

Internet Engineering Task Force (IETF)
Request for Comments: 8036
Category: Standards Track
ISSN: 2070-1721

N. Cam-Winget, Ed.
Cisco Systems
J. Hui
Nest
D. Popa
Itron, Inc
January 2017

Applicability Statement for
the Routing Protocol for Low-Power and Lossy Networks (RPL) in
Advanced Metering Infrastructure (AMI) Networks

Abstract

This document discusses the applicability of the Routing Protocol for Low-Power and Lossy Networks (RPL) in Advanced Metering Infrastructure (AMI) networks.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8036>.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	3
1.2.	Required Reading	3
1.3.	Out-of-Scope Requirements	4
2.	Routing Protocol for LLNs (RPL)	4
3.	Description of AMI Networks for Electric Meters	4
3.1.	Deployment Scenarios	5
4.	Smart Grid Traffic Description	7
4.1.	Smart Grid Traffic Characteristics	7
4.2.	Smart Grid Traffic QoS Requirements	8
4.3.	RPL Applicability per Smart Grid Traffic Characteristics	9
5.	Layer-2 Applicability	9
5.1.	IEEE Wireless Technology	9
5.2.	IEEE Power Line Communication (PLC) Technology	9
6.	Using RPL to Meet Functional Requirements	10
7.	RPL Profile	11
7.1.	RPL Features	11
7.1.1.	RPL Instances	11
7.1.2.	DAO Policy	11
7.1.3.	Path Metrics	11
7.1.4.	Objective Function	12
7.1.5.	DODAG Repair	12
7.1.6.	Multicast	12
7.1.7.	Security	13
7.2.	Description of Layer-2 Features	13
7.2.1.	IEEE 1901.2 PHY and MAC Sub-layer Features	13
7.2.2.	IEEE 802.15.4 (Amendments G and E) PHY and MAC Features	14
7.2.3.	IEEE MAC Sub-layer Security Features	15
7.3.	6LowPAN Options	17
7.4.	Recommended Configuration Defaults and Ranges	17
7.4.1.	Trickle Parameters	17
7.4.2.	Other Parameters	18
8.	Manageability Considerations	18
9.	Security Considerations	19
9.1.	Security Considerations during Initial Deployment	20
9.2.	Security Considerations during Incremental Deployment	20
9.3.	Security Considerations Based on RPL's Threat Analysis	20
10.	Privacy Considerations	21
11.	References	21
11.1.	Normative References	21
11.2.	Informative references	22
	Acknowledgements	24
	Authors' Addresses	24

1. Introduction

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

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

1.1. Requirements Language

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

1.2. Required Reading

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

A NAN can use wireless communication technology, which is based on the IEEE 802.15.4 standard family: more specifically, the Physical Layer (PHY) amendment [IEEE.802.15.4g] and the Media Access Control (MAC) sub-layer amendment [IEEE.802.15.4e], which are adapted to smart grid networks.

NAN can also use Power Line Communication (PLC) technology as an alternative to wireless communications. Several standards for PLC technology have emerged, such as [IEEE.1901.2].

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

1.3. Out-of-Scope Requirements

The following are outside the scope of this document:

- o Applicability statement for RPL [RFC6550] in AMI networks composed of battery-powered devices (i.e., gas/water meters).
- o Applicability statement for RPL in AMI networks composed of a mix of devices powered by alternating current (i.e., electric meters) and battery-powered meters (i.e., gas/water meters).
- o Applicability statement for RPL storing mode of operation in AMI networks.

2. Routing Protocol for LLNs (RPL)

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

RPL builds a Directed Acyclic Graph (DAG) routing structure rooted at an 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. RPL's ability to support multiple different metrics and constraints at the same time enables it to run efficiently in heterogeneous networks composed of nodes and links with vastly different characteristics [RFC6551].

This document describes the applicability of RPL non-storing mode (as defined in [RFC6550]) to AMI deployments. The Routing Requirements for Urban Low-Power and Lossy Networks [RFC5548] are applicable to AMI networks as well. The terminology used in this document is defined in [RFC7102].

3. Description of AMI Networks for Electric Meters

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

sensors to monitor electric grid quality and to support applications such as electric vehicle charging.

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

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

3.1. Deployment Scenarios

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

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

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

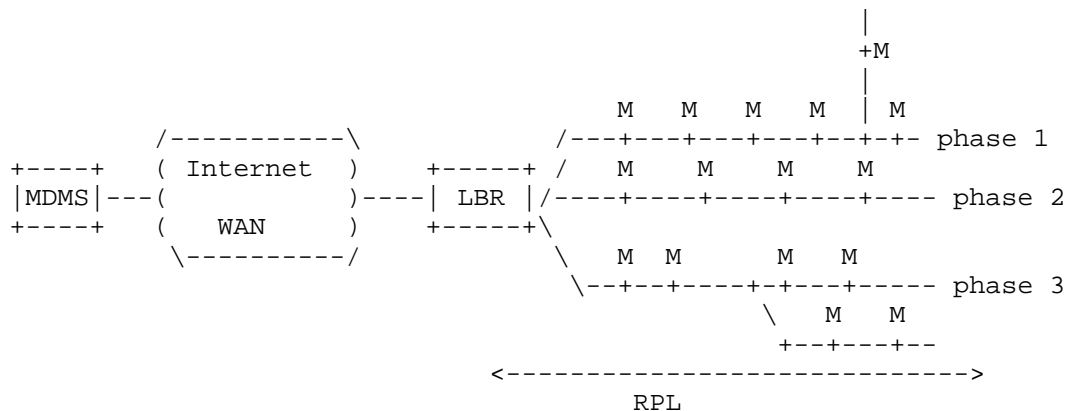


Figure 1: Typical NAN Topology

A typical AMI network architecture (see Figure 1) is composed of a Meter Data Management System (MDMS) connected through an IP network to an LBR, which can be located in the power substation or somewhere else in the field. The power substation connects the households and buildings. The physical topology of the electrical grid is a tree structure, either due to the three different power phases coming through the substation or just to the electrical network topology. Meters (represented by a M in the previous figure) can also participate in a HAN. The scope of this document is the communication between the LBR and the meters, i.e., the NAN segment.

Node density can vary significantly. For example, apartment buildings in urban centers may have hundreds of meters in close proximity, whereas rural areas may have sparse node distributions and may 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 based on Worldwide Interoperability for Microwave Access (WiMAX), and optical fiber. 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.

As a function of the technology used to exchange information, the logical network topology will not necessarily match the electric grid topology. If meters exchange information through radio technologies such as in the IEEE 802.15.4 family, then the topology is a meshed

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

4. Smart Grid Traffic Description

4.1. Smart Grid Traffic Characteristics

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

Distribution Automation (DA) applications typically involve a small number of devices that communicate with each other in a P2P fashion and may or may not be in close physical proximity. DA applications typically have more stringent latency requirements than SMD applications.

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

4.2. Smart Grid Traffic QoS Requirements

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

- A) Meter-initiated traffic (Meter-to-Head-End - M2HE)
- B) Head-end-initiated traffic (Head-End-to-Meter - HE2M)
 - B1) request is sent in P2P to a specific meter
 - B2) request is sent in multicast to a subset of meters
 - B3) request is sent in multicast to all meters

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

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

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

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

4.3. RPL Applicability per Smart Grid Traffic Characteristics

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

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

5. Layer-2 Applicability

5.1. IEEE Wireless Technology

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

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

5.2. IEEE Power Line Communication (PLC) Technology

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

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

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

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

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

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

6. Using RPL to Meet Functional Requirements

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

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

RPL supports the following features:

- o Scalability: Large-scale networks characterized by highly directed traffic flows between each smart meter and the head-end servers in the utility network. To this end, RPL builds a Directed Acyclic Graph (DAG) rooted at each LBR.
- o Zero-touch configuration: This is done through in-band methods for configuring RPL variables using DIO (DODAG Information Object) messages and DIO message options [RFC6550].
- o The use of links with time-varying quality characteristics: This is accomplished by allowing the use of metrics that effectively capture the quality of a path (e.g., Expected Transmission Count (ETX)) and by limiting the impact of changing local conditions by discovering and maintaining multiple DAG parents (and by using local repair mechanisms when DAG links break).

7. RPL Profile

7.1. RPL Features

7.1.1. RPL Instances

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

7.1.2. DAO Policy

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

7.1.3. Path Metrics

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

Additional metrics may be defined in companion RFCs.

7.1.4. Objective Function

RPL relies on an Objective Function for selecting parents and computing path costs and rank. This objective function is decoupled from the core RPL mechanisms and also from the metrics in use in the network. Two objective functions for RPL have been defined at the time of this writing, Objective Function 0 (OF0) [RFC6552] and Minimum Rank with Hysteresis Objective Function (MRHOF) [RFC6719], both of which define the selection of a preferred parent and backup parents and are suitable for AMI deployments.

Additional objective functions may be defined in companion RFCs.

7.1.5. DODAG Repair

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

7.1.6. Multicast

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

7.1.7. Security

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

7.2. Description of Layer-2 Features

7.2.1. IEEE 1901.2 PHY and MAC Sub-layer Features

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

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

The IEEE 1901.2 MAC frame format shares most in common with the IEEE 802.15.4 MAC frame format [IEEE.802.15.4]. A few exceptions are described below.

- o The IEEE 1901.2 MAC frame is obtained by prepending a Segment Control Field to the IEEE 802.15.4 MAC header. One function of the Segment Control Field is to signal the use of the MAC sub-layer segmentation and reassembly.

- o IEEE 1901.2 MAC frames use only the 802.15.4 MAC addresses with a length of 16 and 64 bits.
- o The IEEE 1901.2 MAC sub-layer endorses the concept of Information Elements, as defined in [IEEE.802.15.4e]. The format and use of Information Elements are not relevant to the RPL applicability statement.

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

The IEEE 1901.2 PHY MTU size is variable and dependent on the PHY settings in use (e.g., bandwidth, modulation, tones, etc). As quoted from the IEEE 1901.2 specification:

For CENELEC A/B, if MSDU size is more than 247 octets for robust OFDM (ROBO) and Super-ROBO modulations or more than 239 octets for all other modulations, the MAC layer shall divide the MSDU into multiple segments as described in 5.3.7. For FCC and ARIB, if the MSDU size meets one of the following conditions: a) For ROBO and Super-ROBO modulations, the MSDU size is more than 247 octets but less than 494 octets, b) For all other modulations, the MSDU size is more than 239 octets but less than 478 octets.

7.2.2. IEEE 802.15.4 (Amendments G and E) PHY and MAC Features

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

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

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

- o Timeslotslotted Channel Hopping (TSCH): specifically designed for application domains such as process automation

- o Low-Latency Deterministic Networks (LLDN): for application domains such as factory automation.
- o Deterministic and Synchronous Multi-channel Extension (DSME): for general industrial and commercial application domains that includes channel diversity to increase network robustness.
- o Asynchronous Multi-channel Adaptation (AMCA): for large infrastructure application domains.

The MAC addressing scheme supports short (16-bit) addresses along with extended (64-bit) addresses. These addresses are assigned in different ways and are specified by specific standards organizations. Information Elements, Enhanced Beacons, and frame version 2, as defined in IEEE 802.15.4e, MUST be supported.

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

7.2.3. IEEE MAC Sub-layer Security Features

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

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

The 802.15.4 security layer is handled at the media access control layer, below the 6LowPAN (IPv6 over Low-Power Wireless Personal Area Network) layer. The application specifies its security requirements by setting the appropriate control parameters into the radio/PLC stack. IEEE 802.15.4 defines four packet types: beacon frames, data frames, acknowledgment frames, and command frames for the media access control layer. The 802.15.4 specification does not support security for acknowledgment frames; data frames, beacon frames, and command frames can support integrity protection and confidentiality protection for the frames' data field. An application has a choice of security suites that control the type of security protection that is provided for the transmitted MAC frame. Each security suite offers a different set of security properties and guarantees, and

ultimately offers different MAC frame formats. The 802.15.4 specification defines eight different security suites, outlined below. We can broadly classify the suites by the properties that they offer: no security, encryption only (AES-CTR), authentication only (AES-CBC-MAC), and encryption and authentication (AES-CCM). Each category that supports authentication comes in three variants depending on the size of the Message Authentication Code that it offers. The MAC can be either 4, 8, or 16 bytes long. Additionally, for each suite that offers encryption, the recipient can optionally enable replay protection.

- o Null = No security
- o AES-CTR = Encryption only, CTR mode
- o AES-CBC-MAC-128 = No encryption, 128-bit MAC
- o AES-CBC-MAC-64 = No encryption, 64-bit MAC
- o AES-CCM-128 = Encryption and 128-bit MAC
- o AES-CCM-64 = Encryption and 64-bit MAC
- o AES-CCM-32 = Encryption and 32-bit MAC

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

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

A cryptographic Message Authentication Code (or keyed MIC) is used to authenticate messages. While longer MICs lead to improved resiliency

of the code, they also make the packet size larger and thus take up bandwidth in the network. In constrained environments such as metering infrastructures, an optimum balance between security requirements and network throughput must be found.

7.3. 6LowPAN Options

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

7.4. Recommended Configuration Defaults and Ranges

7.4.1. Trickle Parameters

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

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

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

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

7.4.2. Other Parameters

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

8. Manageability Considerations

Network manageability is a critical aspect of smart grid network deployment and operation. With millions of devices participating in the smart grid network, many requiring real-time reachability, automatic configuration, and lightweight-network health monitoring and management are crucial for achieving network availability and efficient operation. RPL enables automatic and consistent configuration of RPL routers through parameters specified by the DODAG root and disseminated through DIO packets. The use of Trickle for scheduling DIO transmissions ensures lightweight yet timely propagation of important network and parameter updates and allows network operators to choose the trade-off point with which they are comfortable with respect to overhead vs. reliability and timeliness of network updates. The metrics in use in the network along with the Trickle Timer parameters used to control the frequency and redundancy

of network updates can be dynamically varied by the root during the lifetime of the network. To that end, all DIO messages SHOULD contain a Metric Container option for disseminating the metrics and metric values used for DODAG setup. In addition, DIO messages SHOULD contain a DODAG Configuration option for disseminating the Trickle Timer parameters throughout the network. The possibility of dynamically updating the metrics in use in the network as well as the frequency of network updates allows deployment characteristics (e.g., network density) to be discovered during network bring-up and to be used to tailor network parameters once the network is operational rather than having to rely on precise pre-configuration. This also allows the network parameters and the overall routing protocol behavior to evolve during the lifetime of the network. RPL specifies a number of variables and events that can be tracked for purposes of network fault and performance monitoring of RPL routers. Depending on the memory and processing capabilities of each smart grid device, various subsets of these can be employed in the field.

9. Security Considerations

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

In contrast to other types of LLNs, in smart grid networks both centralized administrative control and access to a permanent secure infrastructure are available. As a result, smart grid networks are deployed with security mechanisms such as link-layer, transport-layer, and/or application-layer security mechanisms; while it is best practice to secure all layers, using RPL's secure mode may not be necessary. Failure to protect any of these layers can result in various attacks; a lack of strong authentication of devices in the infrastructure can lead to uncontrolled and unauthorized access. Similarly, failure to protect the communication layers can enable passive (in wireless mediums) attacks as well as man-in-the-middle and active attacks.

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

9.1. Security Considerations during Initial Deployment

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

9.2. Security Considerations during Incremental Deployment

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

9.3. Security Considerations Based on RPL's Threat Analysis

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

Like any secure network infrastructure, an AMI deployment's ability to address node impersonation and active man-in-the-middle attacks rely on a mutual authentication and authorization process. To enable strong mutual authentication, all nodes, from smart meters to nodes in the infrastructure, must have a credential. The credential may be bootstrapped at the time the node is manufactured but must be appropriately managed and classified through the authorization process. The management and authorization process ensures that the nodes are properly authenticated and behaving or 'acting' in their assigned roles.

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

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

10. Privacy Considerations

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

11. References

11.1. Normative References

[IEEE.1901.2]

IEEE, "IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", IEEE 1901.2-2013, DOI 10.1109/ieeestd.2013.6679210, December 2013, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6679208>>.

[IEEE.802.15.4]

IEEE, "IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANS)", IEEE 802.15.4-2011, DOI 10.1109/ieeestd.2011.6012487, September 2011, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6012485>>.

[IEEE.802.15.4e]

IEEE, "IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANS) Amendment 1: MAC sublayer", IEEE 802.15.4e-2012, DOI 10.1109/ieeestd.2012.6185525, April 2012, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6185523>>.

[IEEE.802.15.4g]

IEEE, "IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANS) Amendment 3: Physical Layer (PHY) Specifications for Low-Data-Rate, Wireless, Smart Metering Utility Networks", IEEE 802.15.4g-2012, DOI 10.1109/ieeestd.2012.6190698, April 2012, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6190696>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [surveySG] Gungor, V., Sahin, D., Kocak, T., Ergut, S., Buccella, C., Cecati, C., and G. Hancke, "A Survey on Smart Grid Potential Applications and Communication Requirements", IEEE Transactions on Industrial Informatics Volume 9, Issue 1, pp. 28-42, DOI 10.1109/TII.2012.2218253, February 2013.

11.2. Informative references

- [DOEVCC] "Voluntary Code of Conduct (VCC) Final Concepts and Principles", January 2015, <http://energy.gov/sites/prod/files/2015/01/f19/VCC%20Concepts%20and%20Principles%202015_01_08%20FINAL.pdf>.
- [EUPR] "Information for investors and data controllers", June 2016, <<https://ec.europa.eu/energy/node/1748>>.
- [IEEE.802.11] IEEE, "IEEE Standard for Information technology-- Telecommunications and information exchange between systems Local and metropolitan area networks--Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE 802.11-2012, DOI 10.1109/ieeestd.2012.6178212, March 2012, <<https://standards.ieee.org/getieee802/download/802.11-2012.pdf>>.
- [PRIVACY] Thaler, D., "Privacy Considerations for IPv6 Adaptation Layer Mechanisms", Work in Progress, draft-ietf-6lo-privacy-considerations-04, October 2016.
- [RFC5548] Dohler, M., Ed., Watteyne, T., Ed., Winter, T., Ed., and D. Barthel, Ed., "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, DOI 10.17487/RFC5548, May 2009, <<http://www.rfc-editor.org/info/rfc5548>>.

- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<http://www.rfc-editor.org/info/rfc6206>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<http://www.rfc-editor.org/info/rfc6551>>.
- [RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, DOI 10.17487/RFC6552, March 2012, <<http://www.rfc-editor.org/info/rfc6552>>.
- [RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, DOI 10.17487/RFC6719, September 2012, <<http://www.rfc-editor.org/info/rfc6719>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<http://www.rfc-editor.org/info/rfc7416>>.
- [RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <<http://www.rfc-editor.org/info/rfc7731>>.

Acknowledgements

Matthew Gillmore, Laurent Toutain, Ruben Salazar, and Kazuya Monden were contributors and noted as authors in earlier versions of this document. The authors would also like to acknowledge the review, feedback, and comments of Jari Arkko, Dominique Barthel, Cedric Chauvenet, Yuichi Igarashi, Philip Levis, Jeorjeta Jetcheva, Nicolas Dejean, and JP Vasseur.

Authors' Addresses

Nancy Cam-Winget (editor)
Cisco Systems
3550 Cisco Way
San Jose, CA 95134
United States of America

Email: ncamwing@cisco.com

Jonathan Hui
Nest
3400 Hillview Ave
Palo Alto, CA 94304
United States of America

Email: jonhui@nestlabs.com

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

Email: daniel.popa@itron.com