

Roll
Internet-Draft
Intended status: Informational
Expires: April 9, 2015

A. Brandt
Sigma Designs
E. Baccelli
INRIA
R. Cragie
ARM
P. van der Stok
Consultant
October 6, 2014

**Applicability Statement: The use of the RPL protocol suite in Home
Automation and Building Control
draft-ietf-roll-applicability-home-building-05**

Abstract

The purpose of this document is to provide guidance in the selection and use of protocols from the RPL protocol suite to implement the features required for control 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 April 9, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Relationship to other documents	4
1.2.	Requirements language	4
1.3.	Terminology	4
1.4.	Required Reading	5
1.5.	Out of scope requirements	5
2.	Deployment Scenario	5
2.1.	Network Topologies	6
2.2.	Traffic Characteristics	7
2.2.1.	General	8
2.2.2.	Source-sink (SS) communication paradigm	8
2.2.3.	Publish-subscribe (PS, or pub/sub)) communication paradigm	9
2.2.4.	Peer-to-peer (P2P) communication paradigm	9
2.2.5.	Peer-to-multipeer (P2MP) communication paradigm	9
2.2.6.	Additional considerations: Duocast and N-cast	10
2.2.7.	RPL applicability per communication paradigm	10
2.3.	Layer-2 applicability	11
3.	Using RPL to meet Functional Requirements	11
4.	RPL Profile	12
4.1.	RPL Features	13
4.1.1.	RPL Instances	13
4.1.2.	Storing vs. Non-Storing Mode	13
4.1.3.	DAO Policy	13
4.1.4.	Path Metrics	14
4.1.5.	Objective Function	14
4.1.6.	DODAG Repair	14
4.1.7.	Multicast	14
4.1.8.	Security	15
4.1.9.	P2P communications	15
4.1.10.	IPv6 address configuration	15
4.2.	Layer 2 features	15
4.2.1.	Specifics about layer-2	16
4.2.2.	Services provided at layer-2	16
4.2.3.	6LowPAN options assumed	16
4.2.4.	MLE and other things	16
4.3.	Recommended Configuration Defaults and Ranges	16
4.3.1.	Trickle parameters	16
4.3.2.	Other Parameters	16
5.	MPL Profile	17
5.1.	Recommended configuration Defaults and Ranges	17

5.1.1.	Real-Time optimizations	17
5.1.2.	Trickle parameters	17
5.1.3.	Other parameters	18
6.	Manageability Considerations	18
7.	Security Considerations	19
7.1.	Security considerations during initial deployment	19
7.2.	Security Considerations during incremental deployment	20
7.3.	Security Considerations for P2P uses	20
7.4.	MPL routing	20
7.5.	RPL Security features	20
8.	Other related protocols	21
9.	IANA Considerations	21
10.	Acknowledgements	21
11.	Changelog	21
12.	References	23
12.1.	Normative References	23
12.2.	Informative References	25
Appendix A.	RPL shortcomings in home and building deployments	26
A.1.	Risk of undesired long P2P routes	26
A.1.1.	Traffic concentration at the root	27
A.1.2.	Excessive battery consumption in source nodes	27
A.2.	Risk of delayed route repair	27
A.2.1.	Broken service	27
Appendix B.	Communication failures	28
Authors' Addresses	29

[1.](#) Introduction

The primary purpose of this document is to give guidance in the use of the RPL protocol suite in two application domains:

- o Home automation
- o Building automation

The guidance is based on 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](#)] respectively. The Advanced Metering Infrastructure is also considered where appropriate. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most likely network structures. An abstract set of distinct communication paradigms is then used to frame the applicability domains.

Home automation and building automation application domains share a substantial number of properties:

- o In both domains, the network can be disconnected from the ISP and must still continue to provide control to the occupants of the home/building. Routing needs to be possible independent of the existence of a border router
- o Both domains are subject to unreliable links but require instant and very reliable reactions. This has impact on routing because of timeliness and multipath routing.

The differences between the two application domains mostly appear in commissioning, maintenance and the user interface, which do not typically affect routing. Therefore, the focus of this applicability document is on reliability, timeliness, and local routing.

1.1. Relationship to other documents

The ROLL working group has specified a set of routing protocols for Lossy and Low- resource Networks (LLN) [[RFC6550](#)]. This applicability text describes a subset of those protocols and the conditions under which the subset is appropriate and provides recommendations and requirements for the accompanying parameter value ranges.

In addition, an extension document has been produced specifically to provide a solution for reactive discovery of point-to-point routes in LLNs [[RFC6997](#)]. The present applicability document provides recommendations and requirements for the accompanying parameter value ranges.

A common set of security threats are described in [[I-D.ietf-roll-security-threats](#)]. The applicability statements complement the security threats document by describing preferred security settings and solutions within the applicability statement conditions. This applicability statement recommends more light weight security solutions and specify the conditions under which these solutions are appropriate.

1.2. 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](#)].

1.3. Terminology

This document uses terminology from [[RFC6997](#)], [[I-D.ietf-roll-trickle-mcast](#)], [[RFC7102](#)], [[IEEE802.15.4](#)], and [[RFC6550](#)].

1.4. Required Reading

Applicable requirements are described in [[RFC5826](#)] and [[RFC5867](#)]. A survey of the application field is described in [[BCsurvey](#)].

1.5. Out of scope requirements

The considered network diameter is limited to a maximum 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

The use of communications networks in buildings is essential to satisfy energy saving regulations. Environmental conditions of buildings can be adapted to suit the comfort of the individuals present inside. Consequently when no one is present, energy consumption can be reduced. Cost is the main driving factor behind deployment of wireless networking in buildings, especially in the case of retrofitting, where wireless connectivity saves costs incurred due to cabling and building modifications.

A typical home automation network is comprised of 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 typically organized in sub-networks. Each sub-network in a building automation deployment typically contains tens to hundreds of nodes, and for critical operations may operate independently from the other sub-networks.

The main purpose of the home or building automation network is to provide control over light and heating/cooling resources. User intervention via wall controllers is combined with movement, light and temperature sensors to enable automatic adjustment of window blinds, reduction of room temperature, etc. In general, the sensors and actuators in a home or building typically have fixed physical locations and will remain in the same home or building automation network.

People expect an immediate and reliable response to their presence or actions. For example, a light not switching on after entry into a room may lead to confusion and a profound dissatisfaction with the lighting product.

Monitoring of functional correctness is at least as important. Devices typically communicate their status regularly and send alarm messages notifying a malfunction of equipment or network.

In building control, the infrastructure of the building management network can be shared with the security/access, the IP telephony, and the fire/alarm networks. This approach has a positive impact on the operation and cost of the network; however, care should be taken to ensure that the availability of the building management network does not become compromised beyond the ability for critical functions to perform adequately.

In homes, the entertainment network for audio/video streaming and gaming has different requirements, where the most important requirement is the need for high bandwidth not typically needed for home or building control. It is therefore expected that the entertainment network in the home will mostly be separate from the control network, which also lessens the impact on availability of the control network

2.1. Network Topologies

In general, the home automation network or building control network consists of wired and wireless sub-networks. In large buildings especially, the wireless sub-networks can be connected to an IP backbone network where all infrastructure services are located, such as DNS, automation servers, etc.

The wireless sub-network can be configured according to any of the following topologies:

- o A stand-alone network of 10-100 nodes without border router. This typically occurs in the home with a stand-alone control network, in low cost buildings, and during installation of high end control systems in buildings.
- o A connected network with one border router. This configuration will happen in homes where home appliances are controlled from outside the home, possibly via a smart phone, and in many building control scenarios.
- o A connected network with multiple border routers. This will typically happen in installations of large buildings.

Many of the nodes are battery-powered and may be sleeping nodes which wake up according to clock signals or external events.

In a building control network, for large installation with multiple border routers, sub-networks often overlap both geographically and from a wireless coverage perspective. Due to two purposes of the network, (i) direct control and (ii) monitoring, there may exist two types of routing topologies in a given sub-network: (i) a tree-shaped collection of routes spanning from a central building controller via the border router, on to destination nodes in the sub-network; and/or (ii) a flat, un-directed collection of intra-network routes between functionally related nodes in the sub-network.

The majority of nodes in home and building automation networks are typically class 0 devices [[RFC7228](#)], such as individual wall switches. Only a few nodes (such as multi-purpose remote controls) are more expensive Class 1 devices, which can afford more memory capacity.

2.2. Traffic Characteristics

Traffic may enter the network originating from a central controller or it may originate from an intra-network node. 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, where firmware updates enter the network and logs leave the network. Group communication is used for service discovery or to control groups of nodes, such as light fixtures.

Often, there is a direct physical relation between a controlling sensor and the controlled equipment. For example the temperature sensor and room controller are located in the same room sharing the same climate conditions. Consequently, the bulk of senders and receivers are separated by a distance that allows one-hop direct path communication. A graph of the communication will show several fully connected subsets of nodes. However, due to interference, multipath fading, reflection and other transmission mechanisms, the one-hop direct path may be temporally disconnected. For reliability purposes, it is therefore essential that alternative n-hop communication routes exist for quick error recovery. (See [Appendix B](#) for motivation.)

Looking over time periods of a day, the networks are very lightly loaded. However, bursts of traffic can be generated by e.g. incessant pushing of the button of a remote control, the occurrence of a defect, and other unforeseen events. Under those conditions, the timeliness must nevertheless be maintained. Therefore, measures are necessary to remove any unnecessary traffic. Short routes are preferred. Long multi-hop routes via the border router, should be avoided whenever possible.

Group communication is essential for lighting control. For example, once the presence of a person is detected in a given room, lighting control applies to that room only and no other lights should be dimmed, or switched on/off. In many cases, this means that a multicast message with a 1-hop and 2-hop radius would suffice to control the required lights. The same argument holds for HVAC and other climate control devices. To reduce network load, it is advisable that messages to the lights in a room are not distributed any further in the mesh than necessary based on intended receivers.

An example of an office surface is shown in [[office-light](#)], and the current use of wireless lighting control products is shown in [[occuswitch](#)].

2.2.1. General

Whilst air conditioning and other environmental-control applications may accept response delays of tens of seconds or longer, 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 may activate the controls again, thus causing a sequence of commands such as `Light{on,off,on,off,...}` or `Volume{up,up,up,up,up,...}`. In addition the repetitive sending of commands creates an unnecessary loading of the network, which in turn increases the bad responsiveness of the network.

2.2.2. Source-sink (SS) communication paradigm

This paradigm translates to many sources sending messages to the same sink, sometimes reachable via the border router. As such, source-sink (SS) traffic can be present in home and building networks. The traffic may be generated by environmental sensors (often present in a wireless sub-network) which push periodic readings to a central server. The readings may be used for pure logging, or more often, processed to adjust light, heating and ventilation. Alarm sensors may also generate SS style traffic. The central server in a home automation network will be connected mostly to a wired network segment of the home network, although it is likely that cloud services will also be used. The central server in a building automation network may be connected to a backbone or be placed outside the building.

With regards to message latency, most SS transmissions can tolerate worst-case delays measured in tens of seconds. Fire detectors, however, represent an exception; For example, special provisions with

respect to the location of the Fire detectors and the smoke dampers need to be put in place to meet the stringent delay requirements.

2.2.3. Publish-subscribe (PS, or pub/sub)) communication paradigm

This paradigm translates to a number of devices expressing their interest for a service provided by a server device. For example, a server device can be a sensor delivering temperature readings on the basis of delivery criteria, like changes in acquisition value or age of the latest acquisition. In building automation networks, this paradigm may be closely related to the SS paradigm given that servers, which are connected to the backbone or outside the building, can subscribe to data collectors that are present at strategic places in the building automation network. The use of PS will probably differ significantly from installation to installation.

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

This paradigm translates to a device transferring data to another device often connected to the same sub-network. Peer-to-peer (P2P) traffic is a common traffic type in home automation networks. Most building automation networks rely on P2P traffic, described in the next paragraph. Other building automation networks rely on P2P control traffic between controls and a local controller box for advanced scene and group control. A local controller box can be further connected to service control boxes, thus generating more SS or PS 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 stringent 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 should be delivered within a few hundred milliseconds - even when connections fail momentarily.

2.2.5. Peer-to-multiplepeer (P2MP) communication paradigm

This paradigm translates to a device sending a message as many times as there are destination devices. Peer-to-multiplepeer (P2MP) traffic is common in home and building automation networks. Often, a thermostat in a living room responds to temperature changes by sending temperature acquisitions to several fans and valves consecutively. This paradigm is also closely related to the PS paradigm in the case where a single server device has multiple subscribers.

2.2.6. Additional considerations: Duocast and N-cast

This paradigm translates to a device sending a message to many destinations in one network transfer invocation. Multicast is well suited for lighting where a presence sensor sends a presence message to a set of lighting devices. Multicast increases the probability that the message is delivered within the strict time constraints. The recommended multicast algorithm (e.g. [[I-D.ietf-roll-trickle-mcast](#)]) assures that messages are delivered to ALL intended destinations.

2.2.7. RPL applicability per communication paradigm

In the case of the SS paradigm applied to a wireless sub-network to a server reachable via a border router, the use of RPL [[RFC6550](#)] in non-storing mode is appropriate. Given the low resources of the devices, source routing will be used from the border router to the destination in the wireless sub-network for messages generated outside the mesh network. No specific timing constraints are associated with the SS type messages so network repair does not violate the operational constraints. When no SS traffic takes place, it is good practice to load only RPL code enabling P2P mode of operation [[RFC6997](#)] to reduce the code size and satisfy memory requirements.

P2P-RPL [[RFC6997](#)] is required for all P2P and P2MP traffic taking place between nodes within a wireless sub-network (excluding the border router) to assure responsiveness. Source and destination devices are typically physically close based on room layout. Consequently, most P2P and P2MP traffic is 1-hop or 2-hop traffic. [Appendix A](#) explains why P2P-RPL is preferable to RPL for this type of communication. [Appendix B](#) explains why reliability measures such as multi-path routing are necessary even when 1-hop communication dominates.

Additional advantages of P2P-RPL for home and building automation networks are, for example:

- o Individual wall switches are typically inexpensive class 0 devices [[RFC7228](#)] 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. P2P-RPL 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.
- o The reactive discovery features of P2P-RPL ensure that commands are normally delivered within the 250 msec time window. When

connectivity needs to be restored, discovery is typically completed within seconds. In most cases, an alternative (earlier discovered) route will work and route rediscovery is not necessary.

- o Broadcast storms typically associated with route discovery for AODV are less disruptive for P2P-RPL. P2P-RPL has a "STOP" bit which is set by the target of a route discovery to notify all other nodes that no more DIOs should be forwarded for this temporary DAG. Something looking like a broadcast storm may happen when no target is responding however, in this case, the Trickle suppression mechanism kicks in, limiting the number of DIO forwards in dense networks.

Due to the limited memory of the majority of devices, P2P-RPL is preferably deployed with source routing in non-storing mode as explained in [Section 4.1.2](#).

Multicast with MPL [[I-D.ietf-roll-trickle-mcast](#)] is preferably deployed for N-cast over the wireless network. Configuration constraints that are necessary to meet reliability and timeliness with MPL are discussed in [Section 4.1.7](#).

[2.3. Layer-2 applicability](#)

This document applies to [[IEEE802.15.4](#)] and [[G.9959](#)] which are adapted to IPv6 by the adaption layers [[RFC4944](#)] and [[I-D.ietf-6lo-lowpanz](#)]. Other layer-2 technologies, accompanied by an "IP over Foo" specification, are also relevant provided there is no frame size issue, and there are link layer acknowledgements.

The above mentioned 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 to achieve this even while using source routing.

Dropped packets 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 are compulsory when [[IEEE802.15.4](#)] or [[G.9959](#)] is used with RPL and P2P-RPL.

[3. Using RPL to meet Functional Requirements](#)

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 signalling. Furthermore, P2P paths provided by RPL are not satisfactory in all cases because they involve too many intermediate nodes before reaching the destination.

P2P-RPL [[RFC6997](#)] is necessary in home automation and building control networks, as point-to-point style traffic is substantial and route repair needs to be completed within seconds. P2P-RPL provides a reactive mechanism for quick, efficient and root-independent route discovery/repair. The use of P2P-RPL 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 root of the tree.

When more reliability is required, P2P-RPL enables the establishment of multiple independent paths. For 1-hop destinations this means that one 1-hop communication and a second 2-hop communication take place via a neighbouring node. Such a pair of redundant communication paths can be achieved by using MPL where the source is a MPL forwarder, while a second MPL forwarder is 1 hop away from both the source and the destination node. When the source multicasts the message, it may be received by both the destination and the 2nd forwarder. The 2nd forwarder forwards the message to the destination, thus providing two routes from sender to destination.

To provide more reliability with multiple paths, P2P-RPL can maintain two independent P2P source routes per destination, at the source. Good practice is to use the paths alternately to assess their existence. When one P2P path has failed (possibly only temporarily), as described in [Appendix B](#), the alternative P2P path can be used without discarding the failed path. The failed P2P path, unless proven to work again, can be safely discarded after a timeout (typically 15 minutes). A new route discovery is done when the number of P2P paths is exhausted due to persistent link failures.

4. RPL Profile

P2P-RPL is necessary in home automation and building control networks. Its reactive discovery allows for low application response times even when on-the-fly route repair is needed. Non-storing mode is preferable to reduce memory consumption in repeaters with constrained memory when source routing is used.

4.1.1. RPL Features

An important constraint on the application of RPL is the presence of sleeping nodes.

For example, in a stand-alone network, the master node (or coordinator) providing the logical layer-2 identifier and unique node identifiers to connected nodes may be a remote control which returns to sleep once new nodes have been added. Due to the absence of the border router, there may be no global routable prefixes at all. Likewise, there may be no authoritative always-on root node since there is no border router to host this function.

In a network with a border router and many sleeping nodes, there may be battery powered sensors and wall controllers configured to contact other nodes in response to events and then return to sleep. Such nodes may never detect the announcement of new prefixes via multicast.

In each of the above mentioned constrained deployments, a link layer node (e.g. coordinator or master) assumes the role as authoritative root node, transmitting singlecast RAs with a ULA prefix information option to nodes during the joining process to prepare the nodes for a later operational phase, where a border router is added.

A border router is designed to be aware of sleeping nodes in order to support the distribution of updated global prefixes to such sleeping nodes.

4.1.1.1. RPL Instances

When operating P2P-RPL on a stand-alone basis, there is no authoritative root node maintaining a permanent RPL DODAG. A node necessarily joins at least one RPL instance, as a new, temporary instance is created during each P2P-RPL route discovery operation. A node can be designed to join multiple RPL instances.

4.1.1.2. Storing vs. Non-Storing Mode

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

4.1.1.3. DAO Policy

A node can be designed to join multiple RPL instances; in that case DAO policies may be needed. However, DAO policy is out of scope for this applicability statement.

4.1.4. Path Metrics

OF0 is the recommended metric. [[RFC6551](#)] provides other options. Using other objective functions than OF0 may affect inter-operability.

4.1.5. Objective Function

OF0 is the recommended Objective Function. Other Objective Functions should be used only when dictated by circumstances.

4.1.6. DODAG Repair

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

For SS traffic, local repair is sufficient. The accompanying process is known as poisoning and is described in [Section 8.2.2.5 of \[RFC6550\]](#). Given that the majority of nodes in the building do not physically move around, creating new DODAGs should not happen frequently.

4.1.7. Multicast

Commercial lighting deployments may have a need for multicast to distribute commands to a group of lights in a timely fashion. Several mechanisms exist for achieving such functionality; [[I-D.ietf-roll-trickle-mcast](#)] is the generally accepted protocol for home and building deployments. This section relies heavily on the conclusions of [[RT-MPL](#)].

The density of forwarders and the frequency of message generation are important aspects to obtain timeliness during control operations. A high frequency of message generation can be expected when a remote control button is incessantly pressed, or when alarm situations arise.

Guaranteeing timeliness is intimately related to the density of the MPL routers. In ideal circumstances the message is propagated as a single wave through the network, such that the maximum delay is related to the number of hops times the smallest repetition interval of MPL. Each forwarder that receives the message passes the message on to the next hop by repeating the message. When several copies of a message reach the forwarder, it is specified that the copy need not be repeated. Repetition of the message can be inhibited by a small value of k . To assure timeliness, the value of k should be chosen high enough to make sure that messages are repeated at the first arrival of the message in the forwarder. However, a network that is

too dense leads to a saturation of the medium that can only be prevented by selecting a low value of k . Consequently, timeliness is assured by choosing a relatively high value of k but assuring at the same time a low enough density of forwarders to reduce the risk of medium saturation. Depending on the reliability of the network channels, it is advisable to choose the network such that at least 2 forwarders per hop repeat messages to the same set of destinations.

There are no rules about selecting forwarders for MPL. In buildings with central management tools, the forwarders can be selected, but in the home is not possible to automatically configure the forwarder topology at the time of writing this document.

4.1.8. Security

In order to support low-cost devices and devices running on a battery, RPL uses either unsecured messages or secured messages. If RPL is used with unsecured messages, link layer security is a minimum security requirement (see [Section 7](#)). If RPL is used with secured messages, the following RPL security parameter values are recommended:

- o $T = '0'$: Do not use timestamp in the Counter Field.
- o $Algorithm = '0'$: Use CCM with AES-128
- o $KIM = '10'$: Use group key, Key Source present, Key Index present
- o $LVL = 0$: Use MAC-32

4.1.9. P2P communications

[RFC6997] is recommended to accommodate P2P traffic, which is typically substantial in home and building automation networks.

4.1.10. IPv6 address configuration

Assigned IP addresses follow IETF standards to be routable and unique within the routing domain.

4.2. Layer 2 features

No particular requirements exist for layer 2 but for the ones cited in the IP over Foo RFCs. (See [Section 2.3](#))

[4.2.1.](#) **Specifics about layer-2**

Not applicable

[4.2.2.](#) **Services provided at layer-2**

Not applicable

[4.2.3.](#) **6LowPAN options assumed**

Not applicable

[4.2.4.](#) **MLE and other things**

Not applicable

[4.3.](#) **Recommended Configuration Defaults and Ranges**

The following sections describe the recommended parameter values for P2P-RPL and Trickle.

[4.3.1.](#) **Trickle parameters**

Trickle is used to distribute network parameter values to all nodes without stringent time restrictions. The recommended Trickle parameter values are:

- o DIOIntervalMin 4 = 16 ms
- o DIOIntervalDoublings 14
- o DIORedundancyConstant 1

[4.3.2.](#) **Other Parameters**

This section discusses the P2P-RPL parameters.

P2P-RPL [[RFC6997](#)] provides the features requested by [[RFC5826](#)] and [[RFC5867](#)]. P2P-RPL uses a subset of the frame formats and features defined for RPL [[RFC6550](#)] but may be combined with RPL frame flows in advanced deployments.

The recommended parameter values for P2P-RPL are:

- o MinHopRankIncrease 1
- o MaxRankIncrease 0

- o MaxRank 6
- o Objective function: OF0

5. MPL Profile

MPL is used to distribute values to groups of devices. Using MPL, based on the Trickle algorithm, timeliness should also be guaranteed. A deadline of 200 ms needs to be met when human action is followed by an immediately observable action such as switching on lights. The deadline needs to be met in a building where the number of hops from seed to destination varies between 1 and 10.

5.1. Recommended configuration Defaults and Ranges

5.1.1. Real-Time optimizations

When the network is heavily loaded, MAC delays contribute significantly to the end to end delays when MPL intervals between 10 to 100 ms are used to meet the 200 ms deadline. It is possible to set the number of buffers in the MAC to 1 and set the number of Back-off repetitions to 1. The number of MPL repetitions compensates for the reduced probability of transmission per MAC invocation [[RT-MPL](#)].

In addition, end to end delays and message losses are reduced, by adding a real-time layer between MPL and MAC to throw away too late messages and favour the most recent ones.

5.1.2. Trickle parameters

This section proposes values for the Trickle parameters used by MPL for the distribution of packets that need to meet a 200 ms deadline. The probability of meeting the deadline is increased by (1) choosing a small I_{min} value, (2) reducing the number of MPL intervals thus reducing the load, and (3) reducing the number of MPL forwarders to also reduce the load.

The consequence of this approach is that the value of k can be larger than 1 because network load reduction is already guaranteed by the network configuration.

Under the condition that the density of MPL repeaters can be limited, it is possible to choose low MPL repeat intervals (I_{min}) connected to k values such that $k > 1$. The minimum value of k is related to:

- o Value of I_{min} . The length of I_{min} determines the number of packets that can be received within the listening period of I_{min} .

- o Number of repeaters receiving the broadcast message from the same forwarder or seed. These repeaters repeat within the same I_{min} interval, thus increasing the c counter.

Within the first MPL interval a limited number, q , of messages can be transmitted. Assuming a 3 ms transmission interval, q is given by $q = I_{min}/3$. Assuming that at most q message copies can reach a given forwarder within the first repeat interval of length I_{min} , the related MPL parameter values are suggested in the following sections.

5.1.2.1. I_{min}

The recommended value is $I_{min} = 10 - 50$ ms.

When I_{min} is chosen much smaller, the interference between the copies leads to significant losses given that q is much smaller than the number of repeated packets. With much larger intervals the probability that the deadline will be met decreases with increasing hop count.

5.1.2.2. I_{max}

The recommended value is $I_{max} = 100 - 400$ ms.

The value of I_{max} is less important than the value of $max_expiration$. Given an I_{min} value of 10 ms, the 3rd MPL interval has a value of $10 \cdot 2 \cdot 2 = 40$ ms. When I_{min} has a value of 40 ms, the 3rd interval has a value of 160 ms. Given that more than 3 intervals are unnecessary, the I_{max} does not contribute much to the performance.

5.1.3. Other parameters

Other parameters are the k parameter and the $max_expiration$ parameter.

$k > q$ (see condition above). Under this condition and for small I_{min} , a value of $k=2$ or $k=3$ is usually sufficient to minimize the losses of packets in the first repeat interval.

$max_expiration = 2 - 4$. Higher values lead to more network load while generating copies which will probably not meet their deadline.

6. Manageability Considerations

Manageability is out of scope for home network scenarios. In building automation scenarios, central control could be applied based on MIBs.

7. Security Considerations

This section refers to the security considerations of [[RFC6997](#)], [[RFC6550](#)], [[I-D.ietf-roll-trickle-mcast](#)], and the counter measures discussed in sections 6 and 7 of [[I-D.ietf-roll-security-threats](#)].

Communications network security is based on providing integrity protection and encryption to messages. This can be applied at various layers in the network protocol stack based on using various credentials and a network identity.

The credentials which are relevant in the case of RPL are: (i) the credential used at the link layer in the case where link layer security is applied (see [Section 7.1](#)) or (ii) the credential used for securing RPL messages. In both cases, the assumption is that the credential is a shared key. Therefore, a mechanism is required which allows secure distribution of a shared key and configuration of network identity. Both can rely on: (i) pre-installation using an out-of-band method, (ii) delivered securely when a device is introduced into the network or (iii) delivered securely by a trusted neighbouring device. The shared key **MUST** be stored in a secure fashion which makes it difficult to be read by an unauthorized party.

Securely delivering a key requires a delivery mechanism that has data origin authentication, confidentiality and integrity protection. On reception of the delivered key, freshness of the delivered key needs to be ensured. Securely storing a key requires a storage mechanism that has confidentiality and integrity protection and is only accessible by an authorized party.

The network security domain is typically distinct from the application security domains within the network, of which there may be more than one. For this reason, end-to-end security between applications is recommended by using DTLS [[RFC6347](#)] or TLS [[RFC5246](#)].

7.1. Security considerations during initial deployment

Wireless mesh networks are typically secured at the link layer in order to prevent unauthorized parties from accessing the information exchanged over the links. It is good practice to create a network of nodes which share the same keys for link layer security and exclude nodes sending unsecured messages. With per-message data origin authentication, it is possible to prevent unauthorized nodes joining the mesh.

At initial deployment the network is secured by consecutively securing nodes at the link layer, thus building a network of secured nodes. The Protocol for carrying Authentication for Network Access

(PANA) Relay Element [[RFC6345](#)] in conjunction with PANA [[RFC5191](#)] provides a framework for network access and delivery of common link keys. A new DTLS-based protocol is proposed in [[I-D.kumar-dice-dtls-relay](#)].

For building control an installer will probably use an installation tool that establishes a secure communication path with the joining node. In the home, nodes can be visually inspected by the home owner and simple measures like pushing buttons simultaneously on joint and joining devices is probably sufficient.

This recommendation is in line with the countermeasures described in section 6.1.1 of [[I-D.ietf-roll-security-threats](#)]

[7.2.](#) Security Considerations during incremental deployment

When nodes are lost, no additional security measures are needed, the network remains secure as before by not allowing the addition of new nodes. New nodes can be added by using the same protocols used for initial deployment. Some protocols may need a state change to a subset of the secured nodes, other protocols only need the presence of a "trusted" installation node [[RFC6345](#)], [[RFC5191](#)], or [[I-D.kumar-dice-dtls-relay](#)].

[7.3.](#) Security Considerations for P2P uses

Refer to the security considerations of [[RFC6997](#)]. Many initiatives are under way to provide lighter weight security such as: [[I-D.ietf-dice-profile](#)] and [[I-D.keoh-dice-multicast-security](#)]

[7.4.](#) MPL routing

The routing of MPL is determined by the enabling of the interfaces for specified Multicast addresses. The specification of these addresses can be done via a CoAP application as specified in [[I-D.ietf-core-groupcomm](#)]. An alternative is the creation of a MPL MIB and use of SNMPv3 [[RFC3411](#)] or equivalent techniques to specify the Multicast addresses in the MIB. The application of security measures for the specification of the multicast addresses assures that the routing of MPL packets is secured.

[7.5.](#) RPL Security features

This section follows the structure of [section 7](#), "RPL security features" of [[I-D.ietf-roll-security-threats](#)], where a thorough analysis of security threats and proposed counter measures relevant to RPL and MPL are done.

In accordance with section 7.1 of [[I-D.ietf-roll-security-threats](#)], "Confidentiality features", a secured RPL protocol implements payload protection, as explained in [Section 7](#) of this document. The attributes key-length and life-time of the keys depend on operational conditions, maintenance and installation procedures.

[Section 7.1](#) and [Section 7.2](#) of this document recommend link-layer measures to assure integrity in accordance with section 7.2 of [[I-D.ietf-roll-security-threats](#)], "Integrity features".

The provision of multiple paths recommended in [section 7.3](#) "Availability features" of [[I-D.ietf-roll-security-threats](#)] is also recommended from a reliability point of view. Randomly choosing paths is a possibility.

Key management discussed in [section 7.4](#), "Key Management" of [[I-D.ietf-roll-security-threats](#)], is not standardized and discussions continue.

[Section 7.5](#), "Considerations on Matching Application Domain Needs" of [[I-D.ietf-roll-security-threats](#)] applies as such.

[8.](#) Other related protocols

Application and transport protocols used in home and building automation domains are expected to mostly consist in CoAP over UDP, or equivalents. Typically, UDP is used for IP transport to keep down the application response time and bandwidth overhead. CoAP is used at the application layer to reduce memory footprint and bandwidth requirements.

[9.](#) IANA Considerations

No considerations for IANA pertain to this document.

[10.](#) Acknowledgements

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

[11.](#) Changelog

Changes from version 0 to version 1.

- o Adapted section structure to template.

- o Standardized the reference syntax.
- o [Section 2.2](#), moved everything concerning algorithms to [section 2.2.7](#), and adapted text in 2.2.1-2.2.6.
- o Added MPL parameter text to [section 4.1.7](#) and [section 4.3.1](#).
- o Replaced all TODO sections with text.
- o Consistent use of border router, monitoring, home- and building network.
- o Reformulated security aspects with references to other publications.
- o MPL and RPL parameter values introduced.

Changes from version 1 to version 2.

- o Clarified common characteristics of control in home and building.
- o Clarified failure behaviour of point to point communication in appendix.
- o Changed examples, more hvac and less lighting.
- o Clarified network topologies.
- o replaced reference to smart_object paper by reference to I-D.roll-security-threats
- o Added a concise definition of secure delivery and secure storage
- o text about securing network with PANA

Changes from version 2 to version 3.

- o Changed security section to follow the structure of security threats draft.
- o Added text to DODAG repair sub-section

Changes from version 3 to version 4.

- o Renumbered sections and moved text to conform to applicability template
- o Extended MPL parameter value text

- o Added references to building control products

Changes from version 4 to version 5.

- o Large editing effort to streamline text
- o Rearranged Normative and Informative references
- o Replaced [RFC2119](#) terminology by non-normative terminology
- o Rearranged text of [section 7](#), 7.1, and 7.2 to agree with the intention of [section 7.2](#)

[12.](#) References

[12.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.
- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", [RFC 5191](#), May 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", [RFC 5548](#), May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", [RFC 5673](#), October 2009.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5826](#), April 2010.
- [RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5867](#), June 2010.

- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), September 2011.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", [RFC 6347](#), January 2012.
- [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.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", [RFC 6554](#), March 2012.
- [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.
- [RFC6998] Goyal, M., Baccelli, E., Brandt, A., and J. Martocci, "A Mechanism to Measure the Routing Metrics along a Point-to-Point Route in a Low-Power and Lossy Network", [RFC 6998](#), August 2013.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", [RFC 7102](#), January 2014.
- [I-D.ietf-roll-trickle-mcast]
Hui, J. and R. Kelsey, "Multicast Protocol for Low power and Lossy Networks (MPL)", [draft-ietf-roll-trickle-mcast-09](#) (work in progress), April 2014.
- [I-D.ietf-roll-security-threats]
Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, "A Security Threat Analysis for Routing Protocol for Low-power and lossy networks (RPL)", [draft-ietf-roll-security-threats-11](#) (work in progress), October 2014.

[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>.

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

12.2. Informative References

[RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, [RFC 3411](#), December 2002.

[RFC6345] Duffy, P., Chakrabarti, S., Cragie, R., Ohba, Y., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA) Relay Element", [RFC 6345](#), August 2011.

[RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", [RFC 7228](#), May 2014.

[I-D.ietf-6lo-lowpanz]

Brandt, A. and J. Buron, "Transmission of IPv6 packets over ITU-T G.9959 Networks", [draft-ietf-6lo-lowpanz-07](#) (work in progress), September 2014.

[I-D.ietf-dice-profile]

Tschöfenig, H., "A Datagram Transport Layer Security (DTLS) 1.2 Profile for the Internet of Things", [draft-ietf-dice-profile-04](#) (work in progress), August 2014.

[I-D.keoh-dice-multicast-security]

Keoh, S., Kumar, S., Garcia-Morchon, O., Dijk, E., and A. Rahman, "DTLS-based Multicast Security in Constrained Environments", [draft-keoh-dice-multicast-security-08](#) (work in progress), July 2014.

[I-D.kumar-dice-dtls-relay]

Kumar, S., Keoh, S., and O. Garcia-Morchon, "DTLS Relay for Constrained Environments", [draft-kumar-dice-dtls-relay-01](#) (work in progress), April 2014.

[I-D.ietf-core-groupcomm]

Rahman, A. and E. Dijk, "Group Communication for CoAP", [draft-ietf-core-groupcomm-25](#) (work in progress), September 2014.

- [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.
- [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, January 2012.
- [RT-MPL] van der Stok, P., "Real-Time multicast for wireless mesh networks using MPL", White paper, <http://www.vanderstok.org/papers/Real-time-MPL.pdf>, April 2014.
- [occuswitch] Lighting, Philips., "OccuSwitch wireless", Brochure, http://www.philipslightingcontrols.com/assets/cms/uploads/files/osw/MK_OSWNETBROC_5.pdf, May 2012.
- [office-light] Clanton and Associates, ., "A Life Cycle Cost Evaluation of Multiple Lighting Control Strategies", Wireless Lighting Control, http://www.daintree.net/wp-content/uploads/2014/02/clanton_lighting_control_report_0411.pdf, February 2014.
- [RTN2011] Holtman, K. and P. van der Stok, "Real-time routing for low-latency 802.15.4 control networks", International Workshop on Real-Time Networks; Euromicro Conference on Real-Time Systems, July 2011.
- [MEAS] Holtman, K., "Connectivity loss in large scale IEEE 802.15.4 network", Private Communication, November 2013.
- [BCsurvey] Kastner, W., Neugschwandtner, G., Soucek, S., and H. Newman, "Communication Systems for Building Automation and Control", Proceedings of the IEEE Vol 93, No 6, June 2005.

[Appendix A.](#) RPL shortcomings in home and building deployments

[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 [[RFC6998](#)]. 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 sub-tree 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 DIO/DAO update will eventually repair the broken routes, however this may not occur in a timely manner appropriate to the application. 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. There is no mechanism in RPL to initiate restoration of 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 and/or 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.

Appendix B. Communication failures

Measurements on the connectivity between neighbouring nodes are discussed in [\[RTN2011\]](#) and [\[MEAS\]](#).

The work is motivated by the measurements in literature which affirm that the range of an antenna is not circle symmetric but that the signal strength of a given level follows an intricate pattern around the antenna, and there may be holes within the area delineated by an iso-strength line. It is reported that communication is not symmetric: reception of messages from node A by node B does not imply reception of messages from node B by node A. The quality of the signal fluctuates over time, and also the height of the antenna within a room can have consequences for the range. As function of the distance from the source, three regions are generally recognized: (1) a clear region with excellent signal quality, (2) a region with fluctuating signal quality, (3) a region without reception. In the text below it is shown that installation of meshes with neighbours in the clear region is not sufficient.

[\[RTN2011\]](#) extends existing work by:

- o Observations over periods of at least a week,
- o Testing links that are in the clear region,
- o Observation in an office building during working hours,
- o Concentrating on one-hop and two-hop routes.

Eight nodes were distributed over a surface of 30m². All nodes are at one hop distance from each other and are situated in the clear region of each other. Each node sends messages to each of its neighbours, and repeats the message until it arrives. The latency of the message was measured over periods of at least a week. It is noticed that latencies longer than a second occurred without apparent reasons, but only during working days and never in the weekends. Bad periods could last for minutes. By sending messages via two paths:

(1) one hop path directly, and (2) two hop path via a randomly chosen neighbour, the probability of delays larger than 100 ms decreased significantly.

The conclusion is that even for 1-hop communication between not too distant "Line of Sight" nodes, there are periods of low reception in which communication deadlines of 200 ms are exceeded. It pays to send a second message over a 2-hop path to increase the reliability of timely message transfer.

[MEAS] confirms that temporary bad reception by close neighbours can occur within other types of areas. Nodes were installed on the ceiling in a grid with a distance of 30-50 cm between nodes. 200 nodes were distributed over an area of 10m x 5m. It clearly transpired that with increasing distance the probability of reception decreases. At the same time a few nodes furthest away from the sender had a high probability of message reception, while some close neighbours of the sender did not receive messages. The patterns of clear reception nodes evolved over time.

The conclusion is that even for direct neighbours reception can temporarily be bad during periods of several minutes. For a reliable and timely communication it is imperative to have at least two communication paths available (e.g. two hop paths next to the 1-hop path for direct neighbours).

Authors' Addresses

Anders Brandt
Sigma Designs

Email: abr@sdesigns.dk

Emmanuel Baccelli
INRIA

Email: Emmanuel.Baccelli@inria.fr

Robert Cragie
ARM

Email: robert.cragie@gridmerge.com

Peter van der Stok
Consultant

Email: consultancy@vanderstok.org