RPL applicability in industrial networks
draft-ietf-roll-rpl-industrial-applicability-01

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 March 11, 2014.

Copyright Notice

Table of Contents

1. Introduction

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

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

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

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 Reading

1.3. 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 those
   requirements.

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

      P1: Construction or major modification phase
      P2: Planned startup phase
      P3: Normal operation phase
      P4: Planned shutdown phase

   followed eventually by an (at least theoretical)

      P5: Plant decommissioning phase.

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

      P6: Post-emergency recovery and repair phase.

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

   Extra LLN devices and even extra LLN subnets may be employed during
   Planned startup (P2) and Planned shutdown (P4) phases, in support of
   the task of transitioning the plant or plant area between operational
   and shutdown states. The extra devices typically provide extra
   monitoring as the plant transitions infrequent activity states. (In
   many continuous process plants, up to 2x extra staff are employed at
   monitoring and control workstations during these two phases,
   precisely because the plant is undergoing extraordinary behavior as

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

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

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

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

Additionally, it could be envisioned to include RPL in the normal
operation provided that a new Objective Function is defined that
actually interacts with the PCE is 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
   * Class 1: Closed loop regulatory control - Often a critical
      function
   * Class 2: Closed loop supervisory control - Usually non-critical
      function
   * Class 3: Open loop control - Operator takes action and controls
      the actuator (human in the loop)
o Monitoring
   * Class 4: Alerting - Short-term operational effect (for example
      event-based maintenance)
   * Class 5: Logging and downloading / uploading - No immediate
      operational consequence (e.g., history collection, sequence-of-
      events, preventive maintenance)

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

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

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

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

It appears from the above sections that whether and the way RPL can

   be applied for a given flow depends both on the deployment scenario
   and on the class of application / traffic. At a high level, this can
   be summarized by the following matrix:

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

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

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.

```
                  ---+------------------------
                     |            Plant Network
                     |
               +-----+
               |     | Gateway
               |     |
```

```
              +-----+
                 |
                 |       Backbone
        +--------------------+------------------+
        |                    |                  |
     +-----+             +-----+            +-----+
     |     | LLN border  |     | LLN border |     | LLN border
  o  |     | router      |     | router     |     | router
     +-----+             +-----+            +-----+
    o                o                    o                    o
       o    o    o         o   o   o    o         o   o   o  o
                               LLN
```
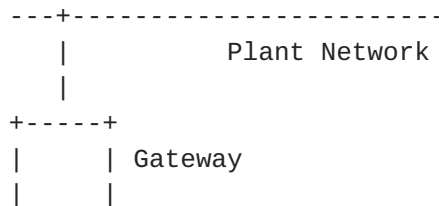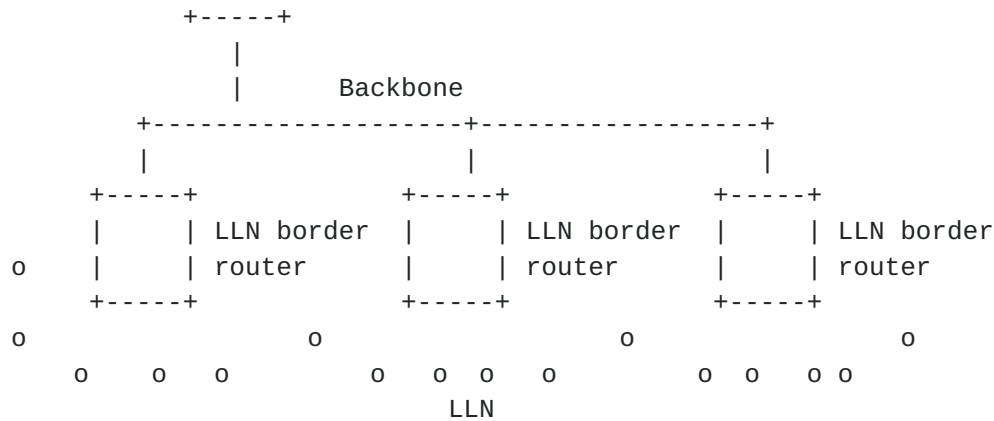
  o : stationary wireless field device, seldom acting as an LLN router
  For factory automation networks, the basic communications cycle for
  control is typically much faster, on the order of 100 Hz or more. In
  this case the LLN itself may be based on high-data-rate IEEE 802.11
  or a 100 Mbit/s or faster optical link, and the higher-rate network
  used by the LBRs to connect the LLN to superior automation equipment
  typically might be based on fiber-optic IEEE 802.3, with multiple
  LBRs around the periphery of the factory area, so that most high-rate
  communications again requires only a single wireless LLN hop.
  Multi-hop LLN routing is used within the LLN portion of such networks
  to provide backup communications paths when primary single-hop LLN
  paths fail, or for lower repetition rate communications where longer
  LLN transit times and higher variance are not an issue. Typically,
  the majority of devices in an IACS can tolerate such higher-delay
  higher-variance paths, so routing choices often are driven by energy
  considerations for the affected devices, rather than simply by IACS
  performance requirements, as illustrated in Figure 3.

```
              ---+------------------------
                 |               Plant Network
                 |
           +-----+
           |     | Gateway
           |     |
           +-----+
              |
              |          Backbone
        +--------------------+------------------+
        |                    |                  |
     +-----+              +-----+            +-----+
     |     | Backbone     |     | Backbone   |     | Backbone
     |     | router       |     | router     |     | router
     +-----+              +-----+            +-----+
        o     o   o    o      o   o  o    o    o   o  o   o o
      o o   o  o      o o  o o    o  o  o    o    o   o  o  o o
     o  o o  o o    o     o    o  o  o  o    M    o  o  o o o
     o   o M o  o  o      o  o    o  o  o    o  o    o  o   o
       o    o o        o        o  o         o       o o
           o              o           o           o      o
                              LLN
```
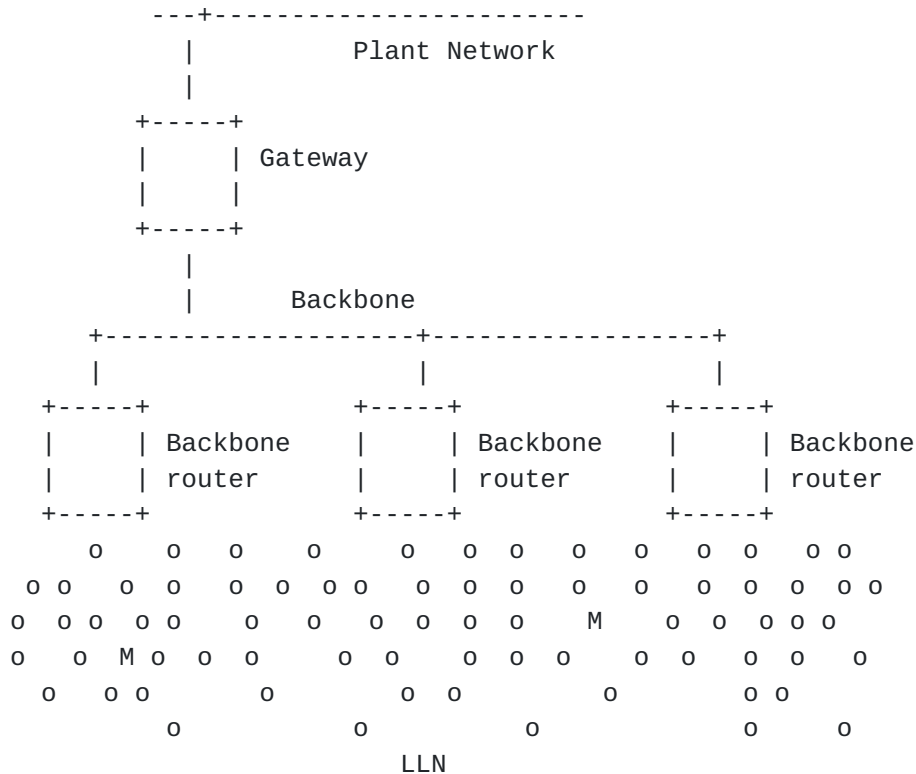
  o : stationary wireless field device, often acting as an LLN router
  M : mobile wireless device
Two decades of experience with digital fieldbuses has shown that four
communications paradigms dominate in IACS:
SS: Source-sink
PS: Publish-subscribe
P2P: Peer-to-peer
P2MP: Peer-to-multipeer

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
multicast messages if their work assignment covers process unit N,
and not otherwise.

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

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

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

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

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

Timeliness of message delivery is a critical aspect of PS
communication. Individual messages can be lost without significant
impact on the controlled physical process, but typically a sequence
of four consecutive lost messages will trigger fallback behavior of
the control algorithms, which is considered a system failure by most
system owner/operators. (In general, and unless a local catastrophic
event such as a major explosion or a tornado occurs in the plant,
invocation of more than one instance of such fallback handling per
year, per plant, is considered unacceptable.)
Message loss, delay and jitter in delivery of PS messaging is a
relative matter. PS messaging is used for transfer of process
measurements and associated status from sensors to control
computation elements, from control computation elements to actuators,
and of current commanded position and status from actuators back to
control computation elements. The actual time interval of interest
is that which starts with sensing of the physical process (which
necessarily occurs before the sensed value can be sent in the first
message) and which ends when the computed control correction is
applied to the physical process by the appropriate actuator (which
cannot occur until after the second message containing the computed
control output has been received by that actuator). With rare
exception, the control algorithms used with PS messaging in the
process automation industries - those managing continuous material
flows - rely on fixed-period sampling, computation and transfer of
outputs, while those in the factory automation industries - those
managing discrete manufacturing operations - rely on bounded delay
between sampling of inputs, control computation and transfer of
outputs to physical actuators that affect the controlled process.
Deliberately manipulated message delay and jitter in delivery of PS
messaging has the potential to destabilize control loops. It is the
responsibility of conveyed higher-level protocols to protect against
such potential security attacks by detecting overly delayed or
jittered messages at delivery, converting them into instances of
message loss. Thus network and data-link protocols such as IPv6 and
Ethernet need not themselves address such issues, although their
selection and employment should take the existence (or lack) of such
higher-layer protection mechanisms, and the resulting consequences
due to excessive delay and jitter, into consideration in their
parameterization.

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

In some system architectures, message streams that use PS to convey
current process measurements and status are compressed at the source
through a 2-dimensional winnowing process that compares
1) the process measurement values and status of the about-to-be-sent
   message with that of the last actually-sent message, and
2) the current time vs. the queueing time for the last actually-sent
   message.
If the interval since that last-sent message is less than a
predefined maximum time, and the status is unchanged, and the process
measurement(s) conveyed in the message is within predefined
deadband(s) of the last-sent measurement value(s), then transmission
of the new message is suppressed. Often this suppression takes the
form of not queuing the new message for transmission, but in some
protocols a brief placeholder message indicating "no significant
change" is queued in its stead.

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.

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

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

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

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

Note: Reliable group download protocols, such as the no-longer-
published IEEE 802.1E (ISO/IEC 15802-4) system load protocol, and
reliable multicast protocols based on the guidance of [RFC2887], are
instructive in how P2MP can be used for initial bulk download,
followed by either P2MP or P2P selective retransmissions for missed
download segments.

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

   Hz loops to be supported in large wireless systems.
   When available, duocast/N-cast typically is used only for one-hop PS
   traffic on Class 1 and Class 0 control loops. It may also be
   employed for rapid, reliable one-hop delivery of Class 0 and
   sometimes Class 1 process alarms and device alerts, which use the SS
   paradigm. Because it requires scheduling of multiple receivers that
   are prepared to acknowledge the received message during the
   transaction, in general it is not appropriate for the other types of
   traffic in such systems - P2P and P2MP - and is not needed for other
   classes of control loops or other types of traffic, which do not have
   such stringent reporting requirements.
   Note: Although there are known patent applications for duocast and
   N-cast, at the time of this writing the patent assignee, Honeywell
   International, has offered to permit cost-free RAND use in those
   industrial wireless standards that have chosen to employee 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.
      Storing mode with multicast support: defined in [RFC6550], Mode-of-
                                            operation (MOP) of 3. This mode
                                            improves MOP 2 by adding the
                                            capability to register multicast
                                            groups and perform multicast
                                            forwarding along the instance
                                            DODAG (or a spanning subtree
                                            within the DODAG).
      Reactive: defined in [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:

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

2.2. Layer 2 applicability.
   To be completed.
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:

   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, including all three
of the following situations:

1: <to be added>
2: <to be added>
3: <to be added>

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.

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 tradeoffs 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
   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 Computaton 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-roll-forwarding-frags]
      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 higly 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 reliabiliy 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 Computaton Engine
      present in the System Manager. This is due to the deterministic,
      centralized nature of wireless LLNs.

Autonomous instance: An autonomous instance requires limited to no
    configuration. It, 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
    autonomic 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 construcstion of the autonomous instance do not need
    approval from the Path Computation Engine present in the 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 packets
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.

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
fielde 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 accomodate 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
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 wireles 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.

   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. Multicast

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

   <to be added>

4.2. Layer-two features

4.2.1. Need layer-2 expert here.

4.2.2. Security functions provided by 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

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.

DIOIntervalDoublings: Industrial automation network deployments SHOULD set DIOIntervalDoublings such that the Trickle Imax 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 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.

4.3.2. **Other Parameters**

<to be added>

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.

6. **Security Considerations**

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. 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 deployment

6.2. Security Considerations during incremental deployment

7. Other Related Protocols

8. IANA Considerations

This specification has no requirement on IANA.

9. Acknowledgements

10. References

10.1. 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]

                Vasseur, J., "Terminology in Low power And Lossy
                Networks", Internet-Draft draft-ietf-roll-terminology-12,
                March 2013.
   [RFC2887]   Handley, M., Floyd, S., Whetten, B., Kermode, R.,
                Vicisano, L. and M. Luby, "The Reliable Multicast Design
                Space for Bulk Data Transfer", RFC 2887, August 2000.
   [RFC5548]   Dohler, M., Watteyne, T., Winter, T. and D. Barthel,
                "Routing Requirements for Urban Low-Power and Lossy
                Networks", RFC 5548, May 2009.
   [RFC5826]   Brandt, A., Buron, J. and G. Porcu, "Home Automation
                Routing Requirements in Low-Power and Lossy Networks", RFC
                5826, April 2010.
   [RFC5867]   Martocci, J., De Mil, P., Riou, N. and W. Vermeylen,
                "Building Automation Routing Requirements in Low-Power and
                Lossy Networks", RFC 5867, June 2010.
   [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.
   [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.
   [RFC6552]   Thubert, P., "Objective Function Zero for the Routing
                Protocol for Low-Power and Lossy Networks (RPL)", RFC
                6552, March 2012.
   [RFC6719]   Gnawali, O. and P. Levis, "The Minimum Rank with
                Hysteresis Objective Function", RFC 6719, September 2012.
   [RFC6997]   Goyal, M., Baccelli, E., Philipp, M., Brandt, A. and J.
                Martocci, "Reactive Discovery of Point-to-Point Routes in
                Low-Power and Lossy Networks", RFC 6997, August 2013.
   [I-D.thubert-roll-asymlink]
                Thubert, P., "RPL adaptation for asymmetrical links",
                Internet-Draft draft-thubert-roll-asymlink-02, December
                2011.
   [I-D.thubert-roll-forwarding-frags]
                Thubert, P. and J. Hui, "LLN Fragment Forwarding and
                Recovery", Internet-Draft draft-thubert-roll-forwarding-
                frags-01, February 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.

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>.

Authors' Addresses

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

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

Pascal Thubert
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
MOUGINS - Sophia Antipolis, 06254
FRANCE

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

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

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