

ROLL
Internet-Draft
Intended status: Standards Track
Expires: January 26, 2012

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

Applicability Statement for the Routing Protocol for Low Power and Lossy
Networks (RPL) in AMI Networks
draft-ietf-roll-applicability-ami-01

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 January 26, 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	4
1.3.	Routing Protocol for LLNs (RPL)	4
1.4.	Requirements Language	5
2.	Deployment Scenarios	5
2.1.	Network Topology	5
2.2.	Traffic Characteristics	6
2.2.1.	Meter Data Management	6
2.2.2.	Distribution Automation	7
2.2.3.	Emerging Applications	7
3.	Using RPL to Meet Functional Requirements	7
4.	RPL Profile	8
4.1.	RPL Features	8
4.1.1.	RPL Instances	8
4.1.2.	Storing vs. Non-Storing Mode	8
4.1.3.	DAO Policy	8
4.1.4.	Path Metrics	9
4.1.5.	Objective Function	9
4.1.6.	DODAG Repair	9
4.1.7.	Multicast	10
4.1.8.	Security	10
4.2.	RPL Options	10
4.3.	Recommended Configuration Defaults and Ranges	10
5.	Manageability Considerations	11
6.	Security Considerations	11
7.	Other Related Protocols	12
8.	IANA Considerations	12
9.	Security Considerations	12
10.	Acknowledgements	12
11.	References	12
11.1.	Informative References	12
11.2.	Normative References	13
	Authors' Addresses	13

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 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. This results 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 network through a network aggregation point (NAP). These kinds of networks increase coverage and reduce installation cost, time and complexity, as well as operational costs, as compared to single-hop wireless networks, relying on a wireline or cellular backhaul. Each electric meter mesh typically has on the order of several thousand wireless endpoints, with densities varying based on the area and the terrain. 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 one or two network neighbors. Paths in the mesh between a network device and the nearest aggregation point may be composed of several hops or even tens of hops.

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 (i.e. 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, such meters may not have the luxury of relying on a powered routing infrastructure but must communicate through other energy-constrained devices (i.e., through other gas and water meters) to reach a NAP. In some cases, battery-powered meters need to communicate directly with a sparsely deployed network infrastructure, requiring them to use high transmit power levels (and thus more energy) in order to achieve the necessary range to reach the infrastructure. In all of these types of networks, the routing protocol must operate with energy consumption in mind.

RPL is 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].

1.3. Routing Protocol for LLNs (RPL)

RPL provides routing functionality for mesh networks composed of a large number of resource-constrained devices, interconnected by low power and lossy links, and communicating with the external network infrastructure through a common aggregation point (e.g., a border router).

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

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 WiFi). In addition to serving as sources and destinations of packets, many network elements typically also forward packets to reduce the need for dedicated network infrastructure and the associated deployment and operational costs.

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. These routing domains are connected to the larger IP infrastructure through one or more LLN Border Routers (LBRs), which provide Wide Area Network (WAN) connectivity through various traditional network technologies, e.g., Ethernet, Cellular, private WAN.

Powered from the main line, electric meters have less energy constraints than battery powered devices and can afford the additional resources required to route packets. In mixed environments, electric meters provide the routing topology while gas and water meters operate as leaf nodes. However, 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 form their own routing topology, albeit with energy consumption in mind.

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.

2.2. Traffic Characteristics

2.2.1. Meter Data Management

Meter Data Management (MDM) applications typically require every smart meter to communicate with a few head-end servers deployed in a utility data center. As a result, all smart metering traffic typically goes through the LBRs, with the vast majority of traffic flowing from smart meter devices to the head-end servers, i.e., in a Multipoint-to-Point (MP2P) fashion.

Smart meters may generate 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 events (e.g., power outages, leak detections). Such traffic is typically unicast since it is sent to a single head-end server.

Head-end servers generate traffic to configure smart metering devices or initiate queries, and use unicast and multicast to efficiently communicate with a single device (i.e., Point-to-Point (P2P) communication) 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 (e.g., a meter read request, or small configuration change) or many consecutive large packets (e.g., a firmware upgrade across one or even thousands of devices).

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

2.2.2. Distribution Automation

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

DA applications typically have more stringent latency requirements than MDM 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 MDM 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 MDM and DA traffic.

4.1.2. Storing vs. Non-Storing Mode

In most scenarios, electric meters are powered by the electric grid they are monitoring and are not energy-constrained. Instead, the capabilities of an electric meter are primarily determined by cost. As a result, different AMI deployments can vary significantly in terms of the memory, computation, and communication trade-offs that they embody. 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 route tables to support downward routing, non-storing mode is preferred. However, when nodes are capable of adequately storing such routing tables, storing mode can lead to reduced overhead and shorter route repair latency.

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 can exhibit significant packet loss. 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 cannot rely on a powered infrastructure, energy constraints may require simpler metrics that do not require as much energy to compute. In particular, hop count and link quality level may be more suitable in such deployments. Other possible metrics to use may be vendor-specific or defined at a later time in companion RFCs.

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 basic 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 a basic AMI deployment. Neither of these supports multiple metrics that might be required in heterogeneous networks (i.e. networks composed of devices with different energy constraints). A new objective function can be defined to meet this requirement.

4.1.6. DODAG Repair

To effectively handle time-varying link characteristics and availability, AMI deployments SHOULD utilize the local repair mechanisms in RPL.

The first mechanism for local repair when a node loses connectivity to its parents is to detach from a DODAG then re-attach 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 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. 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. 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.

The third local repair mechanism enables loop detection, and is implemented by including the rank value of the transmitting node in packets forwarded towards the root (in the packet's RPL Packet Information option [I-D.ietf-6man-rpl-option]). Note that loop detection is not needed when sending packets using strict source routing.

4.1.7. Multicast

RPL defines multicast support for its storing mode of operation. 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 technologies used within AMI deployments typically provide security mechanisms to ensure confidentiality, integrity, and freshness. As a result, AMI deployments may not need to implement RPL's security mechanisms and could rely on link-layer security features.

4.2. RPL Options

4.3. Recommended Configuration Defaults and Ranges

- o AMI deployments can involve densities of hundreds of devices within communication range. As a result, such networks SHOULD set the DIOIntervalMin to 16 or more, resulting in a Trickle Imin of 1 minute or more. In networks with low-energy consumption requirements, DIOIntervalMin SHOULD be set to a higher value.
- o AMI deployments SHOULD set DIOIntervalDoublings to a value that gives a Trickle Imax of 2 hours or more. In networks with low-

energy consumption requirements, `DIOIntervalDoublings` SHOULD be set to a value that results in a Trickle `Imax` of several (e.g., 2) days.

- o AMI deployments SHOULD set `DIORedundancyConstant` to a value of 10 or more.
- 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.

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.

The CoRE Working Group is developing lightweight resource management mechanisms for LLNs that are applicable to smart grid RPL networks as well.

6. Security Considerations

Smart grid networks are subject to stringent security requirements as they are considered a critical national 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 security mechanisms are typically in place and using RPL's secure mode is not necessary. Smart grid networks are often secured at other layers as well, including end-to-end at the application layer.

7. Other Related Protocols

This document contains no other related protocols.

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

This memo includes no security considerations.

10. Acknowledgements

The authors would like to acknowledge the review, feedback, and comments from Dominique Barthel.

11. References

11.1. 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-03 (work in progress), March 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-05 (work in
progress), March 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.

11.2. 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

Email: daniel.popa@itron.com

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