Roll   D. Popa
Internet-Draft   M. Gillmore
Intended status: Standards Track   Itron, Inc
Expires: January 23, 2015   L. Toutain
                                                      Telecom Bretagne
                                                                J. Hui
                                                                 Cisco
                                                              R. Ruben
                                                            Landis+Gyr
                                                             K. Monden
                      Hitachi, Ltd., Yokohama Research Laboratory
                                                         July 22, 2014

        Applicability Statement for the Routing Protocol for Low Power and Lossy
                         Networks (RPL) in AMI Networks
                      draft-ietf-roll-applicability-ami-09

Abstract

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

Status of This Memo

Copyright Notice

Table of Contents

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
   devices. They are typically inter-connected using some combination
   of wireless and power-line communications, 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", "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.

   NAN can use wireless communication technology in which case is using,
   from the IEEE 802.15.4 standard family, the IEEE 802.15.4g PHY Layer
   amendment and IEEE 802.15.4e MAC sub-layer amendment, specifically
   adapted to smart grid networks.

   NAN can also use PLC (Power Line Communication) technology as an
   alternative to wireless communications. Several standards for PLC
   technology have emerged, such as IEEE P1901.2.

   NAN can further use a mix of wireless and PLC technologies to
   increase the network coverage ratio, 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 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 AC powered devices (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, 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
a LBR (LLN Border Router), 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. 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-13].

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 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 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 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.

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 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 wireless and wired (e.g., IEEE 802.15.4g,
   IEEE 802.15.4e, IEEE 1901.2, 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.

```
                                                    |
                                                   +M
                                                    |
                                        M   M   M   M | M
            /-----------\             /---+---+---+---+--+-+- phase 1
 +----+    ( Internet  )    +----+ /   M   M    M    M
 |MDMS|---(            )----|LBR |/----+----+----+----+----+---- phase 2
 +----+    (    WAN    )    +----+\
             \----------/          \   M  M       M   M
                                    \--+--+----+-+---+----- phase 3
                                        \   M   M
                                         +--+---+--
                         <---------------------------->
                                       RPL
```
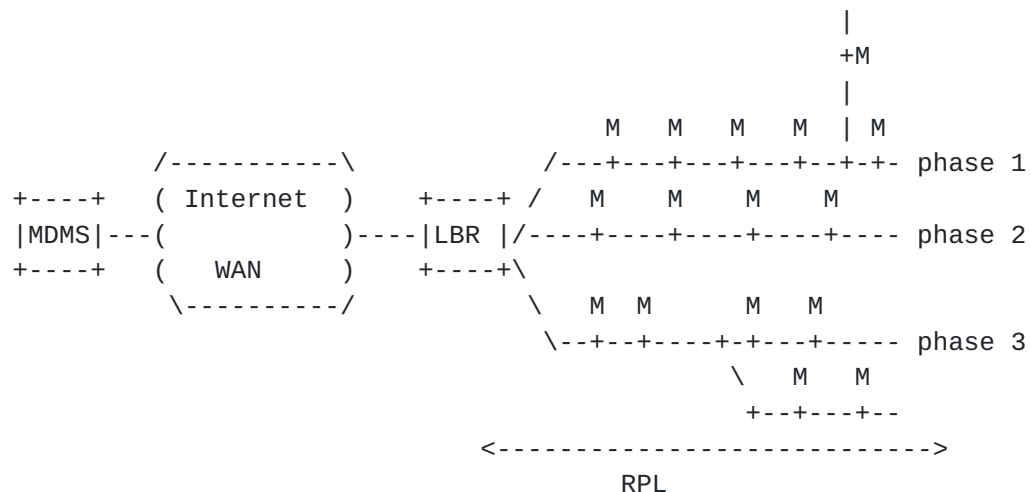
              Figure 1: Typical NAN Topology

A typical AMI network architecture (see figure Figure 1) is composed
of a MDMS (Meter Data Management System) connected through IP network
to a 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 3 different power phases coming through
the sub-station or just to the electrical network topology. Meters
(represented by a M in the previous figure) can also participate to a
HAN (Home-Area Network). 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
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 WiMAX-based WAN, 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 the technology used to exchange information, the
logical network topology will not necessarily match the eletric grid
topology. If meters exchange information through radio technologies
such as in IEEE 802.15.4 family, the topology is a meshed network,
where nodes belonging to the same DODAG can be connected to the grid
through different substations. If narrowband PLC technology is used,
it will follow more or less the physical tree structure since
diaphony may allow one phase to communicate with the other. This is
particularly true near the LBR. Some mixt topology can also be
observed, since some LBR may be strategically installed in the field
to avoid all the communications going through a single LBR.
Nethertheless, 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, deployed
in the utility company data center. Head-end servers generate data
traffic to configure smart meter 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, 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, 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
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.

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.

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 point-to-point 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 decomposed into several
classes :

C1) Highly 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 Power Quality Events, Meter Service
     Connection and Disconnection require 98%+ packet delivery under
     10s. Payload size < 150 bytes

C3) Normal Priority traffic for System Events including Faults,
     Configuration and Security require 98%+ packet delivery under 30s.
     Payload size < 200 bytes

C4) Low Priority traffic for Recurrent Meter Reading require 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 with in 7 days. Average firmware
     update is 1 MB.

4.3. RPL applicability per Smart Grid Traffic Characteristics

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

Group communication model used in smart grid requires 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.

Although not currently used in metering applications, support of peer-
to-peer communications between DODAG nodes is identified as a key
feature to be supported in smart grid networks.

5. Layer 2 applicability

5.1. IEEE Wireless Technology

IEEE Std. 802.15.4g-2012 and IEEE 802.15.4e-2012 amendments to
802.15.2-2011 standard have been specifically developed for smart
grid networks. They are the most common PHY and MAC layers used for
wireless AMI network. 802.15.4g specifies multiple mode of operation
(FSK, OQPSK and OFDM modulations) with speeds from 50kbps to 600kbps,
and allows for transport of a full IPv6 packet (i.e., 1280 octets)
without the need for upper layer segmentation and re-assembly.

IEEE Std. 802.15.4e-2012 is an amendment to IEEE Std 802.15.4-2011
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 802.15.4e does not change the link-layer
security scheme defined in 802.15.4-2011 (and 802.15.4-2006).

5.2. IEEE PowerLine Communication (PLC) technology

   The IEEE Std 1901.2-2013 standard specifies communications for low
   frequency (less than 500 kHz) narrowband power line devices via
   alternating current and direct current electric power lines. IEEE
   Std 1901.2-2013 standard supports indoor and outdoor communications
   over low voltage line (line between transformer and meter, less than
   1000 V), through transformer low-voltage to medium-voltage (1000 V up
   to 72 kV) and through transformer 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 sub-set of
   802.15.4-2006 and 802.15.4e-2012 MAC sub-layer features.

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

   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 standard specifies link-layer security features that
   fully endorse 802.15.4-2006 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]:

   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 messages, and DIO message
      options.

   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 SDM and DA
   traffic.

7.1.2. 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. In this
document we only focus on non-storing mode of operation.

7.1.3. 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.

7.1.4. 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 [RFC6551].

Additional metrics may be defined in companion RFCs.

7.1.5. 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 [RFC6552] and 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.

7.1.6. 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. The first local repair mechanism consists of a node
detaching from a DODAG and then re-attaching to the same or to a
different DODAG at a later time. While detached, a node advertises
an infinite rank value so that its children can select a different
parent. This process is known as poisoning and is described in
Section 8.2.2.5 of [RFC6550]. While RPL provides an option to form a
local DODAG, doing so in AMI for electric meters 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.

7.1.7. Multicast

Multicast support for RPL in non-storing mode are being developed in
companion RFCs (see [draft-ietf-roll-trickle-mcast-06]).

7.1.8. 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.

7.1.9. Peer-to-Peer communications

Basic peer-to-peer capabilities are already defined in the RPL
[RFC6550]. Additional mechanisms for peer-to-peer communications are
proposed in companion RFCs (see [RFC6997]).

7.2. Description of Layer-two 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 sub-carrier can vary
significantly IEEE 1901.2 specifies two complementary mechanisms
allowing to fine-tune 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,...). The second (fine-
grained) maps the sub-carriers which 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 sub-carriers. 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.

IEEE 1901.2 MAC frame format shares most in common with the IEEE 802.15.4-2006 MAC frame format [IEEE802.15.4], with a few exceptions described below.

o IEEE 1901.2 MAC frame is obtained by prepending a Segment Control Field to the IEEE 802.15.4-2006 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 uses only the 802.15.4 MAC addresses with a length of 16 and 64 bits.

o IEEE 1901.2 MAC sub-layer endorses the concept of Information Elements, as defined in IEEE 802.15.4e-2012 [IEEE802.15.4e]. The format and use of Information Elements are not relevant to 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 maximum IEEE 1901.2 PHY frame payload is 512 bytes. The maximum IEEE 1901.2 MAC frame payload is 1280 bytes, which supports the IPv6 minimum MTU requirement.

When there is a mistmatch between the PHY frame payload size and the size of a MAC frame carrying an IPv6 packet, IEEE 1901.2 specifies a MAC sub-layer segmentation and re-assembly mechanism that provides a reliable one-hop transfer of that MAC frame segments.

7.2.2. IEEE 802.15.4 (g + 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, 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.

The IEEE Std 802.15.4g defines the following modulations: MR-FSK (Multi- Rate FSK), MR-OFDM (Multi-Rate OFDM) and MR-O-QPSK (Multi-Rate O-QPSK). The (over-the-air) bit rates for these modulations range from 4.8 to 600kbps for MR-FSK, from 50 to 600kbps for MR-OFDM and from 6.25 to 500kbps 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: Timetimeslotslotted channel hopping (TSCH), specifically designed for application domains such as process automation, Low latency deterministic networks (LLDN), for application domains such as factory automation, Deterministic and synchronous multi-channel extension (DSME), for general industrial and commercial application domains that includes Channel diversity to increase network robustness, and Asynchronous multi-channel adaptation (AMCA), for large infrastructure application domains.

The MAC addressing scheme supports short (16-bits) addresses along with extended (64-bits) addresses.

Information Elements, Enhanced Beacons and frame version 2, as defined in 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 able to carry a full IPv6 packet of 1280 octets without upper layer fragmentation and re-assembly.

7.2.3. IEEE MAC sub-layer Security Features

Since IEEE 1901.2 standard is based on the 802.15.4-2006 MAC sub-layer and fully endorses the security scheme defined in 802.15.4-2006, we only focus on description of 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. The 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 6LowPAN layer. The application specifies its security

requirements by setting the appropriate control parameters into the
radio/PLC stack. The 802.15.4 defines four packet types: beacon
frames, data frames, acknowledgments frame, and command frames for
the media access control layer. The 802.15.4 specification does not
support security for acknowledgement frames; data frames, beacon
frames and command frames can support integrity protection and
confidentiality protection for the frames's 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.
The 802.15.4 specification defines eight different security suites,
outlined bellow. 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 MAC (Message Authentication
Control) 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.

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 (MAC) 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 can not 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 can not
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 MAC is used to authenticate messages. While longer
MACs lead to improved resiliency of the code, they also make 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 1901.2 and 802.15.4(g+e) can utilize all
of the IPv6 Header Compression schemes specified in [RFC6282]
Section 3 and all of the IPv6 Next Header compression schemes
specified in [RFC6282] Section 4, if reducing over the air/wire
overhead is a requirement. However, since the link-layer MTU in both
wireless and PLC links supports the transmission of a full IPv6
packet, the use of 6LowPAN fragmentation is NOT RECOMMENDED.

7.4. Recommended Configuration Defaults and Ranges

7.4.1. Trickle Parameters

Trickle was designed to be density-aware and perform well in networks
characterized by a wide range of node densities. The combination of
DIO packet suppression and adaptive timers for sending updates allows
Trickle to perform well in both sparse and dense environments. 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. 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.

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 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, 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.

9.1. Security Considerations during initial deployment

During the manufacturing process, the meters are loaded with the
appropriate security credentials (keys, certificates). These
security credentials are unique per device and only known by the
device itself and the head-end security appliances. The
manufacturing security credentials are then used by the devices to
authenticate with the system and 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 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.

10. Other Related Protocols

This section is intentionally left blank.

11. IANA Considerations

This memo includes no request to IANA.

12. Acknowledgements

The authors would 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.

13. References

13.1. Informative References

[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",
          Feb 2013.

[IEEE802.15.4]
          IEEE SA, "IEEE Standard for Information technology-- Local
          and metropolitan area networks-- Specific requirements--
          Part 15.4: Wireless Medium Access Control (MAC) and
          Physical Layer (PHY) Specifications for Low Rate Wireless
          Personal Area Networks (WPANs)", September 2006.

[IEEE802.15.4e]
          IEEE SA, "IEEE Standard for Local and metropolitan area
          networks--Part 15.4: Low-Rate Wireless Personal Area
          Networks (LR-WPANs) Amendment 1: MAC sublayer", April
          2012.

[IEEE1901.2]
          IEEE SA, "IEEE Standard for Low-Frequency (less than 500
          kHz) Narrowband Power Line Communications for Smart Grid
          Applications", December 2013.

13.2. Normative references

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

[RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel,
          "Routing Requirements for Urban Low-Power and Lossy
          Networks", RFC 5548, May 2009.

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

[RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R.,
          Levis, P., Pister, K., Struik, R., Vasseur, JP., and R.
          Alexander, "RPL: IPv6 Routing Protocol for Low-Power and
          Lossy Networks", RFC 6550, March 2012.

[RFC6551] Vasseur, JP., Kim, M., Pister, K., Dejean, N., and D.
          Barthel, "Routing Metrics Used for Path Calculation in
          Low-Power and Lossy Networks", RFC 6551, March 2012.

[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.

[RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation
          Routing Requirements in Low-Power and Lossy Networks", RFC
          5826, April 2010.

   [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney,
              "Industrial Routing Requirements in Low-Power and Lossy
              Networks", RFC 5673, October 2009.

   [RFC6997] Goyal, M., Baccelli, E., Philipp, M., Brandt, A., and J.
              Martocci, "Reactive Discovery of Point-to-Point Routes in
              Low-Power and Lossy Networks", RFC 6997, August 2013.

   [RFC6552] Thubert, P., "Objective Function Zero for the Routing
              Protocol for Low-Power and Lossy Networks (RPL)", RFC
              6552, March 2012.

   [RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with
              Hysteresis Objective Function", RFC 6719, September 2012.

Authors' Addresses

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

   Email: daniel.popa@itron.com

   Matthew Gillmore
   Itron, Inc
   2111 N Molter Rd.
   Liberty Lake, WA 99019
   USA

   Email: matthew.gillmore@itron.com

   Laurent Toutain
   Telecom Bretagne
   2 rue de la Chataigneraie
   Cesson Sevigne 35510
   FR

   Email: laurent.toutain@telecom-bretagne.eu

   Jonathan Hui
   Cisco
   170 West Tasman Drive
   San Jose, CA 95134
   USA

   Email: johui@cisco.com

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

   Email: ruben.salazar@landisgyr.com

   Kazuya Monden
   Hitachi, Ltd., Yokohama Research Laboratory
   292, Yoshida-cho, Totsuka-ku, Yokohama-shi
   Kanagawa-ken   244-0817
   Japan

   Email: kazuya.monden.vw@hitachi.com