

roll
Internet-Draft
Intended status: Informational
Expires: August 10, 2015

P. van der Stok
Consultant
R. Cragie
Gridmerge
February 6, 2015

Forwarder policy for multicast with admin-local scope in the Multicast
Protocol for Low power and Lossy Networks (MPL)
draft-ietf-roll-admin-local-policy-03

Abstract

The purpose of this document is to specify an automated policy for the routing of Multicast Protocol for Low power and Lossy Networks (MPL) multicast messages with admin-local scope in a border router.

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 10, 2015.

Copyright Notice

Copyright (c) 2015 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	2
1.1. Requirements Language	4
1.2. Terminology and Acronyms	4
2. Network identifier	4
2.1. IEEE 802.15.4	4
2.2. IEEE 802.11	5
2.3. ITU-T G.9959	5
2.4. BLUETOOTH Low Energy	5
3. MPL4 router	5
3.1. MPL interface parameters	5
3.2. Determination of MPL4 zone	6
4. Admin-Local policy	7
4.1. Legal multicast messages	7
4.2. Forwarding legal packets	7
4.2.1. MPL message	8
4.2.2. Multicast messages without MPL option	8
4.3. Encryption rules	9
5. MPL domains and zones	9
6. Default parameter values	10
7. Security Considerations	10
8. IANA Considerations	12
9. Acknowledgements	12
10. Change log	12
11. References	13
11.1. Normative References	13
11.2. Informative References	14
Authors' Addresses	15

1. Introduction

Multicast scopes are defined in [RFC4291]. The [RFC7346] extends the scope definition with the text:

"Interface-Local, Link-Local, and Realm-Local scope boundaries are automatically derived from physical connectivity or other, non-multicast related configuration. Global scope has no boundary. The boundaries of all other non-reserved scopes of Admin-Local or larger are administratively configured."

The admin-local scope must therefore be administratively configured. In this document "administratively configured" does not imply actions by a human beyond installing the here specified protocol. "Administratively configured" means an automatic derivation as described in this document.

This draft describes an automated policy for the Multicast Protocol for Low power and Lossy Networks (MPL) [[I-D.ietf-roll-trickle-mcast] forwarding of multicast messages with admin-local scope within a border router that lies between a network running MPL and some other network. This wish is in line with the autonomous networking ideas presented in [I-D.irtf-nmrg-an-gap-analysis].

The realm-local multicast address is currently used by MPL to propagate the multicast message to all receivers and forwarders within a mesh network. The multicast propagation is limited to a mesh network with a common layer-2. For example, a LoWPAN is defined by an IEEE 802.15.4 layer-2 mesh network, composed of all connected nodes sharing the same Personal Area Network (PAN) ID [RFC4944].

The network concept differs between mesh network technologies. This document maps a general network identifier to the specific network identifier of existing mesh technologies.

In current and projected deployments, there is a requirement to propagate a multicast message beyond the boundaries of the mesh network it originated in independent of the mesh technology.

Consider the case where propagation over two mesh networks is required. In one example, each mesh network has a border router and the two border routers are connected with an Ethernet link. In another example each mesh network is connected to its own network interface connected to the same border router. In both cases, an admin-local multicast message originating in one network needs to propagate into the other mesh network. The boundary of the admin-local scope is administratively configured.

This document describes an "MPL4 router" that forwards MPL messages with a multicast address with admin-local scope to all interfaces connected to links that connect to other MPL enabled interfaces. The MPL4 router enables all its interfaces for MPL messages and allocates an additional variable MPL_BLOCKED that permits(forbids) the forwarding of MPL messages.

The MPL4 router uses the following technique to establish over which links MPL4 messages must be forwarded. The MPL4 router listens on its interfaces for arrival of MPL4 messages. When MPL4 messages arrive over an interface, the MPL4 router includes this interface to the set of interfaces over which incoming MPL4 messages are forwarded. Regularly, the MPL4 router sends MPL4 messages over its interfaces to provoke the return of MPL4 messages to maintain or remove the interfaces in/from the set of forwarding interfaces.

It is expected that the private network of an organization, building, or home, is connected to the Internet via the edge routers provided by an ISP. The intention is that MPL messages with multicast addresses of admin-local scope are freely forwarded within the private network, but are never forwarded outside the private network by edge routers.

1.1. Requirements Language

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

1.2. Terminology and Acronyms

This document uses terminology defined in [I-D.ietf-roll-trickle-mcast] and [RFC7346]. In addition, the following terms are used in this document:

- o MPL4 refers to MPL with admin-local scope 4.
- o MPL4 message: an MPL DATA message with a destination multicast address of scope 4.
- o MPL4 zone: a convex zone of interconnected interfaces over which MPL messages with admin-local scope propagate. A MPL4 zone is bounded by a zone as defined in [RFC4007].
- o MPL4 router: automatically determines the MPL4 zone in which MPL messages with admin-local scope can be propagated.

2. Network identifier

Links may have the concept of a channel, for example in wireless networks such a channel is associated with a communication frequency. Additionally, for some link technologies, several networks can coexist using the same channel. For these link technologies, a network identifier exists. The network identifier is determined by the link technology specification. When no network identifier exists for a given link, the network identifier has the value "any".

2.1. IEEE 802.15.4

IPv6 over IEEE 802.15.4 is described in [RFC4944]. A LoWPAN is composed of the nodes connected by an IEEE 802.15.4 mesh sharing the same PAN ID. The PAN ID identifies a network in the IEEE 802.15.4 mesh. Several networks with different PAN IDs can coexist on the same channel [IEEE802.15.4]. The PAN ID of an interface is defined

when the interface is enabled. The value of the network identifier of an IEEE 802.15.4 link is the value of the PAN ID.

2.2. IEEE 802.11

IP over IEEE 802.11 is described in [RFC5416]. The Service Set Identifier (SSID) identifies a network in the IEEE 802.11 link. Several networks with different SSIDs can coexist on the same channel [IEEE802.11]. The SSID of an interface is defined when the interface is switched on. The value of the network identifier of a IEEE 802.11 link is the value of the SSID.

2.3. ITU-T G.9959

IPv6 over ITU-T G.9959 is specified in [I-D.ietf-6lo-lowpanz]. The HomeID identifies a network of connected nodes [G.9959]. Several HomeIDs can coexist within communication range, but nodes adhering to a network with a given HomeID cannot communicate with nodes adhering to a network with a different HomeID. The value of the network identifier of a G.9959 link is the value of the HomeID.

2.4. BLUETOOTH Low Energy

IPv6 over BLUETOOTH Low Energy (BTLE) is specified in [I-D.ietf-6lo-btle]. The medium is specified in [btle]. BTLE does not know the concept of multiple networks in one channel. The value of the network identifier of a BTLE link is "any".

3. MPL4 router

The concept of an MPL4 router serves to automatically determine the MPL4 zone in which MPL messages with a scope 4 multicast address can propagate. The MPL4 router periodically executes an algorithm that determines the presence of MPL interfaces on the links connected to its interfaces. When no MPL interfaces are present on a given link, the corresponding MPL interface is signalled as not being part of the MPL4 zone.

3.1. MPL interface parameters

One parameter is associated with every MPL interface in the MPL4 router, and two parameters are associated with the behaviour of the MPL4 router as a whole.

- o `MPL_BLOCKED`: Boolean value that indicates whether the associated interface belongs to the MPL4 zone.

- o MPL_CHECK_INT: integer that indicates the time interval between successive activations of the MPL4 router algorithm in seconds.
- o MPL_TO: integer that indicates the interval in which MPL messages are expected to be received in seconds.

3.2. Determination of MPL4 zone

All interfaces of the MPL4 router MUST be associated with following parameters coming from MPL protocol [I-D.ietf-roll-trickle-mcast]: PROACTIVE_FORWARDING, DATA_MESSAGE_IMIN, DATA_MESSAGE_IMAX, DATA_MESSAGE_K, DATA_MESSAGE_TIMER_EXPIRATIONS. At start-up of the MPL4 router, the parameters associated with all interfaces are assigned the following values: PROACTIVE_FORWARDING = true, MPL_BLOCKED = false. All interfaces MUST subscribe to the multicast addresses ALL_MPL_FORWARDERS scope 3 and scope 4.

The MPL4 router executes the following algorithm for each interface:

- o With a frequency determined by the value of MPL_CHECK_INT, the MPL4 router sends an MPL4 message on each interface with a header that includes the MPL option [I-D.ietf-roll-trickle-mcast] and is sent to multicast address ALL_MPL_FORWARDERS with scope 4.
- o When within an interval determined by the value of MPL_TO no MPL message is received, the value of MPL_BLOCKED is set to true.
- o At reception of an MPL4 message with a multicast address with scope 4, the value of MPL_BLOCKED of the receiving interface is set to false.

This protocol leads to a state where for each interface MPL_BLOCKED is set to false if and only if MPL enabled interfaces are connected to the link associated with the interface. When an MPL message is submitted to an MPL-enabled interface -called A- in the MPL router, the Trickle algorithm [RFC6206] is activated to send the MPL message. The MPL4 message with multicast address ALL_MPL_FORWARDERS scope 4 is accepted by every interface connected to the link that has subscribed to ALL_MPL_FORWARDERS with scope 4. On acceptance of the MPL4 message by an interface -called B-, the MPL4 message is returned with Trickle over interface B. Consequently, the MPL4 message is received by the originating interface A, after which MPL_BLOCKED is set to false.

When a new node is connected to the link, it can immediately send an MPL4 message, or can wait for the reception of an MPL4 message to announce its intention to be part of the MPL4 zone.

4. Admin-Local policy

The section starts with specifying what multicast messages arriving at an interface are legal. It continues with a description of forwarding legal admin-local multicast messages over other MPL interfaces.

The policy for forwarding admin-local multicast messages automatically to a MPL interface is specified as function of the state of the MPL interface and the multicast message. The state of the multicast message is determined by the presence of the MPL option [I-D.ietf-roll-trickle-mcast] and the destination multicast address. The state of the MPL interface is determined by the subscribed multicast addresses, the zone index [RFC4007], and the values of the PROACTIVE_FORWARDING parameter and the MPL_BLOCKED parameter of the MPL interface.

When zone is undefined or not enabled, all interfaces have the same zone index.

4.1. Legal multicast messages

Multicast messages can be created within the node by an application or can arrive at an interface.

A multicast message created at a source (MPL seed) is legal when it conforms to the properties described in section 9.1 of [I-D.ietf-roll-trickle-mcast].

A multicast message received at a given interface is legal when:

- o The message carries an MPL option (MPL message) and the incoming MPL interface is subscribed to the destination multicast address.
- o The message does not carry an MPL option, the multicast address is unequal to ALL_MPL_FORWARDERS scope 4 or scope 3, and the interface has expressed interest to receive messages with the specified multicast address via MLD [RFC3810] or via IGMP [RFC3376]. The message was sent on according to PIM-DM [RFC3973] or according to PIM-SM [RFC4601].

Illegal multicast messages are discarded.

4.2. Forwarding legal packets

A legal multicast message received at a given interface is assigned the network identifier of the interface of the incoming link . A

message that is created within the node is assigned the network identifier "any".

Two types of legal multicast messages are considered: (1) MPL messages, and (2) multicast messages which do not carry the MPL option.

4.2.1. MPL message

MPL messages are forwarded on MPL interfaces using the Trickle parameter values assigned to the MPL interface according to the following rules:

- o Link-local (scope 2) MPL messages are not forwarded.
- o Realm-local (scope 3) MPL messages are forwarded on all MPL interfaces that are subscribed to the same multicast address, have the same zone index, and have PROACTIVE-FORWARDING set to true, and the assigned network identifier of the multicast message is identical to the network identifier of the MPL interface, or the assigned network identifier of the multicast message is "any".
- o Admin-local (scope 4) MPL messages are forwarded on all MPL interfaces that are subscribed to the same multicast address, have the same zone index, have PROACTIVE-FORWARDING set to true, and have MPL_BLOCKED set to false.
- o MPL messages with a multicast scope of 5 or higher MUST encapsulate a message with the same multicast address without MPL option. The decapsulated message can be forwarded over an interface when the interface is subscribed with MLD to the same multicast address.

4.2.2. Multicast messages without MPL option

Multicast messages without MPL option are forwarded on MPL interfaces according to the following rules:

- o Link-local (scope 2) messages or realm-local (scope 3) multicast messages are not forwarded.
- o Admin-local (scope 4) multicast messages are encapsulated with a header carrying the MPL option and are forwarded on all MPL interfaces that are subscribed to the multicast address, have the same zone index, have PROACTIVE_FORWARDING set to true, and have MPL_BLOCKED set to false.

- o Multicast messages with a multicast scope of 5 or higher are encapsulated with a header carrying the MPL option and are forwarded on all MPL interfaces that are subscribed to the multicast address, have PROACTIVE_FORWARDING set to true, and have MPL_BLOCKED set to false. In addition these messages follow the Multicast forwarding rules as specified by PIM [RFC3973], [RFC4601] according to group specifications enabled by MLD [RFC3810] or IGMP [RFC3376].

4.3. Encryption rules

An incoming message protected at layer-2 MUST be subsequently re-protected at layer-2 at all outgoing interfaces. Incoming messages are integrity checked and optionally decrypted at the incoming interface at layer-2 using the keys and protection algorithm appropriate to the incoming interface's network and re-protected at the outgoing interface using the keys and protection algorithm appropriate to the outgoing interface's network. It may be necessary to assess the relative levels of protection on the respective interfaces and apply policy rules, for example to avoid downgrading security where one network has a lower level of security than another.

An incoming MPL4 messages which is not protected at layer-2 MUST NOT be re-protected at layer-2 at all outgoing interfaces.

5. MPL domains and zones

An MPL domain is a scope zone in which MPL interfaces subscribe to the same MPL Domain Address [I-D.ietf-roll-trickle-mcast]. In accordance with [RFC4007] a zone boundary passes through a node. For example, a small LLN node usually has one MPL mesh interface which is enabled to the ALL_MPL_FORWARDERS multicast address with a scope value of 3 (realm-local) [RFC7346]. The node interface belongs to the zone and the corresponding zone boundary does not pass through this node. In the border router with MPL interfaces enabled to the multicast address ALL_MPL_FORWARDERS with scope value 3, the zone includes usually this single interface and excludes all other interfaces. A notable exception is provided by a node where MPL interfaces of the same technology share the same network identifier. These interfaces belong to the same MPL4 zone when the interfaces share the same zone index.

In an MPL4 router, every MPL interface subscribes to the admin_local ALL_MPL_FORWARDERS multicast address next to the realm-local ALL_MPL_FORWARDERS address.

Every interface that belongs to an MPL domain that extends over border routers MUST be subscribed to the admin-local ALL_MPL_FORWARDERS address.

The MPL4 zone corresponding with the MPL multicast address ALL_MPL_FORWARDERS with scope 4 (Admin-local) applies to border routers with multiple interfaces, of which at least one interface is MPL enabled and is subscribed to multicast address ALL_MPL_FORWARDERS with scope 4. In a border router, all MPL enabled interfaces which subscribe to the ALL_MPL_FORWARDERS address with scope 4 and for which MPL_BLOCKED is false belong to the same MPL4 zone when the interfaces share the same zone index.

MPL4 messages remain bounded within a zone as defined in [RFC4007]. Consequently, MPL4 messages cannot be routed between interfaces belonging to different zones. When the concept of zone is unknown or disabled in a router, all interfaces belong to the same zone. For example, consider a router with 5 interfaces where interfaces A and B belong to zone 1 and interfaces C,D, and E belong to zone 2. MPL4 messages can be routed freely between interfaces A and B, and freely between C,D, and E. However, a MPL4 message MUST NOT be routed from Interface A to interface D.

6. Default parameter values

Three parameters are created in this draft. Their values are related to the Trickle timer intervals.

MPL_TO = DATA_MESSAGE_IMAX times 2. Which leaves the time to receive the second response message.

MPL_CHECK_INT = 5 minutes. Which means that a reaction to network malfunctioning happens within 5 minutes.

MPL_BLOCKED = true. Which means that the interface has not received MPL-enabled messages to include the interface to the MPL4 zone.

7. Security Considerations

The security considerations of [I-D.ietf-roll-trickle-mcast] also apply to MPL4 routers.

The sending of MPL4 messages by a malicious node can have unwanted consequences explained with the following example. It is not unusual for a wired (e.g. ethernet) link to be used between two floors or sections of an LLN, as radio propagation through reinforced concrete is generally poor. The MPL4 zone can thus envelop multiple routers, meshes and links. It is possible that a malicious node connects to a

wired link, on which no MPL enabled nodes are foreseen. In this example configuration, the malicious node can send MPL4 messages to the MPL4 router interfaces. When nothing is done, the MPL4 routers will consequently distribute MPL4 messages from one mesh over the wired link to the next mesh, although the wired link was not expected to transport MPL4 messages.

To understand the consequences of this unwanted behaviour, the following cases should be distinguished:

- o The source mesh uses layer-2 encryption.
- o The MPL4 router can be managed.

The four possible combinations are discussed below:

Layer-2 unsecured, Router unmanaged: In this case MPL4 messages are freely distributed over meshes and links which are interconnected by MPL4 routers within a zone. The MPL enabled (malicious) nodes can read all MPL4 messages and distribute MPL4 messages over a network limited by a zone. This situation can be acceptable for an isolated network, within a clearly defined space, where the connection of nodes can be tightly controlled. A completely wired LLN -- such as is seen in BACnet -- is an example of an unencrypted LLN which would be considered physically secure.

Layer-2 secured, Router unmanaged: In this case MPL4 messages are freely distributed over meshes and links, which are interconnected by MPL4 routers within a zone. Following the rules of Section 4.3, the MPL4 enabled (malicious) nodes can not read the MPL4 messages and MPL4 messages sent by the malicious node are not accepted by other nodes. This situation is acceptable for a home network or managed network extending over precisely one zone, occupying a clearly defined physical space, where ease of installation is important. In such a network, the presence of the malicious node is not different from any other malicious node, which tries to send messages over layer-2 protected links. Because the network occupies exactly one zone, the MPL4 message distribution cannot be extended outside the network.

Layer-2 unsecured, Router managed: In this case the distribution of MPL4 messages over MPL4 router interfaces can be limited to those interfaces, which a manager enabled for MPL and a set of multicast addresses. The malicious node cannot extend the distribution of MPL4 messages over unwanted interfaces. It is important that the handling of the interfaces by the manager is protected. However, MPL4 messages sent over the mesh can be interpreted by malicious nodes and malicious messages can be injected into the set of

meshes and links which are connected by the MPL4 routers for which the manager enabled the interfaces. This situation can be practical for interconnected links and meshes, which are connected to a LAN over a limited period, for example during installation of the interconnected meshes and links.

Layer-2 secured, Router managed: In this case the distribution of MPL4 messages over MPL4 router interfaces can be limited to those interfaces, which a manager enabled for MPL and a set of multicast addresses. Following the rules of Section 4.3, the malicious node cannot extend the distribution of MPL4 messages over unwanted interfaces and MPL4 messages sent by the malicious node are not accepted by other nodes. It is important that the handling of the interfaces by the manager is protected. The MPL enabled (malicious) nodes can not read the MPL4 messages and MPL4 messages sent by the malicious node are not accepted by other nodes. Dependent on the number of managed interfaces, the network can progressively pass from auto-configured to fully administratively controlled.

8. IANA Considerations

No considerations for IANA are formulated in this document.

9. Acknowledgements

This document reflects discussions and remarks from several individuals including (in alphabetical order): Scott Bradner, Esko Dijk, Adrian Farrel, Matthew Gillmore, Joel Halpern, Steve Hanna, Michael Richardson, and Pascal Thubert.

10. Change log

When published as a RFC, this section needs to be removed.

Version 03 - version 01

- o Explained MPL acronym
- o Added relation of MPL4 zone to zone as defined in [RFC4007]
- o Added a section on encryption rules
- o Revised and clarified the security considerations

Version 00 - version 01

- o Default parameter values declared

- o Security section extended
- o scope 5 of higher messages specified
- o messages with address ALL_MPL_FORWARDERS are not allowed from outside zone

Changes from personal version to WG version-00.

- o Aligned terminology with MPL terminology
[I-D.ietf-roll-trickle-mcast]
- o Text on MPL4 router included

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, March 2005.
- [RFC5416] Calhoun, P., Montemurro, M., and D. Stanley, "Control and Provisioning of Wireless Access Points (CAPWAP) Protocol Binding for IEEE 802.11", RFC 5416, March 2009.
- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, March 2011.
- [RFC7346] Droms, R., "IPv6 Multicast Address Scopes", RFC 7346, August 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-11 (work in progress), November 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>.

[IEEE802.11]

"IEEE 802.11 - Telecommunications and information exchange between systems Local and metropolitan area networks -- Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", <IEEE Standard 802.11>.

[G.9959]

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

[btle]

"BLUETOOTH Specification Version 4.0", <BLUETOOTH low energy>.

11.2. Informative References

[RFC3973]

Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.

[RFC4601]

Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[I-D.irtf-nmrg-an-gap-analysis]

Jiang, S., Carpenter, B., and M. Behringer, "Gap Analysis for Autonomic Networking", draft-irtf-nmrg-an-gap-analysis-03 (work in progress), December 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-08 (work in progress), October 2014.

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

Nieminen, J., Savolainen, T., Isomaki, M., Patil, B.,
Shelby, Z., and C. Gomez, "Transmission of IPv6 Packets
over BLUETOOTH(R) Low Energy", draft-ietf-6lo-btle-07
(work in progress), January 2015.

Authors' Addresses

Peter van der Stok
Consultant

Email: consultancy@vanderstok.org

Robert Cragie
Gridmerge

Email: robert.cragie@gridmerge.com

Roll
Internet-Draft
Intended status: Standards Track
Expires: April 9, 2017

N. Cam-Winget, Ed.
Cisco Systems
J. Hui
Nest
D. Popa
Itron, Inc
October 6, 2016

Applicability Statement for the Routing Protocol for Low Power and Lossy
Networks (RPL) in AMI Networks
draft-ietf-roll-applicability-ami-15

Abstract

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

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, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
1.2. Required Reading	3
1.3. Out of scope requirements	3
2. Routing Protocol for LLNs (RPL)	4
3. Description of AMI networks for electric meters	4
3.1. Deployment Scenarios	5
4. Smart Grid Traffic Description	7
4.1. Smart Grid Traffic Characteristics	7
4.2. Smart Grid Traffic QoS Requirements	8
4.3. RPL applicability per Smart Grid Traffic Characteristics	9
5. Layer 2 applicability	9
5.1. IEEE Wireless Technology	9
5.2. IEEE PowerLine Communication (PLC) technology	9
6. Using RPL to Meet Functional Requirements	10
7. RPL Profile	11
7.1. RPL Features	11
7.1.1. RPL Instances	11
7.1.2. DAO Policy	11
7.1.3. Path Metrics	11
7.1.4. Objective Function	11
7.1.5. DODAG Repair	12
7.1.6. Multicast	12
7.1.7. Security	12
7.2. Description of Layer-two features	13
7.2.1. IEEE 1901.2 PHY and MAC sub-layer features	13
7.2.2. IEEE 802.15.4 (g + e) PHY and MAC features	14
7.2.3. IEEE MAC sub-layer Security Features	15
7.3. 6LowPAN Options	16
7.4. Recommended Configuration Defaults and Ranges	17
7.4.1. Trickle Parameters	17
7.4.2. Other Parameters	18
8. Manageability Considerations	18
9. Security Considerations	19
9.1. Security Considerations during initial deployment	19
9.2. Security Considerations during incremental deployment	19
9.3. Security Considerations based on RPL's Threat Analysis	20
10. Privacy Considerations	20
11. IANA Considerations	20
12. Acknowledgements	20
13. References	21
13.1. Normative References	21
13.2. Informative references	22

Authors' Addresses	23
--------------------	----

1. Introduction

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

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

1.1. Requirements Language

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

1.2. Required Reading

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

NAN can use wireless communication technology in which case is using, from the IEEE 802.15.4 standard family, the [IEEE802.15.4g] PHY Layer amendment and [IEEE802.15.4e] MAC sub-layer amendment, specifically adapted to smart grid networks.

NAN can also use PLC (Power Line Communication) technology as an alternative to wireless communications. Several standards for PLC technology have emerged, such as [IEEE1901.2].

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

1.3. Out of scope requirements

The following are outside the scope of this document:

- o Applicability statement for RPL (Routing Protocol for Low Power and Lossy Networks) [RFC6550] in AMI networks composed of battery-powered devices (i.e., gas/water meters).
- o Applicability statement for RPL in AMI networks composed of a mix of AC powered devices (i.e., electric meters) and battery-powered meters (i.e., gas/water meters).
- o Applicability statement for RPL storing mode of operation in AMI networks.

2. Routing Protocol for LLNs (RPL)

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

RPL builds a Directed Acyclic Graph (DAG) routing structure rooted at a LBR (LLN Border Router), ensures loop-free routing, and provides support for alternate routes, as well as, for a wide range of routing metrics and policies.

RPL was designed to operate in energy-constrained environments and includes energy-saving mechanisms (e.g., Trickle timers) and energy-aware metrics. RPL's ability to support multiple different metrics and constraints at the same time enables it to run efficiently in heterogeneous networks composed of nodes and links with vastly different characteristics [RFC6551].

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

3. Description of AMI networks for electric meters

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

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

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

3.1. Deployment Scenarios

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

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

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

Figure 1: Typical NAN Topology

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

4. Smart Grid Traffic Description

4.1. Smart Grid Traffic Characteristics

In current AMI deployments, metering applications typically require all smart meters to communicate with a few head-end servers, deployed in the utility company data center. Head-end servers generate data traffic to configure smart data reading or initiate queries, and use unicast and multicast to efficiently communicate with a single device (i.e. Point-to-Point (P2P) communications) or groups of devices respectively (i.e., Point-to-Multipoint (P2MP) communication). The head-end server may send a single small packet at a time to the meters (e.g., a meter read request, a small configuration change, service switch command) or a series of large packets (e.g., a firmware download across one or even thousands of devices). The frequency of large file transfers (e.g., firmware download of all metering devices) is typically much lower than the frequency of sending configuration messages or queries. Each smart meter generates Smart Metering Data (SMD) traffic according to a schedule (e.g., periodic meter reads), in response to on-demand queries (e.g., on-demand meter reads), or in response to some local event (e.g., power outage, leak detection). Such traffic is typically destined to a single head-end server. The SMD traffic is thus highly asymmetric, where the majority of the traffic volume generated by the smart meters typically goes through the LBRs, and is directed from the smart meter devices to the head-end servers, in a MP2P (Mesh Peer to Peer)fashion. Current SMD traffic patterns are fairly uniform and well-understood. The traffic generated by the head-end server and destined to metering devices is dominated by periodic meter reads, while traffic generated by the metering devices is typically uniformly spread over some periodic read time-window.

Smart metering applications typically do not have hard real-time constraints, but they are often subject to bounded latency and stringent reliability service level agreements.

Distribution Automation (DA) applications typically involve a small number of devices that communicate with each other in a Point-to-

Point (P2P) fashion, and may or may not be in close physical proximity. DA applications typically have more stringent latency requirements than SMD applications.

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

4.2. Smart Grid Traffic QoS Requirements

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

- A) Meter-initiated traffic (Meter-to-head-end - M2HE)
 - B1) request is sent in point-to-point to a specific meter
 - B2) request is sent in multicast to a subset of meters
 - B3) request is sent in multicast to all meters

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

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

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

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

4.3. RPL applicability per Smart Grid Traffic Characteristics

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

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

5. Layer 2 applicability

5.1. IEEE Wireless Technology

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

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

5.2. IEEE PowerLine Communication (PLC) technology

The IEEE Std. 1901.2 standard specifies communications for low frequency (less than 500 kHz) narrowband power line devices via alternating current and direct current electric power lines. IEEE Std P1901.2 standard supports indoor and outdoor communications over low voltage line (line between transformer and meter, less than 1000 V), through transformer low-voltage to medium-voltage (1000 V up to 72 kV) and through transformer medium-voltage to low-voltage power

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

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

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

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

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

6. Using RPL to Meet Functional Requirements

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

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

RPL supports the following features:

- o Scalability: Large-scale networks characterized by highly directed traffic flows between each smart meter and the head-end servers in

the utility network. To this end, RPL builds a Directed Acyclic Graph (DAG) rooted at each LBR.

- o Zero-touch configuration: This is done through in-band methods for configuring RPL variables using DIO (DODAG Information Object) messages, and DIO message options [RFC6550].
- o The use of links with time-varying quality characteristics: This is accomplished by allowing the use of metrics that effectively capture the quality of a path (e.g., Expected Transmission Count (ETX)) and by limiting the impact of changing local conditions by discovering and maintaining multiple DAG parents, and by using local repair mechanisms when DAG links break.

7. RPL Profile

7.1. RPL Features

7.1.1. RPL Instances

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

7.1.2. DAO Policy

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

7.1.3. Path Metrics

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

Additional metrics may be defined in companion RFCs.

7.1.4. Objective Function

RPL relies on an Objective Function for selecting parents and computing path costs and rank. This objective function is decoupled from the core RPL mechanisms and also from the metrics in use in the network. Two objective functions for RPL have been defined at the

time of this writing, OF0 [RFC6552] and MRHOF [RFC6719], both of which define the selection of a preferred parent and backup parents, and are suitable for AMI deployments.

Additional objective functions may be defined in companion RFCs.

7.1.5. DODAG Repair

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

7.1.6. Multicast

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

7.1.7. Security

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

7.2. Description of Layer-two features

7.2.1. IEEE 1901.2 PHY and MAC sub-layer features

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

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

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

- o IEEE 1901.2 MAC frame is obtained by prepending a Segment Control Field to the IEEE 802.15.4 MAC header. One function of the Segment Control Field is to signal the use of the MAC sub-layer segmentation and reassembly.
- o IEEE 1901.2 MAC frames uses only the 802.15.4 MAC addresses with a length of 16 and 64 bits.
- o IEEE 1901.2 MAC sub-layer endorses the concept of Information Elements, as defined in [IEEE802.15.4e]. The format and use of Information Elements are not relevant to RPL applicability statement.

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

The IEEE 1901.2 PHY MTU size is variable and dependent on the PHY settings in use (e.g. bandwidth, modulation, tones, etc). As quoted from the IEEE 1901.2 specification: For CENELEC A/B, if MSDU size is more than 247 octets for robust OFDM (ROBO) and Super-ROBO modulations or more than 239 octets for all other modulations, the MAC layer shall divide the MSDU into multiple segments as described in 5.3.7. For FCC and ARIB, if the MSDU size meets one of the following conditions: a) For ROBO and Super-ROBO modulations, the MSDU size is more than 247 octets but less than 494 octets, b) For all other modulations, the MSDU size is more than 239 octets but less than 478 octets.

7.2.2. IEEE 802.15.4 (g + e) PHY and MAC features

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

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

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

- o Timeslotslotted channel hopping (TSCH): specifically designed for application domains such as process automation
- o Low latency deterministic networks (LLDN): for application domains such as factory automation
- o Deterministic and synchronous multi-channel extension (DSME): for general industrial and commercial application domains that includes Channel diversity to increase network robustness
- o Asynchronous multi-channel adaptation (AMCA): for large infrastructure application domains.

The MAC addressing scheme supports short (16-bits) addresses along with extended (64-bits) addresses. These addresses are assigned in different ways and is specified by specific standards organizations.

Information Elements, Enhanced Beacons and frame version 2, as defined in 802.15.4e, MUST be supported.

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

7.2.3. IEEE MAC sub-layer Security Features

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

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

The 802.15.4 security layer is handled at the media access control layer, below 6LoWPAN layer. The application specifies its security requirements by setting the appropriate control parameters into the radio/PLC stack. The 802.15.4 defines four packet types: beacon frames, data frames, acknowledgments frame, and command frames for the media access control layer. The 802.15.4 specification does not support security for acknowledgement frames; data frames, beacon frames and command frames can support integrity protection and confidentiality protection for the frames's data field. An application has a choice of security suites that control the type of security protection that is provided for the transmitted MAC frame. Each security suite offers a different set of security properties and guarantees, and ultimately different MAC frame formats. The 802.15.4 specification defines eight different security suites, outlined below. We can broadly classify the suites by the properties that they offer: no security, encryption only (AES-CTR), authentication only (AES-CBC-MAC), and encryption and authentication (AES-CCM). Each category that supports authentication comes in three variants depending on the size of the MAC (Message Authentication Control) that it offers. The MAC can be either 4, 8, or 16 bytes long. Additionally, for each suite that offers encryption, the recipient can optionally enable replay protection.

- o Null = No security.
- o AES-CTR = Encryption only, CTR mode.

- o AES-CBC-MAC-128 = No encryption, 128-bit MAC.
- o AES-CBC-MAC-64 = No encryption, 64-bit MAC.
- o AES-CCM-128 = Encryption and 128-bit MAC.
- o AES-CCM-64 = Encryption and 64-bit MAC.
- o AES-CCM-32 = Encryption and 32-bit MAC.

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

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

A cryptographic MAC is used to authenticate messages. While longer MACs lead to improved resiliency of the code, they also make packet size larger and thus take up bandwidth in the network. In constrained environments such as metering infrastructures, an optimum balance between security requirements and network throughput must be found.

7.3. 6LowPAN Options

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

7.4. Recommended Configuration Defaults and Ranges

7.4.1. Trickle Parameters

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

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

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

DIORedundancyConstant: AMI deployments SHOULD set DIORedundancyConstant to a value of at least 10. This is due to the larger chosen value for DIOIntervalMin and the proportional

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

7.4.2. Other Parameters

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

8. Manageability Considerations

Network manageability is a critical aspect of smart grid network deployment and operation. With millions of devices participating in the smart grid network, many requiring real-time reachability, automatic configuration, and lightweight network health monitoring and management are crucial for achieving network availability and efficient operation. RPL enables automatic and consistent configuration of RPL routers through parameters specified by the DODAG root and disseminated through DIO packets. The use of Trickle for scheduling DIO transmissions ensures lightweight yet timely propagation of important network and parameter updates and allows network operators to choose the trade-off point they are comfortable with respect to overhead vs. reliability and timeliness of network updates. The metrics in use in the network along with the Trickle Timer parameters used to control the frequency and redundancy of network updates can be dynamically varied by the root during the lifetime of the network. To that end, all DIO messages SHOULD contain a Metric Container option for disseminating the metrics and metric values used for DODAG setup. In addition, DIO messages SHOULD contain a DODAG Configuration option for disseminating the Trickle Timer parameters throughout the network. The possibility of dynamically updating the metrics in use in the network as well as the frequency of network updates allows deployment characteristics (e.g., network density) to be discovered during network bring-up and to be used to tailor network parameters once the network is operational rather than having to rely on precise pre-configuration. This also allows the network parameters and the overall routing protocol behavior to evolve during the lifetime of the network. RPL specifies

a number of variables and events that can be tracked for purposes of network fault and performance monitoring of RPL routers. Depending on the memory and processing capabilities of each smart grid device, various subsets of these can be employed in the field.

9. Security Considerations

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

In contrast to other types of LLNs, in smart grid networks centralized administrative control and access to a permanent secure infrastructure is available. As a result, smart grid networks are deployed with security mechanisms such as link-layer, transport layer and/or application-layer security mechanisms; while it is best practice to secure all layers, using RPL's secure mode may not be necessary. Failure to protect any of these layers can result in various attacks; without strong authentication of devices in the infrastructure can lead to uncontrolled and unauthorized access. Similarly, failure to protect the communication layers can enable passive (in wireless mediums) attacks as well as man-in-the-middle and active attacks.

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

9.1. Security Considerations during initial deployment

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

9.2. Security Considerations during incremental deployment

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

9.3. Security Considerations based on RPL's Threat Analysis

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

Like any secure network infrastructure, an AMI deployment's ability to address node impersonation, active man-in-the-middle attacks relies on mutual authentication and authorization process. The credential management from the manufacturing imprint of security credentials of smart meters to all credentials of nodes in the infrastructure, all credentials must be appropriately managed and classified through the authorization process to ensure beyond the identity of the nodes, that the nodes are behaving or 'acting' in their assigned roles.

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

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

10. Privacy Considerations

Privacy of information flowing through smart grid networks are subject to consideration. An evolving set of recommendations and requirements are being defined by different groups and consortiums; for example, the U.S. Department of Energy issued a document [DOEVCC] defining a process and set of recommendations to address privacy issues. As this document describes the applicability of RPL, the privacy considerations as defined in [I-D.ietf-6lo-privacy-considerations] and [EUPR] apply to this document and to AMI deployments.

11. IANA Considerations

This memo includes no request to IANA.

12. Acknowledgements

Matthew Gillmore, Laurent Toutain, Ruben Salazar, and Kazuya Monden were contributors and noted as authors in earlier drafts. The authors would also like to acknowledge the review, feedback, and

comments of Jari Arkko, Dominique Barthel, Cedric Chauvenet, Yuichi Igarashi, Philip Levis, Jeorjeta Jetcheva, Nicolas Dejean, and JP Vasseur.

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [IEEE802.15.4]
"IEEE Standard for Information technology-- Local and metropolitan area networks-- Specific requirements-- Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPANs)", IEEE Standard 802.15.4, September 2006.
- [IEEE802.15.4e]
"IEEE Standard for Local and metropolitan area networks-- Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer", IEEE Standard 802.15.4e, April 2012.
- [IEEE802.15.4g]
"IEEE Standard for Local and metropolitan area networks-- Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 3: Physical Layer (PHY) Specifications for Low-Data-Rate, Wireless, Smart Metering Utility Networks", IEEE Standard 802.15.4g, November 2012.
- [IEEE1901.2]
"IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", IEEE Standard 1901.2, December 2013.

[surveySG]

Gungor, V., Sahin, D., Kocak, T., Ergut, S., Buccella, C., Cecati, C., and G. Hancke, "A Survey on Smart Grid Potential Applications and Communication Requirements", Feb 2013.

13.2. Informative references

- [RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <<http://www.rfc-editor.org/info/rfc7731>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [RFC5548] Dohler, M., Ed., Watteyne, T., Ed., Winter, T., Ed., and D. Barthel, Ed., "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, DOI 10.17487/RFC5548, May 2009, <<http://www.rfc-editor.org/info/rfc5548>>.
- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<http://www.rfc-editor.org/info/rfc6206>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<http://www.rfc-editor.org/info/rfc6551>>.
- [RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, DOI 10.17487/RFC6719, September 2012, <<http://www.rfc-editor.org/info/rfc6719>>.
- [RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, DOI 10.17487/RFC6552, March 2012, <<http://www.rfc-editor.org/info/rfc6552>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.

[RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<http://www.rfc-editor.org/info/rfc7416>>.

[I-D.ietf-6lo-privacy-considerations] Thaler, D., "Privacy Considerations for IPv6 over Networks of Resource-Constrained Nodes", draft-ietf-6lo-privacy-considerations-03 (work in progress), September 2016.

[DOEVCC] "Voluntary Code of Conduct (VCC) Final Concepts and Principles", Jan 2015, <http://energy.gov/sites/prod/files/2015/01/f19/VCC%20Concepts%20and%20Principles%202015_01_08%20FINAL.pdf>.

[EUPR] "Information for investors and data controllers", Jun 2016, <<https://ec.europa.eu/energy/node/1748>>.

Authors' Addresses

Nancy Cam-Winget (editor)
Cisco Systems
3550 Cisco Way
San Jose, CA 95134
US

Email: ncamwing@cisco.com

Jonathan Hui
Nest
3400 Hillview Ave
Palo Alto, CA 94304
USA

Email: jonhui@nestlabs.com

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

Email: daniel.popa@itron.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 4, 2016

M. Richardson
SSW
May 3, 2016

ROLL Applicability Statement Template
draft-ietf-roll-applicability-template-09

Abstract

This document is a template applicability statement for the Routing over Low-power and Lossy Networks (ROLL) WG. This document is not for publication, but rather is to be used as a template.

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 November 4, 2016.

Copyright Notice

Copyright (c) 2016 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	3
1.2. Requirements Language	3
1.3. Terminology	4
1.4. Required Reading	4
1.5. Out of scope requirements	4
2. Deployment Scenario	4
2.1. Network Topologies	4
2.2. Traffic Characteristics	4
2.2.1. General	4
2.2.2. Source-sink (SS) communication paradigm	4
2.2.3. Publish-subscribe (PS, or pub/sub) communication paradigm	4
2.2.4. Peer-to-peer (P2P) communication paradigm	5
2.2.5. Peer-to-multipeer (P2MP) communication paradigm	5
2.2.6. Additional considerations: Duocast and N-cast	5
2.2.7. RPL applicability per communication paradigm	5
2.3. Layer-2 applicability.	5
3. Using RPL to Meet Functional Requirements	5
4. RPL Profile	5
4.1. RPL Features	5
4.1.1. RPL Instances	5
4.1.2. Storing vs. Non-Storing Mode	5
4.1.3. DAO Policy	5
4.1.4. Path Metrics	5
4.1.5. Objective Function	5
4.1.6. DODAG Repair	6
4.1.7. Multicast	6
4.1.8. Security	6
4.1.9. P2P communications	6
4.1.10. IPv6 address configuration	6
4.2. Layer-2 features	6
4.2.1. Specifics about layer-2	6
4.2.2. Services provided at layer-2	6
4.2.3. 6LowPAN options assumed.	6
4.2.4. MLE and other things	6
4.3. Recommended Configuration Defaults and Ranges	6
4.3.1. Trickle Parameters	6
4.3.2. Other Parameters	6
5. MPL Profile	7
5.1. Recommended Configuration Defaults and Ranges	8
5.1.1. Trickle Parameters	8
5.1.2. Other Parameters	8
6. Manageability Considerations	8
7. Security Considerations	8
7.1. Security Considerations during initial deployment	8

7.2. Security Considerations during incremental deployment . .	8
7.3. Security Considerations for P2P uses	8
8. Other Related Protocols	9
9. IANA Considerations	9
10. Acknowledgements	9
11. References	9
11.1. Normative References	9
11.2. Informative References	9
Author's Address	10

1. Introduction

This document describes a series of questions which should be answered. This document is intended to remain as a Internet Draft.

The idea is that current and future Applicability statements will use the table of contents provided. The goal is that all applicability statements will have to cover the listed items as a minimum.

1.1. Relationship to other documents

EDITORIAL: The following should appear in all applicability statements:

ROLL has specified a set of routing protocols for Lossy and Low-resource Networks (LLN) [RFC6550]. This applicability text describes a subset of these protocols and the conditions which make the subset the correct choice. The text recommends and motivates the accompanying parameter value ranges. Multiple applicability domains are recognized including: Building and Home, and Advanced Metering Infrastructure. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most probable network structures. Each applicability statement identifies the distinguishing properties according to a common set of subjects described in as many sections.

A common set of security threats are described in [RFC7416]. The applicability statements complement the security threats document by describing preferred security settings and solutions within the applicability statement conditions. This applicability statements may recommend more light weight security solutions and specify the conditions under which these solutions are appropriate.

1.2. Requirements Language

(RFC2119 reference)

1.3. Terminology

A reference to draft-ietf-roll-terminology is appropriate. A reference to layer-2 specific terminology and/or inclusion of any terms that are normatively referenced is appropriate here.

1.4. Required Reading

References/Overview of requirements documents, both IETF and industry group. (two pages maximum. This text should be (very) technical, should be aimed at IETF *participants*, not industry group participants, and should explain this industries' specific issues)

1.5. Out of scope requirements

This should list other documents (if any) which deal with situations where things are not in scope for this document.

(For instance, the AMI document tries to cover both line-powered urban metering networks, and energy-constrained metering networks, and also tries to deal with rural requirements. This should be three or four documents, so this section should list the limits of what this document covers)

2. Deployment Scenario

2.1. Network Topologies

describe a single scenario, with possibly multiple topologies that a single utility would employ.

2.2. Traffic Characteristics

Explain what kind of traffic is being transmitted, where it is initiated, and what kinds of protocols (CoAP, multicast, HTTPS, etc.) are being used. Explain what assumptions are being made about authentication and authorization in those protocols.

2.2.1. General

2.2.2. Source-sink (SS) communication paradigm

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

- 2.2.4. Peer-to-peer (P2P) communication paradigm
- 2.2.5. Peer-to-multipeer (P2MP) communication paradigm
- 2.2.6. Additional considerations: Duocast and N-cast
- 2.2.7. RPL applicability per communication paradigm
- 2.3. Layer-2 applicability.

Explain what layer-2 technologies this statement applies to, and if there are options, they should be listed generally here, and specifically in section 4.2.

3. Using RPL to Meet Functional Requirements

This should explain in general terms how RPL is going to be used in this network topology. If trees that are multiple layers deep are expected, then this should be described so that the fan out is understood. Some sample topologies (from simulations) should be explained, perhaps with image references from other publications.

This section should tell an *implementer* in a lab, having a simulation tool or a building/city/etc. to use as a testbed, how to construct an LLN of sufficient complexity (but not too much) to validate an implementation.

4. RPL Profile

This section should list the various features of RPL plus other layers of the LLN, and how they will be used.

4.1. RPL Features

- 4.1.1. RPL Instances
- 4.1.2. Storing vs. Non-Storing Mode
- 4.1.3. DAO Policy
- 4.1.4. Path Metrics
- 4.1.5. Objective Function

4.1.6. DODAG Repair

4.1.7. Multicast

4.1.8. Security

4.1.9. P2P communications

4.1.10. IPv6 address configuration

4.2. Layer-2 features

4.2.1. Specifics about layer-2

this section should detail the specific layer-2 network technology that this document applies to. A class of technologies is generally not acceptable.

4.2.2. Services provided at layer-2

4.2.3. 6LowPAN options assumed.

4.2.4. MLE and other things

4.3. Recommended Configuration Defaults and Ranges

4.3.1. Trickle Parameters

This section is intended to document the specific value (or ranges) appropriate for this kind of deployment. This includes trickle specific parameters such as those of RFC6550, section 8.3.1: Imin (DIOIntervalMin), Imax (DIOIntervalDoublings), and k (DIORedundancyConstant). While it is not necessary to hard code these parameters into RPL nodes, as they are announced as part of the DIO message, it is important for researchers who are trying to validate the convergence properties of the resulting deployment to understand what values have been selected.

4.3.2. Other Parameters

There are additional values which are present in the DODAG Configuration option. The purpose of this section is to: a) document what values are configured, b) if a default value is used, if it is appropriate for this deployment.

These values include: MaxRankIncrease, MinHopRankIncrease, the Objective Code Point to use, Default Lifetime, Lifetime Units...

In addition, the kinds of metrics which will be used (RFC6551) needs to be specified. If Objective Function 0 (RFC6552) is used, then it specifies a number of values, but also needs definitions of the `stretch_of_rank`, and `rank_factor`.

If MRHOF (RFC6719) is used, then section 5 of this document requires selection of: `MAX_LINK_METRIC`, `MAX_PATH_COST`, `PARENT_SWITCH_THRESHOLD`, `PARENT_SET_SIZE`, and `ALLOW_FLOATING_ROOT`.

5. MPL Profile

This section should list the various features of MPL. In considering the parameters, a number of questions come up:

- 1) What are the maximum and minimum 1-hop MPL router neighbours of all the MPL routers?
- 2) what is the arrival rate of new packets that need repetition in a MPL router
- 3) Is there a deadline associated with the packets
- 4) What is the shortest number of hops of the longest path between sources and destinations
- 5) What are the values of the MAC: back-off values, retries, buffer size.
- 6) What is the background load of other non MPL applications.
- 7) arrival probability of 1-hop packets

As the corresponding design space is incredibly large, probably only a limited subset of the design space is viable.

Here is an example scenario:

- o 5 neighbours
- o once every 100 ms (rate at sources is once every 300-500 ms)
- o yes, 200 ms
- o 5 hops, with mostly 1 hop
- o no buffer, retry 1, back-off 2
- o absent

- o 100-80%

leading to $k=3-5$, $I_{min}=30-70$ ms, repeat = 2, I_{max} n/a.

It is critical operational boundary conditions together with appropriate MPL parameter values are published in this applicability statements. All applicability statements together may give a good hint which MPL parameters and boundary conditions to choose.

5.1. Recommended Configuration Defaults and Ranges

5.1.1. Trickle Parameters

5.1.1.1. I_{min}

5.1.1.2. I_{max}

5.1.2. Other Parameters

5.1.2.1. Hot Limit

6. Manageability Considerations

7. Security Considerations

7.1. Security Considerations during initial deployment

(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>)

7.2. Security Considerations during incremental deployment

(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)

7.3. Security Considerations for P2P uses

(When layer-3 RPL security is used, P2P DODAGs are ephemeral, and may have different security needs.)

8. Other Related Protocols

9. IANA Considerations

10. Acknowledgements

This document was created from a number source applicatbility templates, including draft-ietf-roll-applicability-ami-06.txt, draft-phinney-rpl-industrial-applicability-00.txt.

The document has benefitted from advance review by the IETF Security Directorate.

A number of edits were contributed from Peter van der Stok, including the MPL considerations/calculations

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<http://www.rfc-editor.org/info/rfc7416>>.

11.2. Informative References

- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<http://www.rfc-editor.org/info/rfc6206>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.

Author's Address

Michael C. Richardson
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z 5V7
CA

Email: mcr+ietf@sandelman.ca
URI: <http://www.sandelman.ca/>

roll
Internet-Draft
Intended status: Standards Track
Expires: May 5, 2016

Y. Doi
TOSHIBA Corporation
M. Gillmore
Itron, Inc
November 2, 2015

MPL Parameter Configuration Option for DHCPv6
draft-ietf-roll-mpl-parameter-configuration-08

Abstract

This document defines a way to configure a parameter set for MPL (Multicast Protocol for Low power and Lossy Networks) via a DHCPv6 option. MPL has a set of parameters to control its behavior, and the parameter set is often configured as a network-wide parameter because the parameter set should be identical for each MPL forwarder in an MPL domain. Using the MPL Parameter Configuration Option defined in this document, a network can easily be configured with a single set of MPL parameters.

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 May 5, 2016.

Copyright Notice

Copyright (c) 2015 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	2
2. MPL Parameter Configuration Option	3
2.1. MPL Parameter Configuration Option Format	3
2.2. DHCPv6 Client Behavior	5
2.3. MPL Forwarder Behavior	6
2.4. DHCPv6 Server Behavior	7
2.5. DHCPv6 Relay Behavior	7
2.6. Operational Considerations	7
3. IANA Considerations	8
4. Security Considerations	8
5. References	9
5.1. Normative References	9
5.2. Informative References	9
Appendix A. Update History (TO EDITORS: this section is intended to be removed before this document becomes an RFC) .	10
Authors' Addresses	11

1. Introduction

Multicast Protocol for Low power and Lossy Networks (MPL) [I-D.ietf-roll-trickle-mcast] defines a protocol to make a multicast network among low-power and lossy networks, e.g., wireless mesh networks. MPL has a set of parameters to control an MPL domain. The parameters control the trade-off between end-to-end delay and network utilization. In most environments, the default parameters are acceptable. However, in some environments, the parameter set must be configured carefully in order to meet the requirements of each environment. According to the MPL document section 5.4, each parameter in the set should be the same for all nodes within an MPL domain, but the MPL document does not define a method to configure the MPL parameter set.

Some managed wireless mesh networks may have a DHCP server to configure network parameters. MPL parameter sets shall be considered as a part of network parameters (nodes in an MPL domain should use an identical parameter set). And a parameter set is required to configure an MPL domain.

This document defines the way to distribute parameter sets for MPL forwarders as a DHCPv6 [RFC3315] option. This document is intended to follow [RFC7227] the guideline.

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].

2. MPL Parameter Configuration Option

As stated in Section 5.4 of [I-D.ietf-roll-trickle-mcast], there are the following 10 parameters per MPL domain. An MPL domain is defined by an MPL domain address, as described in Section 2 of [I-D.ietf-roll-trickle-mcast].

- PROACTIVE_FORWARDING
- SEED_SET_ENTRY_LIFETIME
- DATA_MESSAGE_IMIN
- DATA_MESSAGE_IMAX
- DATA_MESSAGE_K
- DATA_MESSAGE_TIMER_EXPIRATIONS
- CONTROL_MESSAGE_IMIN
- CONTROL_MESSAGE_IMAX
- CONTROL_MESSAGE_K
- CONTROL_MESSAGE_TIMER_EXPIRATIONS

One network may have multiple MPL domains with different configurations. To configure more than one MPL domain via DHCP, there may be more than one MPL Parameter Configuration Option given to DHCP clients by a DHCP server.

2.1. MPL Parameter Configuration Option Format

To distribute a configuration of an MPL domain or a default value for all MPL domains (wildcard) under the network managed by the DHCP server, this document defines a DHCPv6 option format as follows.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
OPTION_MPL_PARAMETERS																				option_len																			
P										Z										TUNIT										SE_LIFETIME									

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|   DM_K   |           DM_IMIN           |   DM_IMAX   |
+-----+-----+-----+-----+-----+-----+-----+
|           DM_T_EXP           |   C_K   |   C_IMIN   | >
+-----+-----+-----+-----+-----+-----+-----+
>(cont'ed) |   C_IMAX   |           C_T_EXP           |
+-----+-----+-----+-----+-----+-----+-----+

```

```

(if option_len = 32 )
+-----+-----+-----+-----+-----+-----+-----+-----+
|           MPL Domain Address   (128bits)           | >
+-----+-----+-----+-----+-----+-----+-----+-----+
>           (cont'ed)           | >
+-----+-----+-----+-----+-----+-----+-----+-----+
>           (cont'ed)           | >
+-----+-----+-----+-----+-----+-----+-----+-----+
>           (cont'ed)           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

OPTION_MPL_PARAMETERS: DHCPv6 option identifier (not yet assigned).

option_len: Length of the option, which is 16 if no MPL domain address is present, or 32 if there is an MPL domain address.

P (1 bit): A flag to indicate PROACTIVE_FORWARDING. The flag is set if PROACTIVE_FORWARDING is true.

Z (7 bits): Reserved for future use. Servers MUST set them to zero. Clients SHOULD ignore the bits set.

TUNIT (unsigned 8-bit integer): Unit time of timer parameters (SE_LIFETIME, and *_IMIN) in this option. 0 and 0xff are reserved and MUST NOT be used.

SE_LIFETIME (unsigned 16-bit integer): SEED_SET_ENTRY_LIFETIME/TUNIT in milliseconds. 0 and 0xffff are reserved and MUST NOT be used.

DM_K (unsigned 8-bit integer): DATA_MESSAGE_K.

DM_IMIN (unsigned 16-bit integer): DATA_MESSAGE_IMIN/TUNIT in milliseconds. 0 and 0xffff are reserved and MUST NOT be used.

DM_IMAX (unsigned 8-bit integer): DATA_MESSAGE_IMAX. The actual maximum timeout is described as a number of doublings of DATA_MESSAGE_IMIN, as described in [RFC6206] Section 4.1. 0 and 0xff are reserved and MUST NOT be used.

DM_T_EXP (unsigned 16-bit integer): DATA_MESSAGE_TIMER_EXPIRATIONS.
0 and 0xffff are reserved and MUST NOT be used.

C_K (unsigned 8-bit integer): CONTROL_MESSAGE_K.

C_IMIN (unsigned 16-bit integer): CONTROL_MESSAGE_IMIN/TUNIT in
milliseconds. 0 and 0xffff are reserved and MUST NOT be used.

C_IMAX (unsigned 8-bit integer): CONTROL_MESSAGE_IMAX. The actual
maximum timeout is described as a number of doublings of
CONTROL_MESSAGE_IMIN. 0 and 0xff are reserved and MUST NOT be
used.

C_T_EXP (unsigned 16-bit integer): CONTROL_MESSAGE_TIMER_EXPIRATIONS
. 0 and 0xffff are reserved and MUST NOT be used.

Note that the time values (SEED_SET_ENTRY_LIFETIME,
DATA_MESSAGE_IMIN, and CONTROL_MESSAGE_IMIN) in MPL are defined in
TUNIT milliseconds precision in MPL Parameter Configuration Options.
For example, if TUNIT is 20 and the data message interval minimum
(DATA_MESSAGE_IMIN) is 1000ms, then DM_IMIN shall be set to 50.

For maximum interval size (*_IMAX), [RFC6206] defines them as
follows:

The maximum interval size, I_{max}, is described as a number of
doublings of the minimum interval size (the base-2 log(max/min)).
For example, a protocol might define I_{max} as 16. If the minimum
interval is 100 ms, then the amount of time specified by I_{max} is
100 ms * 65,536, i.e., 6,553.6 seconds or approximately 109
minutes.

Because minimum interval size in the MPL Parameter Configuration
Options is described as TUNIT millisecond precision, corresponding
maximum interval size is also in TUNIT precision. For example, if
TUNIT is 10 and C_IMIN is 50, the minimum interval size of the
trickle timer for control messages is 500ms. In this case, the
maximum interval size of the trickle timer is 32 seconds (500ms *
2⁶) if C_IMAX is 6.

2.2. DHCPv6 Client Behavior

Clients MAY request the MPL Parameter Configuration Option, as
described in [RFC3315], sections 17.1.1, 18.1.1, 18.1.3, 18.1.4,
18.1.5, and 22.7. As a convenience to the reader, we mention here
that the client includes requested option codes in the Option Request
Option.

Clients MUST support multiple MPL Parameter Configuration Option, as stated in section 2.

If a DHCPv6 client with an MPL forwarder configured by the MPL Parameter Configuration Option is unable to receive a valid response from a server within T2 of the last valid DHCPv6 message sent from the server (if stateful) or twice the Information Refresh Time (if stateless), it MUST suspend the MPL forwarders of the MPL domains configured by the option. MPL forwarders configured by other methods such as static configuration file MUST NOT be suspended.

Clients MUST ignore all MPL Parameter Configuration Options if the options in a DHCPv6 message contains any invalid value (e.g., it uses reserved all-0 or all-1 values in parameters). In this case, the message is considered not received in MPL context and the condition described in the previous paragraph applies.

2.3. MPL Forwarder Behavior

If a DHCPv6 client requests and receives the MPL Parameter Configuration Option, the node SHOULD join the MPL domain given by the option and act as an MPL forwarder. Note that there may be cases in which a node may fail to join a domain (or domains) due to local resource constraints. Each joining node SHOULD configure its MPL forwarder with the given parameter set for the MPL domain. Each MPL domain is defined by an MPL Domain Address given by an MPL Parameter Configuration Option. As defined in Section 2 of [I-D.ietf-roll-trickle-mcast], an MPL Domain Address is an IPv6 multicast address associated to a set of MPL network interfaces in an MPL Domain.

The priority of MPL Parameter Configurations applied to an MPL Domain is as follows (high to low):

- o Specific MPL Parameter Configuration to the MPL Domain (option_len=32)
- o Wildcard MPL Parameter Configuration (option_len=16)
- o Default configuration given in the MPL specification.

Priority of other configurations such as manual configuration given on a node is not defined in the document.

There MUST be no more than one MPL Parameter Configuration Option for an MPL domain or the wildcard. Thus, the order of DHCPv6 options in the packet has no effect on precedence.

A node MUST leave an MPL domain if it receives an updated and all-valid MPL Parameter Configuration Options without a configuration for the MPL domain, unless it has overriding manual configuration on the MPL domain. In other words, if a node is configured to work as a MPL Forwarder for a MPL domain regardless of DHCPv6 Options, the node MAY stay on the MPL domain even if it receives an MPL Parameter Configuration Option without configuration for the MPL domain.

MPL parameters may be updated occasionally. With stateful DHCPv6, updates can be done when the renewal timer expires. Information Refresh Time Option [RFC4242] shall be used to keep each forwarder updated.

To reduce periodic update traffic, a node may try to use a very long interval between updates. In this case, reconfigure messages may be used to keep forwarder parameter sets synchronized.

2.4. DHCPv6 Server Behavior

Sections 17.2.2 and 18.2 of [RFC3315] govern server operation in regards to option assignment. As a convenience to the reader, we mention here that the server will send the MPL Parameter Configuration Option only if it was configured with specific values for the MPL Parameter Configuration Option and the client requested it.

Servers MUST ignore an incoming MPL Parameter Configuration Option. Servers MUST support multiple MPL Parameter Configuration Option, as stated in section 2.

2.5. DHCPv6 Relay Behavior

It's never appropriate for a relay agent to add options to a message heading toward the client, and relay agents don't actually construct Relay-Reply messages anyway. There are no additional requirements for relays.

2.6. Operational Considerations

This draft introduces dynamic update of MPL parameters. Because the update process is not synchronized, nodes may have inconsistent parameter sets.

[RFC6206] section 6 describe various problems that happens if the trickle timers do not match between communicating nodes. To keep the timers synchronized, it is RECOMMENDED not to update the parameters of an MPL domain too often. A reasonable update rate would be once per expected information refresh time interval, such as T1 in [RFC3315] or Information Refresh Time in [RFC4242].

Inconsistent parameter sets may reduce performance. On the other hand, this situation will work as long as both new and old parameter sets are reasonable parameter sets for a given communication load. As the motivations for parameter update include update of the environment, node density, or communication load, operators of MPL networks shall be aware of unupdated nodes and make sure old and new parameter sets are reasonable for the expected refresh intervals.

3. IANA Considerations

IANA is requested to assign one option code for OPTION_MPL_PARAMETERS from the "DHCP Option Codes" table of the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Registry (<http://www.iana.org/assignments/dhcpv6-parameters>).

4. Security Considerations

There are detailed discussion on security threats on DHCPv6 in Section 23 of RFC3315 [RFC3315], Section 23 of RFC7227 [RFC7227], and Section 13 of [I-D.ietf-roll-trickle-mcast].

In addition, a forged MPL parameter configuration may cause excessive layer-2 broadcasting. Implementations should set reasonable bounds for each parameter. For example, not too high DM/C_K, not too low DM/C_IMIN, etc. These bounds may be implementation dependent or may be derived from MAC/PHY specifications. DHCPv6 server and client implementations need to take care in setting reasonable bounds for each parameter in order to avoid overloading the network.

The DHCP server or the network itself should be trusted by some means such as DHCPv6 authentications described in Section 21 of RFC3315 [RFC3315]. However, ROLL environment may expect less computing resource, and DHCPv6 authentication may not be available. In such cases, other methods to protect integrity between DHCPv6 servers and clients should be applied to a ROLL network. Some ROLL specification such as ZigBee IP [ZigBeeIP] expects RFC5191 [RFC5191] to authenticate joining nodes and all nodes in the network can be trusted. To protect attacks from outside of the network, DHCPv6 packets SHOULD be filtered on the border router between the ROLL network and the Internet, except for the packets between the ROLL network and a remote DHCPv6 server or DHCPv6 relays configured to manage the network.

5. References

5.1. Normative References

- [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-12 (work in progress), June 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.
- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<http://www.rfc-editor.org/info/rfc6206>>.
- [RFC7227] Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", BCP 187, RFC 7227, May 2014.

5.2. Informative References

- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008.

[ZigBeeIP]

ZigBee Alliance, "ZigBee IP Specification", Mar 2014.

Appendix A. Update History (TO EDITORS: this section is intended to be removed before this document becomes an RFC)

Updates on draft-ietf-roll-mpl-configuration-07 to draft-ietf-roll-mpl-configuration-08:

- o clarified when to leave (SHOULD->MUST)
- o moved Trickle parameter considerations on appendix to operational considerations
- o even clarified some texts

Updates on draft-ietf-roll-mpl-configuration-06 to draft-ietf-roll-mpl-configuration-07:

- o clearly stated multiple option support is mandatory (#171)
- o operational consideration now refers RFC6206 and some texts are moved to section 2.2 (#171)
- o added more per-section reference to I-D.ietf-roll-trickle-mcast (#171)
- o field 'Z' clarified (#171, #172)
- o fixed other nits (#171)
- o clarified use of TUNIT, *_IMIN, and *_IMAX with reference to RFC6206 (#172)

Updates on draft-ietf-roll-mpl-configuration-05 to draft-ietf-roll-mpl-configuration-06:

- o added description on manual (external) configurations

Updates on draft-ietf-roll-mpl-configuration-04 to draft-ietf-roll-mpl-configuration-05:

- o fixed *_IMAX definition as RFC6206 defines
- o fixed *_EXP definition as draft-ietf-roll-trickle-mcast defines
- o added references to RFC3315 and RFC7227 in security considerations section

- o added a paragraph on security consideration according to secdir review
- o fixed some nits and updated references

Updates on draft-ietf-roll-mpl-configuration-03 to draft-ietf-roll-mpl-configuration-04:

- o References updated (Non-normative -> Informative)
- o IANA section is updated to make clear request of option ID
- o Reserved numbers are clearly denoted

Updates on draft-ietf-roll-mpl-configuration-02 to draft-ietf-roll-mpl-configuration-03:

- o References updated
- o Removed reference for DHCPv6 stateless reconfiguration as it has expired

Updates on draft-ietf-roll-mpl-configuration-01 to draft-ietf-roll-mpl-configuration-02:

- o Short unsigned floating point is dropped (#159)
- o Packed value is removed and now every value has its own byte(s) (#159)

Updates on draft-ietf-roll-mpl-configuration-00 to draft-ietf-roll-mpl-configuration-01:

- o Operational considerations (normative) and appendix considerations (non-normative) are added (Issue #157)
- o More control on nodes / allow constrained nodes to ignore the configuration: "the node s/SHOULD/MAY/ join the MPL domain given by the option" (Issue #158)

Updates on draft-doi-roll-mpl-configuration-05 to draft-ietf-roll-mpl-configuration-00:

- o I-D renamed.

Authors' Addresses

Yusuke Doi
TOSHIBA Corporation
Komukai Toshiba Cho 1
Saiwai-Ku
Kawasaki, Kanagawa 2128582
JAPAN

Phone: +81-45-342-7230
Email: yusuke.doi@toshiba.co.jp

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

Email: matthew.gillmore@itron.com

Routing Over Low-Power and Lossy Networks
Internet-Draft
Intended status: Informational
Expires: April 05, 2015

T. Tsao
R. Alexander
Cooper Power Systems
M. Dohler
CTTC
V. Daza
A. Lozano
Universitat Pompeu Fabra
M. Richardson, Ed.
Sandelman Software Works
October 02, 2014

A Security Threat Analysis for Routing Protocol for Low-power and lossy
networks (RPL)
draft-ietf-roll-security-threats-11

Abstract

This document presents a security threat analysis for the Routing Protocol for Low-power and lossy networks (RPL, ROLL). The development builds upon previous work on routing security and adapts the assessments to the issues and constraints specific to low-power and lossy networks. A systematic approach is used in defining and evaluating the security threats. Applicable countermeasures are application specific and are addressed in relevant applicability statements.

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 05, 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
2. Relationship to other documents	4
3. Terminology	4
4. Considerations on RPL Security	5
4.1. Routing Assets and Points of Access	6
4.2. The ISO 7498-2 Security Reference Model	8
4.3. Issues Specific to or Amplified in LLNs	10
4.4. RPL Security Objectives	12
5. Threat Sources	13
6. Threats and Attacks	13
6.1. Threats due to failures to Authenticate	14
6.1.1. Node Impersonation	14
6.1.2. Dummy Node	14
6.1.3. Node Resource Spam	14
6.2. Threats due to failure to keep routing information confidential	15
6.2.1. Routing Exchange Exposure	15
6.2.2. Routing Information (Routes and Network Topology) Exposure	15
6.3. Threats and Attacks on Integrity	16
6.3.1. Routing Information Manipulation	16
6.3.2. Node Identity Misappropriation	17
6.4. Threats and Attacks on Availability	17
6.4.1. Routing Exchange Interference or Disruption	17
6.4.2. Network Traffic Forwarding Disruption	17
6.4.3. Communications Resource Disruption	19
6.4.4. Node Resource Exhaustion	19
7. Countermeasures	20
7.1. Confidentiality Attack Countermeasures	20
7.1.1. Countering Deliberate Exposure Attacks	20
7.1.2. Countering Passive Wiretapping Attacks	21
7.1.3. Countering Traffic Analysis	22
7.1.4. Countering Remote Device Access Attacks	23

7.2.	Integrity Attack Countermeasures	23
7.2.1.	Countering Unauthorized Modification Attacks	23
7.2.2.	Countering Overclaiming and Misclaiming Attacks	24
7.2.3.	Countering Identity (including Sybil) Attacks	24
7.2.4.	Countering Routing Information Replay Attacks	25
7.2.5.	Countering Byzantine Routing Information Attacks	25
7.3.	Availability Attack Countermeasures	26
7.3.1.	Countering HELLO Flood Attacks and ACK Spoofing Attacks	26
7.3.2.	Countering Overload Attacks	27
7.3.3.	Countering Selective Forwarding Attacks	28
7.3.4.	Countering Sinkhole Attacks	29
7.3.5.	Countering Wormhole Attacks	30
8.	RPL Security Features	31
8.1.	Confidentiality Features	31
8.2.	Integrity Features	32
8.3.	Availability Features	33
8.4.	Key Management	33
9.	IANA Considerations	34
10.	Security Considerations	34
11.	Acknowledgments	34
12.	References	34
12.1.	Normative References	34
12.2.	Informative References	35
	Authors' Addresses	38

1. Introduction

In recent times, networked electronic devices have found an increasing number of applications in various fields. Yet, for reasons ranging from operational application to economics, these wired and wireless devices are often supplied with minimum physical resources; the constraints include those on computational resources (RAM, clock speed, storage), communication resources (duty cycle, packet size, etc.), but also form factors that may rule out user access interfaces (e.g., the housing of a small stick-on switch), or simply safety considerations (e.g., with gas meters). As a consequence, the resulting networks are more prone to loss of traffic and other vulnerabilities. The proliferation of these low-power and lossy networks (LLNs), however, are drawing efforts to examine and address their potential networking challenges. Securing the establishment and maintenance of network connectivity among these deployed devices becomes one of these key challenges.

This document presents a threat analysis for securing the Routing Protocol for LLNs (RPL). The process requires two steps. First, the analysis will be used to identify pertinent security issues. The second step is to identify necessary countermeasures to secure RPL.

As there are multiple ways to solve the problem and the specific tradeoffs are deployment specific, the specific countermeasure to be used is detailed in applicability statements.

This document uses [ISO.7498-2.1988] model, which describes Authentication, Access Control, Data Confidentiality, Data Integrity, and Non-Repudiation security services and to which Availability is added. As explained below, Non-Repudiation does not apply to routing protocols.

Many of the issues in this document were also covered in The IAB Smart Object Workshop [RFC6574], and The IAB Smart Object Security Workshop [I-D.gilger-smart-object-security-workshop].

All of this document concerns itself with securing the control plane traffic. As such it does not address authorization or authentication of application traffic. RPL uses multicast as part of its protocol, and therefore mechanisms which RPL uses to secure this traffic might also be applicable to MPL control traffic as well: the important part is that the threats are similar.

2. Relationship to other documents

ROLL has specified a set of routing protocols for Lossy and Low-resource Networks (LLN) [RFC6550]. A number of applicability texts describes a subset of these protocols and the conditions which make the subset the correct choice. The text recommends and motivates the accompanying parameter value ranges. Multiple applicability domains are recognized including: Building and Home, and Advanced Metering Infrastructure. The applicability domains distinguish themselves in the way they are operated, their performance requirements, and the most probable network structures. Each applicability statement identifies the distinguishing properties according to a common set of subjects described in as many sections.

The common set of security threats herein are referred to by the applicability statements, and that series of documents describes the preferred security settings and solutions within the applicability statement conditions. This applicability statements may recommend more light weight security solutions and specify the conditions under which these solutions are appropriate.

3. Terminology

This document adopts the terminology defined in [RFC6550], in [RFC4949], and in [RFC7102].

The terms control plane and forwarding plane are used consistently with section 1 of [RFC6192].

The term DODAG is from [RFC6550].

EAP-TLS is defined in [RFC5216].

PANA is defined in [RFC5191].

CCM mode is defined in [RFC3610].

[RFC7102] introduces the term Sleepy Node, referring to a node which may sometimes go into a low-power state, suspending protocol communications

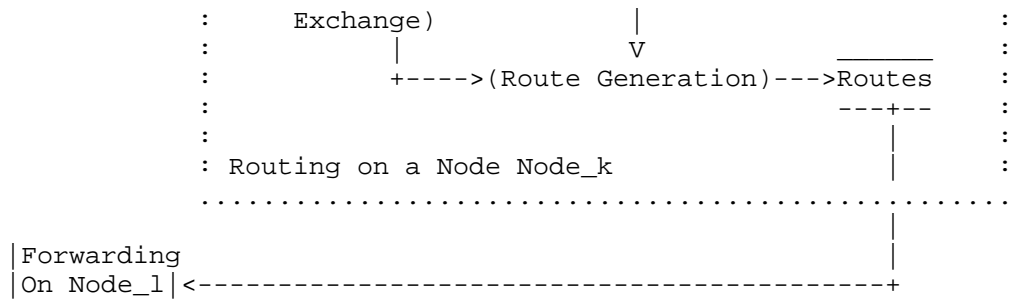
The terms SSID, ESSID and PAN refer to network identifiers, defined in [IEEE.802.11] and [IEEE.802.15.4].

Although this is not a protocol specification, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] in order to clarify and emphasize the guidance and directions to implementers and deployers of LLN nodes that utilize RPL.

4. Considerations on RPL Security

Routing security, in essence, ensures that the routing protocol operates correctly. It entails implementing measures to ensure controlled state changes on devices and network elements, both based on external inputs (received via communications) or internal inputs (physical security of device itself and parameters maintained by the device, including, e.g., clock). State changes would thereby involve not only authorization of injector's actions, authentication of injectors, and potentially confidentiality of routing data, but also proper order of state changes through timeliness, since seriously delayed state changes, such as commands or updates of routing tables, may negatively impact system operation. A security assessment can therefore begin with a focus on the assets [RFC4949] that may be the target of the state changes and the access points in terms of interfaces and protocol exchanges through which such changes may occur. In the case of routing security, the focus is directed towards the elements associated with the establishment and maintenance of network connectivity.

This section sets the stage for the development of the analysis by applying the systematic approach proposed in [Myagmar2005] to the routing security, while also drawing references from other reviews



Notation:

(Proc) A process Proc

topology A structure storing neighbor adjacency (parent/child)

routes A structure storing the forwarding information base (FIB)

|Node_n| An external entity Node_n

-----> Data flow

Figure 1: Data Flow Diagram of a Generic Routing Process

It is seen from Figure 1 that

o Assets include

- * routing and/or topology information;
- * route generation process;
- * communication channel resources (bandwidth);
- * node resources (computing capacity, memory, and remaining energy);
- * node identifiers (including node identity and ascribed attributes such as relative or absolute node location).

o Points of access include

- * neighbor discovery;
- * route/topology exchange;
- * node physical interfaces (including access to data storage).

A focus on the above list of assets and points of access enables a more directed assessment of routing security; for example, it is readily understood that some routing attacks are in the form of attempts to misrepresent routing topology. Indeed, the intention of the security threat analysis is to be comprehensive. Hence, some of the discussion which follows is associated with assets and points of access that are not directly related to routing protocol design but nonetheless provided for reference since they do have direct consequences on the security of routing.

4.2. The ISO 7498-2 Security Reference Model

At the conceptual level, security within an information system in general and applied to RPL in particular is concerned with the primary issues of authentication, access control, data confidentiality, data integrity, and non-repudiation. In the context of RPL:

Authentication

Authentication involves the mutual authentication of the routing peers prior to exchanging route information (i.e., peer authentication) as well as ensuring that the source of the route data is from the peer (i.e., data origin authentication). [RFC5548] points out that LLNs can be drained by unauthenticated peers before configuration. [RFC5673] requires availability of open and untrusted side channels for new joiners, and it requires strong and automated authentication so that networks can automatically accept or reject new joiners.

Access Control

Access Control provides protection against unauthorized use of the asset, and deals with the authorization of a node.

Confidentiality

Confidentiality involves the protection of routing information as well as routing neighbor maintenance exchanges so that only authorized and intended network entities may view or access it. Because LLNs are most commonly found on a publicly accessible shared medium, e.g., air or wiring in a building, and sometimes formed ad hoc, confidentiality also extends to the neighbor state and database information within the routing device since the deployment of the network creates the potential for unauthorized access to the physical devices themselves.

Integrity

Integrity entails the protection of routing information and routing neighbor maintenance exchanges, as well as derived information maintained in the database, from unauthorized modification, insertions, deletions or replays. to be addressed beyond the routing protocol.

Non-repudiation

Non-repudiation is the assurance that the transmission and/or reception of a message cannot later be denied. The service of non-repudiation applies after-the-fact and thus relies on the logging or other capture of on-going message exchanges and signatures. Routing protocols typically do not have a notion of repudiation, so non-repudiation services are not required. Further, with the LLN application domains as described in [RFC5867] and [RFC5548], proactive measures are much more critical than retrospective protections. Finally, given the significant practical limits to on-going routing transaction logging and storage and individual device digital signature verification for each exchange, non-repudiation in the context of routing is an unsupportable burden that bears no further considered as an RPL security issue.

It is recognized that, besides those security issues captured in the ISO 7498-2 model, availability, is a security requirement:

Availability

Availability ensures that routing information exchanges and forwarding services need to be available when they are required for the functioning of the serving network. Availability will apply to maintaining efficient and correct operation of routing and neighbor discovery exchanges (including needed information) and forwarding services so as not to impair or limit the network's central traffic flow function

It should be emphasized here that for RPL security the above requirements must be complemented by the proper security policies and enforcement mechanisms to ensure that security objectives are met by a given RPL implementation.

4.3. Issues Specific to or Amplified in LLNs

The requirements work detailed in Urban Requirements ([RFC5548]), Industrial Requirements ([RFC5673]), Home Automation ([RFC5826], and Building Automation ([RFC5867]) have identified specific issues and constraints of routing in LLNs. The following is a list of observations from those requirements and evaluation of their impact on routing security considerations.

Limited energy, memory, and processing node resources

As a consequence of these constraints, there is an even more critical need than usual for a careful study of trade-offs on which and what level of security services are to be afforded during the system design process. The chosen security mechanisms also needs to work within these constraints. Synchronization of security states with sleepy nodes [RFC7102] is yet another issue. A non-rechargeable battery powered node may well be limited in energy for it's lifetime: once exhausted, it may well never function again.

Large scale of rolled out network

The possibly numerous nodes to be deployed make manual on-site configuration unlikely. For example, an urban deployment can see several hundreds of thousands of nodes being installed by many installers with a low level of expertise. Nodes may be installed and not activated for many years, and additional nodes may be added later on, which may be from old inventory. The lifetime of the network is measured in decades, and this complicates the operation of key management.

Autonomous operations

Self-forming and self-organizing are commonly prescribed requirements of LLNs. In other words, a routing protocol designed for LLNs needs to contain elements of ad hoc networking and in most cases cannot rely on manual configuration for initialization or local filtering rules. Network topology/ownership changes, partitioning or merging, as well as node replacement, can all contribute to complicating the operations of key management.

Highly directional traffic

Some types of LLNs see a high percentage of their total traffic traverse between the nodes and the LLN Border Routers (LBRs)

where the LLNs connect to non-LLNs. The special routing status of and the greater volume of traffic near the LBRs have routing security consequences as a higher valued attack target. In fact, when Point-to-MultiPoint (P2MP) and MultiPoint-to-Point (MP2P) traffic represents a majority of the traffic, routing attacks consisting of advertising incorrect preferred routes can cause serious damage.

While it might seem that nodes higher up in the acyclic graph (i.e. those with lower rank) should be secured in a stronger fashion, it is not in general easy to predict which nodes will occupy those positions until after deployment. Issues of redundancy and inventory control suggests that any node might wind up in such a sensitive attack position, so all nodes to be capable of being fully secured.

In addition, even if it were possible to predict which nodes will occupy positions of lower rank and provision them with stronger security mechanisms, in the absense of a strong authorization model, any node could advertise an incorrect preferred route.

Unattended locations and limited physical security

Many applications have the nodes deployed in unattended or remote locations; furthermore, the nodes themselves are often built with minimal physical protection. These constraints lower the barrier of accessing the data or security material stored on the nodes through physical means.

Support for mobility

On the one hand, only a limited number of applications require the support of mobile nodes, e.g., a home LLN that includes nodes on wearable health care devices or an industry LLN that includes nodes on cranes and vehicles. On the other hand, if a routing protocol is indeed used in such applications, it will clearly need to have corresponding security mechanisms.

Additionally nodes may appear to move from one side of a wall to another without any actual motion involved, the result of changes to electromagnetic properties, such as opening and closing of a metal door.

Support for multicast and anycast

Support for multicast and anycast is called out chiefly for large-scale networks. Since application of these routing mechanisms in autonomous operations of many nodes is new, the consequence on security requires careful consideration.

The above list considers how an LLN's physical constraints, size, operations, and variety of application areas may impact security. However, it is the combinations of these factors that particularly stress the security concerns. For instance, securing routing for a large number of autonomous devices that are left in unattended locations with limited physical security presents challenges that are not found in the common circumstance of administered networked routers. The following subsection sets up the security objectives for the routing protocol designed by the ROLL WG.

4.4. RPL Security Objectives

This subsection applies the ISO 7498-2 model to routing assets and access points, taking into account the LLN issues, to develop a set of RPL security objectives.

Since the fundamental function of a routing protocol is to build routes for forwarding packets, it is essential to ensure that:

- o routing/topology information integrity remains intact during transfer and in storage;
- o routing/topology information is used by authorized entities;
- o routing/topology information is available when needed.

In conjunction, it is necessary to be assured that

- o authorized peers authenticate themselves during the routing neighbor discovery process;
- o the routing/topology information received is generated according to the protocol design.

However, when trust cannot be fully vested through authentication of the principals alone, i.e., concerns of insider attack, assurance of the truthfulness and timeliness of the received routing/topology information is necessary. With regard to confidentiality, protecting the routing/topology information from unauthorized exposure may be desirable in certain cases but is in itself less pertinent in general to the routing function.

One of the main problems of synchronizing security states of sleepy nodes, as listed in the last subsection, lies in difficulties in authentication; these nodes may not have received in time the most recent update of security material. Similarly, the issues of minimal manual configuration, prolonged rollout and delayed addition of nodes, and network topology changes also complicate key management.

Hence, routing in LLNs needs to bootstrap the authentication process and allow for flexible expiration scheme of authentication credentials.

The vulnerability brought forth by some special-function nodes, e.g., LBRs, requires the assurance, particularly in a security context,

- o of the availability of communication channels and node resources;
- o that the neighbor discovery process operates without undermining routing availability.

There are other factors which are not part of RPL but directly affecting its function. These factors include weaker barrier of accessing the data or security material stored on the nodes through physical means; therefore, the internal and external interfaces of a node need to be adequate for guarding the integrity, and possibly the confidentiality, of stored information, as well as the integrity of routing and route generation processes.

Each individual system's use and environment will dictate how the above objectives are applied, including the choices of security services as well as the strengths of the mechanisms that must be implemented. The next two sections take a closer look at how the RPL security objectives may be compromised and how those potential compromises can be countered.

5. Threat Sources

[RFC4593] provides a detailed review of the threat sources: outsiders and byzantine. RPL has the same threat sources.

6. Threats and Attacks

This section outlines general categories of threats under the ISO 7498-2 model and highlights the specific attacks in each of these categories for RPL. As defined in [RFC4949], a threat is "a potential for violation of security, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm."

An attack is "an assault on system security that derives from an intelligent threat, i.e., an intelligent act that is a deliberate attempt (especially in the sense of a method or technique) to evade security services and violate the security policy of a system."

The subsequent subsections consider the threats and the attacks that can cause security breaches under the ISO 7498-2 model to the routing

assets and via the routing points of access identified in Section 4.1. The assessment steps through the security concerns of each routing asset and looks at the attacks that can exploit routing points of access. The threats and attacks identified are based on the routing model analysis and associated review of the existing literature. The source of the attacks is assumed to be from either inside or outside attackers. While some attackers inside the network will be using compromised nodes, and therefore are only able to do what an ordinary node can ("node-equivalent"), other attacks may not be limited in memory, CPU, power consumption or long term storage. Moore's law favours the attacker with access to the latest capabilities, while the defenders will remain in place for years to decades.

6.1. Threats due to failures to Authenticate

6.1.1. Node Impersonation

If an attacker can join a network using any identity, then it may be able to assume the role of a legitimate (and existing) node. It may be able to report false readings (in metering applications), or provide inappropriate control messages (in control systems involving actuators) if the security of the application is implied by the security of the routing system.

Even in systems where there is application layer security, the ability to impersonate a node would permit an attacker to direct traffic to itself. This may permit various on-path attacks which would otherwise be difficult, such as replaying, delaying, or duplicating (application) control messages.

6.1.2. Dummy Node

If an attacker can join a network using any identity, then it can pretend to be a legitimate node, receiving any service legitimate nodes receive. It may also be able to report false readings (in metering applications), or provide inappropriate authorizations (in control systems involving actuators), or perform any other attacks that are facilitated by being able to direct traffic towards itself.

6.1.3. Node Resource Spam

If an attacker can join a network with any identity, then it can continuously do so with new (random) identities. This act may drain down the resources of the network (battery, RAM, bandwidth). This may cause legitimate nodes of the network to be unable to communicate.

6.2. Threats due to failure to keep routing information confidential

The assessment in Section 4.2 indicates that there are attacks against the confidentiality of routing information at all points of access. This threat may result in disclosure, as described in Section 3.1.2 of [RFC4593], and may involve a disclosure of routing information.

6.2.1. Routing Exchange Exposure

Routing exchanges include both routing information as well as information associated with the establishment and maintenance of neighbor state information. As indicated in Section 4.1, the associated routing information assets may also include device specific resource information, such as available memory, remaining power, etc., that may be metrics of the routing protocol.

The routing exchanges will contain reachability information, which would identify the relative importance of different nodes in the network. Nodes higher up in the DODAG, to which more streams of information flow, would be more interesting targets for other attacks, and routing exchange exposures can identify them.

6.2.2. Routing Information (Routes and Network Topology) Exposure

Routes (which may be maintained in the form of the protocol forwarding table) and neighbor topology information are the products of the routing process that are stored within the node device databases.

The exposure of this information will allow attackers to gain direct access to the configuration and connectivity of the network thereby exposing routing to targeted attacks on key nodes or links. Since routes and neighbor topology information is stored within the node device, attacks on the confidentiality of the information will apply to the physical device including specified and unspecified internal and external interfaces.

The forms of attack that allow unauthorized access or disclosure of the routing information will include:

- o Physical device compromise;
- o Remote device access attacks (including those occurring through remote network management or software/field upgrade interfaces).

Both of these attack vectors are considered a device specific issue, and are out of scope for RPL to defend against. In some

applications, physical device compromise may be a real threat and it may be necessary to provide for other devices to securely detect a compromised device and react quickly to exclude it.

6.3. Threats and Attacks on Integrity

The assessment in Section 4.2 indicates that information and identity assets are exposed to integrity threats from all points of access. In other words, the integrity threat space is defined by the potential for exploitation introduced by access to assets available through routing exchanges and the on-device storage.

6.3.1. Routing Information Manipulation

Manipulation of routing information that range from neighbor states to derived routes will allow unauthorized sources to influence the operation and convergence of the routing protocols and ultimately impact the forwarding decisions made in the network.

Manipulation of topology and reachability information will allow unauthorized sources to influence the nodes with which routing information is exchanged and updated. The consequence of manipulating routing exchanges can thus lead to sub-optimality and fragmentation or partitioning of the network by restricting the universe of routers with which associations can be established and maintained.

A sub-optimal network may use too much power and/or may congest some routes leading to premature failure of a node, and a denial of service on the entire network.

In addition, being able to attract network traffic can make a blackhole attack more damaging.

The forms of attack that allow manipulation to compromise the content and validity of routing information include

- o Falsification, including overclaiming and misclaiming (claiming routes to devices which the device can not in fact reach);
- o Routing information replay;
- o Byzantine (internal) attacks that permit corruption of routing information in the node even where the node continues to be a validated entity within the network (see, for example, [RFC4593] for further discussions on Byzantine attacks);
- o Physical device compromise or remote device access attacks.

6.3.2. Node Identity Misappropriation

Falsification or misappropriation of node identity between routing participants opens the door for other attacks; it can also cause incorrect routing relationships to form and/or topologies to emerge. Routing attacks may also be mounted through less sophisticated node identity misappropriation in which the valid information broadcast or exchanged by a node is replayed without modification. The receipt of seemingly valid information that is however no longer current can result in routing disruption, and instability (including failure to converge). Without measures to authenticate the routing participants and to ensure the freshness and validity of the received information the protocol operation can be compromised. The forms of attack that misuse node identity include

- o Identity attacks, including Sybil attacks (see [Sybil2002]) in which a malicious node illegitimately assumes multiple identities;
- o Routing information replay.

6.4. Threats and Attacks on Availability

The assessment in Section 4.2 indicates that the process and resources assets are exposed to threats against availability; attacks in this category may exploit directly or indirectly information exchange or forwarding (see [RFC4732] for a general discussion).

6.4.1. Routing Exchange Interference or Disruption

Interference is the threat action and disruption is threat consequence that allows attackers to influence the operation and convergence of the routing protocols by impeding the routing information exchange.

The forms of attack that allow interference or disruption of routing exchange include:

- o Routing information replay;
- o ACK spoofing;
- o Overload attacks. (Section 7.3.2)

In addition, attacks may also be directly conducted at the physical layer in the form of jamming or interfering.

6.4.2. Network Traffic Forwarding Disruption

The disruption of the network traffic forwarding capability will undermine the central function of network routers and the ability to handle user traffic. This affects the availability of the network because of the potential to impair the primary capability of the network.

In addition to physical layer obstructions, the forms of attack that allows disruption of network traffic forwarding include [Karlof2003]

- o Selective forwarding attacks;

```
|Node_1|--(msg1|msg2|msg3)-->|Attacker|--(msg1|msg3)-->|Node_2|
```

Figure 2: Selective forwarding example

- o Wormhole attacks;

```
|Node_1|-----Unreachable-----x|Node_2|
|                                     ^
|                                     |
|                                     Private Link
|-->|Attacker_1|=====|Attacker_2|--'
```

Figure 3: Wormhole Attacks

- o Sinkhole attacks.




Figure 4: sinkhole attack example

These attacks are generally done to both control plane and forwarding plane traffic. A system that prevents control plane traffic (RPL messages) from being diverted in these ways will also prevent actual data from being diverted.

6.4.3. Communications Resource Disruption

Attacks mounted against the communication channel resource assets needed by the routing protocol can be used as a means of disrupting its operation. However, while various forms of Denial of Service (DoS) attacks on the underlying transport subsystem will affect routing protocol exchanges and operation (for example physical layer RF jamming in a wireless network or link layer attacks), these attacks cannot be countered by the routing protocol. As such, the threats to the underlying transport network that supports routing is considered beyond the scope of the current document. Nonetheless, attacks on the subsystem will affect routing operation and so must be directly addressed within the underlying subsystem and its implemented protocol layers.

6.4.4. Node Resource Exhaustion

A potential threat consequence can arise from attempts to overload the node resource asset by initiating exchanges that can lead to the exhaustion of processing, memory, or energy resources. The establishment and maintenance of routing neighbors opens the routing process to engagement and potential acceptance of multiple neighboring peers. Association information must be stored for each peer entity and for the wireless network operation provisions made to

periodically update and reassess the associations. An introduced proliferation of apparent routing peers can therefore have a negative impact on node resources.

Node resources may also be unduly consumed by attackers attempting uncontrolled topology peering or routing exchanges, routing replays, or the generating of other data traffic floods. Beyond the disruption of communications channel resources, these consequences may be able to exhaust node resources only where the engagements are able to proceed with the peer routing entities. Routing operation and network forwarding functions can thus be adversely impacted by node resources exhaustion that stems from attacks that include:

- o Identity (including Sybil) attacks (see [Sybil2002]);
- o Routing information replay attacks;
- o HELLO-type flood attacks;
- o Overload attacks. (Section 7.3.2)

7. Countermeasures

By recognizing the characteristics of LLNs that may impact routing, this analysis provides the basis for understanding the capabilities within RPL used to deter the identified attacks and mitigate the threats. The following subsections consider such countermeasures by grouping the attacks according to the classification of the ISO 7498-2 model so that associations with the necessary security services are more readily visible.

7.1. Confidentiality Attack Countermeasures

Attacks to disclosure routing information may be mounted at the level of the routing information assets, at the points of access associated with routing exchanges between nodes, or through device interface access. To gain access to routing/topology information, the attacker may rely on a compromised node that deliberately exposes the information during the routing exchange process, may rely on passive wiretapping or traffic analysis, or may attempt access through a component or device interface of a tampered routing node.

7.1.1. Countering Deliberate Exposure Attacks

A deliberate exposure attack is one in which an entity that is party to the routing process or topology exchange allows the routing/topology information or generated route information to be exposed to an unauthorized entity.

For instance, due to mis-configuration or inappropriate enabling of a diagnostic interface, an entity might be copying ("bridging") traffic from a secured ESSID/PAN to an unsecured interface.

A prerequisite to countering this attack is to ensure that the communicating nodes are authenticated prior to data encryption applied in the routing exchange. The authentication ensures the LLN starts with trusted nodes, but it does not provide an indication of whether the node has been compromised.

Reputation systems could be used to help when some nodes may sleep for extended periods of times. It is also unclear if resulting dataset would even fit into constrained devices.

To mitigate the risk of deliberate exposure, the process that communicating nodes use to establish session keys must be peer-to-peer (i.e., between the routing initiating and responding nodes). As is pointed out in [RFC4107], automatic key management is critical for good security. This helps ensure that neither node is exchanging routing information with another peer without the knowledge of both communicating peers. For a deliberate exposure attack to succeed, the compromised node will need to be more overt and take independent actions in order to disclose the routing information to 3rd party.

Note that the same measures which apply to securing routing/topology exchanges between operational nodes must also extend to field tools and other devices used in a deployed network where such devices can be configured to participate in routing exchanges.

7.1.2. Countering Passive Wiretapping Attacks

A passive wiretap attack seeks to breach routing confidentiality through passive, direct analysis and processing of the information exchanges between nodes.

Passive wiretap attacks can be directly countered through the use of data encryption for all routing exchanges. Only when a validated and authenticated node association is completed will routing exchange be allowed to proceed using established session keys and an agreed encryption algorithm. The mandatory to implement CCM mode AES-128 method, is described in [RFC3610], and is believed to be secure against a brute force attack by even the most well-equipped adversary.

The significant challenge for RPL is in the provisioning of the key, which in some modes of RFC6550 is used network-wide. RFC6550 does not solve this problem, and it is the subject of significant future work: see, for instance: [AceCharterProposal], [SolaceProposal], [SmartObjectSecurityWorkshop].

A number of deployments, such as [ZigBeeIP] specify no layer-3/RPL encryption or authentication and rely upon similiar security at layer-2. These networks are immune to outside wiretapping attacks, but are vulnerable to passive (and active) routing attacks through compromises of nodes. (see Section 8.2).

Section 10.9 of [RFC6550] specifies AES-128 in CCM mode with a 32-bit MAC.

Section 5.6 Zigbee IP [ZigBeeIP] specifies use of CCM, with PANA and EAP-TLS for key management.

7.1.3. Countering Traffic Analysis

Traffic analysis provides an indirect means of subverting confidentiality and gaining access to routing information by allowing an attacker to indirectly map the connectivity or flow patterns (including link-load) of the network from which other attacks can be mounted. The traffic analysis attack on an LLN, especially one founded on shared medium, is passive and relies on the ability to read the immutable source/destination layer-2 and/or layer-3 routing information that must remain unencrypted to permit network routing.

One way in which passive traffic analysis attacks can be muted is through the support of load balancing that allows traffic to a given destination to be sent along diverse routing paths. RPL does not generally support multi-path routing within a single DODAG. Multiple DODAGs are supported in the protocol, and an implementation could make use of that. RPL does not have any inherent or standard way to guarantee that the different DODAGs would have significantly diverse paths. Having the diverse DODAGs routed at different border routers might work in some instances, and this could be combined with a multipath technology like MPTCP ([RFC6824]). It is unlikely that it will be affordable in many LLNs, as few deployments will have memory space for more than a few sets of DODAG tables.

Another approach to countering passive traffic analysis could be for nodes to maintain constant amount of traffic to different destinations through the generation of arbitrary traffic flows; the drawback of course would be the consequent overhead and energy expenditure.

The only means of fully countering a traffic analysis attack is through the use of tunneling (encapsulation) where encryption is applied across the entirety of the original packet source/destination addresses. Deployments which use layer-2 security that includes encryption already do this for all traffic.

7.1.4. Countering Remote Device Access Attacks

Where LLN nodes are deployed in the field, measures are introduced to allow for remote retrieval of routing data and for software or field upgrades. These paths create the potential for a device to be remotely accessed across the network or through a provided field tool. In the case of network management a node can be directly requested to provide routing tables and neighbor information.

To ensure confidentiality of the node routing information against attacks through remote access, any local or remote device requesting routing information must be authenticated, and must be authorized for that access. Since remote access is not invoked as part of a routing protocol, security of routing information stored on the node against remote access will not be addressable as part of the routing protocol.

7.2. Integrity Attack Countermeasures

Integrity attack countermeasures address routing information manipulation, as well as node identity and routing information misuse. Manipulation can occur in the form of falsification attack and physical compromise. To be effective, the following development considers the two aspects of falsification, namely, the unauthorized modifications and the overclaiming and misclaiming content. The countering of physical compromise was considered in the previous section and is not repeated here. With regard to misuse, there are two types of attacks to be deterred, identity attacks and replay attacks.

7.2.1. Countering Unauthorized Modification Attacks

Unauthorized modifications may occur in the form of altering the message being transferred or the data stored. Therefore, it is necessary to ensure that only authorized nodes can change the portion of the information that is allowed to be mutable, while the integrity of the rest of the information is protected, e.g., through well-studied cryptographic mechanisms.

Unauthorized modifications may also occur in the form of insertion or deletion of messages during protocol changes. Therefore, the protocol needs to ensure the integrity of the sequence of the exchange sequence.

The countermeasure to unauthorized modifications needs to:

- o implement access control on storage;
- o provide data integrity service to transferred messages and stored data;
- o include sequence number under integrity protection.

7.2.2. Countering Overclaiming and Misclaiming Attacks

Both overclaiming and misclaiming aim to introduce false routes or a false topology that would not occur otherwise, while there are not necessarily unauthorized modifications to the routing messages or information. In order to counter overclaiming, the capability to determine unreasonable routes or topology is required.

The counter to overclaiming and misclaiming may employ:

- o comparison with historical routing/topology data;
- o designs which restrict realizable network topologies.

RPL includes no specific mechanisms in the protocol to counter overclaims or misclaims. An implementation could have specific heuristics implemented locally.

7.2.3. Countering Identity (including Sybil) Attacks

Identity attacks, sometimes simply called spoofing, seek to gain or damage assets whose access is controlled through identity. In routing, an identity attacker can illegitimately participate in routing exchanges, distribute false routing information, or cause an invalid outcome of a routing process.

A perpetrator of Sybil attacks assumes multiple identities. The result is not only an amplification of the damage to routing, but extension to new areas, e.g., where geographic distribution is explicitly or implicitly an asset to an application running on the LLN, for example, the LBR in a P2MP or MP2P LLN.

RPL includes specific public key based authentication at layer-3 that provide for authorization. Many deployments use layer-2 security

that includes admission controls at layer-2 using mechanisms such as PANA.

7.2.4. Countering Routing Information Replay Attacks

In many routing protocols, message replay can result in false topology and/or routes. This is often countered with some kind of counter to ensure the freshness of the message. Replay of a current, literal RPL message are in general idempotent to the topology. An older (lower DODAGVersionNumber) message, if replayed would be rejected as being stale. The trickle algorithm further dampens the effect of any such replay, as if the message was current, then it would contain the same information as before, and it would cause no network changes.

Replays may well occur in some radio technologies (not very likely, 802.15.4) as a result of echos or reflections, and so some replays must be assumed to occur naturally.

Note that for there to be no affect at all, the replay must be done with the same apparent power for all nodes receiving the replay. A change in apparent power might change the metrics through changes to the ETX and therefore might affect the routing even though the contents of the packet were never changed. Any replay which appears to be different should be analyzed as a Selective Forwarding Attack, Sinkhole Attack or Wormhole Attack.

7.2.5. Countering Byzantine Routing Information Attacks

Where a node is captured or compromised but continues to operate for a period with valid network security credentials, the potential exists for routing information to be manipulated. This compromise of the routing information could thus exist in spite of security countermeasures that operate between the peer routing devices.

Consistent with the end-to-end principle of communications, such an attack can only be fully addressed through measures operating directly between the routing entities themselves or by means of external entities able to access and independently analyze the routing information. Verification of the authenticity and liveness of the routing entities can therefore only provide a limited counter against internal (Byzantine) node attacks.

For link state routing protocols where information is flooded with, for example, areas (OSPF [RFC2328]) or levels (ISIS [RFC7142]), countermeasures can be directly applied by the routing entities through the processing and comparison of link state information received from different peers. By comparing the link information

from multiple sources decisions can be made by a routing node or external entity with regard to routing information validity; see Chapter 2 of [Perlman1988] for a discussion on flooding attacks.

For distance vector protocols, such as RPL, where information is aggregated at each routing node it is not possible for nodes to directly detect Byzantine information manipulation attacks from the routing information exchange. In such cases, the routing protocol must include and support indirect communications exchanges between non-adjacent routing peers to provide a secondary channel for performing routing information validation. S-RIP [Wan2004] is an example of the implementation of this type of dedicated routing protocol security where the correctness of aggregate distance vector information can only be validated by initiating confirmation exchanges directly between nodes that are not routing neighbors.

RPL does not provide any direct mechanisms like S-RIP. It does listen to multiple parents, and may switch parents if it begins to suspect that it is being lied to.

7.3. Availability Attack Countermeasures

As alluded to before, availability requires that routing information exchanges and forwarding mechanisms be available when needed so as to guarantee proper functioning of the network. This may, e.g., include the correct operation of routing information and neighbor state information exchanges, among others. We will highlight the key features of the security threats along with typical countermeasures to prevent or at least mitigate them. We will also note that an availability attack may be facilitated by an identity attack as well as a replay attack, as was addressed in Section 7.2.3 and Section 7.2.4, respectively.

7.3.1. Countering HELLO Flood Attacks and ACK Spoofing Attacks

HELLO Flood [Karlof2003],[I-D.suhopark-hello-wsn] and ACK Spoofing attacks are different but highly related forms of attacking an LLN. They essentially lead nodes to believe that suitable routes are available even though they are not and hence constitute a serious availability attack.

A HELLO attack mounted against RPL would involve sending out (or replaying) DIO messages by the attacker. Lower power LLN nodes might then attempt to join the DODAG at a lower rank than they would otherwise.

The most effective method from [I-D.suhopark-hello-wsn] is the verify bidirectionality. A number of layer-2 links are arranged in

controller/spoke arrangements, and continuously are validating connectivity at layer 2.

In addition, in order to calculate metrics, the ETX must be computed, and this involves, in general, sending a number of messages between nodes which are believed to be adjacent.

[I-D.kelsey-intarea-mesh-link-establishment] is one such protocol.

In order to join the DODAG, a DAO message is sent upwards. In RPL the DAO is acknowledged by the DAO-ACK message. This clearly checks bidirectionality at the control plane.

As discussed in section 5.1, [I-D.suhopark-hello-wsn] a receiver with a sensitive receiver could well hear the DAOs, and even send DAO-ACKs as well. Such a node is a form of wormhole attack.

These attacks are also all easily defended against using either layer-2 or layer-3 authentication. Such an attack could only be made against a completely open network (such as might be used for provisioning new nodes), or by a compromised node.

7.3.2. Countering Overload Attacks

Overload attacks are a form of DoS attack in that a malicious node overloads the network with irrelevant traffic, thereby draining the nodes' energy store more quickly, when the nodes rely on batteries or energy scavenging. It thus significantly shortens the lifetime of networks of energy-constrained nodes and constitutes another serious availability attack.

With energy being one of the most precious assets of LLNs, targeting its availability is a fairly obvious attack. Another way of depleting the energy of an LLN node is to have the malicious node overload the network with irrelevant traffic. This impacts availability since certain routes get congested which:

- o renders them useless for affected nodes and data can hence not be delivered;
- o makes routes longer as shortest path algorithms work with the congested network;
- o depletes battery and energy scavenging nodes more quickly and thus shortens the network's availability at large.

Overload attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o introduce quotas on the traffic rate each node is allowed to send;
- o isolate nodes which send traffic above a certain threshold based on system operation characteristics;
- o allow only trusted data to be received and forwarded.

As for the first one, a simple approach to minimize the harmful impact of an overload attack is to introduce traffic quotas. This prevents a malicious node from injecting a large amount of traffic into the network, even though it does not prevent said node from injecting irrelevant traffic at all. Another method is to isolate nodes from the network at the network layer once it has been detected that more traffic is injected into the network than allowed by a prior set or dynamically adjusted threshold. Finally, if communication is sufficiently secured, only trusted nodes can receive and forward traffic which also lowers the risk of an overload attack.

Receiving nodes that validate signatures and sending nodes that encrypt messages need to be cautious of cryptographic processing usage when validating signatures and encrypting messages. Where feasible, certificates should be validated prior to use of the associated keys to counter potential resource overloading attacks. The associated design decision needs to also consider that the validation process requires resources and thus itself could be exploited for attacks. Alternatively, resource management limits can be placed on routing security processing events (see the comment in Section 6, paragraph 4, of [RFC5751]).

7.3.3. Countering Selective Forwarding Attacks

Selective forwarding attacks are a form of DoS attack which impacts the availability of the generated routing paths.

A selective forwarding attack may be done by a node involved with the routing process, or it may be done by what otherwise appears to be a passive antenna or other RF feature or device, but is in fact an active (and selective) device. An RF antenna/repeater which is not selective, is not a threat.

An insider malicious node basically blends neatly in with the network but then may decide to forward and/or manipulate certain packets. If all packets are dropped, then this attacker is also often referred to as a "black hole". Such a form of attack is particularly dangerous if coupled with sinkhole attacks since inherently a large amount of traffic is attracted to the malicious node and thereby causing significant damage. In a shared medium, an outside malicious node would selectively jam overheard data flows, where the thus caused collisions incur selective forwarding.

Selective Forwarding attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o multipath routing of the same message over disjoint paths;
- o dynamically selecting the next hop from a set of candidates.

The first measure basically guarantees that if a message gets lost on a particular routing path due to a malicious selective forwarding attack, there will be another route which successfully delivers the data. Such a method is inherently suboptimal from an energy consumption point of view; it is also suboptimal from a network utilization perspective. The second method basically involves a constantly changing routing topology in that next-hop routers are chosen from a dynamic set in the hope that the number of malicious nodes in this set is negligible. A routing protocol that allows for disjoint routing paths may also be useful.

7.3.4. Countering Sinkhole Attacks

In sinkhole attacks, the malicious node manages to attract a lot of traffic mainly by advertising the availability of high-quality links even though there are none [Karlof2003]. It hence constitutes a serious attack on availability.

The malicious node creates a sinkhole by attracting a large amount of, if not all, traffic from surrounding neighbors by advertising in and outwards links of superior quality. Affected nodes hence eagerly route their traffic via the malicious node which, if coupled with other attacks such as selective forwarding, may lead to serious availability and security breaches. Such an attack can only be executed by an inside malicious node and is generally very difficult to detect. An ongoing attack has a profound impact on the network topology and essentially becomes a problem of flow control.

Sinkhole attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o use geographical insights for flow control;
- o isolate nodes which receive traffic above a certain threshold;
- o dynamically pick up next hop from set of candidates;
- o allow only trusted data to be received and forwarded.

A canary node could periodically call home (using a cryptographic process), with the home system noting if it fails to call in. This provides detection of a problem, but does not mitigate it, and it may have significant energy consequences for the LLN.

Some LLNs may provide for geolocation services, often derived from solving triangulation equations from radio delay calculations, such calculations could in theory be subverted by a sinkhole that transmitted at precisely the right power in a node to node fashion.

While geographic knowledge could help assure that traffic always went in the physical direction desired, it would not assure that the traffic was taking the most efficient route, as the lowest cost real route might not match the physical topology; such as when different parts of an LLN are connected by high-speed wired networks.

7.3.5. Countering Wormhole Attacks

In wormhole attacks at least two malicious nodes claim to have a short path between themselves [Karlof2003]. This changes the availability of certain routing paths and hence constitutes a serious security breach.

Essentially, two malicious insider nodes use another, more powerful, transmitter to communicate with each other and thereby distort the would-be-agreed routing path. This distortion could involve shortcutting and hence paralyzing a large part of the network; it could also involve tunneling the information to another region of the network where there are, e.g., more malicious nodes available to aid the intrusion or where messages are replayed, etc.

In conjunction with selective forwarding, wormhole attacks can create race conditions which impact topology maintenance, routing protocols as well as any security suits built on "time of check" and "time of use".

A pure wormhole attack is nearly impossible to detect. A wormhole which is used in order to subsequently mount another kind of attack would be defeated by defeating the other attack. A perfect wormhole, in which there is nothing adverse that occurs to the traffic, would

be difficult to call an attack. The worst thing that a benign wormhole can do in such a situation is to cease to operate (become unstable), causing the network to have to recalculate routes.

A highly unstable wormhole is no different than a radio opaque (i.e. metal) door that opens and closes a lot. RPL includes hysteresis in its objective functions [RFC6719] in an attempt to deal with frequent changes to the ETX between nodes.

8. RPL Security Features

The assessments and analysis in Section 6 examined all areas of threats and attacks that could impact routing, and the countermeasures presented in Section 7 were reached without confining the consideration to means only available to routing. This section puts the results into perspective; dealing with those threats which are endemic to this field, those which have been mitigated through RPL protocol design, and those which require specific decisions to be made as part of provisioning a network.

The first part of this section, Section 8.1 to Section 8.3, is a description of RPL security features that address specific threats. The second part of this section, Section 8.4, discusses issues of provisioning of security aspects that may impact routing but that also require considerations beyond the routing protocol, as well as potential approaches.

RPL employs multicast and so these alternative communications modes MUST be secured with the same routing security services specified in this section. Furthermore, irrespective of the modes of communication, nodes MUST provide adequate physical tamper resistance commensurate with the particular application domain environment to ensure the confidentiality, integrity, and availability of stored routing information.

8.1. Confidentiality Features

With regard to confidentiality, protecting the routing/topology information from unauthorized disclosure is not directly essential to maintaining the routing function. Breaches of confidentiality may lead to other attacks or the focusing of an attacker's resources (see Section 6.2) but does not of itself directly undermine the operation of the routing function. However, to protect against, and reduce consequences from other more direct attacks, routing information should be protected. Thus, to secure RPL:

- o implement payload encryption using layer-3 mechanisms described in [RFC6550];

- o or: implement layer-2 confidentiality;

Where confidentiality is incorporated into the routing exchanges, encryption algorithms and key lengths need to be specified in accordance with the level of protection dictated by the routing protocol and the associated application domain transport network. For most networks, this means use of AES128 in CCM mode, but this needs to be specified clearly in the applicability statement.

In terms of the life time of the keys, the opportunity to periodically change the encryption key increases the offered level of security for any given implementation. However, where strong cryptography is employed, physical, procedural, and logical data access protection considerations may have more significant impact on cryptoperiod selection than algorithm and key size factors. Nevertheless, in general, shorter cryptoperiods, during which a single key is applied, will enhance security.

Given the mandatory protocol requirement to implement routing node authentication as part of routing integrity (see Section 8.2), key exchanges may be coordinated as part of the integrity verification process. This provides an opportunity to increase the frequency of key exchange and shorten the cryptoperiod as a complement to the key length and encryption algorithm required for a given application domain.

8.2. Integrity Features

The integrity of routing information provides the basis for ensuring that the function of the routing protocol is achieved and maintained. To protect integrity, RPL must either run using only the Secure versions of the messages, or must run over a layer-2 that uses channel binding between node identity and transmissions.

Some layer-2 security mechanisms use a single key for the entire network, and these networks can not provide significant amount of integrity protection, as any node that has that key may impersonate any other node. This mode of operation is likely acceptable when an entire deployment is under the control of a single administrative entity.

Other layer-2 security mechanisms form a unique session key for every pair of nodes that needs to communicate; this is often called a per-link key. Such networks can provide a strong degree of origin authentication and integrity on unicast messages.

However, some RPL messages are broadcast, and even when per-node layer-2 security mechanisms are used, the integrity and origin

authentication of broadcast messages can not be as trusted due to the proliferation of the key used to secure them.

RPL has two specific options which are broadcast in RPL Control Messages: the DODAG Information Object (DIO), and the DODAG Information Solicitation (DIS). The purpose of the DIS is to cause potential parents to reply with a DIO, so the integrity of the DIS is not of great concern. The DIS may also be unicast.

The DIO is a critical piece of routing and carries many critical parameters. RPL provides for asymmetric authentication at layer 3 of the RPL Control Message carrying the DIO and this may be warranted in some deployments. A node could, if it felt that the DIO that it had received was suspicious, send a unicast DIS message to the node in question, and that node would reply with a unicast DIS. Those messages could be protected with the per-link key.

8.3. Availability Features

Availability of routing information is linked to system and network availability which in the case of LLNs require a broader security view beyond the requirements of the routing entities. Where availability of the network is compromised, routing information availability will be accordingly affected. However, to specifically assist in protecting routing availability, nodes:

- o MAY restrict neighborhood cardinality;
- o MAY use multiple paths;
- o MAY use multiple destinations;
- o MAY choose randomly if multiple paths are available;
- o MAY set quotas to limit transmit or receive volume;
- o MAY use geographic information for flow control.

8.4. Key Management

The functioning of the routing security services requires keys and credentials. Therefore, even though not directly a RPL security requirement, an LLN MUST have a process for initial key and credential configuration, as well as secure storage within the associated devices. Anti-tampering SHOULD be a consideration in physical design. Beyond initial credential configuration, an LLN is also encouraged to have automatic procedures for the revocation and replacement of the maintained security credentials.

While RPL has secure modes, but some modes are impractical without use of public key cryptography believed to be too expensive by many. RPL layer-3 security will often depend upon existing LLN layer-2 security mechanisms, which provides for node authentication, but little in the way of node authorization.

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

The analysis presented in this document provides security analysis and design guidelines with a scope limited to RPL. Security services are identified as requirements for securing RPL. The specific mechanisms to be used to deal with each threat is specified in link-layer and deployment specific applicability statements.

11. Acknowledgments

The authors would like to acknowledge the review and comments from Rene Struik and JP Vasseur. The authors would also like to acknowledge the guidance and input provided by the RPL Chairs, David Culler, and JP Vasseur, and the Area Director Adrian Farrel.

This document started out as a combined threat and solutions document. As a result of a series of security reviews performed by Steve Kent, the document was split up by RPL co-Chair Michael Richardson and security Area Director Sean Turner as it went through the IETF publication process. The solutions to the threats are application and layer-2 specific, and have therefore been moved to the relevant applicability statements.

Ines Robles and Robert Cragie kept track of the many issues that were raised during the development of this document

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.
- [RFC4107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", BCP 107, RFC 4107, June 2005.
- [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.

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

[RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, January 2014.

[ZigBeeIP]
ZigBee Public Document 15-002r00, "ZigBee IP Specification", 2013.

12.2. Informative References

[AceCharterProposal]
Li, Kepeng., Ed., "Authentication and Authorization for Constrained Environment Charter (work-in-progress)", December 2013, <http://trac.tools.ietf.org/wg/core/trac/wiki/ACE_charter>.

[I-D.gilger-smart-object-security-workshop]
Gilger, J. and H. Tschofenig, "Report from the 'Smart Object Security Workshop', March 23, 2012, Paris, France", draft-gilger-smart-object-security-workshop-02 (work in progress), October 2013.

[I-D.kelsey-intarea-mesh-link-establishment]
Kelsey, R., "Mesh Link Establishment", draft-kelsey-intarea-mesh-link-establishment-05 (work in progress), February 2013.

[I-D.suhopark-hello-wsn]
Park, S., "Routing Security in Sensor Network: HELLO Flood Attack and Defense", draft-suhopark-hello-wsn-00 (work in progress), December 2005.

[IEEE.802.11]
, "Draft Standard for Information Technology - Telecommunications and information exchange between systems - Local and metropolitan area networks Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications ", IEEE 802.11-REVma, 2006.

[IEEE.802.15.4]
, "Information technology - Telecommunications and information exchange between systems - Local and

metropolitan area networks - Specific requirements - Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANs) ", IEEE Std 802.15.4-2006, June 2006, <<http://standards.ieee.org/getieee802/802.15.html>>.

[ISO.7498-2.1988]

International Organization for Standardization,
"Information Processing Systems - Open Systems
Interconnection Reference Model - Security Architecture",
ISO Standard 7498-2, 1988.

[Karlof2003]

Karlof, C. and D. Wagner, "Secure routing in wireless sensor networks: attacks and countermeasures", Elsevier AdHoc Networks Journal, Special Issue on Sensor Network Applications and Protocols, 1(2):293-315, September 2003, <<http://nest.cs.berkeley.edu/papers/sensor-route-security.pdf>>.

[Myagmar2005]

Myagmar, S., Lee, A.J., and W. Yurcik, "Threat Modeling as a Basis for Security Requirements", in Proceedings of the Symposium on Requirements Engineering for Information Security (SREIS'05), Paris, France, pp. 94-102, Aug 29, 2005.

[Perlman1988]

Perlman, N., "Network Layer Protocols with Byzantine Robustness", MIT LCS Tech Report, 429, 1988.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, September 2003.

[RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, October 2006.

[RFC4732] Handley, M., Rescorla, E., IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 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.

- [RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, March 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.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [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.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, March 2011.
- [RFC6574] Tschofenig, H. and J. Arkko, "Report from the Smart Object Workshop", RFC 6574, April 2012.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.
- [RFC7142] Shand, M. and L. Ginsberg, "Reclassification of RFC 1142 to Historic", RFC 7142, February 2014.
- [SmartObjectSecurityWorkshop] Klausen, T., Ed., "Workshop on Smart Object Security", March 2012, <<http://www.lix.polytechnique.fr/hipercom/SmartObjectSecurity>>.
- [SolaceProposal] Bormann, C., Ed., "Notes from the SOLACE ad-hoc at IETF85 (work-in-progress)", November 2012, <<http://www.ietf.org/mail-archive/web/solace/current/msg00015.html>>.
- [Sybil2002]

Douceur, J., "The Sybil Attack", First International Workshop on Peer-to-Peer Systems , March 2002.

[Wan2004] Wan, T., Kranakis, E., and PC. van Oorschot, "S-RIP: A Secure Distance Vector Routing Protocol", in Proceedings of the 2nd International Conference on Applied Cryptography and Network Security, Yellow Mountain, China, pp. 103-119, Jun. 8-11 2004.

[Yourdon1979]
Yourdon, E. and L. Constantine, "Structured Design",
Yourdon Press, New York, Chapter 10, pp. 187-222, 1979.

Authors' Addresses

Tzeta Tsao
Cooper Power Systems
910 Clopper Rd. Suite 201S
Gaithersburg, Maryland 20878
USA

Email: tzeta.tsao@cooperindustries.com

Roger K. Alexander
Cooper Power Systems
910 Clopper Rd. Suite 201S
Gaithersburg, Maryland 20878
USA

Email: roger.alexander@cooperindustries.com

Mischa Dohler
CTTC
Parc Mediterrani de la Tecnologia, Av. Canal Olímpic S/N
Castelldefels, Barcelona 08860
Spain

Email: mischa.dohler@cttc.es

Vanesa Daza
Universitat Pompeu Fabra
P/ Circumval.lacio 8, Oficina 308
Barcelona 08003
Spain

Email: vanesa.daza@upf.edu

Angel Lozano
Universitat Pompeu Fabra
P/ Circumval.lacio 8, Oficina 309
Barcelona 08003
Spain

Email: angel.lozano@upf.edu

Michael Richardson (ed) (editor)
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z5V7
Canada

Email: mcr+ietf@sandelman.ca

6MAN
Internet-Draft
Intended status: Standards Track
Expires: February 24, 2015

P. Thubert, Ed.
Cisco
August 25, 2014

The IPv6 Flow Label within a LLN domain
draft-thubert-6man-flow-label-for-rpl-05

Abstract

This document presents how the Flow Label can be used inside a LLN domain such as a RPL domain or an ISA100.11a D-subnet, and provides updated rules for a domain Border Router to set and reset the Flow Label when forwarding between inside the domain and the larger Internet in both direction. Rules for routers inside the domain are also provided.

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 February 24, 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	2
2. Terminology	3

3. Requirements for LLN Flows	3
4. On Compatibility With Existing Standards	4
5. Updated Rules	5
6. Security Considerations	6
7. IANA Considerations	6
8. Acknowledgements	6
9. References	7
9.1. Normative References	7
9.2. Informative References	7
Author's Address	8

1. Introduction

The design of Lowpower Lossy Networks (LLNs) is generally focussed on saving energy, which is typically the most constrained resource of all. Other classical constraints, such as memory capacity, frame size, as well as the duty cycling of the LLN devices, derive from that primary concern.

In isolated devices, energy is typically available from batteries that are expected to last for years, or scavenged from the environment in very limited quantities. Any protocol that is intended for use in LLNs must be designed with the primary concern of saving energy as a strict requirement.

The IEEE802.15.4 [IEEE802154] was designed to offer the Physical (PHY) and Medium Access Control (MAC) layers for low-cost, low-speed, low-power Wireless Personal Area Networks (WPANs), which are a wireless form of LLNs.

With the traditional IEEE802.15.4 PHY, frames are limited to 127 octets. In order to adapt IPv6 [RFC2460] over IEEE802.15.4, 6LoWPAN [RFC4944] introduced a fragmentation mechanism under IP, which in turn causes even more energy spending and other issues as discussed in LLN Fragment Forwarding and Recovery [I-D.thubert-6lo-forwarding-fragments].

The IEEE802.15.4e Task Group further defined the TimeSlotted Channel Hopping [I-D.ietf-6tisch-tsch] (TSCH) mode of operation as an update to the MAC specification in order to address Time Sensitive applications.

The 6TiSCH architecture [I-D.ietf-6tisch-architecture] specifies the operation of IPv6 over IEEE802.15.4e TSCH networks attached and synchronized by backbone routers. 6TiSCH was created to simplify the adoption of IETF technology by other Standard Defining Organizations (SDOs), in particular in the Industrial Automation space, which already relies on variations of IEEE802.15.4e TSCH for Wireless Sensor Networking.

The ISA100.11a [ISA100.11a] specification provides an example of such an industrial WSN standard, using a precursor to IEEE802.15.4e over the classical IEEE802.14.5 PHY. In that case, after security is applied, roughly 80 octets are available per frame for IP and Payload. In order to 1) avoid fragmentation and 2) conserve energy, the ISA100 WG in charge of that specification did scrutinize the use of every bit in the frame and rejected any perceived waste.

The challenge to obtain the adoption of IPv6 in the original standard was thus to save all possible bits in the frames, including the UDP checksum which was an interesting discussion on its own. This work was actually one of the roots for the 6LoWPAN Header Compression [RFC6282] work, which goes down to the individual bits to save space in the frames for actual data, and allowed ISA100.11a to adopt IPv6.

ISA100.11a (now IEC62734) uses IPv6 over UDP, and conforms to a number of other IETF RFCs including the IPv6 Flow Label Specification [RFC3697] that was the reference at the time the standard was elaborated, but fails to conform to the newer IPv6 Flow Label Specification [RFC6437] that obsoleted it.

The bone of contention is the use of the Flow Label as an index called a contract ID, and the capability for the Backbone Router, that is the Border Router of a ISA100.11a WSN (also called a D-subnet), to modify the Flow Label. There is work at ROLL that indicates that RPL nodes may benefit from similar abilities to also transport flow-related information in the Flow Label.

This document adds an exception to the rules in [RFC6437], for application within a well-defined LLN domain, whereby the Border Routers would be in a position to ensure that from an external viewpoint, the domain complies to the new Flow Label specification even though the internal use of the Flow Label does not.

2. Terminology

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].

This document uses Terminology defined in Terminology in Low power And Lossy Networks [RFC7102], as well as [RFC6550] and [RFC6553].

3. Requirements for LLN Flows

In Industrial Automation and Control Systems (IACS) [RFC5673], a packet loss is usually acceptable but jitter and latency must be strictly controlled as they can play a critical role in the interpretation of the measured information. Sensory systems are often distributed, and the control information can in fact be originated from multiple sources and aggregated. In such cases, related packets from multiple sources should not be load-balanced

along their path in the Internet.

In a typical LLN application, the bulk of the traffic consists of small chunks of data (in the order few bytes to a few tens of bytes) at a time. 4Hz is a typical loop frequency in Process Control, though it can be a lot slower than that in, say, environmental monitoring. The granularity of traffic from a single source is too small to make a lot of sense in load balancing application.

As a result, it can be a requirement for related measurements from multiple sources to be treated as a single flow following a same path over the Internet so as to experience similar jitter and latency. The traditional tuple of source, destination and ports might then not be the proper indication to isolate a consistent flow. On the other hand, the flow integrity can be preserved in a simple manner if the setting of the Flow Label in the IPv6 header of packets outgoing a LLN domain, is centralized to the Border Router, such as the root of a RPL DODAG structure, or an ISA100.11a Backbone Router, as opposed to distributed across the actual sources.

Considering that the goal for setting the Flow Label as prescribed in the IPv6 Flow Label Specification [RFC6437] is to improve load balancing in the core of the Internet, it is unlikely that LLN devices will consume energy to generate and then transmit a Flow Label to serve outside interests and the Flow Label is generally left to zero so as to be elided in the 6LoWPAN [RFC6282] compression. So in a general manner the interests of the core are better served if the RPL roots systematically rewrite the flow label rather than if they never do.

For packets coming into the RPL domain from the Internet, the value for setting the Flow Label as prescribed in [RFC6437] is consumed once the packet has traversed the core and reaches the LLN. Then again, there is little value but a high cost for the LLN in spending 20 bits to transport a Flow Label, that was set by a peer or a router in the Internet, over the constrained network to a destination node that has no use of it.

On a PHY layer with super-short frames such as IEEE802.15.4, compliance with those rules will simply not happen, and the rules will become an bone of contention for IPv6 adoption at a time where great progress is happening towards that goal, as illustrated by the activity at 6lo on multiple LLN Link-layers.

4. On Compatibility With Existing Standards

All the packets from all the nodes in a same DODAG that are leaving a RPL domain towards the Internet will transit via a same RPL root. The RPL root segregates the Internet and the RPL domain, which enables the capability to reuse the Flow Label within the RPL domain. The ISA100.11a Backbone Router plays a similar role and interfaces an ISA100.11a WSN D-subnet with a larger IPv6 network.

This specification enables the operation of resetting or reusing the IPv6 Flow Label at the border of a LLN domain. This is a deviation from the IPv6 Flow Label Specification [RFC6437], in that the LLN border router is neither the source nor the first hop router that sets the final Flow Label for use outside the LLN domain.

But if we consider the whole RPL domain as a large virtual host from the standpoint of the rest of the Internet, the interests that lead to [RFC6437], and in particular load balancing in the core of the Internet, are probably better served if the root guarantees that the Flow Label is set in a compliant fashion than if we rely on each individual sensor that may not use it at all, or use it slightly differently such as done in ISA100.11a.

Additionally, LLN flows can be compound flows aggregating information from multiple sources. The Border Router is an ideal place to rewrite the Flow Label to a same value for a same flow across multiple sources, ensuring compliance with the rules defined by [RFC6437] for use outside of the RPL domain and in particular in the core of the Internet.

This document specifies how the Flow Label can be reused within a LLN domain such as a RPL domain and an ISA100.11a D-subnet, in which a Border Router delineates the limit of the domain and may rewrite the Flow Label on all packets. In a RPL domain, it will become acceptable to use the Flow Label as replacement to the RPL option, though whether that operation gets standardized is left to be discussed. That use of the Flow Label within a RPL domain would be an instance of the stateful scenarios as discussed in [RFC6437] where the flow state in the node is indexed by the RPLInstanceID that identifies the routing topology. ISA100.11a would be another instance where the 16bit Contract ID in the Flow Label identifies a state in a node that is specific to a particular flow.

5. Updated Rules

This specification applies to a constrained LLN domain that forms a stub and is connected to the Internet by and only by its Border Routers. In the case of a RPL domain, the RPL root is such a bottleneck for all the traffic between the Internet and the Destination-Oriented Directed Acyclic Graph (DODAG) that it serves. This specification also covers other LLN domains with the same properties of having strict constraints in energy and/or frame size, such as an ISA100.11a [ISA100.11a] Industrial Wireless Sensor Network, but does not generalize to any arbitrary domain. This updates the IPv6 Flow Label Specification [RFC6437], which does not allow any specific rule in any particular domain, and updates it only in the context of constrained LLN domains.

In that context, a LLN domain Border Router MAY rewrite the Flow Label of all packets entering or leaving the RPL domain in both directions, from and towards the Internet, regardless of its original setting. For the limited context of a constrained LLN domain, this updates the IPv6 Flow Label Specification [RFC6437] which stipulates that once it is set, the Flow Label is left unchanged; but the RFC also indicates a violation to the rule can be accepted for compelling reasons related to security. This specification adds that energy-saving is another compelling reason for a violation to the aforementioned rule, though applicable only inside a constrained LLN.

In particular, the Border Router of a LLN domain MAY set the Flow Label of IPv6 packets that exit the LLN domain. It SHOULD do it if the LLN domain operations do not conform [RFC6437], and if it does modify the Flow Label, then it MUST do it in a manner that conforms [RFC6437] from the perspective of a Node outside the LLN.

It results that a Node in a constrained LLN domain MUST NOT assume that the setting of the Flow Label will be preserved end-to-end, and that an intermediate router inside a constrained LLN MAY alter a non-zero Flow Label between the source in the LLN and the LLN Border Router. This does not modify the expectations on end Nodes but extends the updated rules from [RFC6437] to arbitrary routers in the LLN.

For instance, a RPL root MAY reset the Flow Label of IPv6 packets entering the RPL domain to zero for an optimal Header Compression by 6LoWPAN [RFC6282]. A RPL root MAY also reuse the Flow Label towards the LLN for other purposes, such as to carry the RPL Information [RFC6553]. An ISA100.11s Backbone Router MAY reuse the Flow Label to carry local flow information, such as the Contract ID specified in ISA100.11a [ISA100.11a].

6. Security Considerations

Because the flow label is not protected by IPSec, it is expected that Layer-2 security is deployed in the LLN where is specification is applied. This is the actual best practice in LLNs, which serves in particular to avoid forwarding of untrusted packets over the constrained network.

The specification insists that the LLN Node should not expect that the Flow Label is conserved end-to-end and rather reduces the risk of misinterpretation in case of a rewrite by a router in the middle.

7. IANA Considerations

No IANA action is required for this specification.

8. Acknowledgements

The author wishes to thank Brian Carpenter for his in-depth review and constructive approach to the problem resolution.

9. References

9.1. Normative References

- [IEEE802154]
IEEE standard for Information Technology, "IEEE std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks", June 2011.
- [ISA100.11a]
ISA/ANSI, "Wireless Systems for Industrial Automation: Process Control and Related Applications - ISA100.11a-2011 - IEC 62734", 2011, <<http://www.isa.org/Community/SP100WirelessSystemsforAutomation>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2460] Deering, S.E. and R.M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3697] Rajahalme, J., Conta, A., Carpenter, B. and S. Deering, "IPv6 Flow Label Specification", RFC 3697, March 2004.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, September 2011.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S. and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, November 2011.
- [RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP. and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.
- [RFC6552] Thubert, P., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, March 2012.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, March 2012.

9.2. Informative References

- [I-D.ietf-6tisch-architecture]
Thubert, P., Watteyne, T. and R. Assimiti, "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4e", Internet-Draft draft-ietf-6tisch-architecture-01, February 2014.

[I-D.ietf-6tisch-tsch]

Watteyne, T., Palattella, M. and L. Grieco, "Using IEEE802.15.4e TSCH in an LLN context: Overview, Problem Statement and Goals", Internet-Draft draft-ietf-6tisch-tsch-00, November 2013.

[I-D.thubert-6lo-forwarding-fragments]

Thubert, P. and J. Hui, "LLN Fragment Forwarding and Recovery", Internet-Draft draft-thubert-6lo-forwarding-fragments-01, February 2014.

[RFC4944] Montenegro, G., Kushalnagar, N., Hui, J. and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007.

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

[RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, January 2014.

Author's Address

Pascal Thubert, editor
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis, 06410
FRANCE

Phone: +33 4 97 23 26 34
Email: pthubert@cisco.com