

ROLL
Internet-Draft
Intended status: Standards Track
Expires: April 27, 2012

D. Popa
J. Jetcheva
Itron
N. Dejean
Elster SAS
R. Salazar
Landis+Gyr
J. Hui
Cisco
October 25, 2011

Applicability Statement for the Routing Protocol for Low Power and Lossy
Networks (RPL) in AMI Networks
draft-ietf-roll-applicability-ami-04

Abstract

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

Status of this Memo

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

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

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

This Internet-Draft will expire on April 27, 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 Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

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

Table of Contents

1.	Introduction	3
1.1.	Electric Metering	3
1.2.	Gas and Water Metering	3
1.3.	Routing Protocol for LLNs (RPL)	4
1.4.	Requirements Language	5
2.	Deployment Scenarios	5
2.1.	Network Topology	5
2.1.1.	Electric Meter Network	5
2.1.2.	Energy-Constrained Network Infrastructure	6
2.2.	Traffic Characteristics	6
2.2.1.	Smart Metering Data	7
2.2.2.	Distribution Automation Communication	8
2.2.3.	Emerging Applications	8
3.	Using RPL to Meet Functional Requirements	8
4.	RPL Profile	9
4.1.	RPL Features	9
4.1.1.	RPL Instances	9
4.1.2.	Storing vs. Non-Storing Mode	9
4.1.3.	DAO Policy	10
4.1.4.	Path Metrics	10
4.1.5.	Objective Function	10
4.1.6.	DODAG Repair	11
4.1.7.	Multicast	11
4.1.8.	Security	12
4.1.9.	P2P communications	12
4.2.	Recommended Configuration Defaults and Ranges	12
4.2.1.	Trickle Parameters	12
4.2.2.	Other Parameters	13
5.	Manageability Considerations	14
6.	Security Considerations	14
7.	Other Related Protocols	15
8.	IANA Considerations	15
9.	Acknowledgements	15
10.	References	15
10.1.	Informative References	15
10.2.	Normative References	16
	Authors' Addresses	16

1. Introduction

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

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

1.1. Electric Metering

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

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

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

1.2. Gas and Water Metering

While electric meters typically consume electricity from the same electric feed that they are monitoring, gas and water meters

typically run on a modest source of stored energy (e.g., batteries).

In some scenarios, gas and water meters are integrated into the same AMI network as the electric meters and may operate as network endpoints (rather than routers) in order to prolong their own lifetime. In other scenarios, however, such meters may not have the luxury of relying on a fully powered AMI routing infrastructure but must communicate through a dedicated infrastructure to reach a LBR. This infrastructure can be either powered by the electricity grid, by battery-based devices, or ones relying on alternative sources of energy (e.g., solar power).

1.3. Routing Protocol for LLNs (RPL)

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

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

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

This note describes the applicability of RPL (as defined in [I-D.ietf-roll-rpl]) to AMI deployments. RPL was designed to meet the following application requirements:

- o Routing Requirements for Urban Low-Power and Lossy Networks [RFC5548].
- o Industrial Routing Requirements in Low-Power and Lossy Networks [RFC5673].
- o Home Automation Routing Requirements in Low-Power and Lossy Networks [RFC5826].
- o Building Automation Routing Requirements in Low-Power and Lossy Networks [RFC5867].

The Routing Requirements for Urban Low-Power and Lossy Networks are applicable to AMI networks as well.

The terminology used in this document is defined in [I-D.ietf-roll-terminology].

1.4. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Deployment Scenarios

2.1. Network Topology

AMI networks are composed of millions of endpoints distributed across both urban and rural environments. Such endpoints include electric, gas, and water meters, distribution automation elements, and home area network devices. Devices in the network communicate directly with other devices in close proximity using a variety of low-power and/or lossy link technologies that are both wired and wireless (e.g., IEEE 802.15.4, IEEE P1901.2, and 802.11). In addition to serving as sources and destinations of packets, many network elements typically also forward packets and thus form a mesh topology.

2.1.1. Electric Meter Network

In a typical AMI deployment, groups of meters within physical proximity form routing domains, each in the order of a 1,000 to 10,000 meters. Thus, each electric meter mesh typically has several thousand wireless endpoints, with densities varying based on the area and the terrain. For example, apartment buildings in urban centers may have hundreds of meters in close proximity, whereas rural areas may have sparse node distributions and include nodes that only have a small number of network neighbors.

Each routing domain is connected to the larger IP infrastructure through one or more LBRs, which provide Wide Area Network (WAN) connectivity through various traditional network technologies, e.g., Ethernet, cellular, private WAN. Paths in the mesh between a network node and the nearest LBR may be composed of several hops or even several tens of hops.

Powered from the main line, electric meters have less energy constraints than battery powered devices, such as gas and water meters, and can afford the additional resources required to route

packets. In mixed environments, electric meters can provide the routing topology while gas and water meters can operate as leaf nodes.

Electric meter networks may also serve as transit networks for other types of devices, including distribution automation elements (e.g., sensors and actuators), and in-home devices. These other devices may utilize a different link-layer technology than the one used in the meter network.

The routing protocol operating in networks with the topology characteristics described above needs to be able to scale with network size and number of forwarding hops, and have the ability to handle a wide range of network densities.

2.1.2. Energy-Constrained Network Infrastructure

In the absence of a co-located electric meter network, gas and water meters must either connect directly to the larger IP network infrastructure or rely on a dedicated routing infrastructure. Deploying such infrastructures is a challenging task as the routing devices can sometimes only be placed in specific locations and thus do not always have access to a continuous energy source. Battery-operated or energy-harvesting (e.g., equipped with solar panels) routers are thus often used in these kinds of scenarios.

Due to the expected lifetime (10 to 20 years) of such networks and their reliance on alternative sources of energy, energy consumption needs to be taken into account when designing and deploying them. There are a number of challenging trade-offs and considerations that exist in that respect. One such consideration is that managing a higher number of meters per router leads to increased energy consumption. However, increasing the number of routers in the network and thus reducing the number of meters managed by each router increases deployment and maintenance costs. At the same time, the use of a sparser routing infrastructure necessitates the use of higher transmit power levels at nodes in the network, which causes increased energy consumption.

The deployment and operational needs of energy-constrained network infrastructure require the use of routing mechanisms that take into account energy consumption, minimize energy use and prolong network lifetime.

2.2. Traffic Characteristics

2.2.1. Smart Metering Data

In current AMI deployments, metering applications typically require all smart meters to communicate with a few head-end servers, deployed in the utility company data center.

Head-end servers generate data traffic to configure smart metering devices or initiate queries, and use unicast and multicast to efficiently communicate with a single device or groups of devices respectively (i.e., Point-to-Multipoint (P2MP) communication). The head-end server may send a single small packet at a time to the meters (e.g., a meter read request, a small configuration change) or a series of large packets (e.g., a firmware upgrade across one or even thousands of devices). The frequency of large file transfers, e.g., firmware upgrade of all metering devices, is typically much lower than the frequency of sending configuration messages or queries.

Each smart meter generates Smart Metering Data (SMD) traffic according to a schedule (e.g., periodic meter reads), in response to on-demand queries (e.g., on-demand meter reads), or in response to some local event (e.g., power outage, leak detection). Such traffic is typically destined to a single head-end server.

The bulk of the SMD traffic tends to be directed towards the LBR, both in terms of bytes (since reports are typically much larger than queries) and in terms of number of packets, e.g., some reports have to be split into multiple packets due to packet size limitations, periodic reports can be sent without requiring a query to be sent for each one first, unsolicited events like alarms and outage notifications are only generated by the meters and sent towards the LBR. The SMD traffic is thus highly asymmetric, where the majority of the traffic volume generated by the smart meters typically goes through the LBRs, and is directed from the smart meter devices to the head-end servers, in a Multipoint-to-Point (MP2P) fashion.

Current SMD traffic patterns are fairly uniform and well-understood. The traffic generated by the head-end server and destined to metering devices is dominated by periodic meter reads, while traffic generated by the metering devices is typically uniformly spread over some periodic read time-window.

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

From a routing perspective, SMD applications require efficient P2MP communication between the devices in the network and one or more

LBRs. In addition, timely loop resolution and broken link repair are needed to meet latency requirements. Finally, the availability of redundant paths is important for increasing network reliability.

2.2.2. Distribution Automation Communication

Distribution Automation (DA) applications typically involve a small number of devices that communicate with each other in a Point-to-Point (P2P) fashion, and may or may not be in close physical proximity.

DA applications typically have more stringent latency requirements than SMD applications.

2.2.3. Emerging Applications

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

3. Using RPL to Meet Functional Requirements

The functional requirements for most AMI deployments are similar to those listed in [RFC5548]:

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

RPL supports:

- o Large-scale networks characterized by highly directed traffic flows between each smart meter and the head-end servers in the utility network. To this end, RPL builds a Directed Acyclic Graph (DAG) rooted at each LBR.
- o Zero-touch configuration. This is done through in-band methods for configuring RPL variables using DIO messages.
- o The use of links with time-varying quality characteristics. This is accomplished by allowing the use of metrics that effectively capture the quality of a path (e.g., Expected Transmission Count (ETX)) and by limiting the impact of changing local conditions by discovering and maintaining multiple DAG parents, and by using local repair mechanisms when DAG links break.

4. RPL Profile

This section outlines a RPL profile for a representative AMI deployment.

4.1. RPL Features

4.1.1. RPL Instances

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

4.1.2. Storing vs. Non-Storing Mode

In most scenarios, electric meters are powered by the grid they are monitoring and are not energy-constrained. Instead, electric meters have hardware and communication capacity constraints that are primarily determined by cost, and secondarily by power consumption. As a result, different AMI deployments can vary significantly in terms of memory size, computation power and communication capabilities. For this reason, the use of RPL storing or non-storing mode SHOULD be deployment specific.

When meters are memory constrained and cannot adequately store the route tables necessary to support hop-by-hop routing, RPL non-storing mode SHOULD be preferred. On the other hand, when nodes are capable of storing such routing tables, the use of storing mode may lead to reduced overhead and route repair latency. However, in high-density environments, storing routes can be challenging because some nodes

may have to maintain routing information for a large number of descendents. When the routing table size becomes challenging, it is RECOMMENDED that nodes perform route aggregation, similarly to the approach taken by other routing protocols, although the required set of mechanism may differ.

4.1.3. DAO Policy

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

4.1.4. Path Metrics

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

For water- and gas-only networks that do not rely on powered infrastructure, simpler metrics that require less energy to compute would be more appropriate. In particular, a combination of hop count and link quality can satisfy this requirement. As minimizing energy consumption is critical in these types of networks, available node energy should also be used in conjunction with these two metrics. The usage of additional metrics specifically designed for such networks may be defined in companion RFCs, e.g., [I-D.ietf-roll-routing-metrics] .

4.1.5. Objective Function

RPL relies on an Objective Function for selecting parents and computing path costs and rank. This objective function is decoupled from the core RPL mechanisms and also from the metrics in use in the network. Two objective functions for RPL have been defined at the time of this writing, OF0 and MRHOF, both of which define the selection of a preferred parent and backup parents, and are suitable for AMI deployments.

Neither of the currently defined objective functions supports multiple metrics that might be required in heterogeneous networks (e.g., networks composed of devices with different energy constraints) or combination of metrics that might be required for water- and gas-only networks. 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, AMI deployments SHOULD utilize the local repair mechanisms in RPL.

Local repair is triggered by broken link detection and in storing mode by loop detection as well.

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 [I-D.ietf-roll-rpl]. While RPL provides an option to form a local DODAG, doing so in AMI deployments is of little benefit since AMI applications typically communicate through a LBR. After the detached node has made sufficient effort to send notification to its children that it is detached, the node can rejoin the same DODAG with a higher rank value. The configured duration of the poisoning mechanism needs to take into account the disconnection time applications running over the network can tolerate. Note that when joining a different DODAG, the node need not perform poisoning.

The second local repair mechanism controls how much a node can increase its rank within a given DODAG Version (e.g., after detaching from the DODAG as a result of broken link or loop detection). Setting the DAGMaxRankIncrease to a non-zero value enables this mechanism, and setting it to a value of less than infinity limits the cost of count-to-infinity scenarios when they occur, thus controlling the duration of disconnection applications may experience.

4.1.7. Multicast

RPL defines multicast support for its storing mode of operation, where the DODAG structure built for unicast packet dissemination is used for multicast distribution as well. In particular, multicast forwarding state creation is done through DAO messages with multicast target options sent along the DODAG towards the root. Thereafter nodes with forwarding state for a particular group forward multicast packets along the DODAG by copying them to all children from which they have received a DAO with a multicast target option for the group.

Multicast support for RPL in non-storing mode will be defined in companion RFCs.

4.1.8. Security

AMI deployments operate in areas that do not provide any physical security. For this reason, the link layer, transport layer and application layer technologies utilized within AMI networks typically provide security mechanisms to ensure authentication, confidentiality, integrity, and freshness. As a result, AMI deployments may not need to implement RPL's security mechanisms and could rely on link layer and higher layer security features.

4.1.9. P2P communications

Distribution Automation and other emerging applications may require efficient P2P communications. Basic P2P capabilities are already defined in the RPL RFC [I-D.ietf-roll-rpl]. Additional mechanisms for efficient P2P communication are being developed in companion RFCs.

4.2. Recommended Configuration Defaults and Ranges

4.2.1. 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 AMI deployments can vary greatly, from nodes having only one or a handful of neighbors to nodes having several hundred neighbors. In high density environments, relatively low values for `Imin` may cause a short period of congestion when an inconsistency is detected and DIO updates are sent by a large number of neighboring nodes nearly simultaneously. While the Trickle timer will exponentially backoff, some time may elapse before the congestion subsides. While some link layers employ contention mechanisms that attempt to avoid congestion, relying solely on the link layer to avoid congestion caused by a large number of DIO updates can result in increased communication latency for other control and data traffic in the network.

To mitigate this kind of short-term congestion, this document recommends a more conservative set of values for the Trickle parameters than those specified in [RFC6206]. In particular, `DIOIntervalMin` is set to a larger value to avoid periods of congestion in dense environments, and `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 AMI deployment, the Trickle parameters are specified in terms of the link's maximum capacity for transmitting link-local multicast messages. If the link can transmit m link-local multicast packets per second on average, the expected time it takes to transmit a link-local multicast packet is $1/m$ seconds.

DIOIntervalMin: AMI deployments SHOULD set DIOIntervalMin such that the Trickle Imin is at least 50 times as long as it takes to transmit a link-local multicast packet. This value is larger than that recommended in [RFC6206] to avoid congestion in dense urban deployments as described above. In energy-constrained deployments (e.g., in water and gas battery-based routing infrastructure), DIOIntervalMin MAY be set to a value resulting in a Trickle Imin of several (e.g. 2) hours.

DIOIntervalDoublings: AMI deployments SHOULD set DIOIntervalDoublings such that the Trickle Imax is at least 2 hours or more. For very energy constrained deployments (e.g., water and gas battery-based routing infrastructure), DIOIntervalDoublings MAY be set to a value resulting in a Trickle Imax of several (e.g., 2) days.

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

4.2.2. Other Parameters

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

5. Manageability Considerations

Network manageability is a critical aspect of smart grid network deployment and operation. With millions of devices participating in the smart grid network, many requiring real-time reachability, automatic configuration, and lightweight network health monitoring and management are crucial for achieving network availability and efficient operation.

RPL enables automatic and consistent configuration of RPL routers through parameters specified by the DODAG root and disseminated through DIO packets. The use of Trickle for scheduling DIO transmissions ensures lightweight yet timely propagation of important network and parameter updates and allows network operators to choose the trade-off point they are comfortable with respect to overhead vs. reliability and timeliness of network updates.

The metrics in use in the network along with the Trickle Timer parameters used to control the frequency and redundancy of network updates can be dynamically varied by the root during the lifetime of the network. To that end, all DIO messages SHOULD contain a Metric Container option for disseminating the metrics and metric values used for DODAG setup. In addition, DIO messages SHOULD contain a DODAG Configuration option for disseminating the Trickle Timer parameters throughout the network.

The possibility of dynamically updating the metrics in use in the network as well as the frequency of network updates allows deployment characteristics (e.g., network density) to be discovered during network bring-up and to be used to tailor network parameters once the network is operational rather than having to rely on precise pre-configuration. This also allows the network parameters and the overall routing protocol behavior to evolve during the lifetime of the network.

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

6. Security Considerations

Smart grid 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 smart grid 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 using RPL's secure mode is not necessary.

7. Other Related Protocols

This document contains no other related protocols.

8. IANA Considerations

This memo includes no request to IANA.

9. Acknowledgements

The authors would like to acknowledge the review, feedback, and comments of Jari Arkko, Dominique Barthel, Cedric Chauvenet, Philip Levis, and JP Vasseur.

10. References

10.1. Informative References

[I-D.ietf-6man-rpl-option]

Hui, J. and J. Vasseur, "RPL Option for Carrying RPL Information in Data-Plane Datagrams", draft-ietf-6man-rpl-option-04 (work in progress), October 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-routing-metrics]

Vasseur, J., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics used for Path Calculation in Low

Power and Lossy Networks",
draft-ietf-roll-routing-metrics-19 (work in progress),
March 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.

[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.

[RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko,
"The Trickle Algorithm", RFC 6206, March 2011.

10.2. Normative References

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

Authors' Addresses

Daniel Popa
Itron
52 Rue Camille Desmoulins
92448 Issy les Moulineaux
France

Email: daniel.popa@itron.com

Jorjeta Jetcheva
Itron
2111 N Molter Rd.
Liberty Lake, WA 99019
USA

Email: jorjeta.jetcheva@itron.com

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

Email: nicolas.dejean@coronis.com

Ruben Salazar
Landis+Gyr
30000 Mill Creek Ave # 100
Alpharetta, GA 30022

Email: ruben.salazar@landisgyr.com

Jonathan W. Hui
Cisco
170 West Tasman Drive
San Jose, California 95134
USA

Phone: +408 424 1547
Email: jonhui@cisco.com