

ROLL
Internet-Draft
Intended status: Informational
Expires: April 3, 2012

T. Phinney, Ed.
consultant
P. Thubert
Cisco
RA. Assimiti
Nivis
October 1, 2011

RPL applicability in industrial networks
draft-phinney-roll-rpl-industrial-applicability-00

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, while simultaneously increasing 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 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 3, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal

Internet-Draft RPL-industrial-applicability-statement October 2011

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.

Internet-Draft RPL-industrial-applicability-statement October 2011

Table of Contents

1.	Introduction	4
2.	Terminology	6
3.	Overview	7
3.1.	Deployment scenarii	7
3.2.	Applications and Traffic classes	9
3.3.	RPL applicability matrix	10
4.	Characterization of communication flows in IACS wireless networks	11
4.1.	General	11
4.2.	Source-sink (SS) communication paradigm	13
4.3.	Publish-subscribe (PS, or pub/sub) communication paradigm	13
4.4.	Peer-to-peer (P2P) communication paradigm	15
4.5.	Peer-to-multiper (P2MP) communication paradigm	16
4.6.	Additional considerations: Duocast and N-cast	17
4.7.	RPL applicability per communication paradigm	18
5.	RPL profile	21
5.1.	Use for process control	21
5.2.	RPL features	21
5.2.1.	Storing vs. non-storing mode	21
5.2.2.	DAO policy	21
5.2.3.	Path metrics	22
5.2.4.	Objective functions	22
5.2.5.	DODAG repair	22
5.2.6.	Security	22
5.3.	RPL options	22
5.4.	Recommended configuration defaults and ranges	22
5.4.1.	Trickle parameters	22
5.4.2.	Other parameters	23
5.4.3.	Additional configuration recommendations	23
6.	Other related protocols	24
7.	Manageability	25
8.	IANA considerations	26
9.	Security considerations	27

10.	Acknowledgements	28
11.	References	29
11.1.	Normative References	29
11.2.	Informative References	29
11.3.	External Informative References	30
	Authors' Addresses	31

Internet-Draft RPL-industrial-applicability-statement October 2011

[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), even ISA100.11a leaves the routing operation to proprietary methods.

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) [[I-D.ietf-roll-rpl](#)] specification and its point to point extension/optimization [[I-D.ietf-roll-p2p-rpl](#)] 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](#)].

Internet-Draft RPL-industrial-applicability-statement October 2011

[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 [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)

[3.](#) Overview

[3.1.](#) Deployment scenarios

[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 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.

[3.2.](#) Applications and Traffic classes

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

Internet-Draft RPL-industrial-applicability-statement October 2011

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.

3.3. RPL applicability matrix

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


Figure 1: RPL applicability matrix

Internet-Draft RPL-industrial-applicability-statement October 2011

4. Characterization of communication flows in IACS wireless networks

4.1. General

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.




o : stationary wireless field device, seldom acting as an LLN router


Figure 2: High-rate low-delay low-variance IACS topology

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.





- o : stationary wireless field device, often acting as an LLN router
- M : mobile wireless device

Figure 3: Low-rate higher-delay higher-variance IACS topology

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-multiple

4.2. 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.

4.3. 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.

[4.4.](#) 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.

[4.5.](#) Peer-to-multiple (P2MP) communication paradigm

In P2MP, the peer-to-multiple 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-

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.

[4.6](#). 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 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).

[4.7.](#) 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 [[I-D.ietf-roll-rpl](#)], 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 [[I-D.ietf-roll-rpl](#)], 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 [[I-D.ietf-roll-rpl](#)], 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 [[I-D.ietf-roll-rpl](#)] [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 [[I-D.ietf-roll-rpl](#)], 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 [[I-D.ietf-roll-p2p-rpl](#)], 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)

+-----+-----+-----+

Figure 4: RPL applicability per communication paradigm

Phinney, et al.

Expires April 3, 2012

[Page 20]

Internet-Draft RPL-industrial-applicability-statement

October 2011

[5.](#) RPL profile

[5.1.](#) Use for process control

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.

[5.2.](#) RPL features

[5.2.1.](#) Storing vs. non-storing mode

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 safety and non-safety control and monitoring traffic.

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.

In general, devices which communicate with LBRs through a chain of intermediary devices will use storing mode for their upward DODAGs, but will use non-storing mode for downward DODAGs for messaging that they route further into the LLN. However, routers that provide downward forwarding for PS messaging addressed to controllers within the LLN (which is expected to be a rare occurrence) will use storing mode for those forwarding paths, so that timely, destination-constrained forwarding of such recurring messaging does not overload the routing node(s) and their downstream subnets.

[5.2.2.](#) DAO policy

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

<to be added>

[5.2.3.](#) 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, OF0 and MRHOF, both of which define the selection of a preferred parent and backup parents, and are suitable for industrial automation network deployments.

Neither of the currently defined objective functions supports multiple metrics that might be required in heterogeneous industrial automation networks (e.g., networks composed of devices with different energy and timeliness-of-communication constraints). Additional objective functions specifically designed for such networks may be defined in companion RFCs.

[5.2.4.](#) Objective functions

<to be added>

[5.2.5.](#) DODAG repair

[5.2.6.](#) 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.

[5.3.](#) RPL options

[5.4.](#) Recommended configuration defaults and ranges

[5.4.1.](#) 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.

<to be added>

[5.4.2.](#) Other parameters

<to be added>

[5.4.3.](#) Additional configuration recommendations

<to be added>

Internet-Draft RPL-industrial-applicability-statement

October 2011

[6.](#) Other related protocols

<to be added>

Internet-Draft RPL-industrial-applicability-statement October 2011

7. Manageability

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.

<to be added>

[8.](#) IANA considerations

This specification has no requirement on IANA.

9. Security considerations

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.

[10](#). Acknowledgements

<to be added>

[11.](#) References

[11.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

11.2. Informative References

- [I-D.ietf-roll-of0]
Thubert, P., "RPL Objective Function Zero",
[draft-ietf-roll-of0-20](#) (work in progress), September 2011.
- [I-D.ietf-roll-p2p-rpl]
Goyal, M., Baccelli, E., Philipp, M., Brandt, A., Cragie, R., and J. Martocci, "Reactive Discovery of Point-to-Point Routes in Low Power and Lossy Networks",
[draft-ietf-roll-p2p-rpl-04](#) (work in progress), July 2011.
- [I-D.ietf-roll-rpl]
Winter, T., Thubert, P., Brandt, A., Clausen, T., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., and J. Vasseur, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", [draft-ietf-roll-rpl-19](#) (work in progress), March 2011.
- [I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", [draft-ietf-roll-terminology-06](#) (work in progress), September 2011.
- [I-D.tsao-roll-security-framework]
Tsao, T., Alexander, R., Daza, V., and A. Lozano, "A Security Framework for Routing over Low Power and Lossy Networks", [draft-tsao-roll-security-framework-02](#) (work in progress), March 2010.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", [RFC 5548](#), May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", [RFC 5673](#), October 2009.
- [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.

[11.3](#). 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>>.

Internet-Draft RPL-industrial-applicability-statement

October 2011

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
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

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

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

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

