Roll
Internet-Draft
Intended status: Informational
Expires: August 9, 2013

A. Brandt
Sigma Designs
E. Baccelli
INRIA
R. Cragie
Gridmerge

February 5, 2013

Applicability Statement: The use of RPL-P2P in Home and Building Control
draft-brandt-roll-rpl-applicability-home-building-03

Abstract

The purpose of this document is to provide guidance in the use of RPL-P2P to implement the features required in building and home environments.

Status of this Memo

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

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

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

This Internet-Draft will expire on August 9, 2013.

Copyright Notice

described in the Simplified BSD License.

Table of Contents

1. Introduction

Home automation and building control application spaces share a substantial number of properties. The purpose of this document is to give guidance in the use of RPL-P2P to provide the features required by the requirements documents "Home Automation Routing Requirements in Low-Power and Lossy Networks" [RFC5826] and "Building Automation Routing Requirements in Low-Power and Lossy Networks" [RFC5867].

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 RFC 2119.

1.2. Overview of requirements

Applicable requirements are described in [RFC5826] and [RFC5867].

1.3. Out of scope requirements

The considered network diameter is limited to a max diameter of 10 hops and a typical diameter of 5 hops, which captures the most common cases in home automation and building control networks.

This document does not consider the applicability of RPL-related specifications for urban and industrial applications [RFC5548], [RFC5673], which may exhibit significantly larger network diameters.

2. Deployment Scenario

A typical home automation network is less than 100 nodes. Large building deployments may span 10,000 nodes but to ensure uninterrupted service of light and air conditioning systems in individual zones of the building, nodes are organized in subnetworks. Each subnetwork in a building automation deployment is typically less than 200 nodes and rarely more than 500 nodes.

The main purpose of the network is to provide control over light and heating/cooling resources. User intervention may be enabled via wall controllers combined with movement, light and temperature sensors to enable automatic adjustment of window blinds, reduction of room temperature, etc.

Alarm systems are also important applications in home and building networks.

2.1. Network Topologies

The typical home automation network or building control subnetwork is a mesh network with a border router located at a convenient place in the home. In a building control network there may be several redundant border routers. The network often consists in a number of overlapping wireless subnetworks. Two types of routing topologies may exist in each subnetwork (i) a tree-shaped collection of routes spanning from a central building controller via the border router, on to destination nodes in the subnetwork, and/or (ii) a flat, un-directed collection of intra-network routes between arbitrary nodes in the subnetwork.

Nodes in Home and Building automation networks are typically inexpensive devices with extremely low memory capacities, such as individual wall switches. Only a few nodes (such as multi-purpose remote controls for instance) are more expensive devices, which can afford more memory capacity.

2.2. Traffic Characteristics

Traffic may enter the network from a central controller or it may originate from an intra-network node, such as a wall switch. The majority of traffic is light-weight point-to-point control style; e.g. Put-Ack or Get-Response. There are however exceptions. Bulk data transfer is used for firmware update and logging. Multicast is used for service discovery or to control groups of nodes, such as light fixtures. Firmware updates enter the network while logs leave the network.

2.2.1. Human user responsiveness

While airconditioning and other environmental-control applications may accept certain response delays, alarm and light control applications may be regarded as soft real-time systems. A slight delay is acceptable, but the perceived quality of service degrades significantly if response times exceed 250 msec. If the light does not turn on at short notice, a user will activate the controls again, causing a sequence of commands such as Light{on,off,on,off,..} or Volume{up,up,up,up,up,...}.

The reactive discovery features of RPL-P2P ensures that commands are normally delivered within the 250msec time window and when connectivity needs to be restored, it is typically completed within seconds.

2.2.2. Source-sink (SS) communication paradigm

Source-sink (SS) traffic is a common traffic type in home and building networks. The traffic is generated by environmental sensors which push periodic readings to a central server. The readings may be used for pure logging, or more often, to adjust light, heating and ventilation. Alarm sensors also generate SS style traffic.

With regards to message latency, most SS transmissions can tolerate worst-case delays measured in tens of seconds. Alarm sensors, however, represent one exception.

2.2.3. Peer-to-peer (P2P) communication paradigm

Peer-to-peer (P2P) traffic is a common traffic type in home networks. Some building networks also rely on P2P traffic while others send all control traffic to a local controller box for advanced scene and group control; thus generating more SS and P2MP traffic.

P2P traffic is typically generated by remote controls and wall controllers which push control messages directly to light or heat sources. P2P traffic has a strong requirement for low latency since P2P traffic often carries application messages that are invoked by humans. As mentioned in Section 2.2.1 application messages need to be delivered within less than a second - even when a route repair is needed before the message can be delivered.

2.2.4. Peer-to-multipeer (P2MP) communication paradigm

Peer-to-multipeer (P2MP) traffic is common in home and building networks. Often, a wall switch in a living room responds to user activation by sending commands to a number of light sources simultaneously.

Individual wall switches are typically inexpensive devices with extremely low memory capacities. Multi-purpose remote controls for use in a home environment typically have more memory but such devices are asleep when there is no user activity. RPL-P2P reactive discovery allows a node to wake up and find new routes within a few seconds while memory constrained nodes only have to keep routes to relevant targets.

2.3. Link layer applicability

This document applies to [IEEE802.15.4] and [G.9959] which are adapted to IPv6 by the adaption layers [RFC4944] and [I-D.lowpanz].

Due to the limited memory of a majority of devices (such as

individual light-switches) RPL-P2P MUST be used with source routing in non-storing mode. The abovementioned adaptation layers leverage on the compression capabilities of [RFC6554] and [RFC6282]. Header compression allows small IP packets to fit into a single layer 2 frame even when source routing is used. A network diameter limited to 5 hops helps achieving this.

Packet drops are often experienced in the targeted environments. ICMP, UDP and even TCP flows may benefit from link layer unicast acknowledgments and retransmissions. Link layer unicast acknowledgments MUST be enabled when [IEEE802.15.4] or [G.9959] is used with RPL-P2P.

3. Using RPL-P2P to meet requirements

RPL-P2P MUST be used in home and building networks, as P2P traffic is substantial and route repair must be completed within seconds. RPL-P2P provides a reactive mechanism for quick, efficient and root-independent route discovery/repair. The use of RPL-P2P furthermore allows data traffic to avoid having to go through a central region around the root of the tree, and drastically reduces path length [SOFT11] [INTEROP12]. These characteristics are desirable in home and building automation networks because they substantially decrease unnecessary network congestion around the tree’s root.

4. RPL Profile for RPL-P2P

RPL-P2P MUST be used in home and building networks. Non-storing mode allows for constrained memory in repeaters when source routing is used. Reactive discovery allows for low application response times even when on-the-fly route repair is needed.

4.1. RPL Features

4.1.1. RPL Instances

TBD.

4.1.2. Non-Storing Mode

Non-storing mode MUST be used to cope with the extremely constrained memory of a majority of nodes in the network (such as individual light switches).

4.1.3. DAO Policy

TBD.

4.1.4. Path Metrics

TBD.

4.1.5. Objective Function

OF0 MUST be supported and is the RECOMMENDED OF to use. Other Objective Functions MAY be used as well.

4.1.6. DODAG Repair

Since RPL-P2P only creates DODAGs on a temporary basis during route repair, there is no need to repair DODAGs.

4.1.7. Multicast

TBD.

4.1.8. Security

TBD.

4.1.9. P2P communications

RPL-P2P [RPL-P2P] MUST be used to accomodate P2P traffic, which is typically substantial in home and building automation networks.

4.2. Layer 2 features

Security MUST be applied at layer 2 for [IEEE802.15.4] and [G.9959]. Residential light control can accept a lower security level than other contexts (e.g. a nuclear research lab). Safety critical devices like electronic door locks SHOULD employ additional higher-layer security while light and heating devices may be sufficiently protected by a single network key. The border router MAY enforce access policies to limit access to the trusted LLN domain from the LAN.

4.2.1. Security functions provided by layer-2

TBD.

4.2.2. 6LowPAN options assumed

TBD.

4.2.3. MLE and other things

TBD.

4.3. Recommended Configuration Defaults and Ranges

TODO

5. Manageability Considerations

TODO

6. Security Considerations

TODO

6.1. Security Considerations during initial deployment

TODO: (This section explains how nodes get their initial trust anchors, initial network keys. It explains if this happens at the factory, in a deployment truck, if it is done in the field, perhaps like http://www.lix.polytechnique.fr/hipercom/SmartObjectSecurity/papers/CullenJennings.pdf)

6.2. Security Considerations during incremental deployment

TODO: (This section explains how that replaces a failed node takes on the dead nodes' identity, or not. How are nodes retired. How are nodes removed if they are compromised)

7. Other related protocols

Application transport protocols may be CoAP over UDP or equivalents. Typically, UDP is used for IP transport to keep down the application response time and bandwidth overhead.

Several features required by [RFC5826], [RFC5867] challenge the P2P paths provided by RPL. Appendix A reviews these challenges. In some cases, a node may need to spontaneously initiate the discovery of a path towards a desired destination that is neither the root of a DAG, nor a destination originating DAO signaling. Furthermore, P2P paths

provided by RPL are not satisfactory in all cases because they involve too many intermediate nodes before reaching the destination.

RPL-P2P [RPL-P2P] provides the features requested by [RFC5826] and [RFC5867]. RPL-P2P uses a subset of the frame formats and features defined for RPL [RFC6550] but may be combined with RPL frame flows in advanced deployments.

8. IANA Considerations

9. Acknowledgements

This document reflects discussions and remarks from several individuals including (in alphabetical order): Michael Richardson, Mukul Goyal, Jerry Martocci, Charles Perkins, and Zach Shelby

10. References

11. References

11.1. Normative References

[RFC5826] "Home Automation Routing Requirements in Low-Power and Lossy Networks".

[RFC5867] "Building Automation Routing Requirements in Low-Power and Lossy Networks".

[RFC5673] "Industrial Routing Requirements in Low-Power and Lossy Networks".

[RFC5548] "Routing Requirements for Urban Low-Power and Lossy Networks".

[IEEE802.15.4] "IEEE 802.15.4 - Standard for Local and metropolitan area networks -- Part 15.4: Low-Rate Wireless Personal Area Networks", <IEEE Standard 802.15.4>.

[RFC4944] "Transmission of IPv6 Packets over IEEE 802.15.4 Networks".

[G.9959] "ITU-T G.9959 Short range narrow-band digital radiocommunication transceivers - PHY and MAC layer

specifications", <ITU-T G.9959>.

[I-D.lowpanz] Brandt, A., "Transmission of IPv6 Packets over ITU-T G.9959 Networks", <draft-brandt-6man-lowpanz>.

[RFC6282] Hui, J. and Thubert, P., "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC6282 , September 2011.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and Manral, V., "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC6554 , March 2012.

[RFC6550] "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks".

[RPL-P2P] Goyal, M., Baccelli, E., Phillip, M., Brandt, A., and J. Martocci, "Reactive Discovery of Point-to-Point Routes in Low Power and Lossy Networks", draft-ietf-roll-p2p-rpl , May 2012.

11.2. Informative References

[SOFT11] Baccelli, E., Phillip, M., and M. Goyal, "The P2P-RPL Routing Protocol for IPv6 Sensor Networks: Testbed Experiments", Proceedings of the Conference on Software Telecommunications and Computer Networks, Split, Croatia, September 2011., September 2011.

[INTEROP12] Baccelli, E., Phillip, M., Brandt, A., Valev , H., and J. Buron , "Report on P2P-RPL Interoperability Testing", RR-7864 INRIA Research Report RR-7864, Janurary 2012.

Appendix A. RPL shortcomings in home and building deployments

This document reflects discussions and remarks from several individuals including (in alphabetical order): Charles Perkins, Jerry Martocci, Michael Richardson, Mukul Goyal and Zach Shelby.

A.1. Risk of undesired long P2P routes

The DAG, being a tree structure is formed from a root. If nodes residing in different branches have a need for communicating internally, DAG mechanisms provided in RPL [RFC6550] will propagate

traffic towards the root, potentially all the way to the root, and down along another branch. In a typical example two nodes could reach each other via just two router nodes but in unfortunate cases, RPL may send traffic three hops up and three hops down again. This leads to several undesired phenomena described in the following sections

A.1.1. Traffic concentration at the root

If many P2P data flows have to move up towards the root to get down again in another branch there is an increased risk of congestion the nearer to the root of the DAG the data flows. Due to the broadcast nature of RF systems any child node of the root is not just directing RF power downwards its subtree but just as much upwards towards the root; potentially jamming other MP2P traffic leaving the tree or preventing the root of the DAG from sending P2MP traffic into the DAG because the listen-before-talk link-layer protection kicks in.

A.1.2. Excessive battery consumption in source nodes

Battery-powered nodes originating P2P traffic depend on the route length. Long routes cause source nodes to stay awake for longer periods before returning to sleep. Thus, a longer route translates proportionally (more or less) into higher battery consumption.

A.2. Risk of delayed route repair

The RPL DAG mechanism uses DIO and DAO messages to monitor the health of the DAG. In rare occasions, changed radio conditions may render routes unusable just after a destination node has returned a DAO indicating that the destination is reachable. Given enough time, the next Trickle timer-controlled DIODAO update will eventually repair the broken routes. In a worst-case event this is however too late. In an apparently stable DAG, Trickle-timer dynamics may reduce the update rate to a few times every hour. If a user issues an actuator command, e.g. light on in the time interval between the last DAO message was issued the destination module and the time one of the parents sends the next DIO, the destination cannot be reached. Nothing in RPL kicks in to restore connectivity in a reactive fashion. The consequence is a broken service in home and building applications.

A.2.1. Broken service

Experience from the telecom industry shows that if the voice delay exceeds 250ms users start getting confused, frustrated andor annoyed. In the same way, if the light does not turn on within the same period of time, a home control user will activate the controls again,

causing a sequence of commands such as Light{on,off,off,on,off,..} or Volume{up,up,up,up,up,...} Whether the outcome is nothing or some unintended response this is unacceptable. A controlling system must be able to restore connectivity to recover from the error situation. Waiting for an unknown period of time is not an option. While this issue was identified during the P2P analysis it applies just as well to application scenarios where an IP application outside the LLN controls actuators, lights, etc.

Authors' Addresses

Anders Brandt
Sigma Designs

Email: abr@sdesigns.dk

Emmanuel Baccelli
INRIA

Email: Emmanuel.Baccelli@inria.fr

Robert Cragie
Gridmerge

Email: robert.cragie@gridmerge.com