, A., and D. McGrew, "AES Galois Counter Mode (GCM) Cipher Suites for TLS", RFC 5288, DOI 10.17487/RFC5288, August 2008, <<http://www.rfc-editor.org/info/rfc5288>>.
- [RFC5289] Rescorla, E., "TLS Elliptic Curve Cipher Suites with SHA-256/384 and AES Galois Counter Mode (GCM)", RFC 5289, DOI 10.17487/RFC5289, August 2008, <<http://www.rfc-editor.org/info/rfc5289>>.
- [RFC5487] Badra, M., "Pre-Shared Key Cipher Suites for TLS with SHA-256/384 and AES Galois Counter Mode", RFC 5487, DOI 10.17487/RFC5487, March 2009, <<http://www.rfc-editor.org/info/rfc5487>>.

- [RFC5548] Dohler, M., Ed., Watteyne, T., Ed., Winter, T., Ed., and D. Barthel, Ed., "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, DOI 10.17487/RFC5548, May 2009, <<http://www.rfc-editor.org/info/rfc5548>>.
- [RFC5673] Pister, K., Ed., Thubert, P., Ed., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, DOI 10.17487/RFC5673, October 2009, <<http://www.rfc-editor.org/info/rfc5673>>.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, DOI 10.17487/RFC5826, April 2010, <<http://www.rfc-editor.org/info/rfc5826>>.
- [RFC5867] Martocci, J., Ed., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, DOI 10.17487/RFC5867, June 2010, <<http://www.rfc-editor.org/info/rfc5867>>.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, DOI 10.17487/RFC5996, September 2010, <<http://www.rfc-editor.org/info/rfc5996>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6345] Duffy, P., Chakrabarti, S., Cragie, R., Ohba, Y., Ed., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA) Relay Element", RFC 6345, DOI 10.17487/RFC6345, August 2011, <<http://www.rfc-editor.org/info/rfc6345>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc6551>>.

- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<http://www.rfc-editor.org/info/rfc6554>>.
- [RFC6655] McGrew, D. and D. Bailey, "AES-CCM Cipher Suites for Transport Layer Security (TLS)", RFC 6655, DOI 10.17487/RFC6655, July 2012, <<http://www.rfc-editor.org/info/rfc6655>>.
- [RFC6786] Yegin, A. and R. Cragie, "Encrypting the Protocol for Carrying Authentication for Network Access (PANA) Attribute-Value Pairs", RFC 6786, DOI 10.17487/RFC6786, November 2012, <<http://www.rfc-editor.org/info/rfc6786>>.
- [RFC6997] Goyal, M., Ed., Baccelli, E., Philipp, M., Brandt, A., and J. Martocci, "Reactive Discovery of Point-to-Point Routes in Low-Power and Lossy Networks", RFC 6997, DOI 10.17487/RFC6997, August 2013, <<http://www.rfc-editor.org/info/rfc6997>>.
- [RFC6998] Goyal, M., Ed., Baccelli, E., Brandt, A., and J. Martocci, "A Mechanism to Measure the Routing Metrics along a Point-to-Point Route in a Low-Power and Lossy Network", RFC 6998, DOI 10.17487/RFC6998, August 2013, <<http://www.rfc-editor.org/info/rfc6998>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC7251] McGrew, D., Bailey, D., Campagna, M., and R. Dugal, "AES-CCM Elliptic Curve Cryptography (ECC) Cipher Suites for TLS", RFC 7251, DOI 10.17487/RFC7251, June 2014, <<http://www.rfc-editor.org/info/rfc7251>>.
- [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, <<http://www.rfc-editor.org/info/rfc7416>>.
- [I-D.ietf-roll-trickle-mcast]
Hui, J. and R. Kelsey, "Multicast Protocol for Low power and Lossy Networks (MPL)", draft-ietf-roll-trickle-mcast-12 (work in progress), June 2015.

[IEEE802.15.4]

"IEEE 802.15.4 - Standard for Local and metropolitan area networks -- Part 15.4: Low-Rate Wireless Personal Area Networks", <IEEE Standard 802.15.4>.

[G.9959] "ITU-T G.9959 Short range narrow-band digital radiocommunication transceivers - PHY and MAC layer specifications", <ITU-T G.9959>.

12.2. Informative References

- [RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <<http://www.rfc-editor.org/info/rfc3411>>.
- [RFC3561] Perkins, C., Belding-Royer, E., and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", RFC 3561, DOI 10.17487/RFC3561, July 2003, <<http://www.rfc-editor.org/info/rfc3561>>.
- [RFC5889] Baccelli, E., Ed. and M. Townsley, Ed., "IP Addressing Model in Ad Hoc Networks", RFC 5889, DOI 10.17487/RFC5889, September 2010, <<http://www.rfc-editor.org/info/rfc5889>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.
- [RFC7390] Rahman, A., Ed. and E. Dijk, Ed., "Group Communication for the Constrained Application Protocol (CoAP)", RFC 7390, DOI 10.17487/RFC7390, October 2014, <<http://www.rfc-editor.org/info/rfc7390>>.
- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.
- [SOFT11] Baccelli, E., Phillip, M., and M. Goyal, "The P2P-RPL Routing Protocol for IPv6 Sensor Networks: Testbed Experiments", Proceedings of the Conference on Software Telecommunications and Computer Networks, Split, Croatia,, September 2011.

[INTEROP12]

Baccelli, E., Phillip, M., Brandt, A., Valev, H., and J. Buron, "Report on P2P-RPL Interoperability Testing", RR-7864 INRIA Research Report RR-7864, January 2012.

[RT-MPL]

van der Stok, P., "Real-Time multicast for wireless mesh networks using MPL", White paper, <http://www.vanderstok.org/papers/Real-time-MPL.pdf>, April 2014.

[occuswitch]

Lighting, Philips., "OccuSwitch wireless", Brochure, http://www.philipslightingcontrols.com/assets/cms/uploads/files/osw/MK_OSWNETBROC_5.pdf, May 2012.

[office-light]

Clanton and Associates, ., "A Life Cycle Cost Evaluation of Multiple Lighting Control Strategies", Wireless Lighting Control, http://www.daintree.net/wp-content/uploads/2014/02/clanton_lighting_control_report_0411.pdf, February 2014.

[RTN2011]

Holtman, K. and P. van der Stok, "Real-time routing for low-latency 802.15.4 control networks", International Workshop on Real-Time Networks; Euromicro Conference on Real-Time Systems, July 2011.

[MEAS]

Holtman, K., "Connectivity loss in large scale IEEE 802.15.4 network", Private Communication, November 2013.

[BCsurvey]

Kastner, W., Neugschwandtner, G., Soucek, S., and H. Newman, "Communication Systems for Building Automation and Control", Proceedings of the IEEE Vol 93, No 6, June 2005.

[ZigBeeIP]

ZigBee Alliance, ., "ZigBee IP specification", ZigBee document 095023r34, March 2014.

Appendix A. RPL shortcomings in home and building deployments

A.1. Risk of undesired long P2P routes

The DAG, being a tree structure is formed from a root. If nodes residing in different branches have a need for communicating internally, DAG mechanisms provided in RPL [RFC6550] will propagate traffic towards the root, potentially all the way to the root, and

down along another branch [RFC6998]. In a typical example two nodes could reach each other via just two router nodes but in unfortunate cases, RPL may send traffic three hops up and three hops down again. This leads to several undesired phenomena described in the following sections.

A.1.1. Traffic concentration at the root

If many P2P data flows have to move up towards the root to get down again in another branch there is an increased risk of congestion the nearer to the root of the DAG the data flows. Due to the broadcast nature of RF systems any child node of the root is not just directing RF power downwards its sub-tree but just as much upwards towards the root; potentially jamming other MP2P traffic leaving the tree or preventing the root of the DAG from sending P2MP traffic into the DAG because the listen-before-talk link-layer protection kicks in.

A.1.2. Excessive battery consumption in source nodes

Battery-powered nodes originating P2P traffic depend on the route length. Long routes cause source nodes to stay awake for longer periods before returning to sleep. Thus, a longer route translates proportionally (more or less) into higher battery consumption.

A.2. Risk of delayed route repair

The RPL DAG mechanism uses DIO and DAO messages to monitor the health of the DAG. In rare occasions, changed radio conditions may render routes unusable just after a destination node has returned a DAO indicating that the destination is reachable. Given enough time, the next Trickle timer-controlled DIO/DAO update will eventually repair the broken routes, however this may not occur in a timely manner appropriate to the application. In an apparently stable DAG, Trickle-timer dynamics may reduce the update rate to a few times every hour. If a user issues an actuator command, e.g. light on in the time interval between the last DAO message was issued the destination module and the time one of the parents sends the next DIO, the destination cannot be reached. There is no mechanism in RPL to initiate restoration of connectivity in a reactive fashion. The consequence is a broken service in home and building applications.

A.2.1. Broken service

Experience from the telecom industry shows that if the voice delay exceeds 250ms, users start getting confused, frustrated and/or annoyed. In the same way, if the light does not turn on within the same period of time, a home control user will activate the controls again, causing a sequence of commands such as

Light{on,off,off,on,off,...} or Volume{up,up,up,up,up,...}. Whether the outcome is nothing or some unintended response this is unacceptable. A controlling system must be able to restore connectivity to recover from the error situation. Waiting for an unknown period of time is not an option. While this issue was identified during the P2P analysis, it applies just as well to application scenarios where an IP application outside the LLN controls actuators, lights, etc.

Appendix B. Communication failures

Measurements on the connectivity between neighbouring nodes are discussed in [RTN2011] and [MEAS].

The work is motivated by the measurements in literature which affirm that the range of an antenna is not circle symmetric but that the signal strength of a given level follows an intricate pattern around the antenna, and there may be holes within the area delineated by an iso-strength line. It is reported that communication is not symmetric: reception of messages from node A by node B does not imply reception of messages from node B by node A. The quality of the signal fluctuates over time, and also the height of the antenna within a room can have consequences for the range. As function of the distance from the source, three regions are generally recognized: (1) a clear region with excellent signal quality, (2) a region with fluctuating signal quality, (3) a region without reception. In the text below it is shown that installation of meshes with neighbours in the clear region is not sufficient.

[RTN2011] extends existing work by:

- o Observations over periods of at least a week,
- o Testing links that are in the clear region,
- o Observation in an office building during working hours,
- o Concentrating on one-hop and two-hop routes.

Eight nodes were distributed over a surface of 30m². All nodes are at one hop distance from each other and are situated in the clear region of each other. Each node sends messages to each of its neighbours, and repeats the message until it arrives. The latency of the message was measured over periods of at least a week. It is noticed that latencies longer than a second occurred without apparent reasons, but only during working days and never in the weekends. Bad periods could last for minutes. By sending messages via two paths: (1) one hop path directly, and (2) two hop path via a randomly chosen

neighbour, the probability of delays larger than 100 ms decreased significantly.

The conclusion is that even for 1-hop communication between not too distant "Line of Sight" nodes, there are periods of low reception in which communication deadlines of 200 ms are exceeded. It pays to send a second message over a 2-hop path to increase the reliability of timely message transfer.

[MEAS] confirms that temporary bad reception by close neighbours can occur within other types of areas. Nodes were installed on the ceiling in a grid with a distance of 30-50 cm between nodes. 200 nodes were distributed over an area of 10m x 5m. It clearly transpired that with increasing distance the probability of reception decreases. At the same time a few nodes furthest away from the sender had a high probability of message reception, while some close neighbours of the sender did not receive messages. The patterns of clear reception nodes evolved over time.

The conclusion is that even for direct neighbours reception can temporarily be bad during periods of several minutes. For a reliable and timely communication it is imperative to have at least two communication paths available (e.g. two hop paths next to the 1-hop path for direct neighbours).

Authors' Addresses

Anders Brandt
Sigma Designs

Email: anders_Brandt@sigmadesigns.com

Emmanuel Baccelli
INRIA

Email: Emmanuel.Baccelli@inria.fr

Robert Cragie
ARM Ltd.
110 Fulbourn Road
Cambridge CB1 9NJ
UK

Email: robert.cragie@gridmerge.com

Peter van der Stok
Consultant

Email: consultancy@vanderstok.org

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 4, 2016

M. Richardson
SSW
May 3, 2016

ROLL Applicability Statement Template
draft-ietf-roll-applicability-template-09

Abstract

This document is a template applicability statement for the Routing over Low-power and Lossy Networks (ROLL) WG. This document is not for publication, but rather is to be used as a template.

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 November 4, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Relationship to other documents	3
1.2. Requirements Language	3
1.3. Terminology	4
1.4. Required Reading	4
1.5. Out of scope requirements	4
2. Deployment Scenario	4
2.1. Network Topologies	4
2.2. Traffic Characteristics	4
2.2.1. General	4
2.2.2. Source-sink (SS) communication paradigm	4
2.2.3. Publish-subscribe (PS, or pub/sub) communication paradigm	4
2.2.4. Peer-to-peer (P2P) communication paradigm	5
2.2.5. Peer-to-multipeer (P2MP) communication paradigm	5
2.2.6. Additional considerations: Duocast and N-cast	5
2.2.7. RPL applicability per communication paradigm	5
2.3. Layer-2 applicability.	5
3. Using RPL to Meet Functional Requirements	5
4. RPL Profile	5
4.1. RPL Features	5
4.1.1. RPL Instances	5
4.1.2. Storing vs. Non-Storing Mode	5
4.1.3. DAO Policy	5
4.1.4. Path Metrics	5
4.1.5. Objective Function	5
4.1.6. DODAG Repair	6
4.1.7. Multicast	6
4.1.8. Security	6
4.1.9. P2P communications	6
4.1.10. IPv6 address configuration	6
4.2. Layer-2 features	6
4.2.1. Specifics about layer-2	6
4.2.2. Services provided at layer-2	6
4.2.3. 6LowPAN options assumed.	6
4.2.4. MLE and other things	6
4.3. Recommended Configuration Defaults and Ranges	6
4.3.1. Trickle Parameters	6
4.3.2. Other Parameters	6
5. MPL Profile	7
5.1. Recommended Configuration Defaults and Ranges	8
5.1.1. Trickle Parameters	8
5.1.2. Other Parameters	8
6. Manageability Considerations	8
7. Security Considerations	8
7.1. Security Considerations during initial deployment	8

7.2. Security Considerations during incremental deployment . .	8
7.3. Security Considerations for P2P uses	8
8. Other Related Protocols	9
9. IANA Considerations	9
10. Acknowledgements	9
11. References	9
11.1. Normative References	9
11.2. Informative References	9
Author's Address	10

1. Introduction

This document describes a series of questions which should be answered. This document is intended to remain as a Internet Draft.

The idea is that current and future Applicability statements will use the table of contents provided. The goal is that all applicability statements will have to cover the listed items as a minimum.

1.1. Relationship to other documents

EDITORIAL: The following should appear in all applicability statements:

ROLL has specified a set of routing protocols for Lossy and Low-resource Networks (LLN) [RFC6550]. This applicability text describes a subset of these protocols and the conditions which make the subset the correct choice. The text recommends and motivates the accompanying parameter value ranges. Multiple applicability domains are recognized including: Building and Home, and Advanced Metering Infrastructure. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most probable network structures. Each applicability statement identifies the distinguishing properties according to a common set of subjects described in as many sections.

A common set of security threats are described in [RFC7416]. The applicability statements complement the security threats document by describing preferred security settings and solutions within the applicability statement conditions. This applicability statements may recommend more light weight security solutions and specify the conditions under which these solutions are appropriate.

1.2. Requirements Language

(RFC2119 reference)

1.3. Terminology

A reference to draft-ietf-roll-terminology is appropriate. A reference to layer-2 specific terminology and/or inclusion of any terms that are normatively referenced is appropriate here.

1.4. Required Reading

References/Overview of requirements documents, both IETF and industry group. (two pages maximum. This text should be (very) technical, should be aimed at IETF *participants*, not industry group participants, and should explain this industries' specific issues)

1.5. Out of scope requirements

This should list other documents (if any) which deal with situations where things are not in scope for this document.

(For instance, the AMI document tries to cover both line-powered urban metering networks, and energy-constrained metering networks, and also tries to deal with rural requirements. This should be three or four documents, so this section should list the limits of what this document covers)

2. Deployment Scenario

2.1. Network Topologies

describe a single scenario, with possibly multiple topologies that a single utility would employ.

2.2. Traffic Characteristics

Explain what kind of traffic is being transmitted, where it is initiated, and what kinds of protocols (CoAP, multicast, HTTPS, etc.) are being used. Explain what assumptions are being made about authentication and authorization in those protocols.

2.2.1. General

2.2.2. Source-sink (SS) communication paradigm

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

- 2.2.4. Peer-to-peer (P2P) communication paradigm
- 2.2.5. Peer-to-multipeer (P2MP) communication paradigm
- 2.2.6. Additional considerations: Duocast and N-cast
- 2.2.7. RPL applicability per communication paradigm
- 2.3. Layer-2 applicability.

Explain what layer-2 technologies this statement applies to, and if there are options, they should be listed generally here, and specifically in section 4.2.

3. Using RPL to Meet Functional Requirements

This should explain in general terms how RPL is going to be used in this network topology. If trees that are multiple layers deep are expected, then this should be described so that the fan out is understood. Some sample topologies (from simulations) should be explained, perhaps with image references from other publications.

This section should tell an *implementer* in a lab, having a simulation tool or a building/city/etc. to use as a testbed, how to construct an LLN of sufficient complexity (but not too much) to validate an implementation.

4. RPL Profile

This section should list the various features of RPL plus other layers of the LLN, and how they will be used.

4.1. RPL Features

- 4.1.1. RPL Instances
- 4.1.2. Storing vs. Non-Storing Mode
- 4.1.3. DAO Policy
- 4.1.4. Path Metrics
- 4.1.5. Objective Function

4.1.6. DODAG Repair

4.1.7. Multicast

4.1.8. Security

4.1.9. P2P communications

4.1.10. IPv6 address configuration

4.2. Layer-2 features

4.2.1. Specifics about layer-2

this section should detail the specific layer-2 network technology that this document applies to. A class of technologies is generally not acceptable.

4.2.2. Services provided at layer-2

4.2.3. 6LowPAN options assumed.

4.2.4. MLE and other things

4.3. Recommended Configuration Defaults and Ranges

4.3.1. Trickle Parameters

This section is intended to document the specific value (or ranges) appropriate for this kind of deployment. This includes trickle specific parameters such as those of RFC6550, section 8.3.1: Imin (DIOIntervalMin), Imax (DIOIntervalDoublings), and k (DIORedundancyConstant). While it is not necessary to hard code these parameters into RPL nodes, as they are announced as part of the DIO message, it is important for researchers who are trying to validate the convergence properties of the resulting deployment to understand what values have been selected.

4.3.2. Other Parameters

There are additional values which are present in the DODAG Configuration option. The purpose of this section is to: a) document what values are configured, b) if a default value is used, if it is appropriate for this deployment.

These values include: MaxRankIncrease, MinHopRankIncrease, the Objective Code Point to use, Default Lifetime, Lifetime Units...

In addition, the kinds of metrics which will be used (RFC6551) needs to be specified. If Objective Function 0 (RFC6552) is used, then it specifies a number of values, but also needs definitions of the `stretch_of_rank`, and `rank_factor`.

If MRHOF (RFC6719) is used, then section 5 of this document requires selection of: `MAX_LINK_METRIC`, `MAX_PATH_COST`, `PARENT_SWITCH_THRESHOLD`, `PARENT_SET_SIZE`, and `ALLOW_FLOATING_ROOT`.

5. MPL Profile

This section should list the various features of MPL. In considering the parameters, a number of questions come up:

- 1) What are the maximum and minimum 1-hop MPL router neighbours of all the MPL routers?
- 2) what is the arrival rate of new packets that need repetition in a MPL router
- 3) Is there a deadline associated with the packets
- 4) What is the shortest number of hops of the longest path between sources and destinations
- 5) What are the values of the MAC: back-off values, retries, buffer size.
- 6) What is the background load of other non MPL applications.
- 7) arrival probability of 1-hop packets

As the corresponding design space is incredibly large, probably only a limited subset of the design space is viable.

Here is an example scenario:

- o 5 neighbours
- o once every 100 ms (rate at sources is once every 300-500 ms)
- o yes, 200 ms
- o 5 hops, with mostly 1 hop
- o no buffer, retry 1, back-off 2
- o absent

- o 100-80%

leading to $k=3-5$, $I_{min}=30-70$ ms, repeat = 2, I_{max} n/a.

It is critical operational boundary conditions together with appropriate MPL parameter values are published in this applicability statements. All applicability statements together may give a good hint which MPL parameters and boundary conditions to choose.

5.1. Recommended Configuration Defaults and Ranges

5.1.1. Trickle Parameters

5.1.1.1. I_{min}

5.1.1.2. I_{max}

5.1.2. Other Parameters

5.1.2.1. Hot Limit

6. Manageability Considerations

7. Security Considerations

7.1. Security Considerations during initial deployment

(This section explains how nodes get their initial trust anchors, initial network keys. It explains if this happens at the factory, in a deployment truck, if it is done in the field, perhaps like <http://www.lix.polytechnique.fr/hipercom/SmartObjectSecurity/papers/CullenJennings.pdf>)

7.2. Security Considerations during incremental deployment

(This section explains how that replaces a failed node takes on the dead nodes' identity, or not. How are nodes retired. How are nodes removed if they are compromised)

7.3. Security Considerations for P2P uses

(When layer-3 RPL security is used, P2P DODAGs are ephemeral, and may have different security needs.)

8. Other Related Protocols

9. IANA Considerations

10. Acknowledgements

This document was created from a number source applicatbility templates, including draft-ietf-roll-applicability-ami-06.txt, draft-phinney-rpl-industrial-applicability-00.txt.

The document has benefitted from advance review by the IETF Security Directorate.

A number of edits were contributed from Peter van der Stok, including the MPL considerations/calculations

11. References

11.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [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, <<http://www.rfc-editor.org/info/rfc7416>>.

11.2. Informative References

- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<http://www.rfc-editor.org/info/rfc6206>>.
- [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, <<http://www.rfc-editor.org/info/rfc6550>>.

Author's Address

Michael C. Richardson
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z 5V7
CA

Email: mcr+ietf@sandelman.ca
URI: <http://www.sandelman.ca/>

ROLL
-Draft
Informational

T. Phinney, Ed. Internet
consultant Intended status:
P. Thubert Expires: April 22, 2014
cisco

RA. Assimiti
Nivis

October 21, 2013

RPL applicability in industrial networks

draft-ietf-roll-rpl-industrial-applicability-02 Abstract The wide deployment of wireless devices, with their low installed cost (compared to wired devices), will significantly improve the productivity and safety of industrial plants. It will simultaneously increase the efficiency and safety of the plant's workers, by extending and making more timely the information set available about plant operations. The new Routing Protocol for Low Power and Lossy

Networks (RPL) defines a Distance Vector protocol that is designed for such networks. The aim of this document is to analyze the applicability of that routing protocol in industrial LLNs formed of field devices. Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>. Internet-Drafts are draft documents

valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress." This Internet-Draft will expire on April 22, 2014. Copyright Notice Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved. Phinney, Thubert & Assimi Expires April 22, 2014 [Page 1]

Internet-Draft	RPL-industrial-applicability-statement	October 2013	This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.
Table of Contents			
1. Introduction	3	1.	
1. Requirements Language	4	1.2. Required Reading	4
1.3. Out of scope requirements	4	2. Deployment Scenario	
2.1. Network Topologies		2.1.1. Traffic Characteristics	
2.1.2. Topologies	6	2.1.3. Source-sink (SS) communication paradigm	
2.1.4. Publish-subscribe (PS, or pub/sub) communication paradigm	8	2.1.5. Peer-to-peer (P2P) communication paradigm	11
2.1.6. Peer-to-multipeer (P2MP) communication paradigm	14	2.1.7. Additional considerations: Duocast and N-cast	14
2.1.8. RPL applicability per communication paradigm	16	2.2. Layer 2 applicability	18
3. Using RPL to Meet Functional Requirements	18	4. RPL Profile	
4.1. RPL Features	20	4.1.1. RPL Instances	
4.1.2. Storing vs. Non-Storing Mode	20	4.1.3. DAO Policy	22
4.1.4. Path Metrics	23	4.1.5. Objective Function	24
4.1.6. DODAG Repair	24	4.1.7. MPL Profile	25
4.1.8. Security	25	4.1.9. P2P communications	25
4.2. Layer-two features	26	4.3. Recommended Configuration Defaults and Ranges	
4.3.1. Trickle Parameters	26	4.3.2. Other Parameters	27
5. Manageability Considerations	27	6. Security Considerations	28
6.1. Security Considerations during initial deployment	28	6.2. Security Considerations during incremental deployment	28
7. Other Related Protocols	28	8. IANA Considerations	28
9. Acknowledgements	28	10. References	28
10.1. Normative References	28	10.2. Informative References	28
10.3. External Informative References	30		

Information Technology (IT) is already, and increasingly will be applied to Industrial Automation and Control System (IACS) technology in application areas where those IT technologies can be constrained sufficiently by Service Level Agreements (SLA) or other modest change that they are able to meet the operational needs of IACS. When that happens, the IACS benefits from the large intellectual, experiential and training investment that has already occurred in those IT precursors. One can conclude that future reuse of additional IT protocols for IACS will continue to occur due to the significant intellectual, experiential and training economies which result from that reuse. Following that logic, many vendors are already extending or replacing their local field-bus technology with Ethernet and IP-based solutions. Examples of this evolution include CIP EtherNet/IP, Modbus/TCP, Foundation Fieldbus HSE, PROFINET and Invensys/Foxboro FOXnet. At the same time, wireless, low power field devices are being introduced that facilitate a significant increase in the amount of information which industrial users can collect and the number of control points that can be remotely managed. IPv6 appears as a core technology at the conjunction of both trends, as illustrated by the current [ISA100.11a] industrial Wireless Sensor Networking (WSN) specification, where layers 1-4 technologies developed for end uses other than IACS - IEEE 802.15.4 PHY and MAC, 6LoWPAN and IPv6, and UDP - are adapted to IACS use. But due to the lack of open standards for routing in Low power and Lossy Networks (LLN) at the time ISA100.11a was crafted, routing was accomplished at the link layer and is specific to that standard. The IETF ROLL Working Group has defined application-specific routing requirements for a LLN routing protocol, specified in: Routing Requirements for Urban LLNs [RFC5548], Industrial Routing Requirements in LLNs [RFC5673], Home Automation Routing Requirements in LLNs [RFC5826], and Building Automation Routing Requirements in LLNs [RFC5867]. The Routing Protocol for Low Power and Lossy Networks (RPL) [RFC6550] specification and its point-to-point extension/optimization [RFC6997] define a generic Distance Vector protocol that is adapted to a variety of Low Power and Lossy Networks (LLN) types by the application of specific Objective Functions (OFs). RPL forms Destination Oriented Directed Acyclic Graphs (DODAGs) within instances of the protocol, each instance being associated with an Objective Function to form a routing topology. Phinney, Thubert & Assimi Expires April 22, 2014 [Page

Internet-Draft RPL-industrial-applicability-statement October 2013 A field device that belongs to an instance uses the OF to determine which DODAG and which Version of that DODAG the device should join. The device also uses the OF to select a number of routers within the DODAG current and subsequent Versions to serve as parents or as feasible successors. A new Version of the DODAG is periodically reconstructed to enable a global reoptimization of the graph.

A RPL OF states the outcome of the process used by a RPL node to select and optimize routes within a RPL Instance based on the information objects available. The separation of OFs from the core protocol specification allows RPL to be adapted to meet the different optimization criteria required by the wide range of industrial classes of traffic and applications. This document provides information on how RPL can accommodate the industrial requirements for LLNs, in particular as specified in [RFC5673].1.1. Requirements Language The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119]. Additionally, this document uses terminology from [I-D.ietf-roll-terminology], and uses usual terminology from the Process Control and Factory Automation industries, some of which is recapitulated below: FEC: Forward error correction IACS: Industrial automation and control systems RAND: reasonable and non-discriminatory (relative to licensing of patents)1.2. Required Reading1.3. Out of scope requirements This applicability statement does not address requirements related to wireless LLNs employed in factory automation and related applications.2. Deployment Scenario [RFC5673] describes in detail the routing requirements for industrial LLNs. This RFC provides information on the varying deployment scenarios for such LLNs and how RPL assists in meeting thosePhinney, Thubert & AssimiExpires April 22, 2014 [Page 4]

Internet-Draft RPL-industrial-applicability-statement October 2013 requirements. Large industrial plants, or major operating areas within such plants, repeatedly go through four major phases, each of which typically lasts from months to years: P1: Construction or major modification phase P2: Planned startup phase P3: Normal operation phase P4: Planned shutdown phase followed eventually by an (at least theoretical) P5: Plant decommissioning phase. It is also likely, after a major catastrophe at a plant, to have a P6: Post-emergency recovery and repair phase. The deployment scenarios for wireless LLN devices may be different in each of these phases. In particular, during the Construction or major modification phase (P1), LLN devices may be installed months before the intended LLN can become usefully operational (because needed routers and infrastructure devices are not yet installed or active), and there are likely to be many personnel in whom the plant owner/operator has only limited trust, such as subcontractors and others in the plant area who have undergone only a cursory background investigation (if any at all). In general, during this phase, plant instrumentation is not yet operational, so could be removed and replaced by a Trojaned device without much likelihood of physical detection of the substitution. Thus physical security of LLN devices is generally a more significant risk factor during this phase than once the plant is operational, where simple replacement of device electronics is detectable. Extra LLN devices and even extra LLN subnets may be employed during Planned startup (P2) and Planned shutdown (P4) phases, in support of the task of transitioning the plant or plant area between operational and shutdown states. The extra devices typically provide extra monitoring as the plant transitions infrequent activity states. (In many continuous process plants, up to 2x extra staff are employed at monitoring and control workstations during these two phases, precisely because the plant is undergoing extraordinary behavior as it transitions to or from its steady-state operational condition.) Similar transient devices and subnets may be used during an unscheduled Post-emergency recovery and repair phase (P6) of operation, but in that case the extra devices usually are routers substituting for plant LLN devices that have been damaged by the incident (such as a fire, explosion, flood, tornado or hurricane)Phinney, Thubert & AssimiExpires April 22, 2014 [Page 5]

Internet-Draft RPL-industrial-applicability-statement October 2013 that induced the emergency. The Planned startup (P2) and Planned shutdown (P4) phases are similar in many respects, but the LLN environment of the two can be quite different, since the Planned shutdown phase can assume that the stable LLN environment used for Normal operation (P3) is functional during shutdown, whereas that stable environment usually is still being established during startup. The Post-emergency recovery and repair phase (P6) typically operates in an LLN environment that is somewhere between that of the Planned startup (P2) and Normal operation (P3) phases, but with an indeterminate number of temporary routers placed to facilitate communication across and around the area affected by the catastrophe. Smaller industrial plants and sites may go through similar phases, but often commingle the phases because, in those smaller plants, the phases require less planning and structuring of personnel responsibilities and thus permit less formalization and partitioning of the operating scenarios. For example, it is much simpler, and usually requires much less planning, to bring new equipment on a skid into a plant, using a forklift, than to lay temporary railroad track or employ an extended-axle heavy haul tractor-trailer to deliver a multi-ton process vessel, and temporarily deploy and use very large heavy-lift cranes to install it. In the former cases, nearby equipment usually can continue normal operation while the installation proceeds; in the latter case that is almost always impossible, due to safety and other concerns.

The domain of applicability for the RPL protocol may include all phases but the Normal Operation phase, where the bandwidth allocation and the routes are usually optimized by an external Path Computing Engine (PCE), e.g. an ISA100.11a System Manager. Additionally, it could be envisioned to include RPL in the normal operation provided that a new Objective Function is defined that actually interacts with the PCE in order to establish the reference topology, in which case RPL operations would only apply to emergency repair actions. When the reference topology becomes unusable for some failure, and as long as the problem persists.

2.1. Network Topologies

2.1.1. Traffic Characteristics

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

- o Safety
- * Class 0: Emergency action - Always a critical function
- o Control

Phinney, Thubert & Assimi Expires April 22, 2014

Internet-Draft RPL-industrial-applicability-statement October 2013 *

Class 1: Closed loop regulatory control - Often a critical function

* Class 2: Closed loop supervisory control - Usually non-critical function

ction * Class 3: Open loop control - Operator takes action and controls

the actuator (human in the loop) o Monitoring * Class 4: Alerting

- Short-term operational effect (for example event-based maintenance)


* Class 5: Logging and downloading / uploading - No immediate operational consequence (e.g., history collection, sequence-of-events, preventive maintenance)

Safety critical functions effect the basic safety integrity of the plant. These normally dormant functions kick in only when process control systems, or their operators, have failed. By design and by regular interval inspection, they have a well-understood probability of failure on demand in the range of typically once per 10-1000 years. In-time deliveries of messages becomes more relevant as the class number decreases. Note that for a control application, the jitter is just as important as latency and has a potential of destabilizing control algorithms. The domain of applicability for the RPL protocol probably matches the range of classes where industrial users are interested in deploying wireless networks. This domain includes monitoring classes (4 and 5), and the non-critical portions of control classes (2 and 3). RPL might also be considered as an additional repair mechanism in all situations, and independently of the flow classification and the medium type. It appears from the above sections that whether and the way RPL can


Phinney, Thubert & Assimi Expires April 22, 2014 [Page 7]

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

-----+ ? : typically usable for all but higher-rate classes 0,1 PS traffic 2.1.2. Topologies In an IACS, high-rate communications flows (e.g., 1 Hz or 4 Hz for a traditional process automation network) typically are such that only a single wireless LLN hop separates the source device from a LLN Border Router (LBR) to a significantly higher data-rate backbone network, typically based on IEEE 802.3, IEEE 802.11, or IEEE 802.16, as illustrated in Figure 2. Phinney, Thubert & Assimi Expires April 22, 2014



o : stationary wireless field device, seldom acting as an LLN router F
 or factory automation networks, the basic communications cycle for control is
 typically much faster, on the order of 100 Hz or more. In this case the LLN i
 tself may be based on high-data-rate IEEE 802.11 or a 100 Mbit/s or faster opt
 ical link, and the higher-rate network used by the LBRs to connect the LLN to
 superior automation equipment typically might be based on fiber-optic IEEE 802
 .3, with multiple LBRs around the periphery of the factory area, so that most
 high-rate communications again requires only a single wireless LLN hop. Mult
 i-hop LLN routing is used within the LLN portion of such networks to provide b
 ackup communications paths when primary single-hop LLN paths fail, or for lowe
 r repetition rate communications where longer LLN transit times and higher var
 iance are not an issue. Typically, the majority of devices in an IACS can tol
 erate such higher-delay higher-variance paths, so routing choices often are dr
 iven by energy considerations for the affected devices, rather than simply by
 IACS performance requirements, as illustrated in Figure 3. Phinney, Thubert &
 Assimi Expires April 22, 2014 [Page 9]



LLN : stationary wireless field device, often acting as an LLN router M : mobile wireless device Two decades of experience with digital fieldbuses has shown that four communications paradigms dominate in IACS:

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

2.1.3. Source-sink (SS) communication paradigm In SS, the source-sink communication paradigm, each of many devices in one set, S1, sends UDP-like messages, usually infrequently and intermittently, to a second set of devices, S2, determined by a common multicast address. A typical example would be that all devices within a given process unit N are configured to send process alarm messages to the multicast address Receivers_of_process_alarms_for_unit_N. Receiving devices, typically on non-LLN networks accessed via LBRs, are configured to receive such

Phinney, Thubert & Assimi Expires April 22, 2014

Internet-Draft RPL-industrial-applicability-statement October 2013 multi
cast messages if their work assignment covers process unit N, and not otherwis
e. Timeliness of message delivery is a significant aspect of some SS communi
cation. When the SS traffic conveys process alarms or device alerts, there is
often a contractual requirement, and sometimes even a regulatory requirement,
on the maximum end-to-end transit delay of the SS message, including both the
LLN and non-LLN components of that delay. However, there is no requirement o
n relative jitter in the delivery of multiple SS messages from the same source
, and message reordering during transit is irrelevant. Within the LLN, the S
S paradigm simply requires that messages so addressed be forwarded to the resp
onsible LBR (or set of equivalent LBRs) for further forwarding outside the LLN
. Within the LLN such traffic typically is device-to-LBR or device-to-redundan
t-set-of- equivalent-LBRs. In general, SS traffic may be aggregated before
forwarding when both the multicast destination address and other QoS attribute
s are identical. If information on the target delivery times for SS messages
is available to the aggregating forwarding device, that device may intentiona
ly delay forwarding somewhat to facilitate further aggregation, which can sign
ificantly reduce LLN alarm-reporting traffic during major plant upset events.2
.1.4. Publish-subscribe (PS, or pub/sub) communication paradigm In PS, the pu
blish-subscribe communication paradigm, a device sends UDP-like messages, usua
lly periodically or cyclicly (i.e., repetitively but without fixed periodicity
), to a single multicast address derived from or correlated with the device's
own address. A typical example would be that each sensor and actuator device
within a given process unit N is configured to send process state messages t
o the multicast address that designates its specific publications. In essence
the derived multicast address for device D is Receivers_of_publications_by_dev
ice_D. Typically those receivers are in two categories: controllers (C) for co
ntrol loops in which device D participates, and devices accessed via the LLN's
LBRs that monitor and/or accumulate historical information about device D's s
tatus and outputs. If the controller(s) that receive device D's publication
are all outside the LLN and accessed by LBRs, then within the LLN such traff
ic typically is device-to-LBR or device-to-redundant-set-of- equivalent-LBRs.
But if a controller (Cn) is within the LLN, then a number of different LLN-lo
cal traffic patterns may be employed, depending on the capabilities of the und
erlying link technology and on configured performance requirements for such re
porting. Typically in such a case, publication by device D is forwarded up a
DODAG to an LLN router that is also on a downward DODAG to a destination con
troller Cn, then forwarded down that second DODAG to thatPhinney, Thubert & Assi
miExpires April 22, 2014 [Page 11]

Internet-Draft RPL-industrial-applicability-statement October 2013 destination controller Cn. Of course, if the LLN router (or even the LBR) is itself the intended destination controller, which will often be the case, then no downward forwarding occurs. Timeliness of message delivery is a critical aspect of PS communication. Individual messages can be lost without significant impact on the controlled physical process, but typically a sequence of four consecutive lost messages will trigger fallback behavior of the control algorithms, which is considered a system failure by most system owner/operators. (In general, and unless a local catastrophic event such as a major explosion or a tornado occurs in the plant, invocation of more than one instance of such fallback handling per year, per plant, is considered unacceptable.) Message loss, delay and jitter in delivery of PS messaging is a relative matter. PS messaging is used for transfer of process measurements and associated status from sensors to control computation elements, from control computation elements to actuators, and of current commanded position and status from actuators back to control computation elements. The actual time interval of interest is that which starts with sensing of the physical process (which necessarily occurs before the sensed value can be sent in the first message) and which ends when the computed control correction is applied to the physical process by the appropriate actuator (which cannot occur until after the second message containing the computed control output has been received by that actuator). With rare exception, the control algorithms used with PS messaging in the process automation industries - those managing continuous material flows - rely on fixed-period sampling, computation and transfer of outputs, while those in the factory automation industries - those managing discrete manufacturing operations - rely on bounded delay between sampling of inputs, control computation and transfer of outputs to physical actuators that affect the controlled process. Deliberately manipulated message delay and jitter in delivery of PS messaging has the potential to destabilize control loops. It is the responsibility of conveyed higher-level protocols to protect against such potential security attacks by detecting overly delayed or jittered messages at delivery, converting them into instances of message loss. Thus network and data-link protocols such as IPv6 and Ethernet need not themselves address such issues, although their selection and employment should take the existence (or lack) of such higher-layer protection mechanisms, and the resulting consequences due to excessive delay and jitter, into consideration in their parameterization. Phinney, Thubert & Assimi

Expires April 22, 2014 [Page 12]

Internet-Draft RPL-industrial-applicability-statement October 2013

In general, PS traffic within the LLN is not aggregated before forwarding, to minimize message loss and delay in reception by any relevant controller(s) that are outside the LLN. However, if all intended destination controllers are within the LLN, and at least one of those intended controllers also serves as an LLN router on a DODAG to off-LLN destinations that all are not controllers, then the router functions in that device may aggregate PS traffic before forwarding when the required routing and other QoS attributes are identical. If information on the target delivery times for PS messages to non-controller devices is available to the aggregating forwarding device, that device may intentionally delay forwarding somewhat to facilitate further aggregation. In some system architectures, message streams that use PS to convey current process measurements and status are compressed at the source through a 2-dimensional winnowing process that compares 1) the process measurement values and status of the about-to-be-sent message with that of the last actually-sent message, and 2) the current time vs. the queueing time for the last actually-sent message. If the interval since that last-sent message is less than a predefined maximum time, and the status is unchanged, and the process measurement(s) conveyed in the message is within predefined deadband(s) of the last-sent measurement value(s), then transmission of the new message is suppressed. Often this suppression takes the form of not queuing the new message for transmission, but in some protocols a brief placeholder message indicating "no significant change" is queued in its stead.

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

In P2P, the peer-to-peer communication paradigm, a device sends UDP-like or TCP-like messages from one device (D1) to a second device (D2), usually with bidirectional but asymmetric flow of application data, where the amount of data is significantly greater in one direction than the other. Typical examples are transfer of configuration information to or from a process field device, or transfer of captured process diagnostics (e.g., time-stamped noise signatures from a coriolis flowmeter) to an off-LLN higher-level asset management system. Unicast addressing is used in both directions of data flow. In general, specific P2P traffic has only loose timeliness requirements, typically just those required so that response times to human-operator-initiated actions meet human factors requirements. As a consequence, in general, message aggregation is permitted, although few opportunities are likely to present themselves for such aggregation due to the sporadic nature of such messaging to a single destination, and/or due to the large message payloads that often occur in at least one direction of transmission.

Phinney, Thubert & Assimi Expires April 22, 2014

Internet-Draft RPL-industrial-applicability-statement October 2013 2.1.6. Peer-to-multipoint (P2MP) communication paradigm In P2MP, the peer-to-multipoint communication paradigm, a device sends UDP-like messages downward, from one device (D1) to a set of other devices (Dn). Typical examples are bulk downloads to a set of devices that use identical code image segments or identically-structured database segments; group commands to enable device state transitions that are quasi-synchronized across all or part of the local network (e.g., switch to the next set of point-to-point downloaded session keys, or notifying that the network is switching to an emergency repair and recovery mode); etc. Multicast addressing is used in the downward direction of data flow. Devices can be assigned to a number of multicast groups, for instance by device type. Then, if it becomes necessary to reflash all devices of a given type with a new load image, a multicast distribution mechanism can be leveraged to optimize the distribution operation. In general, P2MP traffic has only loose timeline requirements. As a consequence, in general, message aggregation is permitted, although few opportunities are likely to present themselves for such aggregation due to the sporadic nature of such messaging to a single multicast group destination, and/or due to the large message payloads that often occur when P2MP is used for group downloads. However, in general, message aggregation negatively impacts the delivery success rate for each of the aggregated messages, since the probability of error in a received message increases with message length. Together these considerations often lead to a policy of non-aggregation for P2MP messaging. Note: Reliable group download protocols, such as the no-longer-published IEEE 802.1E (ISO/IEC 15802-4) system load protocol, and reliable multicast protocols based on the guidance of [RFC2887], are instructive in how P2MP can be used for initial bulk download, followed by either P2MP or P2P selective retransmissions for missed download segments. 2.1.7. Additional considerations: Duocast and N-cast In industrial automation systems, some traffic is from (relatively) high-rate monitoring and control loops, of Class 0 and Class 1 as described in [RFC5673]. In such systems, the wireless link protocol, which typically uses immediate in-band acknowledgement to confirm delivery (or, on failure, conclude that a retransmission is required), can be adapted to attempt simultaneous delivery to more than one receiving device, with separated, sequenced immediate in-band acknowledgement by each of those intended receivers. (This mechanism is known colloquially as "duocast" (for two intended receivers), or more generically as "N-cast" (for N intended receivers).) Transmission is deemed successful if at least one such immediate acknowledgement is received by the sending device; Phinney, Thubert & Assimi Expires April 22, 2014 [Page 14]

Internet-Draft RPL-industrial-applicability-statement October 2013

otherwise the device queues the message for retransmission, up until the maximum configured number of retries has been attempted. The logic behind duocast/N-cast is very simple: In wireless systems without FEC (forward error correction), the overall rate of success for transactions consisting of an initial transmission and an immediate acknowledgement is typically 95%. In other words, 5% of such transactions fail, either because the initial message of the transaction is not received correctly by the intended receiver, or because the immediate acknowledgment by that receiver is not received correctly by the transaction initiator. In the generalized case of N-cast, where any received acknowledgment serves to complete the transaction, and where the N intended receivers are spatially diverse, physically separated from each other by multiple wavelengths, the probability that all such receivers fail to receive the initial message of the transaction, or that all generated immediate acknowledgements are not received by the transaction initiator, is typically approximately $(5\%)^N$. Thus, for duocast, the expected success rate for a single transaction goes from 95% ($1.0 - 0.05$) to 99.75% ($1.0 - 0.05^2$), to 99.9875% ($1.0 - 0.05^3$) when $N=3$, and even higher when $N>3$. From the above analysis, it is obvious that the primary benefit of N-cast occurs when N goes from $N=1$ (unicast) to $N=2$ (duocast); the reduction in transaction loss rate for increasing $N>2$ is quite small, and for $N>3$ it is infinitesimal. In the typical industrial automation environment of class 1 process control loops, which typically repeat at a 1 Hz or 4 Hz rate, in a very large process plant with thousands of field devices reporting at that rate, the maximum number of transmission retries that must be planned, and for which capacity must be scheduled (within the requisite 250 ms or 1 s interval) is seven (7) retries for unicast PS reporting, but only three (3) retries with duocast PS reporting. (This is determined by the requirement to not miss four successive reports more than once per year, across the entire plant, as such a loss typically triggers fallback behavior in the controlled loop, which is considered a failure of the wireless system by the plant owner/operator.) In practice, the enormous reduction in both planned and used retransmission capacity provided by duocast/N-cast is what enables 4 Hz loops to be supported in large wireless systems. When available, duocast/N-cast typically is used only for one-hop PS traffic on Class 1 and Class 0 control loops. It may also be employed for rapid, reliable one-hop delivery of Class 0 and sometimes Class 1 process alarms and device alerts, which use the SS paradigm. Because it requires scheduling of multiple receivers that are prepared to acknowledge the received message during the transaction, in general it is not appropriate for the other types of traffic in such systems - P2P and P2MP - and is not needed for other classes of control loops or other types of traffic, which do not have such stringent reporting requirements.

Phinney, Thubert & Assimi
Expires April 22, 2014 [Page 15]

Internet-Draft RPL-industrial-applicability-statement October 2013 Note:

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

2.1.8. RPL applicability per communication paradigm To match the requirements above, RPL provides a number of RPL Modes of Operation (MOP):

- No downward route: defined in [RFC6550], section 6.3.1, MOP of 0. This mode allows only upward routing, that is from nodes (devices) that reside inside the RPL network toward the outside via the DODAG root. Non-storing mode: defined in [RFC6550], section 6.3.1, MOP of 1. This mode improves MOP 0 by adding the capability to use source routing from the root towards registered targets within the instance DODAG.
- Storing mode without multicast support: defined in [RFC6550], section 6.3.1, MOP of 2. This mode improves MOP 0 by adding the capability to use stateful routing from the root towards registered targets within the instance DODAG.
- Storing mode with link-scope multicast DAO: defined in [RFC6550] section 9.10, this mode improves MOP 2 by adding the capability to send Destination Advertisements to all nodes over a single Layer 2 link (e.g. a wireless hop) and enables line-of-sight direct communication.

Phinney, Thubert & Assimi Expires April 22, 2014

operation (MOP) of 3. This mode improves MOP 2 by adding the capability to register multicast groups and perform multicast forwarding along the instance DODAG (or a spanning subtree within the DODAG). Reactive: defined in [RFC6997], the reactive mode creates on-demand additional DAGs that are used to reach a given node acting as DODAG root within a certain number of hops. This mode can typically be used for an ad-hoc closed-loop communication. The RPL MOP that can be applied for a given flow depends on the communication paradigm. It must be noted that a DODAG that is used for PS traffic can also be used for SS traffic since the MOP 2 extends the MOP 0, and that a DODAG that is used for P2MP distribution can also be used for downward PS since the MOP 3 extends the MOP 2. On the other hand, an Objective Function (OF) that optimizes metrics for a pure upwards DODAG might differ from the OF that optimizes a mixed upward and downward DODAG. As a result, it can be expected that different RPL instances are installed with different OFs, different channel allocations, etc... that result in different routing and forwarding topologies, sometimes with differing delay vs. energy profiles, optimized separately for the different flows at hand. This can be broadly summarized in the following table:

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

Internet-Draft RPL-industrial-applicability-statement October 2013 2.2.
Layer 2 applicability. Work at the 6TiSCH WG details layer 2 operations for the most commonly used link Layer for industrial operations, the Timeslotted Channel Hopping (TSCH) mode of IEEE802.15.4e [IEEE802154e]. [I-D.wattheyne-6ti-sch-tsch] provides in-depth information on the IEEE802.15.4e [IEEE802154e] TSCH MAC operation whereas the 6TiSCH architecture [I-D.thubert-6tisch-architecture] provides additional information as of how RPL can be used over TSCH. This contrasts with the SmartGrid area where ZigBee IP [ZigBeeIP] ("ZigBee" is a registered trademark of the ZigBee Alliance) defines an application of RPL over a more classical contention-based operation but will not exhibit the deterministic capabilities that industrial control loops require.

3. Using RPL to Meet Functional Requirements The functional requirements for most industrial automation deployments are similar to those listed in [RFC5673]. The routing protocol MUST be capable of supporting the organization of a large number of nodes into regions, usually corresponding to partitions of the automated process, each containing on the order of 30 to 3000 nodes. The routing protocol MUST provide mechanisms to support configuration of the routing protocol itself. The routing protocol MUST provide mechanisms to support instructed configuration of explicit routing, so that in the absence of failure the routing used for selected flow classes is that which has been remotely configured (typically by a centralized configurator). In such circumstances RPL is used for local network repair; for flow classes to which explicit routing has not been assigned; during bootstrapping of the network itself (which is really just an instance of routing without such an externally-imposed assignment). The routing protocol SHOULD support directed flows with different QoS characteristics, typically with different energy vs. delay tradeoffs, for traffic directed to LBRs.

In practice only two such sets of QoS are relevant: Phinney, Thubert & Assimi Expires April 22, 2014 [Page 18]

one that emphasizes energy minimization for energy-constrained nodes at the expense of greater mean transit delay and variance in transit delay; and one that emphasizes minimization of mean transit delay and transit delay variance at the expense of greater energy demand on originating and intermediary energy-constrained nodes, typically used for critical SS traffic (e.e., infrequent and unpredictable safety alarms with legally-mandated maximum reporting delays) and critical PS traffic (e.g., predictable periodic (for process automation) or cyclic (for factory automation) high-speed safety control loops needed to protect life, the environment, and/or critical national infrastructure assets).

In the absence of configured routing, or when such routes have failed, the routing protocol MUST dynamically compute and select effective routes composed of low-power and lossy links. Local network dynamics SHOULD NOT impact the entire network. The routing protocol MUST compute multiple paths when possible. The routing protocol MUST support multicast addressing, including

multicast originating with a LBR or off the LLN, directed to a predefined group within the LLN multicast originating within the LLN, directed to one or more equivalent LBRs, in support of SS traffic

multicast originating within the LLN, directed to one or more equivalent LBRs, in support of PS traffic. The routing protocol SHOULD support and

utilize a large number of highly directed flows to a few LBRs, to handle scalability. The routing protocol SHOULD support formation of groups of field devices in the network. The routing protocol NEED NOT support anycast

addressing because, as of the date of writing of this document, such addressing is not used by automation and control field devices. In general, no

two such devices are equivalent, except perhaps for intermediary LBRs,

so unicast suffices for situations where anycast might otherwise be employed. Phinney, Thubert & Assimi Expires April 22, 2014 [Page 19]

Internet-Draft RPL-industrial-applicability-statement October 2013 RPL supports:

Large-scale networks characterized by highly directed traffic flows between each field device and servers close to the head-end of the automation network. To this end, RPL builds Directed Acyclic Graphs (DAGs) rooted at LBRs. Zero-touch configuration. This is done through in-band methods for configuring RPL variables using DIO messages. The use of links with time-varying availability and quality characteristics. This is accomplished by allowing the use of metrics that effectively capture the quality of a path (e.g., in terms of the mean and maximum impact of use of that path on packet delivery timing and on endpoint energy demands), and by limiting the impact of changing local conditions by discovering and maintaining multiple DAG parents, and by using local repair mechanisms when DAG links break. For wireless installations of small size with undemanding communication requirements, RPL is likely to generate satisfactory routing without any special effort. However, in larger installations or where timeliness considerations do not permit multi-second wireless-subnet transit times, then flow labeling is likely required so that forwarding routers can make informed trade offs between conserving their own energy resources and meeting overall system needs.

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

4.1. RPL Features

4.1.1. RPL Instances RPL allows formation of multiple instances that operate independently of each other. Each instance may use a different objective function and different modes of operation. It is highly recommended that wireless field devices participate in different instances that utilize objective functions that meet different optimization goals. These optimization goals target:

1. Minimizing and ensuring that a guaranteed latency is being met
2. Maximizing the communication reliability of the packets transferred over the wireless media

Phinney, Thubert & Assimi Expires April 22, 2014 [Page 20]

Internet-Draft RPL-industrial-applicability-statement October 2013 3. Minimizing aggregate power consumption for multi-hop LLNs that are composed of battery powered field devices. Some of these optimization goals will have to be met concurrently in a single instance by imposing various constraints.

Each wireless field device should participate in a set composed of a minimum of three instances that meet optimization goals associated with three traffic flows which need to be supported by all industrial LLNs. Management Instance: Wireless industrial networks are highly deterministic in nature, meaning that wireless field devices do not make any decisions locally but are managed by a centralized System Manager that oversees the join process as well as all communication and security settings present in the devices. The management traffic flow is downward traffic and needs to meet strictly enforced latency and reliability requirements in order to ensure proper operation of the wireless LLN. Hence each field device should participate in an instance dedicated to management traffic. All decisions made while constructing this instance will need to be approved by the Path Computation Engine present in the System Manager due to the deterministic, centralized nature of wireless industrial LLNs. Shallow LLNs with a hop count of up to one, accommodate this downward traffic using non-storing mode. Non-storing involves source routing that is detrimental to the packet size. For large transfers such as image download and configuration files, this can be factorized for a large packet. In that case, a method such as [I-D.thubert-6lo-forwarding-fragments] is required over multi-hop networks to forward and recover individual fragments without the overhead of the source route information in each fragment. If the hop count in the wireless LLN grows (LLN becomes deeper) it is highly recommended that the management instance rely on storing mode in order to relay management related packets. Operational Instance: The bulk of the data that is transferred over wireless LLN consists of process automation related payloads. This data is of paramount importance to the smooth operation of the process that is being monitored. Hence data reliability is of paramount importance. It is also important to note that a vast majority of the wireless field devices that operate in industrial LLNs are battery powered. The operational instance should hence ensure high reliability of the data transmitted while also minimizing the aggregate power consumption of the field devices operating in the LLN. All decisions made while constructing this instance will need to be approved by the Path Computation Engine present in the System Manager. This is due to the deterministic, centralized nature of wireless LLNs.

Phinney, Thubert & Asimi
Expires April 22, 2014

Internet-Draft RPL-industrial-applicability-statement October 2013 Autonomous instance: An autonomous instance requires limited to no configuration. Its primary purpose is to serve as a backup for the operational instance in case the operational instance fails. It is also useful in non-production phases of the network, when the plant is installed or dismantled. [I-D.thubert-roll-asymlink] provides rules and mechanisms whereby an instance can be used as a fallback to another upon failure to forward a packet further.

The autonomous instance should always be active and during normal operations it should be maintained through local repair mechanisms. In normal operation global repairs should be sparingly employed in order to conserve batteries. But a global repair is also probably the fastest and most economical technique in the case the network is extensively damaged. It is recommended to rely on automation that will trigger a global repair upon the detection of a large scale incident such as an explosion or a crash. As the name suggests, the autonomous instance is formed without any dependence on the System Manager. Decisions made during the construction of the autonomous instance do not need approval from the Path Computation Engine present in the System Manager. Participation of each wireless field device in at least one instance that hosts a DODAG with a virtual root is highly recommended. Wireless industrial networks are typically composed of multiple LLNs that terminate in a LLN Border Router (LBR). The LBRs communicate with each other and with other entities present on the backbone (such as the Gateway and the System Manager) over a wired or wireless backbone infrastructure. When a device A that operates in LLN 1 sends a packet to a device B that operates in LLN2, the packet egresses LLN1 through LBR1 and ingresses LLN2 through LBR2 after travelling over the backbone infrastructure that connects the LBRs. In order to accommodate this packet flow that travels from one LLN to another, it is highly recommended that wireless field devices participate in at least one instance that has a DODAG with a virtual root.

4.1.2. Storing vs. Non-Storing Mode

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

Internet-Draft RPL-industrial-applicability-statement October 2013

Deeper multi-hop wireless LLNs (hop count > 1) should support storing mode in order to minimize the overhead associated with source routing given the limited header capacity associated with typical physical layers employed in wireless LLNs. Support for storing mode requires additional RAM resources be present in the constrained wireless field devices. Typical wireless LLNs scale to a maximum of one hundred field devices. Hence the appropriate RAM resources for supporting storing mode should be part of the hardware requirements imposed upon wireless field devices during the design phase. The ISA100.11a standard mandates that all LBRs maintain routing tables with enough capacity to accommodate operation in storing mode. The standard also mandates that all wireless field devices maintain routing tables but it does not make any capacity assumptions, allowing for null routing tables. The System Manager should read the routing table capacity of each wireless field router and LBR during their join phase, and determine if support for storing mode in a particular LLN is feasible.

Lack of support for storing mode is also detrimental to battery operated wireless field devices due to the power consumption associated with transporting the hefty headers associated with source routing. Support for storing mode also ensures path redundancy which in turn allows for better prediction of the latency associated with downward traffic flows. Guaranteed latencies are of paramount importance for various traffic flows in wireless industrial LLNs.

4.1.3. DAO Policy Support for both upward and downward traffic flows is a requirement in industrial automation systems. As a result, nodes send DAO messages to establish downward paths from the root to themselves. DAO messages are not acknowledged in wireless industrial LLNs that are composed of battery operated field devices in order to minimize the power consumption overhead associated with path discovery. Given that wireless field devices in LLNs will typically participate in multiple RPL instances and DODAGs, it is highly recommended that both the RPLInstance ID and the DODAGID be included in the DAO.

4.1.4. Path Metrics RPL relies on an Objective Function for selecting parents and computing path costs and rank. This objective function is decoupled from the core RPL mechanisms and also from the metrics in use in the network. Two objective functions for RPL have been defined at the time of this writing, the RPL Objective Function 0 [RFC6552] and the Minimum Rank with Hysteresis Objective Function [RFC6719], both of

Phinney, Thubert & Assimi Expires April 22, 2014

Internet-Draft RPL-industrial-applicability-statement October 2013 which define a selection method for a preferred parent and backup parents, and are suitable for industrial automation network deployments.4.1.5. Objective Function Industrial wireless LLNs are subject to swift variations in terms of the propagation of the wireless signal, variations that can affect the quality of the links between field devices. This is due to the nature of the environment in which they operate which can be characterized as metal jungles that cause wireless propagation distortions, multi-path fading and scattering. Hence support for hysteresis is needed in order to ensure relative link stability which in turn ensures route stability. As mentioned in previous sections of this document, different traffic flows require different optimization goals. Wireless field devices should participate in multiple instances associated with multiple objective functions. Management Instance: Should utilize an objective function that focuses on optimization of latency and data reliability. Operational instance: Should utilize an objective function that focuses on data reliability and minimizing aggregate power consumption for battery operated field devices. Autonomous instance: Should utilize an objective function that optimizes data latency. The primary purpose of the autonomous instance is as a fallback instance in case the operational instance fails. Data latency is hence paramount for ensuring that the wireless field devices can exchange packets in order to repair the operational instance. More complex objective functions are needed that take in consideration multiple constraints and utilize weighted sums of multiple additive and multiplicative metrics. Additional objective functions specifically designed for such networks may be defined in companion RFCs.4.1.6. DODAG Repair To effectively handle time-varying link characteristics and availability, industrial automation network deployments SHOULD utilize the local repair mechanisms in RPL. Local repair is triggered by broken link detection, and in storing mode also by loop detection. Phinney, Thubert & Assimi Expires April 22, 2014 [Page 24]

Internet-Draft RPL-industrial-applicability-statement October 2013

The first local repair mechanism consists of a node detaching from a DODAG and then re-attaching to the same or to a different DODAG at a later time. While detached, a node advertises an infinite rank value so that its children can select a different parent. This process is known as poisoning and is described in Section 8.2.2.5 of [RFC6550]. While RPL provides an option to form a local DODAG, doing so in industrial automation network deployments is of little benefit since applications typically communicate through a LBR. After the detached node has made sufficient effort to send notification to its children that it is detached, the node can rejoin the same DODAG with a higher rank value. The configured duration of the poisoning mechanism needs to take into account the disconnection time applications running over the network can tolerate. Note that when joining a different DODAG, the node need not perform poisoning. The second local repair mechanism controls how much a node can increase its rank within a given DODAG Version (e.g., after detaching from the DODAG as a result of broken link or loop detection). Setting the DAGMaxRankIncrease to a non-zero value enables this mechanism, and setting it to a value of less than infinity limits the cost of count-to-infinity scenarios when they occur, thus controlling the duration of disconnection applications may experience.

4.1.7. MPL Profile The applicability of MPL is left to be determined. There is a potential for Source/Sink flows in order to control the flooding incurred by alarms and alerts.

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

4.1.9. P2P communications There is definitely a need for route optimizations for the close control loops that sustain the automation systems. [I-D.thubert-6tisch-architecture] discusses the applicability of a central routing computation based on a Path Computation Element (PCE), which would be the natural IETF correspondent to the System Managers or Network Managers that can be found in existing industrial standards. The RPL point to point extension/optimization [RFC6997] (experimental) or its standard track successor may be used as well to establish on-demand paths or repair existing ones.

Phinney, Thubert & Assimi Expires April 22, 2014 [Page 25]

Internet-Draft RPL-industrial-applicability-statement October 20134.2. Layer-two features This section defers to work that is taking place at the 6TiSCH WG. In particular [I-D.wang-6tisch-6top] defines the Link Layer Control (LLC) operation that sustain RPL and IPv6 whereas [I-D.vilajosana-6tisch-minimal] specifies a minimal RPL operation based on a static TSCH schedule.4.3. Recommended Configuration Defaults and Ranges4.3.1. Trickle Parameters Trickle was designed to be density-aware and perform well in networks characterized by a wide range of node densities. The combination of DIO packet suppression and adaptive timers for sending updates allows Trickle to perform well in both sparse and dense environments. Node densities in industrial automation network deployments can vary greatly, from nodes having only one or a handful of neighbors to nodes having several hundred neighbors. In high density environments, relatively low values for Imin may cause a short period of congestion when an inconsistency is detected and DIO updates are sent by a large number of neighboring nodes nearly simultaneously. While the Trickle timer will exponentially backoff, some time may elapse before the congestion subsides. Although some link layers employ contention mechanisms that attempt to avoid congestion, relying solely on the link layer to avoid congestion caused by a large number of DIO updates can result in increased communication latency for other control and data traffic in the network. To mitigate this kind of short-term congestion, this document recommends a more conservative set of values for the Trickle parameters than those specified in [RFC6206]. In particular, DIOIntervalMin is set to a larger value to avoid periods of congestion in dense environments, and DIORefundancyConstant is parameterized accordingly as described below. These values are appropriate for the timely distribution of DIO updates in both sparse and dense scenarios while avoiding the short-term congestion that might arise in dense scenarios. Because the actual link capacity depends on the particular link technology used within an industrial automation network deployment, the Trickle parameters are specified in terms of the link's maximum capacity for conveying link-local multicast messages. If the link can convey m link-local multicast packets per second on average, the expected time it takes to transmit a link-local multicast packet is 1/m seconds. DIOIntervalMin: Industrial automation network deployments SHOULD set DIOIntervalMin such that the Trickle Imin is at least 50 times as long as it takes to convey a link-local multicast packet. This value is larger than that recommended in [RFC6206] to avoid congestion in dense plant deployments as described above. Phinney, Thubert & Assimi Expires April 22, 2014 [Page 26]

Internet-Draft RPL-industrial-applicability-statement October 2013

DIOIntervalDoublings: Industrial automation network deployments SHOULD set DIOIntervalDoublings such that the Trickle I_{max} is at least TBD minutes or more.

DIORedundancyConstant: Industrial automation network deployments SHOULD set DIORedundancyConstant to a value of at least 10. This is due to the larger chosen value for DIOIntervalMin and the proportional relationship between I_{min} and k suggested in [RFC6206]. This increase is intended to compensate for the increased communication latency of DIO updates caused by the increase in the DIOIntervalMin value, though the proportional relationship between I_{min} and k suggested in [RFC6206] is not preserved. Instead, DIORedundancyConstant is set to a lower value in order to reduce the number of packet transmissions in dense environments.

4.3.2. Other Parameters None identified at this time. Further work is required to refine this analysis.

5. Manageability Considerations RPL enables automatic and consistent configuration of RPL routers through parameters specified by the DODAG root and disseminated through DIO packets. The use of Trickle for scheduling DIO transmissions ensures lightweight yet timely propagation of important network and parameter updates and allows network operators to choose the trade-off point they are comfortable with respect to overhead vs. reliability and timeliness of network updates. The metrics in use in the network along with the Trickle Timer parameters used to control the frequency and redundancy of network updates can be dynamically varied by the root during the lifetime of the network. To that end, all DIO messages SHOULD contain a Metric Container option for disseminating the metrics and metric values used for DODAG setup. In addition, DIO messages SHOULD contain a DODAG Configuration option for disseminating the Trickle Timer parameters throughout the network. The possibility of dynamically updating the metrics in use in the network as well as the frequency of network updates allows deployment characteristics (e.g., network density) to be discovered during network bring-up and to be used to tailor network parameters once the network is operational rather than having to rely on precise pre-configuration. This also allows the network parameters and the overall routing protocol behavior to evolve during the lifetime of the network. RPL specifies a number of variables and events that can be tracked for purposes of network fault and performance monitoring of RPL routers. Depending on the memory and processing capabilities of each smart grid device, various subsets of these can be employed in the field.

Phinney, Thubert & Assimi Expires April 22, 2014 [Page 27]

Internet-Draft RPL-industrial-applicability-statement October 20136. Security Considerations Industrial automation network deployments typically operate in areas that provide limited physical security (relative to the risk of a attack). For this reason, the link layer, transport layer and application layer technologies utilized within such networks typically provide security mechanisms to ensure authentication, confidentiality, integrity, timeliness and freshness. As a result, such deployments may not need to implement RPL's security mechanisms and could rely on link layer and higher layer security features. This document does not specify operations that could introduce new threats. Security considerations for RPL deployments are to be developed in accordance with recommendations laid out in, for example, [I-D.tsao-roll-security-framework]. Industrial automation networks are subject to stringent security requirements as they are considered a critical infrastructure component. At the same time, since they are composed of large numbers of resource-constrained devices inter-connected with limited-throughput links, many available security mechanisms are not practical for use in such networks. As a result, the choice of security mechanisms is highly dependent on the device and network capabilities characterizing a particular deployment. In contrast to other types of LLNs, in industrial automation networks centralized administrative control and access to a permanent secure infrastructure is available. As a result link-layer, transport-layer and/or application-layer security mechanisms are typically in place and may make use of RPL's secure mode unnecessary.6.1. Security Considerations during initial deployment6.2. Security Considerations during incremental deployment7. Other Related Protocols8. IANA Considerations This specification has no requirement on IANA.9. Acknowledgements10. References10.1. Normative References [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.10.2. Informative References [I-D.ietf-roll-terminology]Phinney, Thubert & AssimiExpires April 22, 2014 [Page 28]

Internet-Draft RPL-industrial-applicability-statement October 2013
 Vasseur, J., "Terminology in Low power And Lossy Networks", I
 nternet-Draft draft-ietf-roll-terminology-12, March 2013. [RFC288
 7] Handley, M., Floyd, S., Whetten, B., Kermode, R., Vicisano, L.
 and M. Luby, "The Reliable Multicast Design Space for Bulk Data Tra
 nsfer", RFC 2887, August 2000. [RFC5548] Dohler, M., Watteyne, T., Winter, T.
 and D. Barthel, "Routing Requirements for Urban Low-Power and Loss
 y Networks", RFC 5548, May 2009. [RFC5826] Brandt, A., Buron, J.
 and G. Porcu, "Home Automation Routing Requirements in Low-Power a
 nd Lossy Networks", RFC 5826, April 2010. [RFC5867] Martocci, J.
 , De Mil, P., Riou, N. and W. Vermeulen, "Building Automation Routi
 ng Requirements in Low-Power and Lossy Networks", RFC 5867, June 20
 10. [RFC5673] Pister, K., Thubert, P., Dwars, S. and T. Phinney,
 "Industrial Routing Requirements in Low-Power and Lossy Networks",
 RFC 5673, October 2009. [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali,
 O. and J. Ko, "The Trickle Algorithm", RFC 6206, March 2011. [RFC
 6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Le
 vis, P., Pister, K., Struik, R., Vasseur, JP. and R. Alexander, "RP
 L: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 655
 0, March 2012. [RFC6552] Thubert, P., "Objective Function Zero for the Routin
 g Protocol for Low-Power and Lossy Networks (RPL)", RFC
 6552, March 2012. [RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank wit
 h Hysteresis Objective Function", RFC 6719, September 2012. [RFC6
 997] Goyal, M., Baccelli, E., Philipp, M., Brandt, A. and J. Marto
 cci, "Reactive Discovery of Point-to-Point Routes in Low-Power and
 Lossy Networks", RFC 6997, August 2013. [I-D.thubert-roll-asymmlink]
 Thubert, P., "RPL adaptation for asymmetrical links", Internet-D
 raft draft-thubert-roll-asymmlink-02, December 2011. [I-D.thubert-
 6lo-forwarding-fragments] Thubert, P. and J. Hui, "LLN Fragment For
 warding and Recovery", Internet-Draft draft-thubert-6lo-forwarding-
 fragments-00, October 2013. [I-D.thubert-6tisch-architecture]Phi
 nney, Thubert & AssimiExpires April 22, 2014 [Page 29]

Internet-Draft RPL-industrial-applicability-statement October 2013

Thubert, P., Assimiti, R. and T. Watteyne, "An Architecture for IPv6 over the TSCH mode of IEEE IEEE802.15.4e", Internet-Draft draft-thubert-6tisch-architecture-00, October 2013. [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", Internet-Draft draft-tsao-roll-security-framework-02, March 2010. [I-D.watteyne-6tisch-tsch] Watteyne, T., "Using IEEE802.15.4e TSCH in an LLN context: Overview, Problem Statement and Goals", Internet-Draft draft-watteyne-6tisch-tsch-00, October 2013. [I-D.wang-6tisch-6top] Wang, Q., Vilajosana, X. and T. Watteyne, "6TiSCH Operation Sublayer (6top)", Internet-Draft draft-wang-6tisch-6top-00, October 2013. [I-D.vilajosana-6tisch-minimal] Vilajosana, X. and K. Pister, "Minimal 6TiSCH Configuration", Internet-Draft draft-vilajosana-6tisch-minimal-00, October 2013.10.3. External Informative References [HART] www.hartcomm.org, "Highway Addressable Remote Transducer, a group of specifications for industrial process and control devices administered by the HART Foundation", . [ISA100.11a] ISA, "ISA100, Wireless Systems for Automation", May 2008, <http://www.isa.org/Community/SP100WirelessSystemsforAutomation>. [ZigBeeIP] ZigBee Public Document 15-002r00, "ZigBee IP Specification", 2013. Authors' Addresses Tom Phinney, editor consultant 5012 W. Torrey Pines Circle Glendale, AZ 85308-3221 USA Phone: +1 602 938 3163 Email: tom.phinney@cox.net Phinney, Thubert & Assimi Expires April 22, 2014 [Page 30]

Internet-Draft RPL-industrial-applicability-statement October 2013 Pasca
l Thubert Cisco Systems, Inc Building D 45 Allee des Ormes - BP1200 MOUG
INS - Sophia Antipolis, 06254 FRANCE Phone: +33 497 23 26 34 Email: pth
ubert@cisco.com Robert Assimiti Nivis 1000 Circle 75 Parkway SE, Ste 300
Atlanta, GA 30339 USA Phone: +1 678 202 6859 Email: robert.assimiti@ni
vis.comPhinney, Thubert & AssimiExpires April 22, 2014 [Page 31]

Routing Over Low-Power and Lossy Networks
Internet-Draft
Intended status: Informational
Expires: April 05, 2015

T. Tsao
R. Alexander
Cooper Power Systems
M. Dohler
CTTC
V. Daza
A. Lozano
Universitat Pompeu Fabra
M. Richardson, Ed.
Sandelman Software Works
October 02, 2014

A Security Threat Analysis for Routing Protocol for Low-power and lossy
networks (RPL)
draft-ietf-roll-security-threats-11

Abstract

This document presents a security threat analysis for the Routing Protocol for Low-power and lossy networks (RPL, ROLL). The development builds upon previous work on routing security and adapts the assessments to the issues and constraints specific to low-power and lossy networks. A systematic approach is used in defining and evaluating the security threats. Applicable countermeasures are application specific and are addressed in relevant applicability statements.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 05, 2015.

Copyright Notice

Copyright (c) 2014 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. Relationship to other documents	4
3. Terminology	4
4. Considerations on RPL Security	5
4.1. Routing Assets and Points of Access	6
4.2. The ISO 7498-2 Security Reference Model	8
4.3. Issues Specific to or Amplified in LLNs	10
4.4. RPL Security Objectives	12
5. Threat Sources	13
6. Threats and Attacks	13
6.1. Threats due to failures to Authenticate	14
6.1.1. Node Impersonation	14
6.1.2. Dummy Node	14
6.1.3. Node Resource Spam	14
6.2. Threats due to failure to keep routing information confidential	15
6.2.1. Routing Exchange Exposure	15
6.2.2. Routing Information (Routes and Network Topology) Exposure	15
6.3. Threats and Attacks on Integrity	16
6.3.1. Routing Information Manipulation	16
6.3.2. Node Identity Misappropriation	17
6.4. Threats and Attacks on Availability	17
6.4.1. Routing Exchange Interference or Disruption	17
6.4.2. Network Traffic Forwarding Disruption	17
6.4.3. Communications Resource Disruption	19
6.4.4. Node Resource Exhaustion	19
7. Countermeasures	20
7.1. Confidentiality Attack Countermeasures	20
7.1.1. Countering Deliberate Exposure Attacks	20
7.1.2. Countering Passive Wiretapping Attacks	21
7.1.3. Countering Traffic Analysis	22
7.1.4. Countering Remote Device Access Attacks	23

7.2.	Integrity Attack Countermeasures	23
7.2.1.	Countering Unauthorized Modification Attacks	23
7.2.2.	Countering Overclaiming and Misclaiming Attacks	24
7.2.3.	Countering Identity (including Sybil) Attacks	24
7.2.4.	Countering Routing Information Replay Attacks	25
7.2.5.	Countering Byzantine Routing Information Attacks	25
7.3.	Availability Attack Countermeasures	26
7.3.1.	Countering HELLO Flood Attacks and ACK Spoofing Attacks	26
7.3.2.	Countering Overload Attacks	27
7.3.3.	Countering Selective Forwarding Attacks	28
7.3.4.	Countering Sinkhole Attacks	29
7.3.5.	Countering Wormhole Attacks	30
8.	RPL Security Features	31
8.1.	Confidentiality Features	31
8.2.	Integrity Features	32
8.3.	Availability Features	33
8.4.	Key Management	33
9.	IANA Considerations	34
10.	Security Considerations	34
11.	Acknowledgments	34
12.	References	34
12.1.	Normative References	34
12.2.	Informative References	35
	Authors' Addresses	38

1. Introduction

In recent times, networked electronic devices have found an increasing number of applications in various fields. Yet, for reasons ranging from operational application to economics, these wired and wireless devices are often supplied with minimum physical resources; the constraints include those on computational resources (RAM, clock speed, storage), communication resources (duty cycle, packet size, etc.), but also form factors that may rule out user access interfaces (e.g., the housing of a small stick-on switch), or simply safety considerations (e.g., with gas meters). As a consequence, the resulting networks are more prone to loss of traffic and other vulnerabilities. The proliferation of these low-power and lossy networks (LLNs), however, are drawing efforts to examine and address their potential networking challenges. Securing the establishment and maintenance of network connectivity among these deployed devices becomes one of these key challenges.

This document presents a threat analysis for securing the Routing Protocol for LLNs (RPL). The process requires two steps. First, the analysis will be used to identify pertinent security issues. The second step is to identify necessary countermeasures to secure RPL.

As there are multiple ways to solve the problem and the specific tradeoffs are deployment specific, the specific countermeasure to be used is detailed in applicability statements.

This document uses [ISO.7498-2.1988] model, which describes Authentication, Access Control, Data Confidentiality, Data Integrity, and Non-Repudiation security services and to which Availability is added. As explained below, Non-Repudiation does not apply to routing protocols.

Many of the issues in this document were also covered in The IAB Smart Object Workshop [RFC6574], and The IAB Smart Object Security Workshop [I-D.gilger-smart-object-security-workshop].

All of this document concerns itself with securing the control plane traffic. As such it does not address authorization or authentication of application traffic. RPL uses multicast as part of its protocol, and therefore mechanisms which RPL uses to secure this traffic might also be applicable to MPL control traffic as well: the important part is that the threats are similar.

2. Relationship to other documents

ROLL has specified a set of routing protocols for Lossy and Low-resource Networks (LLN) [RFC6550]. A number of applicability texts describes a subset of these protocols and the conditions which make the subset the correct choice. The text recommends and motivates the accompanying parameter value ranges. Multiple applicability domains are recognized including: Building and Home, and Advanced Metering Infrastructure. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most probable network structures. Each applicability statement identifies the distinguishing properties according to a common set of subjects described in as many sections.

The common set of security threats herein are referred to by the applicability statements, and that series of documents describes the preferred security settings and solutions within the applicability statement conditions. This applicability statements may recommend more light weight security solutions and specify the conditions under which these solutions are appropriate.

3. Terminology

This document adopts the terminology defined in [RFC6550], in [RFC4949], and in [RFC7102].

The terms control plane and forwarding plane are used consistently with section 1 of [RFC6192].

The term DODAG is from [RFC6550].

EAP-TLS is defined in [RFC5216].

PANA is defined in [RFC5191].

CCM mode is defined in [RFC3610].

[RFC7102] introduces the term Sleepy Node, referring to a node which may sometimes go into a low-power state, suspending protocol communications

The terms SSID, ESSID and PAN refer to network identifiers, defined in [IEEE.802.11] and [IEEE.802.15.4].

Although this is not a protocol specification, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] in order to clarify and emphasize the guidance and directions to implementers and deployers of LLN nodes that utilize RPL.

4. Considerations on RPL Security

Routing security, in essence, ensures that the routing protocol operates correctly. It entails implementing measures to ensure controlled state changes on devices and network elements, both based on external inputs (received via communications) or internal inputs (physical security of device itself and parameters maintained by the device, including, e.g., clock). State changes would thereby involve not only authorization of injector's actions, authentication of injectors, and potentially confidentiality of routing data, but also proper order of state changes through timeliness, since seriously delayed state changes, such as commands or updates of routing tables, may negatively impact system operation. A security assessment can therefore begin with a focus on the assets [RFC4949] that may be the target of the state changes and the access points in terms of interfaces and protocol exchanges through which such changes may occur. In the case of routing security, the focus is directed towards the elements associated with the establishment and maintenance of network connectivity.

This section sets the stage for the development of the analysis by applying the systematic approach proposed in [Myagmar2005] to the routing security, while also drawing references from other reviews


and assessments found in the literature, particularly, [RFC4593] and [Karlof2003] (i.e. selective forwarding, wormhole and sinkhole attacks). The subsequent subsections begin with a focus on the elements of a generic routing process that is used to establish routing assets and points of access to the routing functionality. Next, the [ISO.7498-2.1988] security model is briefly described. Then, consideration is given to issues specific to or amplified in LLNs. This section concludes with the formulation of a set of security objectives for RPL.


4.1. Routing Assets and Points of Access

An asset is an important system resource (including information, process, or physical resource), the access to, corruption or loss of which adversely affects the system. In the control plane context, an asset is information about the network, processes used to manage and manipulate this data, and the physical devices on which this data is stored and manipulated. The corruption or loss of these assets may adversely impact the control plane of the network. Within the same context, a point of access is an interface or protocol that facilitates interaction between control plane assets. Identifying these assets and points of access will provide a basis for enumerating the attack surface of the control plane.

A level-0 data flow diagram [Yourdon1979] is used here to identify the assets and points of access within a generic routing process. The use of a data flow diagram allows for a clear and concise model of the way in which routing nodes interact and process information, and hence provides a context for threats and attacks. The goal of the model is to be as detailed as possible so that corresponding assets, points of access, and process in an individual routing protocol can be readily identified.

Figure 1 shows that nodes participating in the routing process transmit messages to discover neighbors and to exchange routing information; routes are then generated and stored, which may be maintained in the form of the protocol forwarding table. The nodes use the derived routes for making forwarding decisions.





Notation:

(Proc) A process Proc

topology A structure storing neighbor adjacency (parent/child)

routes A structure storing the forwarding information base (FIB)

|Node_n| An external entity Node_n

-----> Data flow

Figure 1: Data Flow Diagram of a Generic Routing Process

It is seen from Figure 1 that

o Assets include

- * routing and/or topology information;
- * route generation process;
- * communication channel resources (bandwidth);
- * node resources (computing capacity, memory, and remaining energy);
- * node identifiers (including node identity and ascribed attributes such as relative or absolute node location).

o Points of access include

- * neighbor discovery;
- * route/topology exchange;
- * node physical interfaces (including access to data storage).

A focus on the above list of assets and points of access enables a more directed assessment of routing security; for example, it is readily understood that some routing attacks are in the form of attempts to misrepresent routing topology. Indeed, the intention of the security threat analysis is to be comprehensive. Hence, some of the discussion which follows is associated with assets and points of access that are not directly related to routing protocol design but nonetheless provided for reference since they do have direct consequences on the security of routing.

4.2. The ISO 7498-2 Security Reference Model

At the conceptual level, security within an information system in general and applied to RPL in particular is concerned with the primary issues of authentication, access control, data confidentiality, data integrity, and non-repudiation. In the context of RPL:

Authentication

Authentication involves the mutual authentication of the routing peers prior to exchanging route information (i.e., peer authentication) as well as ensuring that the source of the route data is from the peer (i.e., data origin authentication). [RFC5548] points out that LLNs can be drained by unauthenticated peers before configuration. [RFC5673] requires availability of open and untrusted side channels for new joiners, and it requires strong and automated authentication so that networks can automatically accept or reject new joiners.

Access Control

Access Control provides protection against unauthorized use of the asset, and deals with the authorization of a node.

Confidentiality

Confidentiality involves the protection of routing information as well as routing neighbor maintenance exchanges so that only authorized and intended network entities may view or access it. Because LLNs are most commonly found on a publicly accessible shared medium, e.g., air or wiring in a building, and sometimes formed ad hoc, confidentiality also extends to the neighbor state and database information within the routing device since the deployment of the network creates the potential for unauthorized access to the physical devices themselves.

Integrity

Integrity entails the protection of routing information and routing neighbor maintenance exchanges, as well as derived information maintained in the database, from unauthorized modification, insertions, deletions or replays. to be addressed beyond the routing protocol.

Non-repudiation

Non-repudiation is the assurance that the transmission and/or reception of a message cannot later be denied. The service of non-repudiation applies after-the-fact and thus relies on the logging or other capture of on-going message exchanges and signatures. Routing protocols typically do not have a notion of repudiation, so non-repudiation services are not required. Further, with the LLN application domains as described in [RFC5867] and [RFC5548], proactive measures are much more critical than retrospective protections. Finally, given the significant practical limits to on-going routing transaction logging and storage and individual device digital signature verification for each exchange, non-repudiation in the context of routing is an unsupportable burden that bears no further considered as an RPL security issue.

It is recognized that, besides those security issues captured in the ISO 7498-2 model, availability, is a security requirement:

Availability

Availability ensures that routing information exchanges and forwarding services need to be available when they are required for the functioning of the serving network. Availability will apply to maintaining efficient and correct operation of routing and neighbor discovery exchanges (including needed information) and forwarding services so as not to impair or limit the network's central traffic flow function

It should be emphasized here that for RPL security the above requirements must be complemented by the proper security policies and enforcement mechanisms to ensure that security objectives are met by a given RPL implementation.

4.3. Issues Specific to or Amplified in LLNs

The requirements work detailed in Urban Requirements ([RFC5548]), Industrial Requirements ([RFC5673]), Home Automation ([RFC5826], and Building Automation ([RFC5867]) have identified specific issues and constraints of routing in LLNs. The following is a list of observations from those requirements and evaluation of their impact on routing security considerations.

Limited energy, memory, and processing node resources

As a consequence of these constraints, there is an even more critical need than usual for a careful study of trade-offs on which and what level of security services are to be afforded during the system design process. The chosen security mechanisms also needs to work within these constraints. Synchronization of security states with sleepy nodes [RFC7102] is yet another issue. A non-rechargeable battery powered node may well be limited in energy for it's lifetime: once exhausted, it may well never function again.

Large scale of rolled out network

The possibly numerous nodes to be deployed make manual on-site configuration unlikely. For example, an urban deployment can see several hundreds of thousands of nodes being installed by many installers with a low level of expertise. Nodes may be installed and not activated for many years, and additional nodes may be added later on, which may be from old inventory. The lifetime of the network is measured in decades, and this complicates the operation of key management.

Autonomous operations

Self-forming and self-organizing are commonly prescribed requirements of LLNs. In other words, a routing protocol designed for LLNs needs to contain elements of ad hoc networking and in most cases cannot rely on manual configuration for initialization or local filtering rules. Network topology/ownership changes, partitioning or merging, as well as node replacement, can all contribute to complicating the operations of key management.

Highly directional traffic

Some types of LLNs see a high percentage of their total traffic traverse between the nodes and the LLN Border Routers (LBRs)

where the LLNs connect to non-LLNs. The special routing status of and the greater volume of traffic near the LBRs have routing security consequences as a higher valued attack target. In fact, when Point-to-MultiPoint (P2MP) and MultiPoint-to-Point (MP2P) traffic represents a majority of the traffic, routing attacks consisting of advertising incorrect preferred routes can cause serious damage.

While it might seem that nodes higher up in the acyclic graph (i.e. those with lower rank) should be secured in a stronger fashion, it is not in general easy to predict which nodes will occupy those positions until after deployment. Issues of redundancy and inventory control suggests that any node might wind up in such a sensitive attack position, so all nodes to be capable of being fully secured.

In addition, even if it were possible to predict which nodes will occupy positions of lower rank and provision them with stronger security mechanisms, in the absense of a strong authorization model, any node could advertise an incorrect preferred route.

Unattended locations and limited physical security

Many applications have the nodes deployed in unattended or remote locations; furthermore, the nodes themselves are often built with minimal physical protection. These constraints lower the barrier of accessing the data or security material stored on the nodes through physical means.

Support for mobility

On the one hand, only a limited number of applications require the support of mobile nodes, e.g., a home LLN that includes nodes on wearable health care devices or an industry LLN that includes nodes on cranes and vehicles. On the other hand, if a routing protocol is indeed used in such applications, it will clearly need to have corresponding security mechanisms.

Additionally nodes may appear to move from one side of a wall to another without any actual motion involved, the result of changes to electromagnetic properties, such as opening and closing of a metal door.

Support for multicast and anycast

Support for multicast and anycast is called out chiefly for large-scale networks. Since application of these routing mechanisms in autonomous operations of many nodes is new, the consequence on security requires careful consideration.

The above list considers how an LLN's physical constraints, size, operations, and variety of application areas may impact security. However, it is the combinations of these factors that particularly stress the security concerns. For instance, securing routing for a large number of autonomous devices that are left in unattended locations with limited physical security presents challenges that are not found in the common circumstance of administered networked routers. The following subsection sets up the security objectives for the routing protocol designed by the ROLL WG.

4.4. RPL Security Objectives

This subsection applies the ISO 7498-2 model to routing assets and access points, taking into account the LLN issues, to develop a set of RPL security objectives.

Since the fundamental function of a routing protocol is to build routes for forwarding packets, it is essential to ensure that:

- o routing/topology information integrity remains intact during transfer and in storage;
- o routing/topology information is used by authorized entities;
- o routing/topology information is available when needed.

In conjunction, it is necessary to be assured that

- o authorized peers authenticate themselves during the routing neighbor discovery process;
- o the routing/topology information received is generated according to the protocol design.

However, when trust cannot be fully vested through authentication of the principals alone, i.e., concerns of insider attack, assurance of the truthfulness and timeliness of the received routing/topology information is necessary. With regard to confidentiality, protecting the routing/topology information from unauthorized exposure may be desirable in certain cases but is in itself less pertinent in general to the routing function.

One of the main problems of synchronizing security states of sleepy nodes, as listed in the last subsection, lies in difficulties in authentication; these nodes may not have received in time the most recent update of security material. Similarly, the issues of minimal manual configuration, prolonged rollout and delayed addition of nodes, and network topology changes also complicate key management.

Hence, routing in LLNs needs to bootstrap the authentication process and allow for flexible expiration scheme of authentication credentials.

The vulnerability brought forth by some special-function nodes, e.g., LBRs, requires the assurance, particularly in a security context,

- o of the availability of communication channels and node resources;
- o that the neighbor discovery process operates without undermining routing availability.

There are other factors which are not part of RPL but directly affecting its function. These factors include weaker barrier of accessing the data or security material stored on the nodes through physical means; therefore, the internal and external interfaces of a node need to be adequate for guarding the integrity, and possibly the confidentiality, of stored information, as well as the integrity of routing and route generation processes.

Each individual system's use and environment will dictate how the above objectives are applied, including the choices of security services as well as the strengths of the mechanisms that must be implemented. The next two sections take a closer look at how the RPL security objectives may be compromised and how those potential compromises can be countered.

5. Threat Sources

[RFC4593] provides a detailed review of the threat sources: outsiders and byzantine. RPL has the same threat sources.

6. Threats and Attacks

This section outlines general categories of threats under the ISO 7498-2 model and highlights the specific attacks in each of these categories for RPL. As defined in [RFC4949], a threat is "a potential for violation of security, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm."

An attack is "an assault on system security that derives from an intelligent threat, i.e., an intelligent act that is a deliberate attempt (especially in the sense of a method or technique) to evade security services and violate the security policy of a system."

The subsequent subsections consider the threats and the attacks that can cause security breaches under the ISO 7498-2 model to the routing

assets and via the routing points of access identified in Section 4.1. The assessment steps through the security concerns of each routing asset and looks at the attacks that can exploit routing points of access. The threats and attacks identified are based on the routing model analysis and associated review of the existing literature. The source of the attacks is assumed to be from either inside or outside attackers. While some attackers inside the network will be using compromised nodes, and therefore are only able to do what an ordinary node can ("node-equivalent"), other attacks may not be limited in memory, CPU, power consumption or long term storage. Moore's law favours the attacker with access to the latest capabilities, while the defenders will remain in place for years to decades.

6.1. Threats due to failures to Authenticate

6.1.1. Node Impersonation

If an attacker can join a network using any identity, then it may be able to assume the role of a legitimate (and existing) node). It may be able to report false readings (in metering applications), or provide inappropriate control messages (in control systems involving actuators) if the security of the application is implied by the security of the routing system.

Even in systems where there is application layer security, the ability to impersonate a node would permit an attacker to direct traffic to itself. This may permit various on-path attacks which would otherwise be difficult, such as replaying, delaying, or duplicating (application) control messages.

6.1.2. Dummy Node

If an attacker can join a network using any identity, then it can pretend to be a legitimate node, receiving any service legitimate nodes receive. It may also be able to report false readings (in metering applications), or provide inappropriate authorizations (in control systems involving actuators), or perform any other attacks that are facilitated by being able to direct traffic towards itself.

6.1.3. Node Resource Spam

If an attacker can join a network with any identity, then it can continuously do so with new (random) identities. This act may drain down the resources of the network (battery, RAM, bandwidth). This may cause legitimate nodes of the network to be unable to communicate.

6.2. Threats due to failure to keep routing information confidential

The assessment in Section 4.2 indicates that there are attacks against the confidentiality of routing information at all points of access. This threat may result in disclosure, as described in Section 3.1.2 of [RFC4593], and may involve a disclosure of routing information.

6.2.1. Routing Exchange Exposure

Routing exchanges include both routing information as well as information associated with the establishment and maintenance of neighbor state information. As indicated in Section 4.1, the associated routing information assets may also include device specific resource information, such as available memory, remaining power, etc., that may be metrics of the routing protocol.

The routing exchanges will contain reachability information, which would identify the relative importance of different nodes in the network. Nodes higher up in the DODAG, to which more streams of information flow, would be more interesting targets for other attacks, and routing exchange exposures can identify them.

6.2.2. Routing Information (Routes and Network Topology) Exposure

Routes (which may be maintained in the form of the protocol forwarding table) and neighbor topology information are the products of the routing process that are stored within the node device databases.

The exposure of this information will allow attackers to gain direct access to the configuration and connectivity of the network thereby exposing routing to targeted attacks on key nodes or links. Since routes and neighbor topology information is stored within the node device, attacks on the confidentiality of the information will apply to the physical device including specified and unspecified internal and external interfaces.

The forms of attack that allow unauthorized access or disclosure of the routing information will include:

- o Physical device compromise;
- o Remote device access attacks (including those occurring through remote network management or software/field upgrade interfaces).

Both of these attack vectors are considered a device specific issue, and are out of scope for RPL to defend against. In some

applications, physical device compromise may be a real threat and it may be necessary to provide for other devices to securely detect a compromised device and react quickly to exclude it.

6.3. Threats and Attacks on Integrity

The assessment in Section 4.2 indicates that information and identity assets are exposed to integrity threats from all points of access. In other words, the integrity threat space is defined by the potential for exploitation introduced by access to assets available through routing exchanges and the on-device storage.

6.3.1. Routing Information Manipulation

Manipulation of routing information that range from neighbor states to derived routes will allow unauthorized sources to influence the operation and convergence of the routing protocols and ultimately impact the forwarding decisions made in the network.

Manipulation of topology and reachability information will allow unauthorized sources to influence the nodes with which routing information is exchanged and updated. The consequence of manipulating routing exchanges can thus lead to sub-optimality and fragmentation or partitioning of the network by restricting the universe of routers with which associations can be established and maintained.

A sub-optimal network may use too much power and/or may congest some routes leading to premature failure of a node, and a denial of service on the entire network.

In addition, being able to attract network traffic can make a blackhole attack more damaging.

The forms of attack that allow manipulation to compromise the content and validity of routing information include

- o Falsification, including overclaiming and misclaiming (claiming routes to devices which the device can not in fact reach);
- o Routing information replay;
- o Byzantine (internal) attacks that permit corruption of routing information in the node even where the node continues to be a validated entity within the network (see, for example, [RFC4593] for further discussions on Byzantine attacks);
- o Physical device compromise or remote device access attacks.

6.3.2. Node Identity Misappropriation

Falsification or misappropriation of node identity between routing participants opens the door for other attacks; it can also cause incorrect routing relationships to form and/or topologies to emerge. Routing attacks may also be mounted through less sophisticated node identity misappropriation in which the valid information broadcast or exchanged by a node is replayed without modification. The receipt of seemingly valid information that is however no longer current can result in routing disruption, and instability (including failure to converge). Without measures to authenticate the routing participants and to ensure the freshness and validity of the received information the protocol operation can be compromised. The forms of attack that misuse node identity include

- o Identity attacks, including Sybil attacks (see [Sybil2002]) in which a malicious node illegitimately assumes multiple identities;
- o Routing information replay.

6.4. Threats and Attacks on Availability

The assessment in Section 4.2 indicates that the process and resources assets are exposed to threats against availability; attacks in this category may exploit directly or indirectly information exchange or forwarding (see [RFC4732] for a general discussion).

6.4.1. Routing Exchange Interference or Disruption

Interference is the threat action and disruption is threat consequence that allows attackers to influence the operation and convergence of the routing protocols by impeding the routing information exchange.

The forms of attack that allow interference or disruption of routing exchange include:

- o Routing information replay;
- o ACK spoofing;
- o Overload attacks. (Section 7.3.2)

In addition, attacks may also be directly conducted at the physical layer in the form of jamming or interfering.

6.4.2. Network Traffic Forwarding Disruption

The disruption of the network traffic forwarding capability will undermine the central function of network routers and the ability to handle user traffic. This affects the availability of the network because of the potential to impair the primary capability of the network.

In addition to physical layer obstructions, the forms of attack that allows disruption of network traffic forwarding include [Karlof2003]

- o Selective forwarding attacks;

```
|Node_1|--(msg1|msg2|msg3)-->|Attacker|--(msg1|msg3)-->|Node_2|
```

Figure 2: Selective forwarding example

- o Wormhole attacks;

```
|Node_1|-----Unreachable-----x|Node_2|
|                                     ^
|                                     |
|                                     Private Link
|                                     |
'-->|Attacker_1|=====>|Attacker_2|--'
```

Figure 3: Wormhole Attacks

- o Sinkhole attacks.




Figure 4: sinkhole attack example

These attacks are generally done to both control plane and forwarding plane traffic. A system that prevents control plane traffic (RPL messages) from being diverted in these ways will also prevent actual data from being diverted.

6.4.3. Communications Resource Disruption

Attacks mounted against the communication channel resource assets needed by the routing protocol can be used as a means of disrupting its operation. However, while various forms of Denial of Service (DoS) attacks on the underlying transport subsystem will affect routing protocol exchanges and operation (for example physical layer RF jamming in a wireless network or link layer attacks), these attacks cannot be countered by the routing protocol. As such, the threats to the underlying transport network that supports routing is considered beyond the scope of the current document. Nonetheless, attacks on the subsystem will affect routing operation and so must be directly addressed within the underlying subsystem and its implemented protocol layers.

6.4.4. Node Resource Exhaustion

A potential threat consequence can arise from attempts to overload the node resource asset by initiating exchanges that can lead to the exhaustion of processing, memory, or energy resources. The establishment and maintenance of routing neighbors opens the routing process to engagement and potential acceptance of multiple neighboring peers. Association information must be stored for each peer entity and for the wireless network operation provisions made to

periodically update and reassess the associations. An introduced proliferation of apparent routing peers can therefore have a negative impact on node resources.

Node resources may also be unduly consumed by attackers attempting uncontrolled topology peering or routing exchanges, routing replays, or the generating of other data traffic floods. Beyond the disruption of communications channel resources, these consequences may be able to exhaust node resources only where the engagements are able to proceed with the peer routing entities. Routing operation and network forwarding functions can thus be adversely impacted by node resources exhaustion that stems from attacks that include:

- o Identity (including Sybil) attacks (see [Sybil2002]);
- o Routing information replay attacks;
- o HELLO-type flood attacks;
- o Overload attacks. (Section 7.3.2)

7. Countermeasures

By recognizing the characteristics of LLNs that may impact routing, this analysis provides the basis for understanding the capabilities within RPL used to deter the identified attacks and mitigate the threats. The following subsections consider such countermeasures by grouping the attacks according to the classification of the ISO 7498-2 model so that associations with the necessary security services are more readily visible.

7.1. Confidentiality Attack Countermeasures

Attacks to disclosure routing information may be mounted at the level of the routing information assets, at the points of access associated with routing exchanges between nodes, or through device interface access. To gain access to routing/topology information, the attacker may rely on a compromised node that deliberately exposes the information during the routing exchange process, may rely on passive wiretapping or traffic analysis, or may attempt access through a component or device interface of a tampered routing node.

7.1.1. Countering Deliberate Exposure Attacks

A deliberate exposure attack is one in which an entity that is party to the routing process or topology exchange allows the routing/topology information or generated route information to be exposed to an unauthorized entity.

For instance, due to mis-configuration or inappropriate enabling of a diagnostic interface, an entity might be copying ("bridging") traffic from a secured ESSID/PAN to an unsecured interface.

A prerequisite to countering this attack is to ensure that the communicating nodes are authenticated prior to data encryption applied in the routing exchange. The authentication ensures the LLN starts with trusted nodes, but it does not provide an indication of whether the node has been compromised.

Reputation systems could be used to help when some nodes may sleep for extended periods of times. It is also unclear if resulting dataset would even fit into constrained devices.

To mitigate the risk of deliberate exposure, the process that communicating nodes use to establish session keys must be peer-to-peer (i.e., between the routing initiating and responding nodes). As is pointed out in [RFC4107], automatic key management is critical for good security. This helps ensure that neither node is exchanging routing information with another peer without the knowledge of both communicating peers. For a deliberate exposure attack to succeed, the compromised node will need to be more overt and take independent actions in order to disclose the routing information to 3rd party.

Note that the same measures which apply to securing routing/topology exchanges between operational nodes must also extend to field tools and other devices used in a deployed network where such devices can be configured to participate in routing exchanges.

7.1.2. Countering Passive Wiretapping Attacks

A passive wiretap attack seeks to breach routing confidentiality through passive, direct analysis and processing of the information exchanges between nodes.

Passive wiretap attacks can be directly countered through the use of data encryption for all routing exchanges. Only when a validated and authenticated node association is completed will routing exchange be allowed to proceed using established session keys and an agreed encryption algorithm. The mandatory to implement CCM mode AES-128 method, is described in [RFC3610], and is believed to be secure against a brute force attack by even the most well-equipped adversary.

The significant challenge for RPL is in the provisioning of the key, which in some modes of RFC6550 is used network-wide. RFC6550 does not solve this problem, and it is the subject of significant future work: see, for instance: [AceCharterProposal], [SolaceProposal], [SmartObjectSecurityWorkshop].

A number of deployments, such as [ZigBeeIP] specify no layer-3/RPL encryption or authentication and rely upon similiar security at layer-2. These networks are immune to outside wiretapping attacks, but are vulnerable to passive (and active) routing attacks through compromises of nodes. (see Section 8.2).

Section 10.9 of [RFC6550] specifies AES-128 in CCM mode with a 32-bit MAC.

Section 5.6 Zigbee IP [ZigBeeIP] specifies use of CCM, with PANA and EAP-TLS for key management.

7.1.3. Countering Traffic Analysis

Traffic analysis provides an indirect means of subverting confidentiality and gaining access to routing information by allowing an attacker to indirectly map the connectivity or flow patterns (including link-load) of the network from which other attacks can be mounted. The traffic analysis attack on an LLN, especially one founded on shared medium, is passive and relies on the ability to read the immutable source/destination layer-2 and/or layer-3 routing information that must remain unencrypted to permit network routing.

One way in which passive traffic analysis attacks can be muted is through the support of load balancing that allows traffic to a given destination to be sent along diverse routing paths. RPL does not generally support multi-path routing within a single DODAG. Multiple DODAGs are supported in the protocol, and an implementation could make use of that. RPL does not have any inherent or standard way to guarantee that the different DODAGs would have significantly diverse paths. Having the diverse DODAGs routed at different border routers might work in some instances, and this could be combined with a multipath technology like MPTCP ([RFC6824]). It is unlikely that it will be affordable in many LLNs, as few deployments will have memory space for more than a few sets of DODAG tables.

Another approach to countering passive traffic analysis could be for nodes to maintain constant amount of traffic to different destinations through the generation of arbitrary traffic flows; the drawback of course would be the consequent overhead and energy expenditure.

The only means of fully countering a traffic analysis attack is through the use of tunneling (encapsulation) where encryption is applied across the entirety of the original packet source/destination addresses. Deployments which use layer-2 security that includes encryption already do this for all traffic.

7.1.4. Countering Remote Device Access Attacks

Where LLN nodes are deployed in the field, measures are introduced to allow for remote retrieval of routing data and for software or field upgrades. These paths create the potential for a device to be remotely accessed across the network or through a provided field tool. In the case of network management a node can be directly requested to provide routing tables and neighbor information.

To ensure confidentiality of the node routing information against attacks through remote access, any local or remote device requesting routing information must be authenticated, and must be authorized for that access. Since remote access is not invoked as part of a routing protocol, security of routing information stored on the node against remote access will not be addressable as part of the routing protocol.

7.2. Integrity Attack Countermeasures

Integrity attack countermeasures address routing information manipulation, as well as node identity and routing information misuse. Manipulation can occur in the form of falsification attack and physical compromise. To be effective, the following development considers the two aspects of falsification, namely, the unauthorized modifications and the overclaiming and misclaiming content. The countering of physical compromise was considered in the previous section and is not repeated here. With regard to misuse, there are two types of attacks to be deterred, identity attacks and replay attacks.

7.2.1. Countering Unauthorized Modification Attacks

Unauthorized modifications may occur in the form of altering the message being transferred or the data stored. Therefore, it is necessary to ensure that only authorized nodes can change the portion of the information that is allowed to be mutable, while the integrity of the rest of the information is protected, e.g., through well-studied cryptographic mechanisms.

Unauthorized modifications may also occur in the form of insertion or deletion of messages during protocol changes. Therefore, the protocol needs to ensure the integrity of the sequence of the exchange sequence.

The countermeasure to unauthorized modifications needs to:

- o implement access control on storage;
- o provide data integrity service to transferred messages and stored data;
- o include sequence number under integrity protection.

7.2.2. Countering Overclaiming and Misclaiming Attacks

Both overclaiming and misclaiming aim to introduce false routes or a false topology that would not occur otherwise, while there are not necessarily unauthorized modifications to the routing messages or information. In order to counter overclaiming, the capability to determine unreasonable routes or topology is required.

The counter to overclaiming and misclaiming may employ:

- o comparison with historical routing/topology data;
- o designs which restrict realizable network topologies.

RPL includes no specific mechanisms in the protocol to counter overclaims or misclaims. An implementation could have specific heuristics implemented locally.

7.2.3. Countering Identity (including Sybil) Attacks

Identity attacks, sometimes simply called spoofing, seek to gain or damage assets whose access is controlled through identity. In routing, an identity attacker can illegitimately participate in routing exchanges, distribute false routing information, or cause an invalid outcome of a routing process.

A perpetrator of Sybil attacks assumes multiple identities. The result is not only an amplification of the damage to routing, but extension to new areas, e.g., where geographic distribution is explicitly or implicitly an asset to an application running on the LLN, for example, the LBR in a P2MP or MP2P LLN.

RPL includes specific public key based authentication at layer-3 that provide for authorization. Many deployments use layer-2 security

that includes admission controls at layer-2 using mechanisms such as PANA.

7.2.4. Countering Routing Information Replay Attacks

In many routing protocols, message replay can result in false topology and/or routes. This is often countered with some kind of counter to ensure the freshness of the message. Replay of a current, literal RPL message are in general idempotent to the topology. An older (lower DODAGVersionNumber) message, if replayed would be rejected as being stale. The trickle algorithm further dampens the effect of any such replay, as if the message was current, then it would contain the same information as before, and it would cause no network changes.

Replays may well occur in some radio technologies (not very likely, 802.15.4) as a result of echos or reflections, and so some replays must be assumed to occur naturally.

Note that for there to be no affect at all, the replay must be done with the same apparent power for all nodes receiving the replay. A change in apparent power might change the metrics through changes to the ETX and therefore might affect the routing even though the contents of the packet were never changed. Any replay which appears to be different should be analyzed as a Selective Forwarding Attack, Sinkhole Attack or Wormhole Attack.

7.2.5. Countering Byzantine Routing Information Attacks

Where a node is captured or compromised but continues to operate for a period with valid network security credentials, the potential exists for routing information to be manipulated. This compromise of the routing information could thus exist in spite of security countermeasures that operate between the peer routing devices.

Consistent with the end-to-end principle of communications, such an attack can only be fully addressed through measures operating directly between the routing entities themselves or by means of external entities able to access and independently analyze the routing information. Verification of the authenticity and liveness of the routing entities can therefore only provide a limited counter against internal (Byzantine) node attacks.

For link state routing protocols where information is flooded with, for example, areas (OSPF [RFC2328]) or levels (ISIS [RFC7142]), countermeasures can be directly applied by the routing entities through the processing and comparison of link state information received from different peers. By comparing the link information

from multiple sources decisions can be made by a routing node or external entity with regard to routing information validity; see Chapter 2 of [Perlman1988] for a discussion on flooding attacks.

For distance vector protocols, such as RPL, where information is aggregated at each routing node it is not possible for nodes to directly detect Byzantine information manipulation attacks from the routing information exchange. In such cases, the routing protocol must include and support indirect communications exchanges between non-adjacent routing peers to provide a secondary channel for performing routing information validation. S-RIP [Wan2004] is an example of the implementation of this type of dedicated routing protocol security where the correctness of aggregate distance vector information can only be validated by initiating confirmation exchanges directly between nodes that are not routing neighbors.

RPL does not provide any direct mechanisms like S-RIP. It does listen to multiple parents, and may switch parents if it begins to suspect that it is being lied to.

7.3. Availability Attack Countermeasures

As alluded to before, availability requires that routing information exchanges and forwarding mechanisms be available when needed so as to guarantee proper functioning of the network. This may, e.g., include the correct operation of routing information and neighbor state information exchanges, among others. We will highlight the key features of the security threats along with typical countermeasures to prevent or at least mitigate them. We will also note that an availability attack may be facilitated by an identity attack as well as a replay attack, as was addressed in Section 7.2.3 and Section 7.2.4, respectively.

7.3.1. Countering HELLO Flood Attacks and ACK Spoofing Attacks

HELLO Flood [Karlof2003],[I-D.suhopark-hello-wsn] and ACK Spoofing attacks are different but highly related forms of attacking an LLN. They essentially lead nodes to believe that suitable routes are available even though they are not and hence constitute a serious availability attack.

A HELLO attack mounted against RPL would involve sending out (or replaying) DIO messages by the attacker. Lower power LLN nodes might then attempt to join the DODAG at a lower rank than they would otherwise.

The most effective method from [I-D.suhopark-hello-wsn] is the verify bidirectionality. A number of layer-2 links are arranged in

controller/spoke arrangements, and continuously are validating connectivity at layer 2.

In addition, in order to calculate metrics, the ETX must be computed, and this involves, in general, sending a number of messages between nodes which are believed to be adjacent.

[I-D.kelsey-intarea-mesh-link-establishment] is one such protocol.

In order to join the DODAG, a DAO message is sent upwards. In RPL the DAO is acknowledged by the DAO-ACK message. This clearly checks bidirectionality at the control plane.

As discussed in section 5.1, [I-D.suhopark-hello-wsn] a receiver with a sensitive receiver could well hear the DAOs, and even send DAO-ACKs as well. Such a node is a form of wormhole attack.

These attacks are also all easily defended against using either layer-2 or layer-3 authentication. Such an attack could only be made against a completely open network (such as might be used for provisioning new nodes), or by a compromised node.

7.3.2. Countering Overload Attacks

Overload attacks are a form of DoS attack in that a malicious node overloads the network with irrelevant traffic, thereby draining the nodes' energy store more quickly, when the nodes rely on batteries or energy scavenging. It thus significantly shortens the lifetime of networks of energy-constrained nodes and constitutes another serious availability attack.

With energy being one of the most precious assets of LLNs, targeting its availability is a fairly obvious attack. Another way of depleting the energy of an LLN node is to have the malicious node overload the network with irrelevant traffic. This impacts availability since certain routes get congested which:

- o renders them useless for affected nodes and data can hence not be delivered;
- o makes routes longer as shortest path algorithms work with the congested network;
- o depletes battery and energy scavenging nodes more quickly and thus shortens the network's availability at large.

Overload attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o introduce quotas on the traffic rate each node is allowed to send;
- o isolate nodes which send traffic above a certain threshold based on system operation characteristics;
- o allow only trusted data to be received and forwarded.

As for the first one, a simple approach to minimize the harmful impact of an overload attack is to introduce traffic quotas. This prevents a malicious node from injecting a large amount of traffic into the network, even though it does not prevent said node from injecting irrelevant traffic at all. Another method is to isolate nodes from the network at the network layer once it has been detected that more traffic is injected into the network than allowed by a prior set or dynamically adjusted threshold. Finally, if communication is sufficiently secured, only trusted nodes can receive and forward traffic which also lowers the risk of an overload attack.

Receiving nodes that validate signatures and sending nodes that encrypt messages need to be cautious of cryptographic processing usage when validating signatures and encrypting messages. Where feasible, certificates should be validated prior to use of the associated keys to counter potential resource overloading attacks. The associated design decision needs to also consider that the validation process requires resources and thus itself could be exploited for attacks. Alternatively, resource management limits can be placed on routing security processing events (see the comment in Section 6, paragraph 4, of [RFC5751]).

7.3.3. Countering Selective Forwarding Attacks

Selective forwarding attacks are a form of DoS attack which impacts the availability of the generated routing paths.

A selective forwarding attack may be done by a node involved with the routing process, or it may be done by what otherwise appears to be a passive antenna or other RF feature or device, but is in fact an active (and selective) device. An RF antenna/repeater which is not selective, is not a threat.

An insider malicious node basically blends neatly in with the network but then may decide to forward and/or manipulate certain packets. If all packets are dropped, then this attacker is also often referred to as a "black hole". Such a form of attack is particularly dangerous if coupled with sinkhole attacks since inherently a large amount of traffic is attracted to the malicious node and thereby causing significant damage. In a shared medium, an outside malicious node would selectively jam overheard data flows, where the thus caused collisions incur selective forwarding.

Selective Forwarding attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o multipath routing of the same message over disjoint paths;
- o dynamically selecting the next hop from a set of candidates.

The first measure basically guarantees that if a message gets lost on a particular routing path due to a malicious selective forwarding attack, there will be another route which successfully delivers the data. Such a method is inherently suboptimal from an energy consumption point of view; it is also suboptimal from a network utilization perspective. The second method basically involves a constantly changing routing topology in that next-hop routers are chosen from a dynamic set in the hope that the number of malicious nodes in this set is negligible. A routing protocol that allows for disjoint routing paths may also be useful.

7.3.4. Countering Sinkhole Attacks

In sinkhole attacks, the malicious node manages to attract a lot of traffic mainly by advertising the availability of high-quality links even though there are none [Karlof2003]. It hence constitutes a serious attack on availability.

The malicious node creates a sinkhole by attracting a large amount of, if not all, traffic from surrounding neighbors by advertising in and outwards links of superior quality. Affected nodes hence eagerly route their traffic via the malicious node which, if coupled with other attacks such as selective forwarding, may lead to serious availability and security breaches. Such an attack can only be executed by an inside malicious node and is generally very difficult to detect. An ongoing attack has a profound impact on the network topology and essentially becomes a problem of flow control.

Sinkhole attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o use geographical insights for flow control;
- o isolate nodes which receive traffic above a certain threshold;
- o dynamically pick up next hop from set of candidates;
- o allow only trusted data to be received and forwarded.

A canary node could periodically call home (using a cryptographic process), with the home system noting if it fails to call in. This provides detection of a problem, but does not mitigate it, and it may have significant energy consequences for the LLN.

Some LLNs may provide for geolocation services, often derived from solving triangulation equations from radio delay calculations, such calculations could in theory be subverted by a sinkhole that transmitted at precisely the right power in a node to node fashion.

While geographic knowledge could help assure that traffic always went in the physical direction desired, it would not assure that the traffic was taking the most efficient route, as the lowest cost real route might not match the physical topology; such as when different parts of an LLN are connected by high-speed wired networks.

7.3.5. Countering Wormhole Attacks

In wormhole attacks at least two malicious nodes claim to have a short path between themselves [Karlof2003]. This changes the availability of certain routing paths and hence constitutes a serious security breach.

Essentially, two malicious insider nodes use another, more powerful, transmitter to communicate with each other and thereby distort the would-be-agreed routing path. This distortion could involve shortcutting and hence paralyzing a large part of the network; it could also involve tunneling the information to another region of the network where there are, e.g., more malicious nodes available to aid the intrusion or where messages are replayed, etc.

In conjunction with selective forwarding, wormhole attacks can create race conditions which impact topology maintenance, routing protocols as well as any security suits built on "time of check" and "time of use".

A pure wormhole attack is nearly impossible to detect. A wormhole which is used in order to subsequently mount another kind of attack would be defeated by defeating the other attack. A perfect wormhole, in which there is nothing adverse that occurs to the traffic, would

be difficult to call an attack. The worst thing that a benign wormhole can do in such a situation is to cease to operate (become unstable), causing the network to have to recalculate routes.

A highly unstable wormhole is no different than a radio opaque (i.e. metal) door that opens and closes a lot. RPL includes hysteresis in its objective functions [RFC6719] in an attempt to deal with frequent changes to the ETX between nodes.

8. RPL Security Features

The assessments and analysis in Section 6 examined all areas of threats and attacks that could impact routing, and the countermeasures presented in Section 7 were reached without confining the consideration to means only available to routing. This section puts the results into perspective; dealing with those threats which are endemic to this field, those which have been mitigated through RPL protocol design, and those which require specific decisions to be made as part of provisioning a network.

The first part of this section, Section 8.1 to Section 8.3, is a description of RPL security features that address specific threats. The second part of this section, Section 8.4, discusses issues of provisioning of security aspects that may impact routing but that also require considerations beyond the routing protocol, as well as potential approaches.

RPL employs multicast and so these alternative communications modes MUST be secured with the same routing security services specified in this section. Furthermore, irrespective of the modes of communication, nodes MUST provide adequate physical tamper resistance commensurate with the particular application domain environment to ensure the confidentiality, integrity, and availability of stored routing information.

8.1. Confidentiality Features

With regard to confidentiality, protecting the routing/topology information from unauthorized disclosure is not directly essential to maintaining the routing function. Breaches of confidentiality may lead to other attacks or the focusing of an attacker's resources (see Section 6.2) but does not of itself directly undermine the operation of the routing function. However, to protect against, and reduce consequences from other more direct attacks, routing information should be protected. Thus, to secure RPL:

- o implement payload encryption using layer-3 mechanisms described in [RFC6550];

- o or: implement layer-2 confidentiality;

Where confidentiality is incorporated into the routing exchanges, encryption algorithms and key lengths need to be specified in accordance with the level of protection dictated by the routing protocol and the associated application domain transport network. For most networks, this means use of AES128 in CCM mode, but this needs to be specified clearly in the applicability statement.

In terms of the life time of the keys, the opportunity to periodically change the encryption key increases the offered level of security for any given implementation. However, where strong cryptography is employed, physical, procedural, and logical data access protection considerations may have more significant impact on cryptoperiod selection than algorithm and key size factors. Nevertheless, in general, shorter cryptoperiods, during which a single key is applied, will enhance security.

Given the mandatory protocol requirement to implement routing node authentication as part of routing integrity (see Section 8.2), key exchanges may be coordinated as part of the integrity verification process. This provides an opportunity to increase the frequency of key exchange and shorten the cryptoperiod as a complement to the key length and encryption algorithm required for a given application domain.

8.2. Integrity Features

The integrity of routing information provides the basis for ensuring that the function of the routing protocol is achieved and maintained. To protect integrity, RPL must either run using only the Secure versions of the messages, or must run over a layer-2 that uses channel binding between node identity and transmissions.

Some layer-2 security mechanisms use a single key for the entire network, and these networks can not provide significant amount of integrity protection, as any node that has that key may impersonate any other node. This mode of operation is likely acceptable when an entire deployment is under the control of a single administrative entity.

Other layer-2 security mechanisms form a unique session key for every pair of nodes that needs to communicate; this is often called a per-link key. Such networks can provide a strong degree of origin authentication and integrity on unicast messages.

However, some RPL messages are broadcast, and even when per-node layer-2 security mechanisms are used, the integrity and origin

authentication of broadcast messages can not be as trusted due to the proliferation of the key used to secure them.

RPL has two specific options which are broadcast in RPL Control Messages: the DODAG Information Object (DIO), and the DODAG Information Solicitation (DIS). The purpose of the DIS is to cause potential parents to reply with a DIO, so the integrity of the DIS is not of great concern. The DIS may also be unicast.

The DIO is a critical piece of routing and carries many critical parameters. RPL provides for asymmetric authentication at layer 3 of the RPL Control Message carrying the DIO and this may be warranted in some deployments. A node could, if it felt that the DIO that it had received was suspicious, send a unicast DIS message to the node in question, and that node would reply with a unicast DIS. Those messages could be protected with the per-link key.

8.3. Availability Features

Availability of routing information is linked to system and network availability which in the case of LLNs require a broader security view beyond the requirements of the routing entities. Where availability of the network is compromised, routing information availability will be accordingly affected. However, to specifically assist in protecting routing availability, nodes:

- o MAY restrict neighborhood cardinality;
- o MAY use multiple paths;
- o MAY use multiple destinations;
- o MAY choose randomly if multiple paths are available;
- o MAY set quotas to limit transmit or receive volume;
- o MAY use geographic information for flow control.

8.4. Key Management

The functioning of the routing security services requires keys and credentials. Therefore, even though not directly a RPL security requirement, an LLN MUST have a process for initial key and credential configuration, as well as secure storage within the associated devices. Anti-tampering SHOULD be a consideration in physical design. Beyond initial credential configuration, an LLN is also encouraged to have automatic procedures for the revocation and replacement of the maintained security credentials.

While RPL has secure modes, but some modes are impractical without use of public key cryptography believed to be too expensive by many. RPL layer-3 security will often depend upon existing LLN layer-2 security mechanisms, which provides for node authentication, but little in the way of node authorization.

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

The analysis presented in this document provides security analysis and design guidelines with a scope limited to RPL. Security services are identified as requirements for securing RPL. The specific mechanisms to be used to deal with each threat is specified in link-layer and deployment specific applicability statements.

11. Acknowledgments

The authors would like to acknowledge the review and comments from Rene Struik and JP Vasseur. The authors would also like to acknowledge the guidance and input provided by the RPL Chairs, David Culler, and JP Vasseur, and the Area Director Adrian Farrel.

This document started out as a combined threat and solutions document. As a result of a series of security reviews performed by Steve Kent, the document was split up by RPL co-Chair Michael Richardson and security Area Director Sean Turner as it went through the IETF publication process. The solutions to the threats are application and layer-2 specific, and have therefore been moved to the relevant applicability statements.

Ines Robles and Robert Cragie kept track of the many issues that were raised during the development of this document

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", BCP 107, RFC 4107, June 2005.
- [RFC6550] Winter, T., Thubert, P., 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, March 2012.

[RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, September 2012.

[RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, January 2014.

[ZigBeeIP]
ZigBee Public Document 15-002r00, "ZigBee IP Specification", 2013.

12.2. Informative References

[AceCharterProposal]
Li, Kepeng., Ed., "Authentication and Authorization for Constrained Environment Charter (work-in-progress)", December 2013, <http://trac.tools.ietf.org/wg/core/trac/wiki/ACE_charter>.

[I-D.gilger-smart-object-security-workshop]
Gilger, J. and H. Tschofenig, "Report from the 'Smart Object Security Workshop', March 23, 2012, Paris, France", draft-gilger-smart-object-security-workshop-02 (work in progress), October 2013.

[I-D.kelsey-intarea-mesh-link-establishment]
Kelsey, R., "Mesh Link Establishment", draft-kelsey-intarea-mesh-link-establishment-05 (work in progress), February 2013.

[I-D.suhopark-hello-wsn]
Park, S., "Routing Security in Sensor Network: HELLO Flood Attack and Defense", draft-suhopark-hello-wsn-00 (work in progress), December 2005.

[IEEE.802.11]
, "Draft Standard for Information Technology - Telecommunications and information exchange between systems - Local and metropolitan area networks Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications ", IEEE 802.11-REVma, 2006.

[IEEE.802.15.4]
, "Information technology - Telecommunications and information exchange between systems - Local and

metropolitan area networks - Specific requirements - Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANs) ", IEEE Std 802.15.4-2006, June 2006, <<http://standards.ieee.org/getieee802/802.15.html>>.

[ISO.7498-2.1988]

International Organization for Standardization,
"Information Processing Systems - Open Systems
Interconnection Reference Model - Security Architecture",
ISO Standard 7498-2, 1988.

[Karlof2003]

Karlof, C. and D. Wagner, "Secure routing in wireless sensor networks: attacks and countermeasures", Elsevier AdHoc Networks Journal, Special Issue on Sensor Network Applications and Protocols, 1(2):293-315, September 2003, <<http://nest.cs.berkeley.edu/papers/sensor-route-security.pdf>>.

[Myagmar2005]

Myagmar, S., Lee, A.J., and W. Yurcik, "Threat Modeling as a Basis for Security Requirements", in Proceedings of the Symposium on Requirements Engineering for Information Security (SREIS'05), Paris, France, pp. 94-102, Aug 29, 2005.

[Perlman1988]

Perlman, N., "Network Layer Protocols with Byzantine Robustness", MIT LCS Tech Report, 429, 1988.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, September 2003.

[RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, October 2006.

[RFC4732] Handley, M., Rescorla, E., IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.

[RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008.

- [RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, March 2008.
- [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.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [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.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, March 2011.
- [RFC6574] Tschofenig, H. and J. Arkko, "Report from the Smart Object Workshop", RFC 6574, April 2012.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.
- [RFC7142] Shand, M. and L. Ginsberg, "Reclassification of RFC 1142 to Historic", RFC 7142, February 2014.
- [SmartObjectSecurityWorkshop]
Klausen, T., Ed., "Workshop on Smart Object Security", March 2012, <<http://www.lix.polytechnique.fr/hipercom/SmartObjectSecurity>>.
- [SolaceProposal]
Bormann, C., Ed., "Notes from the SOLACE ad-hoc at IETF85 (work-in-progress)", November 2012, <<http://www.ietf.org/mail-archive/web/solace/current/msg00015.html>>.
- [Sybil2002]

Douceur, J., "The Sybil Attack", First International Workshop on Peer-to-Peer Systems , March 2002.

[Wan2004] Wan, T., Kranakis, E., and PC. van Oorschot, "S-RIP: A Secure Distance Vector Routing Protocol", in Proceedings of the 2nd International Conference on Applied Cryptography and Network Security, Yellow Mountain, China, pp. 103-119, Jun. 8-11 2004.

[Yourdon1979]
Yourdon, E. and L. Constantine, "Structured Design",
Yourdon Press, New York, Chapter 10, pp. 187-222, 1979.

Authors' Addresses

Tzeta Tsao
Cooper Power Systems
910 Clopper Rd. Suite 201S
Gaithersburg, Maryland 20878
USA

Email: tzeta.tsao@cooperindustries.com

Roger K. Alexander
Cooper Power Systems
910 Clopper Rd. Suite 201S
Gaithersburg, Maryland 20878
USA

Email: roger.alexander@cooperindustries.com

Mischa Dohler
CTTC
Parc Mediterrani de la Tecnologia, Av. Canal Olímpic S/N
Castelldefels, Barcelona 08860
Spain

Email: mischa.dohler@cttc.es

Vanesa Daza
Universitat Pompeu Fabra
P/ Circumval.lacio 8, Oficina 308
Barcelona 08003
Spain

Email: vanesa.daza@upf.edu

Angel Lozano
Universitat Pompeu Fabra
P/ Circumval.lacio 8, Oficina 309
Barcelona 08003
Spain

Email: angel.lozano@upf.edu

Michael Richardson (ed) (editor)
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z5V7
Canada

Email: mcr+ietf@sandelman.ca

Networking Working Group
Internet-Draft
Intended status: Informational
Expires: March 31, 2014

JP. Vasseur
Cisco Systems, Inc
September 30, 2013

Terms used in Routing for Low power And Lossy Networks
draft-ietf-roll-terminology-13.txt

Abstract

The documents provides a glossary of terminology used in routing requirements and solutions for networks referred to as Low power and Lossy Networks (LLN). An LLN is typically composed of many embedded devices with limited power, memory, and processing resources interconnected by a variety of links. There is a wide scope of application areas for LLNs, including industrial monitoring, building automation (e.g. Heating, Ventilating, Air Conditioning, lighting, access control, fire), connected home, healthcare, environmental monitoring, urban sensor networks, energy management, assets tracking, refrigeration.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 31, 2014.

Copyright Notice

Copyright (c) 2013 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	2
2. Terminology	3
3. IANA Considerations	6
4. Security Considerations	7
5. Acknowledgements	7
6. References	7
6.1. Normative References	7
6.2. Informative References	7
Author's Address	7

1. Introduction

The documents provides a glossary of terminology used in routing requirements solutions for networks referred to as Low power and Lossy Networks (LLN).

Low power and Lossy networks (LLNs) are typically composed of many embedded devices with limited power, memory, and processing resources interconnected by a variety of links, such as IEEE 802.15.4, Low Power WiFi. There is a wide scope of application areas for LLNs, including industrial monitoring, building automation (HVAC, lighting, access control, fire), connected home, healthcare, environmental monitoring, urban sensor networks, energy management, assets tracking and refrigeration.

Since these applications are usually highly specific (for example Industrial Automation, Building Automation, ...), it is not uncommon to see a number of disparate terms to describe the same device or functionality. Thus in order to avoid confusion or discrepancies, this document specifies the common terminology to be used in all ROLL Working Group documents. The terms defined in this document are used in [RFC5548],[RFC5673], [RFC5826] and [RFC5867].

Terminology specific to a particular application are out of the scope of this document.

It is expected that all routing requirements documents defining requirements or specifying routing solutions for LLN will use the common terminology specified in this document. This document should be listed as an informative reference.

2. Terminology

Actuator: a field device that controls a set of equipment. For example, an actuator might control and/or modulate the flow of a gas or liquid, control electricity distribution, perform a mechanical operation, ...

AMI: Advanced Metering Infrastructure that makes use of Smart Grid technologies. A canonical Smart Grid application is smart-metering.

Channel: Radio frequency sub-band used to transmit a modulated signal carrying packets.

Channel Hopping: A procedure by which field devices synchronously change channels during operation.

Commissioning Tool: Any physical or logical device temporarily added to the network for the express purpose of setting up the network and device operational parameters. The commissioning tool can also be temporarily added to the LLN for scheduled or unscheduled maintenance.

Closed Loop Control: A procedure whereby a device controller controls an actuator based on input information sensed by one or more field devices.

Controller: A field device that can receive sensor input and automatically change the environment in the facility by manipulating digital or analog actuators.

DA: Distribution Automation, part of Smart Grid. Encompasses technologies for maintenance and management of electrical distribution systems.

Directed Acyclic Graph: A directed graph with no directed cycles (a graph formed by a collection of vertices and directed edges where each edge connects one vertex to another, such that there is no way to start at some vertex *v* and follow a sequence of edges that eventually loops back to the edge *v* again)

Data sink: A device that collects data from nodes in an LLN.

Downstream: Data direction traveling from outside of the LLN (e.g. traffic coming from a LAN, WAN or the Internet) via a LBR, or in general "deeper" in the Directed Acyclic Graph computed by the routing protocol.

Field Device: A field device is a physical device placed in the network's operating environment (e.g. plant, urban or home). Field devices include sensors, actuators as well as routers and Low power and Lossy Network Border Router (LBR). A field device is usually (but not always) a device with constrained CPU, memory footprint, storage capacity, bandwidth and sometimes power (battery operated). At the time of writing, for the sake of illustration, a typical sensor or actuator would have a few KBytes of RAM, a few dozens of KBytes of ROM/Flash memory, a 8/16/32 bit microcontroller and communication capabilities ranging from a few Kbits/s to a few hundreds of Kbits/s. Although it is expected to see continuous improvements of hardware and software technologies, such devices will likely continue to be seen as resource constrained devices compared compared to computers and routers used in the rest of the Internet.

Flash memory: non-volatile memory that can be re-programmed.

FMS: Facility Management System. A global term applied across all the vertical designations within a building including, Heating, Ventilating, and Air Conditioning also referred to as HVAC, Fire, Security, Lighting and Elevator control.

HART: "Highway Addressable Remote Transducer", a group of specifications for industrial process and control devices administered by the HART Foundation (see [HART]). The latest version for the specifications is HART7 which includes the additions for WirelessHART.

HVAC: Heating, Ventilation and Air Conditioning. A term applied to the comfort level of an internal space.

ISA: "International Society of Automation". ISA is an ANSI accredited standards-making society. ISA100 is an ISA committee whose charter includes defining a family of standards for industrial automation. [ISA100.11a] is a working group within ISA100 that is working on a standard for monitoring and non-critical process control applications.

LAN: Local Area Network.

LBR: Low power and Lossy Network Border Router. The LBR is a device that connects the Low power and Lossy Network to another routing domain such as a Local Area Network (LAN), Wide Area Network (WAN) or

the Internet where a possibly different routing protocol is in operation. The LBR acts as a routing device and may possibly host other functions such as data collector or aggregator.

LLN: Low power and Lossy networks (LLNs) are typically composed of many embedded devices with limited power, memory, and processing resources interconnected by a variety of links, such as IEEE 802.15.4 or Low Power WiFi. There is a wide scope of application areas for LLNs, including industrial monitoring, building automation (HVAC, lighting, access control, fire), connected home, healthcare, environmental monitoring, urban sensor networks, energy management, assets tracking and refrigeration..

MP2P: Multipoint-to-Point is used to describe a particular traffic pattern (e.g. MP2P flows collecting information from many nodes flowing upstream towards a collecting sink or an LBR).

MAC: Medium Access Control. Refers to algorithms and procedures used by the data link layer to coordinate use of the physical layer.

Non-sleepy Node: A non-sleepy node is a node that always remains in a fully powered on state (i.e. always awake) where it has the capability to perform communication.

Open Loop Control: A process whereby a plant operator manually manipulates an actuator over the network where the decision is influenced by information sensed by field devices.

PER: Packet Error Rate. A ratio of the number of unusable packets (not received at all, or received in error- even after any applicable error correction has been applied) to the total number of packets that would have been received in the absence of errors.

P2P: Point To Point. This refers to traffic exchanged between two nodes (regardless of the number of hops between the two nodes).

P2MP: Point-to-Multipoint traffic refers to traffic between one node and a set of nodes. This is similar to the P2MP concept in Multicast or MPLS Traffic Engineering ([RFC4461]and [RFC4875]). A common RPL use case involves P2MP flows from or through a DAG root outward towards other nodes contained in the DAG.

RAM: Random Access Memory. The RAM is a volatile memory.

RFID: Radio Frequency IDentification.

ROM: Read Only Memory.

ROLL: Routing Over Low power and Lossy networks.

RPL: An IPv6 Routing Protocol for Low-Power and Lossy Networks that provides a mechanism whereby multipoint-to-point traffic from devices inside the LLN towards a central control point as well as point-to-multipoint traffic from the central control point to the devices inside the LLN are supported. RPL also support point-to-point traffic between any arbitratry node in the LLN.

RPL Domain: A RPL routing domain is a collection of RPL routers under the control of a single administration. The boundaries of routing domains are defined by network management by setting some links to be exterior, or inter-domain, links.

Schedule: An agreed execution, wake-up, transmission, reception, etc., time-table between two or more field devices.

Sensor: A sensor is a device that measures a physical quantity and converts it to an analog or digital signal that can be read by a program or a user. Sensed data can be of many types: electromagnetic (e.g. current, voltage, power, resistance, ...) , mechanical (e.g. pressure, flow, liquid density, humidity, ...), chemical (e.g. oxygen, carbon monoxide, ...), acoustic (e.g. noise, ultrasound), ...

Sleepy Node: A sleepy node is a node that may sometimes go into a sleep mode (i.e. go into a low power state to conserve power) and temporarily suspend protocol communication. When no in a sleep mode, the sleepy node is in a fully powered on state where it has the capability to perform communication.

Smart Grid: A Smart Grid is a broad class of applications to network and automate utility infrastructure.

Timeslot: A Timeslot is a fixed time interval that may be used for the transmission or reception of a packet between two field devices. A timeslot used for communications is associated with a slotted-link

Upstream: Data direction traveling from the LLN via the LBR to outside of the LLN (LAN, WAN, Internet) or general closer to the root of the Directed Acyclic Graph computed by the routing protocol.

WAN: Wide Area Network.

3. IANA Considerations

This document includes no request for IANA action.

4. Security Considerations

Since this document specifies terminology and does not specify new procedure or protocols, it raises no new security issue.

5. Acknowledgements

The authors would like to thank Christian Jacquenet, Tim Winter, Pieter De Mil, David Meyer, Mukul Goyal and Abdussalam Baryun for their valuable feed-back.

6. References

6.1. Informative References

- [HART] HART Communication Foundation (<http://www.hartcomm.org>)
- [RFC4461] Yasukawa, S., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.
- [RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [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. Vermeulen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, June 2010.

Author's Address

JP Vasseur
Cisco Systems, Inc
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Email: jpv@cisco.com

ROLL
Internet-Draft
Intended status: Standards Track
Expires: December 4, 2015

J. Hui
Nest Labs
R. Kelsey
Silicon Labs
June 2, 2015

Multicast Protocol for Low power and Lossy Networks (MPL)
draft-ietf-roll-trickle-mcast-12

Abstract

This document specifies the Multicast Protocol for Low power and Lossy Networks (MPL) that provides IPv6 multicast forwarding in constrained networks. MPL avoids the need to construct or maintain any multicast forwarding topology, disseminating messages to all MPL Forwarders in a MPL Domain.

MPL has two modes of operation. One mode uses the Trickle algorithm to manage control- and data-plane message transmissions, and is applicable for deployments with few multicast sources. The other mode uses classic flooding. By providing both modes and parameterization of the Trickle algorithm, a MPL implementation can be used in a variety of multicast deployments and can trade between dissemination latency and transmission efficiency.

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 December 4, 2015.

Copyright Notice

Copyright (c) 2015 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. Applicability Statement	5
4. MPL Protocol Overview	6
4.1. MPL Domains	6
4.2. Information Base Overview	7
4.3. Protocol Overview	7
4.4. Signaling Overview	9
5. MPL Parameters and Constants	9
5.1. MPL Multicast Addresses	9
5.2. MPL Message Types	10
5.3. MPL Seed Identifiers	10
5.4. MPL Parameters	10
6. Protocol Message Formats	12
6.1. MPL Option	12
6.2. MPL Control Message	14
6.3. MPL Seed Info	15
7. Information Base	16
7.1. Local Interface Set	16
7.2. Domain Set	16
7.3. Seed Set	16
7.4. Buffered Message Set	16
8. MPL Seed Sequence Numbers	17
9. MPL Data Messages	17
9.1. MPL Data Message Generation	17
9.2. MPL Data Message Transmission	18
9.3. MPL Data Message Processing	19
10. MPL Control Messages	20
10.1. MPL Control Message Generation	20
10.2. MPL Control Message Transmission	20
10.3. MPL Control Message Processing	21
11. Acknowledgements	22
12. IANA Considerations	22
12.1. MPL Option Type	22
12.2. MPL ICMPv6 Type	23
12.3. Well-known Multicast Addresses	23

13. Security Considerations	23
14. References	24
14.1. Normative References	24
14.2. Informative References	25
Authors' Addresses	26

1. Introduction

Low power and Lossy Networks (LLNs) typically operate with strict resource constraints in communication, computation, memory, and energy. Such resource constraints may preclude the use of existing IPv6 multicast routing and forwarding mechanisms. Traditional IP multicast delivery typically relies on topology maintenance mechanisms to discover and maintain routes to all subscribers of a multicast group (e.g. [RFC3973] [RFC4601]). However, maintaining such topologies in Low power and Lossy Networks is costly and may not be feasible given the available resources.

Memory constraints may limit devices to maintaining links/routes to one or a few neighbors. For this reason, the Routing Protocol for LLNs (RPL) specifies both storing and non-storing modes [RFC6550]. The latter allows RPL routers to maintain only one or a few default routes towards a LLN Border Router (LBR) and use source routing to forward messages away from the LBR. For the same reasons, a LLN device may not be able to maintain a multicast routing topology when operating with limited memory.

Furthermore, the dynamic properties of wireless networks can make the cost of maintaining a multicast routing topology prohibitively expensive. In wireless environments, topology maintenance may involve selecting a connected dominating set used to forward multicast messages to all nodes in an administrative domain. However, existing mechanisms often require two-hop topology information and the cost of maintaining such information grows polynomially with network density.

This document specifies the Multicast Protocol for Low power and Lossy Networks (MPL), which provides IPv6 multicast forwarding in constrained networks. MPL avoids the need to construct or maintain any multicast routing topology, disseminating multicast messages to all MPL Forwarders in a MPL Domain. By using the Trickle algorithm [RFC6206], MPL requires only small, constant state for each MPL device that initiates disseminations. The Trickle algorithm also allows MPL to be density-aware, allowing the communication rate to scale logarithmically with density.

2. Terminology

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

The following terms are used throughout this document:

MPL Forwarder	- A router that implements MPL. A MPL Forwarder is equipped with at least one MPL Interface.
MPL Interface	- A MPL Forwarder's attachment to a communications medium, over which it transmits and receives MPL Data Messages and MPL Control Messages according to this specification. A MPL Interface is assigned one or more unicast addresses and is subscribed to one or more MPL Domain Addresses.
MPL Domain Address	- A multicast address that identifies the set of MPL Interfaces within a MPL Domain. MPL Data Messages disseminated in a MPL Domain have the associated MPL Domain Address as their destination address.
MPL Domain	- A scope zone, as defined in [RFC4007], in which MPL Interfaces subscribe to the same MPL Domain Address and participate in disseminating MPL Data Messages.
MPL Data Message	- A multicast message that is used to communicate a multicast payload between MPL Forwarders within a MPL domain. A MPL Data Message contains a MPL Option in the IPv6 header and has as its destination address the MPL Domain Address corresponding to the MPL Domain.
MPL Control Message	- A link-local multicast message that is used to communicate information about recently received MPL Data Messages to neighboring MPL Forwarders.
MPL Seed	- A MPL Forwarder that generates MPL Data Messages and serves as an entry point into a MPL Domain.
MPL Seed Identifier	- An unsigned integer that uniquely identifies a MPL Seed within a MPL Domain.

Node - The term "node" is used within this document to refer to a MPL Forwarder.

3. Applicability Statement

MPL is an IPv6 multicast forwarding protocol designed for the communication characteristics and resource constraints of Low-Power and Lossy Networks. By implementing controlled disseminations of multicast messages using the Trickle algorithm, MPL is designed for networks that communicate using low-power and lossy links with widely varying topologies in both the space and time dimensions.

While designed specifically for Low-Power and Lossy Networks, MPL is not limited to use over such networks. MPL may be applicable to any network where no multicast routing state is desired. MPL may also be used in environments where only a subset of links are considered Low-Power and Lossy links.

A host need not be aware that their multicast is supported by MPL as long as its attachment router forwards multicast messages between the MPL Domain and the host. However, a host may choose to implement MPL so that it can take advantage of the broadcast medium inherent in many Low-Power and Lossy Networks and receive multicast messages carried by MPL directly.

MPL is parameterized to support different dissemination techniques. In one parameterization, MPL may utilize the classic flooding method that involves having each device receiving a message rebroadcast the message. In another parameterization, MPL may utilize Trickle's [RFC6206] "polite gossip" method that involves transmission suppression and adaptive timing techniques. [Clausen2013] questions the efficiency of Trickle's "polite gossip" mechanism in some multicast scenarios, so by also including a classic flooding mode of operation MPL aims to be able to perform satisfactorily in a variety of situations.

To support efficient message delivery in networks that have many poor links, MPL supports a reactive forwarding mode that utilizes MPL Control Messages to summarize the current multicast state. The MPL Control Message size grows linearly with the number of simultaneous MPL Seeds in the MPL Domain - 4 octets per MPL Seed. When reactive forwarding is not enabled, MPL Control Messages are not transmitted and the associated overhead is not incurred.

This document does not specify a cryptographic security mechanism for MPL to ensure that MPL messages are not spoofed by anyone with access to the LLN. In general, the basic ability to inject messages into a Low-power and Lossy Network may be used as a denial-of-service attack

regardless of what forwarding protocol is used. For these reasons, Low-power and Lossy Networks typically employ link-layer security mechanisms to mitigate an attacker's ability to inject messages. For example, the IEEE 802.15.4 [IEEE802154] standard specifies frame security mechanisms using AES-128 to support access control, message integrity, message confidentiality, and replay protection. However, if the attack vector includes attackers that have access to the LLN, then MPL SHOULD NOT be used.

4. MPL Protocol Overview

The goal of MPL is to deliver multicast messages to all interfaces that subscribe to the multicast messages' destination address within a MPL Domain.

4.1. MPL Domains

A MPL Domain is a scope zone, as defined in [RFC4007], in which MPL Interfaces subscribe to the same MPL Domain Address and participate in disseminating MPL Data Messages.

When participating in only one MPL Domain, the MPL Domain Address is the ALL_MPL_FORWARDERS multicast address with Realm-Local scope (scop value 3) [RFC7346].

When a MPL Forwarder participates in multiple MPL Domains simultaneously, at most one MPL Domain may be assigned a MPL Domain Address equal to the ALL_MPL_FORWARDERS multicast address. All other MPL Domains MUST be assigned a unique MPL Domain Address that allows the MPL Forwarder to identify each MPL Domain. The MPL Domains SHOULD be configured automatically based on some underlying topology. For example, when using RPL [RFC6550], MPL Domains may be configured based on RPL Instances.

When MPL is used in deployments that use administratively defined scopes that cover, for example, multiple subnets based on different underlying network technologies, Admin-Local scope (scop value 4) or Site-Local scope (scop value 5) SHOULD be used.

A MPL Forwarder MAY participate in additional MPL Domains identified by other multicast addresses. A MPL Interface MUST subscribe to the MPL Domain Addresses for the MPL Domains that it participates in. The assignment of other multicast addresses is out of scope.

For each MPL Domain Address that a MPL Interface subscribes to, the MPL Interface MUST also subscribe to the same MPL Domain Address with Link-Local scope (scop value 2) when reactive forwarding is in use (i.e. when communicating MPL Control Messages).

4.2. Information Base Overview

A node records necessary protocol state in the following information sets:

- o The Local Interface Set records the set of local MPL Interfaces and the unicast addresses assigned to those MPL Interfaces.
- o The Domain Set records the set of MPL Domain Addresses and the local MPL Interfaces that subscribe to those addresses.
- o A Seed Set records information about received MPL Data Messages received from a MPL Seed within a MPL Domain. Each MPL Domain has an associated Seed Set. A Seed Set maintains the minimum sequence number for MPL Data Messages that the MPL Forwarder is willing to receive or has buffered in its Buffered Message Set from a MPL Seed. MPL uses Seed Sets and Buffered Message Sets to determine when to accept a MPL Data Message, process its payload, and retransmit it.
- o A Buffered Message Set records recently received MPL Data Messages from a MPL Seed within a MPL Domain. Each MPL Domain has an associated Buffered Message Set. MPL Data Messages resident in a Buffered Message Set have sequence numbers that are greater than or equal to the minimum threshold maintained in the corresponding Seed Set. MPL uses Buffered Message Sets to store MPL Data Messages that may be transmitted by the MPL Forwarder for forwarding.

4.3. Protocol Overview

MPL achieves its goal by implementing a controlled flood that attempts to disseminate the multicast data message to all interfaces within a MPL Domain. MPL performs the following tasks to disseminate a multicast message:

- o When having a multicast message to forward into a MPL Domain, the MPL Seed generates a MPL Data Message that includes the MPL Domain Address as the IPv6 Destination Address, the MPL Seed Identifier, a newly generated sequence number, and the multicast message. If the multicast destination address is not the MPL Domain Address, IP-in-IP [RFC2473] is used to encapsulate the multicast message in a MPL Data Message, preserving the original IPv6 Destination Address.
- o Upon receiving a MPL Data Message, the MPL Forwarder extracts the MPL Seed and sequence number and determines whether or not the MPL

Data Message was previously received using the MPL Domain's Seed Set and Buffered Message Set.

- * If the sequence number is less than the lower-bound sequence number maintained in the Seed Set or a message with the same sequence number exists within the Buffered Message Set, the MPL Forwarder marks the MPL Data Message as old.
- * Otherwise, the MPL Forwarder marks the MPL Data Message as new.
- o For each newly received MPL Data Message, a MPL Forwarder updates the Seed Set, adds the MPL Data Message into the Buffered Message Set, processes its payload, and multicasts the MPL Data Message a number of times on all MPL Interfaces participating in the same MPL Domain to forward the message.
- o Each MPL Forwarder may periodically link-local multicast MPL Control Messages on MPL Interfaces to communicate information contained in a MPL Domain's Seed Set and Buffered Message Set.
- o Upon receiving a MPL Control Message, a MPL Forwarder determines whether there are any new MPL Data Messages that have yet to be received by the MPL Control Message's source and multicasts those MPL Data Messages.

MPL's configuration parameters allow two forwarding strategies for disseminating MPL Data Messages via MPL Interfaces.

Proactive Forwarding - With proactive forwarding, a MPL Forwarder schedules transmissions of MPL Data Messages using the Trickle algorithm, without any prior indication that neighboring nodes have yet to receive the message. After transmitting the MPL Data Message a limited number of times, the MPL Forwarder may terminate proactive forwarding for the MPL Data Message.

Reactive Forwarding - With reactive forwarding, a MPL Forwarder link-local multicasts MPL Control Messages using the Trickle algorithm [RFC6206]. MPL Forwarders use MPL Control Messages to discover new MPL Data Messages that have not yet been received. When discovering that a neighboring MPL Forwarder has not yet received a MPL Data Message, the MPL Forwarder schedules those MPL Data Messages for transmission using the Trickle algorithm.

Note that the use of proactive and reactive forwarding strategies within the same MPL Domain are not mutually exclusive and may be used simultaneously. For example, upon receiving a new MPL Data Message when both proactive and reactive forwarding techniques are enabled, a MPL Forwarder will proactively retransmit the MPL Data Message a

limited number of times and schedule further transmissions upon receiving MPL Control Messages.

4.4. Signaling Overview

MPL generates and processes the following messages:

MPL Data Message - Generated by a MPL Seed to deliver a multicast message across a MPL Domain. The MPL Data Message's source is an address in the Local Interface Set of the MPL Seed that generated the message and is valid within the MPL Domain. The MPL Data Message's destination is the MPL Domain Address corresponding to the MPL Domain. A MPL Data Message contains:

- * The Seed Identifier of the MPL Seed that generated the MPL Data Message.
- * The sequence number of the MPL Seed that generated the MPL Data Message.
- * The original multicast message.

MPL Control Message - Generated by a MPL Forwarder to communicate information contained in a MPL Domain's Seed Set and Buffered Message Set to neighboring MPL Forwarders. A MPL Control Message contains a list of tuples for each entry in the Seed Set. Each tuple contains:

- * The minimum sequence number maintained in the Seed Set for the MPL Seed.
- * A bit-vector indicating the sequence numbers of MPL Data Messages resident in the Buffered Message Set for the MPL Seed, where the first bit represents a sequence number equal to the minimum threshold maintained in the Seed Set.
- * The length of the bit-vector.

5. MPL Parameters and Constants

This section describes various program and networking parameters and constants used by MPL.

5.1. MPL Multicast Addresses

MPL makes use of MPL Domain Addresses to identify MPL Interfaces of a MPL Domain. By default, MPL Forwarders subscribe to the

ALL_MPL_FORWARDERS multicast address with Realm-Local scope (scop value 3) [RFC7346].

For each MPL Domain Address that a MPL Interface subscribes to, the MPL Interface MUST also subscribe to the MPL Domain Address with Link-Local scope (scop value 2) when reactive forwarding is in use. MPL Forwarders use the link-scoped MPL Domain Address to communicate MPL Control Messages to neighboring (i.e. on-link) MPL Forwarders.

5.2. MPL Message Types

MPL defines an IPv6 Option for carrying a MPL Seed Identifier and a sequence number within a MPL Data Message. The IPv6 Option Type has value 0x6D.

MPL defines an ICMPv6 Message (MPL Control Message) for communicating information contained in a MPL Domain's Seed Set and Buffered Message Set to neighboring MPL Forwarders. The MPL Control Message has ICMPv6 Type MPL_ICMP_TYPE.

5.3. MPL Seed Identifiers

MPL uses MPL Seed Identifiers to uniquely identify MPL Seeds within a MPL Domain. For each MPL Domain that the MPL Forwarder serves as a MPL Seed, the MPL Forwarder MUST have an associated MPL Seed Identifier. A MPL Forwarder MAY use the same MPL Seed Identifier across multiple MPL Domains, but the MPL Seed Identifier MUST be unique within each MPL Domain. The mechanism for assigning and verifying uniqueness of MPL Seed Identifiers is not specified in this document.

5.4. MPL Parameters

PROACTIVE_FORWARDING A boolean value that indicates whether the MPL Forwarder schedules MPL Data Message transmissions after receiving them for the first time. PROACTIVE_FORWARDING has a default value of TRUE. All MPL interfaces on the same link SHOULD be configured with the same value of PROACTIVE_FORWARDING. An implementation MAY choose to vary the value of PROACTIVE_FORWARDING across interfaces on the same link if reactive forwarding is also in use. The mechanism for setting PROACTIVE_FORWARDING is not specified within this document.

SEED_SET_ENTRY_LIFETIME The minimum lifetime for an entry in the Seed Set. SEED_SET_ENTRY_LIFETIME has a default value of 30 minutes. It is RECOMMENDED that all MPL Forwarders use the same value for SEED_SET_ENTRY_LIFETIME for a given MPL Domain and use a default value of 30 minutes. Using a value of

SEED_SET_ENTRY_LIFETIME that is too small can cause the duplicate detection mechanism to fail, resulting in a MPL Forwarder to receive a given MPL Data Message more than once. The mechanism for setting SEED_SET_ENTRY_LIFETIME is not specified within this document.

As specified in [RFC6206], a Trickle timer runs for a defined interval and has three configuration parameters: the minimum interval size Imin, the maximum interval size Imax, and a redundancy constant k.

This specification defines a fourth Trickle configuration parameter, TimerExpirations, which indicates the number of Trickle timer expiration events that occur before terminating the Trickle algorithm for a given MPL Data Message or MPL Control Message.

Each MPL Interface uses the following Trickle parameters for MPL Data Message and MPL Control Message transmissions.

DATA_MESSAGE_IMIN The minimum Trickle timer interval, as defined in [RFC6206], for MPL Data Message transmissions. DATA_MESSAGE_IMIN has a default value of 10 times the expected link-layer latency.

DATA_MESSAGE_IMAX The maximum Trickle timer interval, as defined in [RFC6206], for MPL Data Message transmissions. DATA_MESSAGE_IMAX has a default value equal to DATA_MESSAGE_IMIN.

DATA_MESSAGE_K The redundancy constant, as defined in [RFC6206], for MPL Data Message transmissions. DATA_MESSAGE_K has a default value of 1.

DATA_MESSAGE_TIMER_EXPIRATIONS The number of Trickle timer expirations that occur before terminating the Trickle algorithm's retransmission of a given MPL Data Message. DATA_MESSAGE_TIMER_EXPIRATIONS has a default value of 3.

CONTROL_MESSAGE_IMIN The minimum Trickle timer interval, as defined in [RFC6206], for MPL Control Message transmissions. CONTROL_MESSAGE_IMIN has a default value of 10 times the worst-case link-layer latency.

CONTROL_MESSAGE_IMAX The maximum Trickle timer interval, as defined in [RFC6206], for MPL Control Message transmissions. CONTROL_MESSAGE_IMAX has a default value of 5 minutes.

CONTROL_MESSAGE_K The redundancy constant, as defined in [RFC6206], for MPL Control Message transmissions. CONTROL_MESSAGE_K has a default value of 1.

CONTROL_MESSAGE_TIMER_EXPIRATIONS The number of Trickle timer expirations that occur before terminating the Trickle algorithm for MPL Control Message transmissions.

CONTROL_MESSAGE_TIMER_EXPIRATIONS has a default value of 10.

As described in [RFC6206], if different nodes have different configuration parameters, Trickle may have unintended behaviors. Therefore, it is RECOMMENDED that all MPL Interfaces attached to the same link of a given MPL Domain use the same values for the Trickle Parameters above for a given MPL Domain. The mechanism for setting the Trickle Parameters is not specified within this document.

The default MPL parameters specify a forwarding strategy that utilizes both proactive and reactive techniques. Using these default values, a MPL Forwarder proactively transmits any new MPL Data Messages it receives then uses MPL Control Messages to trigger additional MPL Data Message retransmissions where message drops are detected. Setting DATA_MESSAGE_IMAX to the same as DATA_MESSAGE_IMIN in this case is acceptable since subsequent MPL Data Message retransmissions are triggered by MPL Control Messages, where CONTROL_MESSAGE_IMAX is greater than CONTROL_MESSAGE_IMIN.

6. Protocol Message Formats

Messages generated and processed by a MPL Forwarder are described in this section.

6.1. MPL Option

The MPL Option is carried in MPL Data Messages in an IPv6 Hop-by-Hop Options header, immediately following the IPv6 header. The MPL Option has the following format:

```

      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
      +-----+-----+-----+-----+-----+-----+-----+-----+
      |                                         Option Type | Opt Data Len |
      +-----+-----+-----+-----+-----+-----+-----+-----+
      | S |M|V|  rsv |   sequence   |   seed-id (optional)   |
      +-----+-----+-----+-----+-----+-----+-----+-----+

```

Option Type 0x6D.

Opt Data Len Length of the Option Data field in octets.

S 2-bit unsigned integer. Identifies the length of seed-id. 0 indicates that the seed-id is the

	IPv6 Source Address and not included in the MPL Option. 1 indicates that the seed-id is a 16-bit unsigned integer. 2 indicates that the seed-id is a 64-bit unsigned integer. 3 indicates that the seed-id is a 128-bit unsigned integer.
M	1-bit flag. 1 indicates that the value in sequence is known to be the largest sequence number that was received from the MPL Seed.
V	1-bit flag. 0 indicates that the MPL Option conforms to this specification. MPL Data Messages with a MPL Option in which this flag is 1 MUST be dropped.
rsv	4-bit reserved field. MUST be set to 0 on transmission and ignored on reception.
sequence	8-bit unsigned integer. Identifies relative ordering of MPL Data Messages from the MPL Seed identified by seed-id.
seed-id	Uniquely identifies the MPL Seed that initiated dissemination of the MPL Data Message. The size of seed-id is indicated by the S field.

The Option Data (specifically the M flag) of the MPL Option is updated by MPL Forwarders as the MPL Data Message is forwarded. Nodes that do not understand the MPL Option MUST discard the MPL Data Message. Thus, according to [RFC2460] the three high order bits of the Option Type are set to '011'. The Option Data length is variable.

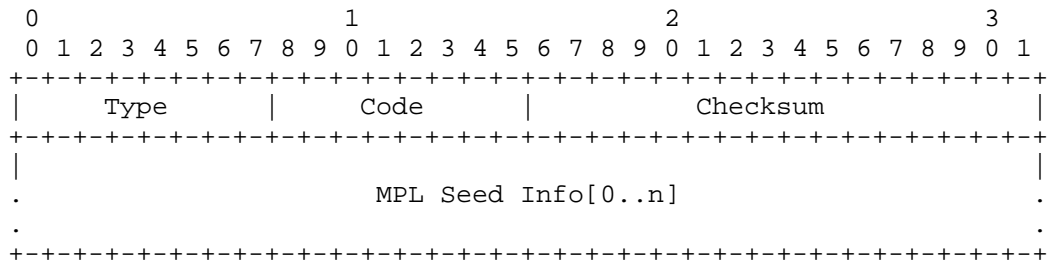
The seed-id uniquely identifies a MPL Seed. When seed-id is 128 bits (S=3), the MPL Seed MAY use an IPv6 address assigned to one of its interfaces that is unique within the MPL Domain. Managing MPL Seed Identifiers is not within scope of this document.

The sequence field establishes a total ordering of MPL Data Messages generated by a MPL Seed for a MPL Domain. The MPL Seed MUST increment the sequence field's value on each new MPL Data Message that it generates for a MPL Domain. Implementations MUST follow the Serial Number Arithmetic as defined in [RFC1982] when incrementing a sequence value or comparing two sequence values.

Future updates to this specification may define additional fields following the seed-id field.

6.2. MPL Control Message

A MPL Forwarder uses ICMPv6 messages to communicate information contained in a MPL Domain's Seed Set and Buffered Message Set to neighboring MPL Forwarders. The MPL Control Message has the following format:



IP Fields:

Source Address An IPv6 address in the AddressSet of the corresponding MPL Interface and MUST be valid within the MPL Domain.

Destination Address The link-scoped MPL Domain Address corresponding to the MPL Domain.

Hop Limit 255

ICMPv6 Fields:

Type MPL_ICMP_TYPE

Code 0

Checksum The ICMP checksum. See [RFC4443].

MPL Seed Info[0..n] List of zero or more MPL Seed Info entries.

The MPL Control Message indicates the sequence numbers of MPL Data Messages that are within the MPL Domain's Buffered Message Set. The MPL Control Message also indicates the sequence numbers of MPL Data Messages that a MPL Forwarder is willing to receive. The MPL Control Message allows neighboring MPL Forwarders to determine whether there are any new MPL Data Messages to exchange.

6.3. MPL Seed Info

A MPL Seed Info encodes the minimum sequence number for an MPL Seed maintained in the MPL Domain's Seed Set. The MPL Seed Info also indicates the sequence numbers of MPL Data Messages generated by the MPL Seed that are stored within the MPL Domain's Buffered Message Set. The MPL Seed Info has the following format:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
| min-seqno | bm-len | S | seed-id (0/2/8/16 octets) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|
| buffered-mpl-messages (variable length)
|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

min-seqno 8-bit unsigned integer. The lower-bound sequence number for the MPL Seed.

bm-len 6-bit unsigned integer. The size of buffered-mpl-messages in octets.

S 2-bit unsigned integer. Identifies the length of seed-id. 0 indicates that the seed-id value is the IPv6 Source Address and not included in the MPL Seed Info. 1 indicates that the seed-id value is a 16-bit unsigned integer. 2 indicates that the seed-id value is a 64-bit unsigned integer. 3 indicates that the seed-id is a 128-bit unsigned integer.

seed-id Variable-length unsigned integer. Indicates the MPL Seed associated with this MPL Seed Info.

buffered-mpl-messages Variable-length bit vector. Identifies the sequence numbers of MPL Data Messages maintained in the corresponding Buffered Message Set for the MPL Seed. The *i*'th bit represents a sequence number of min-seqno + *i*. '0' indicates that the corresponding MPL Data Message does not exist in the Buffered Message Set. '1' indicates that the corresponding MPL Data Message does exist in the Buffered Message Set.

The MPL Seed Info does not have any octet alignment requirement.

7. Information Base

7.1. Local Interface Set

The Local Interface Set records the local MPL Interfaces of a MPL Forwarder. The Local Interface Set consists of Local Interface Tuples, one per MPL Interface: (AddressSet).

AddressSet - a set of unicast addresses assigned to the MPL Interface.

7.2. Domain Set

The Domain Set records the MPL Interfaces that subscribe to each MPL Domain Address. The Domain Set consists of MPL Domain Tuples, one per MPL Domain: (MPLInterfaceSet).

MPLInterfaceSet - a set of MPL Interfaces that subscribe to the MPL Domain Address that identifies the MPL Domain.

7.3. Seed Set

A Seed Set records a sliding window used to determine the sequence numbers of MPL Data Messages that a MPL Forwarder is willing to accept generated by the MPL Seed. A MPL Forwarder maintains a Seed Set for each MPL Domain that it participates in. A Seed Set consists of MPL Seed Tuples: (SeedID, MinSequence, Lifetime).

SeedID - the identifier for the MPL Seed.

MinSequence - a lower-bound sequence number that represents the sequence number of the oldest MPL Data Message the MPL Forwarder is willing to receive or transmit. A MPL Forwarder MUST ignore any MPL Data Message that has sequence value less than than MinSequence.

Lifetime - indicates the minimum remaining lifetime of the Seed Set entry. A MPL Forwarder MUST NOT free a Seed Set entry before the remaining lifetime expires.

7.4. Buffered Message Set

A Buffered Message Set records recently received MPL Data Messages from a MPL Seed within a MPL Domain. A MPL Forwarder uses a Buffered Message Set to buffer MPL Data Messages while the MPL Forwarder is forwarding the MPL Data Messages. A MPL Forwarder maintains a Buffered Message Set for each MPL Domain that it participates in. A

Buffered Message Set consists of Buffered Message Tuples: (SeedID, SequenceNumber, DataMessage).

SeedID - the identifier for the MPL Seed that generated the MPL Data Message.

SequenceNumber - the sequence number for the MPL Data Message.

DataMessage - the MPL Data Message.

All MPL Data Messages within a Buffered Message Set MUST have a sequence number greater than or equal to MinSequence for the corresponding SeedID. When increasing MinSequence for a MPL Seed, the MPL Forwarder MUST delete any MPL Data Messages from the corresponding Buffered Message Set that have sequence numbers less than MinSequence.

8. MPL Seed Sequence Numbers

Each MPL Seed maintains a sequence number for each MPL Domain that it serves. The sequence numbers are included in MPL Data Messages generated by the MPL Seed. The MPL Seed MUST increment the sequence number for each MPL Data Message that it generates for a MPL Domain. Implementations MUST follow the Serial Number Arithmetic as defined in [RFC1982] when incrementing a sequence value or comparing two sequence values. This sequence number is used to establish a total ordering of MPL Data Messages generated by a MPL Seed for a MPL Domain.

9. MPL Data Messages

9.1. MPL Data Message Generation

MPL Data Messages are generated by MPL Seeds when these messages enter the MPL Domain. All MPL Data messages have the following properties:

- o The IPv6 Source Address MUST be an address in the AddressSet of a corresponding MPL Interface and MUST be valid within the MPL Domain.
- o The IPv6 Destination Address MUST be set to the MPL Domain Address corresponding to the MPL Domain.
- o A MPL Data Message MUST contain a MPL Option in its IPv6 Header to identify the MPL Seed that generated the message and the ordering relative to other MPL Data Messages generated by the MPL Seed.

When the destination address is a MPL Domain Address and the source address is in the AddressList of a MPL Interface that belongs to that MPL Domain Address, the application message and the MPL Data Message MAY be identical. In other words, the MPL Data Message may contain a single IPv6 header that includes the MPL Option.

Otherwise, IPv6-in-IPv6 encapsulation MUST be used to satisfy the MPL Data Message requirements listed above [RFC2473]. The complete IPv6-in-IPv6 message forms a MPL Data Message. The outer IPv6 header conforms to the MPL Data Message requirements listed above. The encapsulated IPv6 datagram encodes the multicast data message that is communicated beyond the MPL Domain.

9.2. MPL Data Message Transmission

A MPL Forwarder manages transmission of MPL Data Messages in its Buffered Message Sets using the Trickle algorithm [RFC6206]. A MPL Forwarder MUST use a separate Trickle timer for each MPL Data Message that it is actively forwarding. In accordance with Section 5 of RFC 6206 [RFC6206], this document defines the following:

- o This document defines a "consistent" transmission as receiving a MPL Data Message that has the same MPL Domain Address, seed-id, and sequence value as the MPL Data Message managed by the Trickle timer.
- o This document defines an "inconsistent" transmission as receiving a MPL Data Message that has the same MPL Domain Address, seed-id value, and the M flag set, but has a sequence value less than MPL Data Message managed by the Trickle timer.
- o This document does not define any external "events".
- o This document defines MPL Data Messages as Trickle messages.
- o The actions outside the Trickle algorithm that MPL takes involve managing the MPL Domain's Seed Set and Buffered Message Set.

As specified in [RFC6206], a Trickle timer has three variables: the current interval size *I*, a time within the current interval *t*, and a counter *c*. MPL defines a fourth variable, *e*, which counts the number of Trickle timer expiration events since the Trickle timer was last reset.

After DATA_MESSAGE_TIMER_EXPIRATIONS Trickle timer events, the MPL Forwarder MUST disable the Trickle timer. When a buffered MPL Data Message does not have an associated Trickle timer, the MPL Forwarder MAY delete the message from the Buffered Message Set by advancing

MinSequence of the corresponding MPL Seed in the Seed Set. When the MPL Forwarder no longer buffers any messages for a MPL Seed, the MPL Forwarder MUST NOT increment MinSequence for that MPL Seed.

When transmitting a MPL Data Message, the MPL Forwarder MUST either set the M flag to zero or set it to a level that indicates whether or not the message's sequence number is the largest value that has been received from the MPL Seed.

9.3. MPL Data Message Processing

Upon receiving a MPL Data Message, the MPL Forwarder first processes the MPL Option and updates the Trickle timer associated with the MPL Data Message if one exists.

Upon receiving a MPL Data Message, a MPL Forwarder MUST perform one of the following actions:

- o Accept the message and enter the MPL Data Message in the MPL Domain's Buffered Message Set.
- o Accept the message and update the corresponding MinSequence in the MPL Domain's Seed Set to 1 greater than the message's sequence number.
- o Discard the message without any change to the MPL Information Base.

If a Seed Set entry exists for the MPL Seed, the MPL Forwarder MUST discard the MPL Data Message if its sequence number is less than MinSequence or exists in the Buffered Message Set.

If a Seed Set entry does not exist for the MPL Seed, the MPL Forwarder MUST create a new entry for the MPL Seed before accepting the MPL Data Message.

If memory is limited, a MPL Forwarder SHOULD reclaim memory resources by:

- o Incrementing MinSequence entries in a Seed Set and deleting MPL Data Messages in the corresponding Buffered Message Set that fall below the MinSequence value.
- o Deleting other Seed Set entries that have expired and the corresponding MPL Data Messages in the Buffered Message Set.

If the MPL Forwarder accepts the MPL Data Message, the MPL Forwarder MUST perform the following actions:

- o Reset the Lifetime of the corresponding Seed Set entry to SEED_SET_ENTRY_LIFETIME.
- o If PROACTIVE_FORWARDING is true, the MPL Forwarder MUST initialize and start a Trickle timer for the MPL Data Message.
- o If the MPL Control Message Trickle timer is not running and CONTROL_MESSAGE_TIMER_EXPIRATIONS is non-zero, the MPL Forwarder MUST initialize and start the MPL Control Message Trickle timer.
- o If the MPL Control Message Trickle timer is running, the MPL Forwarder MUST reset the MPL Control Message Trickle timer.

10. MPL Control Messages

10.1. MPL Control Message Generation

A MPL Forwarder generates MPL Control Messages to communicate a MPL Domain's Seed Set and Buffered Message Set to neighboring MPL Forwarders. Each MPL Control Message is generated according to Section 6.2, with a MPL Seed Info for each entry in the MPL Domain's Seed Set. Each MPL Seed Info entry has the following content:

- o S set to the size of the seed-id field in the MPL Seed Info entry.
- o min-seqno set to MinSequence of the MPL Seed.
- o bm-len set to the size of buffered-mpl-messages in octets.
- o seed-id set to the MPL seed identifier.
- o buffered-mpl-messages with each bit representing whether or not a MPL Data Message with the corresponding sequence number exists in the Buffered Message Set. The i'th bit represents a sequence number of min-seqno + i. '0' indicates that the corresponding MPL Data Message does not exist in the Buffered Message Set. '1' indicates that the corresponding MPL Data Message does exist in the Buffered Message Set.

10.2. MPL Control Message Transmission

A MPL Forwarder transmits MPL Control Messages using the Trickle algorithm. A MPL Forwarder maintains a single Trickle timer for each MPL Domain. When CONTROL_MESSAGE_TIMER_EXPIRATIONS is 0, the MPL Forwarder does not execute the Trickle algorithm and does not transmit MPL Control Messages. In accordance with Section 5 of RFC 6206 [RFC6206], this document defines the following:

- o This document defines a "consistent" transmission as receiving a MPL Control Message that results in a determination that neither the receiving nor transmitting node has any new MPL Data Messages to offer.
- o This document defines an "inconsistent" transmission as receiving a MPL Control Message that results in a determination that either the receiving or transmitting node has at least one new MPL Data Message to offer.
- o The Trickle timer is reset in response to external "events." This document defines an "event" as increasing MinSequence of any entry in the corresponding Seed Set or adding a message to the corresponding Buffered Message Set.
- o This document defines a MPL Control Message as a Trickle message.

As specified in [RFC6206], a Trickle timer has three variables: the current interval size I , a time within the current interval t , and a counter c . MPL defines a fourth variable, e , which counts the number of Trickle timer expiration events since the Trickle timer was last reset. After CONTROL_MESSAGE_TIMER_EXPIRATIONS Trickle timer events, the MPL Forwarder MUST disable the Trickle timer.

10.3. MPL Control Message Processing

A MPL Forwarder processes each MPL Control Message that it receives to determine if it has any new MPL Data Messages to receive or offer.

A MPL Forwarder determines if a new MPL Data Message has not been received from a neighboring node if any of the following conditions hold true:

- o The MPL Control Message includes a MPL Seed that does not exist in the MPL Domain's Seed Set.
- o The MPL Control Message indicates that the neighbor has a MPL Data Message in its Buffered Message Set with sequence number greater than MinSequence (i.e. the i -th bit is set to 1 and $\text{min-seqno} + i > \text{MinSequence}$) and is not included in the MPL Domain's Buffered Message Set.

When a MPL Forwarder determines that it has not yet received a MPL Data Message buffered by a neighboring device, the MPL Forwarder MUST reset its Trickle timer associated with MPL Control Message transmissions. If a MPL Control Message Trickle timer is not running, the MPL Forwarder MUST initialize and start a new Trickle timer.

A MPL Forwarder determines if a MPL Data Message in the Buffered Message Set has not yet been received by a neighboring MPL Forwarder if any of the following conditions hold true:

- o The MPL Control Message does not include a MPL Seed for the MPL Data Message.
- o The MPL Data Message's sequence number is greater than or equal to min-seqno and not included in the neighbor's corresponding Buffered Message Set (i.e. the MPL Data Message's sequence number does not have a corresponding bit in buffered-mpl-messages set to 1).

When a MPL Forwarder determines that it has at least one MPL Data Message in its corresponding Buffered Message Set that has not yet been received by a neighbor, the MPL Forwarder MUST reset the MPL Control Message Trickle timer. Additionally, for each of those entries in the Buffered Message Set, the MPL Forwarder MUST reset the Trickle timer and reset e to 0. If a Trickle timer is not associated with the MPL Data Message, the MPL Forwarder MUST initialize and start a new Trickle timer.

11. Acknowledgements

The authors would like to acknowledge the helpful comments of Robert Cragie, Esko Dijk, Ralph Droms, Paul Duffy, Adrian Farrel, Ulrich Herberg, Owen Kirby, Philip Levis, Kerry Lynn, Joseph Reddy, Michael Richardson, Ines Robles, Don Sturek, Dario Tedeschi, and Peter van der Stok, which greatly improved the document.

12. IANA Considerations

This document defines one IPv6 Option, a type that must be allocated from the IPv6 "Destination Options and Hop-by-Hop Options" registry of [RFC2780].

This document defines one ICMPv6 Message, a type that must be allocated from the "ICMPv6 "type" Numbers" registry of [RFC4443].

This document registers a well-known multicast address from the Variable Scope Multicast Address registry.

12.1. MPL Option Type

IANA is requested to allocate an IPv6 Option Type from the IPv6 "Destination Options and Hop-by-Hop Options" registry of [RFC2780], as specified in Table 1 below:

Hex Value	act	chg	rest	Description	Reference
0x6D	01	1	01101	MPL Option	This Document

Table 1: IPv6 Option Type Allocation

12.2. MPL ICMPv6 Type

IANA is requested to allocate an ICMPv6 Type from the "ICMPv6 "type" Numbers" registry of [RFC4443], as specified in Table 2 below:

Type	Name	Reference
TBD	MPL Control Message	This Document

Table 2: IPv6 Option Type Allocation

In this document, the mnemonic MPL_ICMP_TYPE was used to refer to the ICMPv6 Type above, which is TBD by IANA.

12.3. Well-known Multicast Addresses

IANA is requested to allocate an IPv6 multicast address, with Group ID in the range [0x01,0xFF] for 6LoWPAN compression [RFC6282], "ALL_MPL_FORWARDERS" from the "Variable Scope Multicast Addresses" sub-registry of the "IPv6 Multicast Address Space" registry [RFC3307] as specified in Table 3 below:

Address(s)	Description	Reference	Date Registered
FF0X:0:0:0:0:0:FC	ALL_MPL_FORWARDERS	This Document	2013-04-10

Table 3: Variable Scope Multicast Address Allocation

13. Security Considerations

MPL uses sequence numbers to maintain a total ordering of MPL Data Messages from a MPL Seed. The use of sequence numbers allows a denial-of-service attack where an attacker can spoof a message with a sufficiently large sequence number to: (i) flush messages from the

Buffered Message List and (ii) increase the MinSequence value for a MPL Seed in the corresponding Seed Set. In both cases, the side effect allows an attacker to halt the forwarding process of any MPL Data Messages being disseminated and prevents MPL Forwarders from accepting new MPL Data Messages that a MPL Seed generates while the sequence number is less than MinSequence or until the corresponding Seed Set Entry expires. The net effect applies to both proactive and reactive forwarding modes.

In general, the basic ability to inject messages into a Low-power and Lossy Network may be used as a denial-of-service attack regardless of what forwarding protocol is used. Because MPL is a dissemination protocol, the ability to spoof MPL messages allows an attacker to affect an entire MPL Domain. For these reasons, Low-power and Lossy Networks typically employ link-layer security mechanisms to mitigate an attacker's ability to inject messages. For example, the IEEE 802.15.4 [IEEE802154] standard specifies frame security mechanisms using AES-128 to support access control, message integrity, message confidentiality, and replay protection. However, if the attack vector includes attackers that have access to the LLN, then MPL SHOULD NOT be used.

To prevent attackers from injecting packets through a MPL Forwarder, the MPL Forwarder MUST NOT accept or forward MPL Data Messages from a communication interface that does not subscribe to the MPL Domain Address identified in message's destination address.

MPL uses the Trickle algorithm to manage message transmissions and the security considerations described in [RFC6206] apply.

14. References

14.1. Normative References

- [RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.

- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.
- [RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.
- [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, March 2005.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, March 2011.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, September 2011.
- [RFC6550] Winter, T., Thubert, P., 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, March 2012.
- [RFC7346] Droms, R., "IPv6 Multicast Address Scopes", RFC 7346, August 2014.

14.2. Informative References

- [Clausen2013] Clausen, T., Colin de Verdiere, A., and J. Yi, "Performance Analysis of Trickle as a Flooding Mechanism", The 5th IEEE International Conference on Communication Technology (ICCT2013), November 2013.
- [IEEE802154] "IEEE Std. 802.15.4-2006", October 2006.
- [RFC3973] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

Authors' Addresses

Jonathan W. Hui
Nest Labs
3400 Hillview Ave
Palo Alto, California 94304
USA

Phone: +650 253 2770
Email: jonhui@nestlabs.com

Richard Kelsey
Silicon Labs
25 Thomson Place
Boston, Massachusetts 02210
USA

Phone: +617 951 1225
Email: richard.kelsey@silabs.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: April 24, 2014

P. Roux
A. Petrescu
M. Kellil
CEA
October 21, 2013

Preliminary results about MPL performance evaluation
draft-roux-roll-mpl-eval-00.txt

Abstract

This draft presents simulation work and first results related to MPL performance evaluation. The simulation makes it possible to evaluate MPL performances in the context of a large network. The simulated network introduces 500 nodes. The general principles of the simulator are described. Then reference settings are introduced, and evaluation indicators are proposed. Finally preliminary results are presented under the form of a few tables, that show the proposed indicator values depending on some specific parameter which is used as a variable argument. Among various results, the advantage of using reactive mode for MPL is shown in terms of the capability to maintain loss free diffusion in harsh radio conditions.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 24, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Simulated Network Layout	3
4. Simulator Principle	4
5. Radio Simulation Aspects	4
6. Simulation Conditions	5
7. Types of Results	6
7.1. Preliminary Results with Respect to Achievable Data Rate	8
7.2. Preliminary results with respect to the influence of redundancy constants	8
7.3. Preliminary results with respect to the influence of transmit poser	9
8. Security Considerations	10
9. IANA Considerations	10
10. Acknowledgements	10
11. Normative References	10
Appendix A. ChangeLog	10
Authors' Addresses	10

1. Introduction

The RPL protocol is an IPv6 protocol for low-power and lossy networks, defined in [RFC6550].

The [I-D.ietf-roll-trickle-mcast] draft introduces the so called MPL algorithm, with a lot of freedom in terms of possible configuration. The current draft is a preliminary attempt to provide guidance for finding good algorithm settings for MPL.

Simulation makes it possible to address algorithmic capabilities in terms of scalability. A network made of 500 nodes is considered in this document. The proposed objective is to assess performance of the MPL protocol in such large networks, and to compare various MPL settings, in order to understand their influence on performances, so that setting guidelines may be derived.


However, the presented results are still not mature enough to propose solid and motivated guidelines. This draft should be considered as a preliminary attempt in this direction. Depending on the feedback, a more complete set of results may be presented in the coming months.

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 RFC 2119 [RFC2119].

3. Simulated Network Layout

A reference network layout has been assumed for all simulations. It is composed of 500 nodes randomly distributed within a rectangular area, excluding a void sub-area in its center, as shown in Figure 1. The main area width is 20 km and the height is 10 km. The void sub-area is 2 km and its height is 5 km. The layout is built by dropping nodes randomly in the main area. Once a potential node has been dropped at a random location in the main area, it is retained only if it is not located in the void sub-area, and if its closest neighbour among previously defined nodes, is not closer than 20 meters. This process continues until 500 nodes has been retained in the area. Then the closest node to the point of coordinates (x=2km,y=5km) is selected as the seed node.



4. Simulator Principle

The simulator elementary time step has been set in such a way that it takes 4 elementary steps for a sender to send the longest messages, which are supposed not to exceed 128 bytes. The simulator doesn't address anything below this time step. As an example the simulator doesn't actually fill messages with bits.

At each elementary step, the simulator spans each node in the network several times in order to:

- o Check data trickle timers and control trickle timer (if one exists).
- o Treat incoming messages.
- o Take care of (data or control) message transmission (start, continue or complete a message transmission).

5. Radio Simulation Aspects

As said above, only a fragment of a message can be transmitted during an elementary time step of the simulator. The term of fragment, as used in this document, refers to the part of a message which can be transmitted during an elementary simulator step time.

The radio coverage radius is set to 100m, which means that beyond this distance, no fragment can be transmitted between 2 nodes. Below this distance, there is a probability for a transmitted fragment to be received which depends on the simulated power received at the receiver.

The static attenuation in dB is assumed to be equal to $50 + 30 \cdot \log_{10}(d/50)$, with d in meters, which means that we assume 3 as the

path loss exponent. A log normal shadowing with 3 dB as standard deviation is added as a dynamical attenuation. Dynamical attenuation spectrum is limited with a 3 Hz cut off frequency.

A fragment being sent may not be received (or corrupted) at the receiver depending on a probability which is bound to the received level at the receiver at the receiving time. A predefined table defines the probability for a fragment to be received depending on the received power.

For a message to be transmitted between a sender and a receiver it is necessary that all segments involved in the message are received uncorrupted. Otherwise, the message is assumed to be lost (because of the UDP checksum mechanism, it is assumed that a message received corrupted is also a non delivered UDP message). If any segment has been lost then the message is lost as well.

Collisions are simulated as well: if a receiver receives fragments from multiples neighbour nodes under its coverage, then a collision may occur for the received fragment. For a message to be delivered, it is necessary that no collision occurs on any fragment of this message. A collision condition occurs if the received power from a first neighbour is interfered with the received power from one or a set of other neighbours which a total interfering power which is equal at least to half the useful received power from the first neighbour.

6. Simulation Conditions

When not explicitly stated, preliminary results are obtained with the following parameter values:

Parameter	Value
Maximum number of messages in buffer message set	10
Seed message rate	1 per each 4 seconds period
Number of messages sent during one simulation	200
DATA_MESSAGE_IMIN	1.28 second (128 simul. steps)
DATA_MESSAGE_IMAX	10 (max interval = 21 minutes)
DATA_MESSAGE_K	1
DATA_MESSAGE_TIMER_EXPIRATIONS	No expiration
CONTROL_MESSAGE_IMIN	128
CONTROL_MESSAGE_IMAX	10
CONTROL_MESSAGE_K	1
Transmission power	5 mW

Default values assumed for simulation

Table 1

This set of default values should not be seen as a proposed configuration which would be optimum in some sense. It is just meant as a reference point to explore configuration space. A future contribution may propose a more optimized reference configuration, once systematic simulations will have been run.

7. Types of Results

The following performance indicators are exploited:

Indicator	Explanation
Proactive, or reactive data loss ratio	Data message loss ratio observed in each node after simulation, and averaged across all nodes in the network.
Proactive, or reactive data expansion	Total number of data messages transmitted from any node during simulation, divided by the number of nodes in the network, and divided again by the number of data messages sent from the seed.
Reactive control expansion	Total number of control messages transmitted from any node during simulation, divided by the number of nodes in the network, and divided again by the number of data messages sent from the seed.
Reactive expansion	Sum of the reactive data expansion plus the reactive control expansion

Performance indicators

Table 2

With respect to expansion figures, there is no claim about any generic properties for the results which are presented in this document. Given the present definitions, it is expected that the denser is the network (in terms of average number of neighbours per node) the lower expansions rates will be. Therefore expansion rates are not only depending on the routing algorithmic aspects, they also depend on the network layout. Their interest in the context of this document is to compare different configuration settings, rather than to obtain generic performance indicators.

7.1. Preliminary Results with Respect to Achievable Data Rate

Data message generation period at the seed	0.5s	1s	2s	4s
Proactive data loss ratio	0.38	0.15	0.03	0
Proactive data expansion	0.37	0.85	1.37	1.79
Reactive data loss ratio	0.38	0.2	0.02	0
Reactive data expansion	0.36	0.91	1.58	1.93
Reactive control expansion	0.1	0.22	0.44	0.72
Reactive total expansion	0.46	1.13	2.02	2.65

Preliminary results with respect to achievable data rates

Table 3

Table 3 provides results for several message rates at the seed. It can be seen that a saturation phenomenon is observed (introducing significant message loss ratios) when the message rate is too high. Of course, this result is strongly related to the DATA_MESSAGE_IMIN and CONTROL_MESSAGE_IMIN parameters, which have both been set in this case to 1.28 s.

The other observation seems to be that the reactive mode introduces more transmissions in the network (higher total expansion), with no real benefit in terms of achievable message rate given the constraints of negligible message loss ratios.

7.2. Preliminary results with respect to the influence of redundancy constants

DATA_MESSAGE_K	1	2	4
Proactive data expansion	1.79	2.95	4.4

Data expansion in proactive mode, versus redundancy constant

Table 4

DATA_MESSAGE_K	1	2	4
Reactive data expansion	1.93	2.88	4.18
Reactive control expansion	0.72	0.7	0.72
Reactive total expansion	2.65	3.58	4.9

Data and control expansion in reactive mode, versus redundancy constants

Table 5

Table 4 shows the effect of redundancy constant of trickle timers on data expansion, in case of proactive mode. The other configuration parameters are set as described in the simulation conditions section. In particular, the message rate is 1 message for 4s, so that there are not message losses in any simulation. With no surprise, it can be observed that the redundancy constant has a strong influence on expansion ratio.

Table 5 shows the same result with respect to proactive mode. It shows that expansion figures are higher when reactive mode is used, as compared to proactive mode. The influence of introducing vlues that are different of 1 for CONTROL_MESSAGE_K will be studied later.

7.3. Preliminary results with respect to the influence of transmit power

Tx power (mW)	1	2	3	4	5
Proactive data loss ratio	0.56	0.14	0.03	0.01	0
Proactive data expansion	0.8	1.85	1.95	1.88	1.8
Reactive data loss ratio	0.4	0.07	0.01	0	0
Reactive data expansion	1.61	2.55	2.29	2.07	1.92
Reactive control expansion	0.68	0.88	0.81	0.76	0.72
Reactive total expansion	2.29	3.43	3.1	2.83	2.64

Results versus transmit power

Table 6

Table 6 shows the effect of varying the transmit power.

It shows the interest of using reactive mode has it is able to achieve lower message loss ratios in harsh radio environments. Of

course this result is obtained at the cost of higher expansion rates.

8. Security Considerations

9. IANA Considerations

10. Acknowledgements

This work has been supported by the A2NETS project (Autonomic Services in M2M Networks) which is funded by the ITEA 2 program.

11. Normative References

- [I-D.ietf-roll-trickle-mcast]
Hui, J. and R. Kelsey, "Multicast Protocol for Low power and Lossy Networks (MPL)",
draft-ietf-roll-trickle-mcast-05 (work in progress),
August 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC6550] Winter, T., Thubert, P., 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, March 2012.

Appendix A. ChangeLog

The changes are listed in reverse chronological order, most recent changes appearing at the top of the list.

From draft-roux-roll-mpl-eval-00.txt to
draft-authors-ipv6-over-80211p-00.txt:

- o first version.

Authors' Addresses

Pierre Roux
CEA
<http://www.cea.fr>,
France

Phone:
Email: Pierre.Roux@cea.fr

Alexandru Petrescu
CEA
<http://www.cea.fr>,
France

Phone:
Email: Alexandru.Petrescu@cea.fr

Mounir Kellil
CEA
<http://www.cea.fr>,
France

Phone:
Email: Mounir.Kellil@cea.fr

