Building Automation Routing Requirements
in Low-Power and Lossy Networks

Abstract

   The Routing Over Low-Power and Lossy (ROLL) networks Working Group
   has been chartered to work on routing solutions for Low-Power and
   Lossy Networks (LLNs) in various markets: industrial, commercial
   (building), home, and urban networks. Pursuant to this effort, this
   document defines the IPv6 routing requirements for building
   automation.

Status of This Memo

Table of Contents

1. Introduction

   The Routing Over Low-Power and Lossy (ROLL) networks Working Group
   has been chartered to work on routing solutions for Low-Power and
   Lossy Networks (LLNs) in various markets: industrial, commercial
   (building), home, and urban networks. Pursuant to this effort, this
   document defines the IPv6 routing requirements for building
   automation.

   Commercial buildings have been fitted with pneumatic, and
   subsequently electronic, communication routes connecting sensors to
   their controllers for over one hundred years. Recent economic and
   technical advances in wireless communication allow facilities to
   increasingly utilize a wireless solution in lieu of a wired solution,
   thereby reducing installation costs while maintaining highly reliant
   communication.

The cost benefits and ease of installation of wireless sensors allow customers to further instrument their facilities with additional sensors, providing tighter control while yielding increased energy savings.

Wireless solutions will be adapted from their existing wired counterparts in many of the building applications including, but not limited to, heating, ventilation, and air conditioning (HVAC); lighting; physical security; fire; and elevator/lift systems. These devices will be developed to reduce installation costs while increasing installation and retrofit flexibility, as well as increasing the sensing fidelity to improve efficiency and building service quality.

Sensing devices may be battery-less, battery-powered, or mains-powered. Actuators and area controllers will be mains-powered. Due to building code and/or device density (e.g., equipment room), it is envisioned that a mix of wired and wireless sensors and actuators will be deployed within a building.

Building management systems (BMSs) are deployed in a large set of vertical markets including universities, hospitals, government facilities, kindergarten through high school (K-12), pharmaceutical manufacturing facilities, and single-tenant or multi-tenant office buildings. These buildings range in size from 100K-sq.-ft. structures (5-story office buildings), to 1M-sq.-ft. skyscrapers (100-story skyscrapers), to complex government facilities such as the Pentagon. The described topology is meant to be the model to be used in all of these types of environments but clearly must be tailored to the building class, building tenant, and vertical market being served.

Section 3 describes the necessary background to understand the context of building automation including the sensor, actuator, area controller, and zone controller layers of the topology; typical device density; and installation practices.

Section 4 defines the traffic flow of the aforementioned sensors, actuators, and controllers in commercial buildings.

Section 5 defines the full set of IPv6 routing requirements for commercial buildings.

   Appendix A documents important commercial building requirements that
   are out of scope for routing yet will be essential to the final
   acceptance of the protocols used within the building.

   Section 3 and Appendix A are mainly included for educational
   purposes.

   The expressed aim of this document is to provide the set of IPv6
   routing requirements for LLNs in buildings, as described in
   Section 5.

2. Terminology

   For a description of the terminology used in this specification,
   please see [ROLL-TERM].

2.1. 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].

3. Overview of Building Automation Networks

3.1. Introduction

   To understand the network systems requirements of a building
   management system in a commercial building, this document uses a
   framework to describe the basic functions and composition of the
   system. A BMS is a hierarchical system of sensors, actuators,
   controllers, and user interface devices that interoperate to provide
   a safe and comfortable environment while constraining energy costs.

   A BMS is divided functionally across different but interrelated
   building subsystems such as heating, ventilation, and air
   conditioning (HVAC); fire; security; lighting; shutters; and
   elevator/lift control systems, as denoted in Figure 1.

   Much of the makeup of a BMS is optional and installed at the behest
   of the customer. Sensors and actuators have no standalone
   functionality. All other devices support partial or complete
   standalone functionality. These devices can optionally be tethered
   to form a more cohesive system. The customer requirements dictate
   the level of integration within the facility. This architecture
   provides excellent fault tolerance since each node is designed to
   operate in an independent mode if the higher layers are unavailable.

```
               +------+ +-----+ +------+ +------+ +------+ +------+

   Bldg App'ns |      | |     | |      | |      | |      | |      |

               |      | |     | |      | |      | |      | |      |

   Building Cntl |    | |     | |  S   | |  L   | |  S   | |  E   |

               |      | |     | |  E   | |  I   | |  H   | |  L   |

   Area Control |  H  | | F   | |  C   | |  G   | |  U   | |  E   |

               |  V   | | I   | |  U   | |  H   | |  T   | |  V   |

   Zone Control |  A  | | R   | |  R   | |  T   | |  T   | |  A   |

               |  C   | | E   | |  I   | |  I   | |  E   | |  T   |

   Actuators    |     | |     | |  T   | |  N   | |  R   | |  O   |

               |      | |     | |  Y   | |  G   | |  S   | |  R   |

   Sensors      |     | |     | |      | |      | |      | |      |

               +------+ +-----+ +------+ +------+ +------+ +------+
```
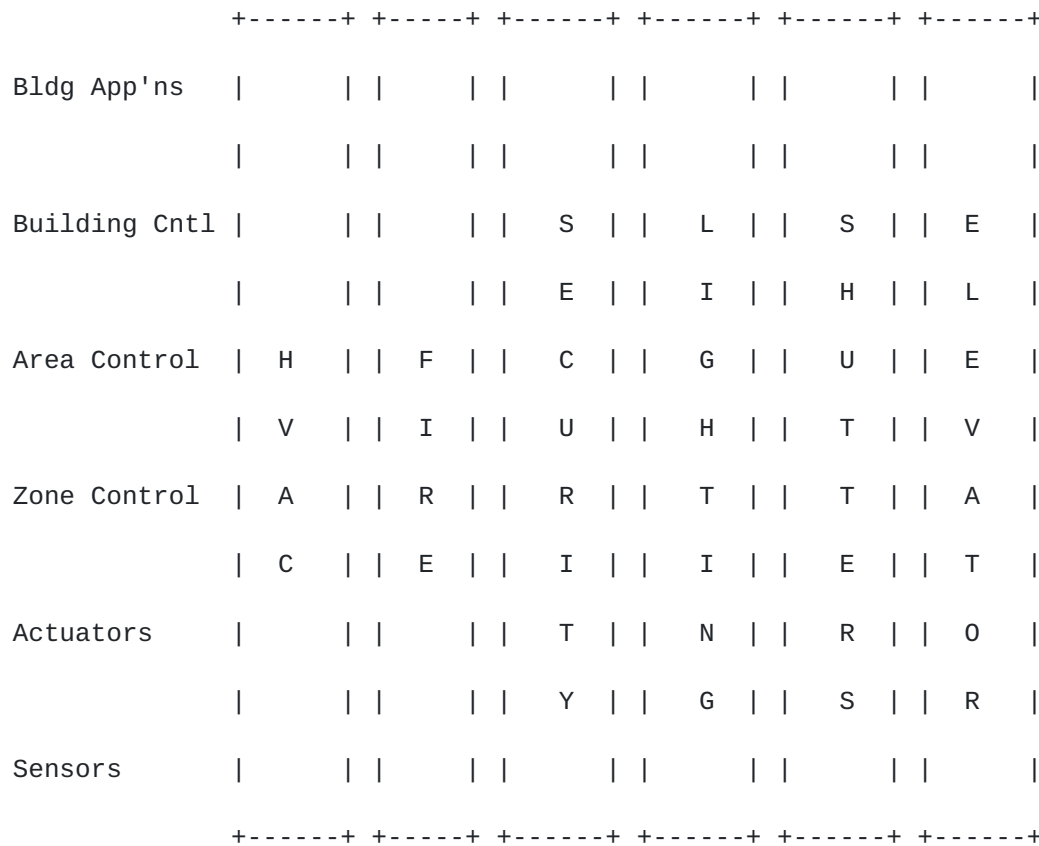
Figure 1: Building Systems and Devices

3.2. Building Systems Equipment

3.2.1. Sensors/Actuators

As Figure 1 indicates, a BMS may be composed of many functional
stacks or silos that are interoperably woven together via building
applications. Each silo has an array of sensors that monitor the
environment and actuators that modify the environment, as determined
by the upper layers of the BMS topology. The sensors typically are
at the edge of the network structure, providing environmental data
for the system. The actuators are the sensors' counterparts,
modifying the characteristics of the system, based on the sensor data
and the applications deployed.

3.2.2. Area Controllers

An area describes a small physical locale within a building,
typically a room. HVAC (temperature and humidity) and lighting (room
lighting, shades, solar loads) vendors oftentimes deploy area
controllers. Area controllers are fed by sensor inputs that monitor

the environmental conditions within the room. Common sensors found
in many rooms that feed the area controllers include temperature,
occupancy, lighting load, solar load, and relative humidity. Sensors
found in specialized rooms (such as chemistry labs) might include air
flow, pressure, and CO_2 and CO particle sensors. Room actuation
includes temperature setpoint, lights, and blinds/curtains.

3.2.3. Zone Controllers

Zone controllers support a similar set of characteristics to area
controllers, albeit for an extended space. A zone is normally a
logical grouping or functional division of a commercial building. A
zone may also coincidentally map to a physical locale such as a
floor.

Zone controllers may have direct sensor inputs (smoke detectors for
fire), controller inputs (room controllers for air handlers in HVAC),
or both (door controllers and tamper sensors for security). Like
area/room controllers, zone controllers are standalone devices that
operate independently or may be attached to the larger network for
more synergistic control.

3.3. Equipment Installation Methods

A BMS is installed very differently from most other IT networks. IT
networks are typically installed as an overlay onto the existing
environment and are installed from the inside out. That is, the
network wiring infrastructure is installed; the switches, routers,
and servers are connected and made operational; and finally, the
endpoints (e.g., PCs, VoIP phones) are added.

BMSs, on the other hand, are installed from the outside in. That is,
the endpoints (thermostats, lights, smoke detectors) are installed in
the spaces first; local control is established in each room and
tested for proper operation. The individual rooms are later lashed
together into a subsystem (e.g., lighting). The individual
subsystems (e.g., lighting, HVAC) then coalesce. Later, the entire
system may be merged onto the enterprise network.

The rationale for this is partly due to the different construction
trades having access to a building under construction at different
times. The sheer size of a building often dictates that even a
single trade may have multiple independent teams working
simultaneously. Furthermore, the HVAC, lighting, and fire systems
must be fully operational before the building can obtain its
occupancy permit. Hence, the BMS must be in place and configured
well before any of the IT servers (DHCP; Authentication,
Authorization, and Accounting (AAA); DNS; etc.) are operational.

This implies that the BMS cannot rely on the availability of the IT
network infrastructure or application servers. Rather, the BMS
installation should be planned to dovetail into the IT system once
the IT system is available for easy migration onto the IT network.
Front-end planning of available switch ports, cable runs, access
point (AP) placement, firewalls, and security policies will
facilitate this adoption.

3.4. Device Density

Device density differs, depending on the application and as dictated
by the local building code requirements. The following subsections
detail typical installation densities for different applications.

3.4.1. HVAC Device Density

HVAC room applications typically have sensors/actuators and
controllers spaced about 50 ft. apart. In most cases, there is a 3:1
ratio of sensors/actuators to controllers. That is, for each room
there is an installed temperature sensor, flow sensor, and damper
actuator for the associated room controller.

HVAC equipment room applications are quite different. An air handler
system may have a single controller with up to 25 sensors and
actuators within 50 ft. of the air handler. A chiller or boiler is
also controlled with a single equipment controller instrumented with
25 sensors and actuators. Each of these devices would be
individually addressed since the devices are mandated or optional as
defined by the specified HVAC application. Air handlers typically
serve one or two floors of the building. Chillers and boilers may be
installed per floor, but many times they service a wing, building, or
the entire complex via a central plant.

These numbers are typical. In special cases, such as clean rooms,
operating rooms, pharmaceutical facilities, and labs, the ratio of
sensors to controllers can increase by a factor of three. Tenant
installations such as malls would opt for packaged units where much
of the sensing and actuation is integrated into the unit; here, a
single device address would serve the entire unit.

3.4.2. Fire Device Density

Fire systems are much more uniformly installed, with smoke detectors
installed about every 50 ft. This is dictated by local building
codes. Fire pull boxes are installed uniformly about every 150 ft.
A fire controller will service a floor or wing. The fireman's fire
panel will service the entire building and typically is installed in
the atrium.

3.4.3. Lighting Device Density

   Lighting is also very uniformly installed, with ballasts installed
   approximately every 10 ft. A lighting panel typically serves 48 to
   64 zones. Wired systems tether many lights together into a single
   zone. Wireless systems configure each fixture independently to
   increase flexibility and reduce installation costs.

3.4.4. Physical Security Device Density

   Security systems are non-uniformly oriented, with heavy density near
   doors and windows and lighter density in the building's interior
   space.

   The recent influx of interior and perimeter camera systems is
   increasing the security footprint. These cameras are atypical
   endpoints requiring up to 1 megabit/second (Mbit/s) data rates per
   camera, as contrasted by the few kbit/s needed by most other BMS
   sensing equipment. Previously, camera systems had been deployed on
   proprietary wired high-speed networks. More recent implementations
   utilize wired or wireless IP cameras integrated into the enterprise
   LAN.

4. Traffic Pattern

   The independent nature of the automation subsystems within a building
   can significantly affect network traffic patterns. Much of the real-
   time sensor environmental data and actuator control stays within the
   local LLN environment, while alarms and other event data will
   percolate to higher layers.

   Each sensor in the LLN unicasts point to point (P2P) about 200 bytes
   of sensor data to its associated controller each minute and expects
   an application acknowledgment unicast returned from the destination.
   Each controller unicasts messages at a nominal rate of 6 kbit/minute
   to peer or supervisory controllers. Thirty percent of each node's
   packets are destined for other nodes within the LLN. Seventy percent
   of each node's packets are destined for an aggregation device
   (multipoint to point (MP2P)) and routed off the LLN. These messages
   also require a unicast acknowledgment from the destination. The
   above values assume direct node-to-node communication; meshing and
   error retransmissions are not considered.

   Multicasts (point to multipoint (P2MP)) to all nodes in the LLN occur
   for node and object discovery when the network is first commissioned.
   This data is typically a one-time bind that is henceforth persisted.
   Lighting systems will also readily use multicasting during normal
   operations to turn banks of lights "on" and "off" simultaneously.

BMSs may be either polled or event-based. Polled data systems will
generate a uniform and constant packet load on the network. Polled
architectures, however, have proven not to be scalable. Today, most
vendors have developed event-based systems that pass data on event.
These systems are highly scalable and generate low data on the
network at quiescence. Unfortunately, the systems will generate a
heavy load on startup since all initial sensor data must migrate to
the controller level. They also will generate a temporary but heavy
load during firmware upgrades. This latter load can normally be
mitigated by performing these downloads during off-peak hours.

Devices will also need to reference peers periodically for sensor
data or to coordinate operation across systems. Normally, though,
data will migrate from the sensor level upwards through the local and
area levels, and then to the supervisory level. Traffic bottlenecks
will typically form at the funnel point from the area controllers to
the supervisory controllers.

Initial system startup after a controlled outage or unexpected power
failure puts tremendous stress on the network and on the routing
algorithms. A BMS is comprised of a myriad of control algorithms at
the room, area, zone, and enterprise layers. When these control
algorithms are at quiescence, the real-time data rate is small, and
the network will not saturate. An overall network traffic load of 6
kbit/s is typical at quiescence. However, upon any power loss, the
control loops and real-time data quickly atrophy. A short power
disruption of only 10 minutes may have a long-term deleterious impact
on the building control systems, taking many hours to regain proper
control. Control applications that cannot handle this level of
disruption (e.g., hospital operating rooms) must be fitted with a
secondary power source.

Power disruptions are unexpected and in most cases will immediately
impact lines-powered devices. Power disruptions, however, are
transparent to battery-powered devices. These devices will continue
to attempt to access the LLN during the outage. Battery-powered
devices designed to buffer data that has not been delivered will
further stress network operations when power returns.

Upon restart, lines-powered devices will naturally dither due to
primary equipment delays or variance in the device self-tests.
However, most lines-powered devices will be ready to access the LLN
network within 10 seconds of power-up. Empirical testing indicates
that routes acquired during startup will tend to be very oblique
since the available neighbor lists are incomplete. This demands an
adaptive routing protocol to allow for route optimization as the
network stabilizes.

5. Building Automation Routing Requirements

Following are the building automation routing requirements for
networks used to integrate building sensor, actuator, and control
products. These requirements are written not presuming any
preordained network topology, physical media (wired), or radio
technology (wireless).

5.1. Device and Network Commissioning

Building control systems typically are installed and tested by
electricians having little computer knowledge and no network
communication knowledge whatsoever. These systems are often
installed during the building construction phase, before the drywall
and ceilings are in place. For new construction projects, the
building enterprise IP network is not in place during installation of
the building control system. For retrofit applications, the
installer will still operate independently from the IP network so as
not to affect network operations during the installation phase.

In traditional wired systems, correct operation of a light
switch/ballast pair was as simple as flipping on the light switch.
In wireless applications, the tradesperson has to assure the same
operation, yet be sure the operation of the light switch is
associated with the proper ballast.

System-level commissioning will later be deployed using a more
computer savvy person with access to a commissioning device (e.g., a
laptop computer). The completely installed and commissioned
enterprise IP network may or may not be in place at this time.
Following are the installation routing requirements.

5.1.1. Zero-Configuration Installation

It MUST be possible to fully commission network devices without
requiring any additional commissioning device (e.g., a laptop). From
the ROLL perspective, "zero configuration" means that a node can
obtain an address and join the network on its own, without human
intervention.

5.1.2. Local Testing

During installation, the room sensors, actuators, and controllers
SHOULD be able to route packets amongst themselves and to any other
device within the LLN, without requiring any additional routing
infrastructure or routing configuration.

5.1.3. Device Replacement

   To eliminate the need to reconfigure the application upon replacing a
   failed device in the LLN, the replaced device must be able to
   advertise the old IP address of the failed device in addition to its
   new IP address. The routing protocols MUST support hosts and routers
   that advertise multiple IPv6 addresses.

5.2. Scalability

   Building control systems are designed for facilities from 50,000 sq.
   ft. to 1M+ sq. ft. The networks that support these systems must
   cost-effectively scale accordingly. In larger facilities,
   installation may occur simultaneously on various wings or floors, yet
   the end system must seamlessly merge. Following are the scalability
   requirements.

5.2.1. Network Domain

   The routing protocol MUST be able to support networks with at least
   2,000 nodes, where 1,000 nodes would act as routers and the other
   1,000 nodes would be hosts. Subnetworks (e.g., rooms, primary
   equipment) within the network must support up to 255 sensors and/or
   actuators.

5.2.2. Peer-to-Peer Communication

   The data domain for commercial BMSs may sprawl across a vast portion
   of the physical domain. For example, a chiller may reside in the
   facility's basement due to its size, yet the associated cooling
   towers will reside on the roof. The cold-water supply and return
   pipes snake through all of the intervening floors. The feedback
   control loops for these systems require data from across the
   facility.

   A network device MUST be able to communicate in an end-to-end manner
   with any other device on the network. Thus, the routing protocol
   MUST provide routes between arbitrary hosts within the appropriate
   administrative domain.

5.3. Mobility

   Most devices are affixed to walls or installed on ceilings within
   buildings. Hence, the mobility requirements for commercial buildings
   are few. However, in wireless environments, location tracking of
   occupants and assets is gaining favor. Asset-tracking applications,
   such as tracking capital equipment (e.g., wheelchairs) in medical

facilities, require monitoring movement with granularity of a minute; however, tracking babies in a pediatric ward would require latencies less than a few seconds.

The following subsections document the mobility requirements in the routing layer for mobile devices. Note, however, that mobility can be implemented at various layers of the system, and the specific requirements depend on the chosen layer. For instance, some devices may not depend on a static IP address and are capable of re-establishing application-level communications when given a new IP address. Alternatively, mobile IP may be used, or the set of routers in a building may give an impression of a building-wide network and allow devices to retain their addresses regardless of where they are, handling routing between the devices in the background.

5.3.1. Mobile Device Requirements

To minimize network dynamics, mobile devices while in motion should not be allowed to act as forwarding devices (routers) for other devices in the LLN. Network configuration should allow devices to be configured as routers or hosts.

5.3.1.1. Device Mobility within the LLN

An LLN typically spans a single floor in a commercial building. Mobile devices may move within this LLN. For example, a wheelchair may be moved from one room on the floor to another room on the same floor.

A mobile LLN device that moves within the confines of the same LLN SHOULD re-establish end-to-end communication with a fixed device also in the LLN within 5 seconds after it ceases movement. The LLN network convergence time should be less than 10 seconds once the mobile device stops moving.

5.3.1.2. Device Mobility across LLNs

A mobile device may move across LLNs, such as a wheelchair being moved to a different floor.

A mobile device that moves outside of its original LLN SHOULD re-establish end-to-end communication with a fixed device also in the new LLN within 10 seconds after the mobile device ceases movement. The network convergence time should be less than 20 seconds once the mobile device stops moving.

5.4. Resource Constrained Devices

Sensing and actuator device processing power and memory may be 4
orders of magnitude less (i.e., 10,000x) than many more traditional
client devices on an IP network. The routing mechanisms must
therefore be tailored to fit these resource constrained devices.

5.4.1. Limited Memory Footprint on Host Devices

The software size requirement for non-routing devices (e.g., sleeping
sensors and actuators) SHOULD be implementable in 8-bit devices with
no more than 128 KB of memory.

5.4.2. Limited Processing Power for Routers

The software size requirements for routing devices (e.g., room
controllers) SHOULD be implementable in 8-bit devices with no more
than 256 KB of flash memory.

5.4.3. Sleeping Devices

Sensing devices will, in some cases, utilize battery power or energy
harvesting techniques for power and will operate mostly in a sleep
mode to maintain power consumption within a modest budget. The
routing protocol MUST take into account device characteristics such
as power budget.

Typically, sensor battery life (2,000 mAh) needs to extend for at
least 5 years when the device is transmitting its data (200 octets)
once per minute over a low-power transceiver (25 mA) and expecting an
application acknowledgment. In this case, the transmitting device
must leave its receiver in a high-powered state, awaiting the return
of the application ACK. To minimize this latency, a highly efficient
routing protocol that minimizes hops, and hence end-to-end
communication, is required. The routing protocol MUST take into
account node properties, such as "low-powered node", that produce
efficient low-latency routes that minimize radio "on" time for these
devices.

Sleeping devices MUST be able to receive inbound data. Messages sent
to battery-powered nodes MUST be buffered by the last-hop router for
a period of at least 20 seconds when the destination node is
currently in its sleep cycle.

5.5. Addressing

Building management systems require different communication schemes
to solicit or post network information. Multicasts or anycasts need
to be used to decipher unresolved references within a device when the
device first joins the network.

As with any network communication, multicasting should be minimized.
This is especially a problem for small embedded devices with limited
network bandwidth. Multicasts are typically used for network joins
and application binding in embedded systems. Routing MUST support
anycast, unicast, and multicast.

5.6. Manageability

As previously noted in Section 3.3, installation of LLN devices
within a BMS follows an "outside-in" work flow. Edge devices are
installed first and tested for communication and application
integrity. These devices are then aggregated into islands, then
LLNs, and later affixed onto the enterprise network.

The need for diagnostics most often occurs during the installation
and commissioning phase, although at times diagnostic information may
be requested during normal operation. Battery-powered wireless
devices typically will have a self-diagnostic mode that can be
initiated via a button press on the device. The device will display
its link status and/or end-to-end connectivity when the button is
pressed. Lines-powered devices will continuously display
communication status via a bank of LEDs, possibly denoting signal
strength and end-to-end application connectivity.

The local diagnostics noted above oftentimes are suitable for
defining room-level networks. However, as these devices aggregate,
system-level diagnostics may need to be executed to ameliorate route
vacillation, excessive hops, communication retries, and/or network
bottlenecks.

In operational networks, due to the mission-critical nature of the
application, the LLN devices will be temporally monitored by the
higher layers to assure that communication integrity is maintained.
Failure to maintain this communication will result in an alarm being
forwarded to the enterprise network from the monitoring node for
analysis and remediation.

In addition to the initial installation and commissioning of the
system, it is equally important for the ongoing maintenance of the
system to be simple and inexpensive. This implies a straightforward
device swap when a failed device is replaced, as noted in Section
5.1.3.

5.6.1. Diagnostics

To improve diagnostics, the routing protocol SHOULD be able to be
placed in and out of "verbose" mode. Verbose mode is a temporary
debugging mode that provides additional communication information
including, at least, the total number of routed packets sent and
received, the number of routing failures (no route available),
neighbor table members, and routing table entries. The data provided
in verbose mode should be sufficient that a network connection graph
could be constructed and maintained by the monitoring node.

Diagnostic data should be kept by the routers continuously and be
available for solicitation at any time by any other node on the
internetwork. Verbose mode will be activated/deactivated via
unicast, multicast, or other means. Devices having available
resources may elect to support verbose mode continuously.

5.6.2. Route Tracking

Route diagnostics SHOULD be supported, providing information such as
route quality, number of hops, and available alternate active routes
with associated costs. Route quality is the relative measure of
"goodness" of the selected source to destination route as compared to
alternate routes. This composite value may be measured as a function
of hop count, signal strength, available power, existing active
routes, or any other criteria deemed by ROLL as the route cost
differentiator.

5.7. Route Selection

Route selection determines reliability and quality of the
communication among the devices by optimizing routes over time and
resolving any nuances developed at system startup when nodes are
asynchronously adding themselves to the network.

5.7.1. Route Cost

The routing protocol MUST support a metric of route quality and
optimize selection according to such metrics within constraints
established for links along the routes. These metrics SHOULD reflect
metrics such as signal strength, available bandwidth, hop count,
energy availability, and communication error rates.

5.7.2. Route Adaptation

   Communication routes MUST be adaptive and converge toward optimality
   of the chosen metric (e.g., signal quality, hop count) in time.

5.7.3. Route Redundancy

   The routing layer SHOULD be configurable to allow secondary and
   tertiary routes to be established and used upon failure of the
   primary route.

5.7.4. Route Discovery Time

   Mission-critical commercial applications (e.g., fire, security)
   require reliable communication and guaranteed end-to-end delivery of
   all messages in a timely fashion. Application-layer time-outs must
   be selected judiciously to cover anomalous conditions such as lost
   packets and/or route discoveries, yet not be set too large to over-
   damp the network response. If route discovery occurs during packet
   transmission time (reactive routing), it SHOULD NOT add more than 120
   ms of latency to the packet delivery time.

5.7.5. Route Preference

   The routing protocol SHOULD allow for the support of manually
   configured static preferred routes.

5.7.6. Real-Time Performance Measures

   A node transmitting a "request with expected reply" to another node
   must send the message to the destination and receive the response in
   not more than 120 ms. This response time should be achievable with 5
   or less hops in each direction. This requirement assumes network
   quiescence and a negligible turnaround time at the destination node.

5.7.7. Prioritized Routing

   Network and application packet routing prioritization must be
   supported to assure that mission-critical applications (e.g., fire
   detection) cannot be deferred while less critical applications access
   the network. The routing protocol MUST be able to provide routes
   with different characteristics, also referred to as Quality of
   Service (QoS) routing.

5.8. Security Requirements

   This section sets forth specific requirements that are placed on any
   protocols developed or used in the ROLL building environment, in
   order to ensure adequate security and retain suitable flexibility of
   use and function of the protocol.

   Due to the variety of buildings and tenants, the BMSs must be
   completely configurable on-site.

   Due to the quantity of the BMS devices (thousands) and their
   inaccessibility (oftentimes above ceilings), security configuration
   over the network is preferred over local configuration.

   Wireless encryption and device authentication security policies need
   to be considered in commercial buildings, while keeping in mind the
   impact on the limited processing capabilities and additional latency
   incurred on the sensors, actuators, and controllers.

   BMSs are typically highly configurable in the field, and hence the
   security policy is most often dictated by the type of building to
   which the BMS is being installed. Single-tenant owner-occupied
   office buildings installing lighting or HVAC control are candidates
   for implementing a low level of security on the LLN, especially when
   the LLN is not connected to an external network. Antithetically,
   military or pharmaceutical facilities require strong security
   policies. As noted in the installation procedures described in
   Sections 3.3 and 5.2, security policies MUST support dynamic
   configuration to allow for a low level of security during the
   installation phase (prior to building occupancy, when it may be
   appropriate to use only diagnostic levels of security), yet to make
   it possible to easily raise the security level network-wide during
   the commissioning phase of the system.

5.8.1. Building Security Use Case

   LLNs for commercial building applications should always implement and
   use encrypted packets. However, depending on the state of the LLN,
   the security keys may either be:

   1) a key obtained from a trust center already operable on the LLN;

   2) a pre-shared static key as defined by the general contractor or
      its designee; or

   3) a well-known default static key.

Unless a node entering the network had previously received its
credentials from the trust center, the entering node will try to
solicit the trust center for the network key. If the trust center is
accessible, the trust center will MAC-authenticate the entering node
and return the security keys. If the trust center is not available,
the entering node will check to determine if it has been given a
network key by an off-band means and use it to access the network.
If no network key has been configured in the device, it will revert
to the default network key and enter the network. If neither of
these keys were valid, the device would signal via a fault LED.

This approach would allow for independent simplified commissioning,
yet centralized authentication. The building owner or building type
would then dictate when the trust center would be deployed. In many
cases, the trust center need not be deployed until all of the local
room commissioning is complete. Yet, at the province of the owner,
the trust center may be deployed from the onset, thereby trading
installation and commissioning flexibility for tighter security.

5.8.2. Authentication

Authentication SHOULD be optional on the LLN. Authentication SHOULD
be fully configurable on-site. Authentication policy and updates
MUST be routable over-the-air. Authentication SHOULD occur upon
joining or rejoining a network. However, once authenticated, devices
SHOULD NOT need to reauthenticate with any other devices in the LLN.
Packets may need authentication at the source and destination nodes;
however, packets routed through intermediate hops should not need
reauthentication at each hop.

These requirements mean that at least one LLN routing protocol
solution specification MUST include support for authentication.

5.8.3. Encryption

5.8.3.1. Encryption Types

Data encryption of packets MUST be supported by all protocol solution
specifications. Support can be provided by use of a network-wide key
and/or an application key. The network key would apply to all
devices in the LLN. The application key would apply to a subset of
devices in the LLN.

The network key and application key would be mutually exclusive. The
routing protocol MUST allow routing a packet encrypted with an
application key through forwarding devices without requiring each
node in the route to have the application key.

Martocci, et al.   Informational   [Page 20]

5.8.3.2. Packet Encryption

The encryption policy MUST support either encryption of the payload
only or of the entire packet. Payload-only encryption would
eliminate the decryption/re-encryption overhead at every hop,
providing more real-time performance.

5.8.4. Disparate Security Policies

Due to the limited resources of an LLN, the security policy defined
within the LLN MUST be able to differ from that of the rest of the IP
network within the facility, yet packets MUST still be able to route
to or through the LLN from/to these networks.

5.8.5. Routing Security Policies to Sleeping Devices

The routing protocol MUST gracefully handle routing temporal security
updates (e.g., dynamic keys) to sleeping devices on their "awake"
cycle to assure that sleeping devices can readily and efficiently
access the network.

6. Security Considerations

The requirements placed on the LLN routing protocol in order to
provide the correct level of security support are presented in
Section 5.8.

LLNs deployed in a building environment may be entirely isolated from
other networks, attached to normal IP networks within the building
yet physically disjoint from the wider Internet, or connected either
directly or through other IP networks to the Internet. Additionally,
even where no wired connectivity exists outside of the building, the
use of wireless infrastructure within the building means that
physical connectivity to the LLN is possible for an attacker.

Therefore, it is important that any routing protocol solution
designed to meet the requirements included in this document addresses
the security features requirements described in Section 5.8.
Implementations of these protocols will be required in the protocol
specifications to provide the level of support indicated in Section
5.8, and will be encouraged to make the support flexibly configurable
to enable an operator to make a judgment of the level of security
that they want to deploy at any time.

As noted in Section 5.8, use/deployment of the different security
features is intended to be optional. This means that, although the
protocols developed must conform to the requirements specified, the
operator is free to determine the level of risk and the trade-offs

against performance. An implementation must not make those choices
on behalf of the operator by avoiding implementing any mandatory-to-
implement security features.

This informational requirements specification introduces no new
security concerns.

7. Acknowledgments

In addition to the authors, JP. Vasseur, David Culler, Ted Humpal,
and Zach Shelby are gratefully acknowledged for their contributions
to this document.

8. References

8.1. Normative References

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

8.2. Informative References

[ROLL-TERM] Vasseur, JP., "Terminology in Low power And Lossy
            Networks", Work in Progress, March 2010.

Appendix A. Additional Building Requirements

   Appendix A contains additional building requirements that were deemed
   out of scope for ROLL, yet provided ancillary substance for the
   reader.

A.1. Additional Commercial Product Requirements

A.1.1. Wired and Wireless Implementations

   Vendors will likely not develop a separate product line for both
   wired and wireless networks. Hence, the solutions set forth must
   support both wired and wireless implementations.

A.1.2. World-Wide Applicability

   Wireless devices must be supportable unlicensed bands.

A.2. Additional Installation and Commissioning Requirements

A.2.1. Unavailability of an IP Network

   Product commissioning must be performed by an application engineer
   prior to the installation of the IP network (e.g., switches, routers,
   DHCP, DNS).

A.3. Additional Network Requirements

A.3.1. TCP/UDP

   Connection-based and connectionless services must be supported.

A.3.2. Interference Mitigation

   The network must automatically detect interference and seamlessly
   switch the channel to improve communication. Channel changes, and
   the nodes' responses to a given channel change, must occur within 60
   seconds.

A.3.3. Packet Reliability

In building automation, it is required that the network meet the
following minimum criteria:

<1% MAC-layer errors on all messages, after no more than three
retries;

<0.1% network-layer errors on all messages, after no more than three
additional retries;

<0.01% application-layer errors on all messages.

Therefore, application-layer messages will fail no more than once
every 100,000 messages.

A.3.4. Merging Commissioned Islands

Subsystems are commissioned by various vendors at various times
during building construction. These subnetworks must seamlessly
merge into networks and networks must seamlessly merge into
internetworks since the end user wants a holistic view of the system.

A.3.5. Adjustable Routing Table Sizes

The routing protocol must allow constrained nodes to hold an
abbreviated set of routes. That is, the protocol should not mandate
that the node routing tables be exhaustive.

A.3.6. Automatic Gain Control

For wireless implementations, the device radios should incorporate
automatic transmit power regulation to maximize packet transfer and
minimize network interference, regardless of network size or density.

A.3.7. Device and Network Integrity

Commercial-building devices must all be periodically scanned to
assure that each device is viable and can communicate data and alarm
information as needed. Routers should maintain previous packet flow
information temporally to minimize overall network overhead.

A.4. Additional Performance Requirements

A.4.1. Data Rate Performance

An effective data rate of 20 kbit/s is the lowest acceptable
operational data rate on the network.

A.4.2. **Firmware Upgrades**

   To support high-speed code downloads, routing should support
   transports that provide parallel downloads to targeted devices, yet
   guarantee packet delivery. In cases where the spatial position of
   the devices requires multiple hops, the algorithm should recurse
   through the network until all targeted devices have been serviced.
   Devices receiving a download may cease normal operation, but upon
   completion of the download must automatically resume normal
   operation.

A.4.3. **Route Persistence**

   To eliminate high network traffic in power-fail or brown-out
   conditions, previously established routes should be remembered and
   invoked prior to establishing new routes for those devices re-
   entering the network.

Authors' Addresses

   Jerry Martocci
   Johnson Controls Inc.
   507 E. Michigan Street
   Milwaukee, WI 53202
   USA
   Phone: +1 414 524 4010
   EMail: jerald.p.martocci@jci.com

   Pieter De Mil
   Ghent University - IBCN
   G. Crommenlaan 8 bus 201
   Ghent 9050
   Belgium
   Phone: +32 9331 4981
   Fax: +32 9331 4899
   EMail: pieter.demil@intec.ugent.be

   Nicolas Riou
   Schneider Electric
   Technopole 38TEC T3
   37 quai Paul Louis Merlin
   38050 Grenoble Cedex 9
   France
   Phone: +33 4 76 57 66 15
   EMail: nicolas.riou@fr.schneider-electric.com

   Wouter Vermeylen
   Arts Centre Vooruit
   Ghent 9000
   Belgium
   EMail: wouter@vooruit.be