6MAN
Internet-Draft
Intended status: Standards Track
Expires: December 11, 2010

J. Hui
Arch Rock Corporation
JP. Vasseur
Cisco Systems, Inc
June 9, 2010

RPL Option for Carrying RPL Information in Data-Plane Datagrams draft-hui-6man-rpl-option-01

Abstract

The RPL protocol requires data-plane datagrams to carry RPL routing information that is processed by RPL routers when forwarding those datagrams. This document describes the RPL option for use within a RPL domain.

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 December 11, 2010.

Copyright Notice

Table of Contents

1. Introduction

RPL is a distance vector IPv6 routing protocol designed for low power and lossy networks [I-D.ietf-roll-rpl]. Such networks are typically constrained in energy and/or channel capacity. To conserve precious resources, a routing protocol must generate control traffic sparingly. However, this is at odds with the need to quickly propagate any new routing information to resolve routing inconsistencies quickly.

To help minimize resource consumption, RPL uses a slow proactive process to construct and maintain a routing topology but a reactive and dynamic approach to resolving routing inconsistencies. In the steady state, RPL maintains the routing topology using a low-rate beaconing process. However, when RPL detects inconsistencies that may prevent proper datagram delivery, RPL temporarily increases the beacon rate to quickly resolve those inconsistencies. Such a dynamic rate of control packets operation is governed by the use of dynamic timers also referred to as "trickle" timers and defined in [I-D.levis-roll-trickle]. By contrast with other routing protocols such as OSPF ([RFC2328]), RPL detects routing inconsistencies using data-path verification, by including routing information within the datagram itself. Data-path verification quickly detects and resolves inconsistencies when routes are needed by the data flow itself. In doing so, repair mechanisms operate only as needed, allowing the control and data planes to operate on similar time scales. The main motivation for data path verification in Low power and Lossy Networks (LLNs) is that control plane traffic should be carefully bounded with respect to the data traffic: there is no need to solve a routing issues (which may be temporary) in the absence of data traffic.

The RPL protocol constructs a DAG that attempts to minimize path costs to the DAG root according to a set of metric and objective functions. There are circumstances where loops may occur, and RPL is designed to use a data-path loop detection method. This is one of the known requirements of RPL and other data-path usage might be defined in the future.

To that end, this document proposes a new IPv6 option called the RPL Option to be carried within the IPv6 Hop-by-Hop header. The RPL Option is for use only within a RPL domain. Routers on the edge of the domain MAY insert the RPL Option into datagrams entering the RPL domain but MUST remove the RPL Option from datagrams exiting the RPL domains, if one exists.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Overview

Datagrams being forwarded within a RPL domain MUST include a RPL Option. For datagrams sourced within a RPL domain, the RPL Option MAY be included in the datagram itself. For datagrams sourced outside a RPL domain, IPv6-in-IPv6 tunneling, as specified in [RFC2473] MUST be used to include a RPL Option. When forwarding the datagram, the router MUST prepend a new IPv6 header and IPv6 Hop-by-Hop Options header containing the RPL Option to the existing datagram. Use of tunneling ensures that the datagram is delivered unmodified and that ICMP errors return to the RPL Option source rather than the source of the original datagram.

To help avoid IP-layer fragmentation, the RPL Option has a maximum size of RPL_OPTION_MAX_SIZE octets and links within a RPL domain SHOULD have a MTU of at least 1280 + 44 (outer IP header, Hop-by-Hop Option header, Option header) + RPL_OPTION_MAX_SIZE + (additional extension headers or options needed within RPL domain).

3. Format of the RPL Option

The RPL option is carried in an IPv6 Hop-by-Hop Options header, immediately following the IPv6 header. The RPL option has the following format:

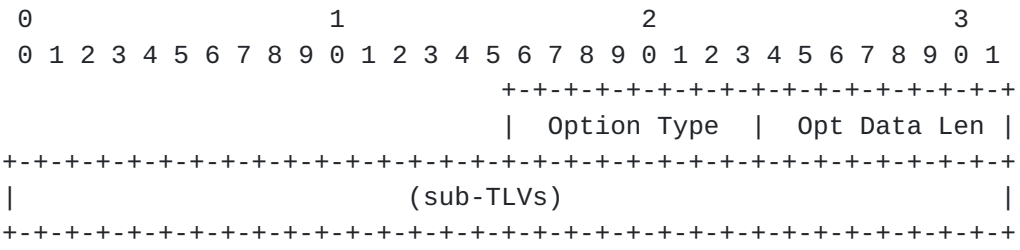

Figure 1: RPL Option

The Opt Data Len MUST NOT exceed RPL_OPTION_MAX_SIZE octets.

The Option Data of the RPL option is expected to change en-route. Nodes that do not understand the RPL option MUST skip over this option and continue processing the header. Thus, according to [RFC2460] the two high order bits of the Option Type must be equal set to zero and the third bit is equal to 1. The RPL Option Data Length is variable.

The action taken by using the RPL Option and the potential set of sub-TLVs carried within the RPL Option MUST be specified by the RFC of the protocol that use that option. No TLVs are currently defined.

4. RPL Router Behavior

Routers MUST include a RPL Option when forwarding datagrams that do not already contain a RPL Option. If one does not already exist, routers MUST use IPv6-in-IPv6 tunneling, as specified in [RFC2473] to include a RPL Option in datagrams that are sourced by other nodes. This ensures that the original datagram is delivered unmodified.

Performing IP-in-IP encapsulation may grow the datagram to a size larger than the IPv6 min MTU of 1280 octets. To help avoid IP-layer fragmentation caused by IP-in-IP encapsulation, links within a RPL domain SHOULD be configured with a MTU of at least 1280 + 44 (outer IP header, Hop-by-Hop Option header, Option header) + RPL_OPTION_MAX_SIZE + (additional extension headers or options needed within RPL domain).

5. RPL Border Router Behavior

RPL Border Routers (referred to as LBRs in [I-D.ietf-roll-terminology]) are responsible for ensuring that a RPL Option is only used within a RPL domain.

For datagrams entering the RPL domain, RPL Border Routers MUST drop received datagrams that contain a RPL Option in the IPv6 Extension headers.

For datagrams exiting the RPL domain, RPL Border Routers MUST remove the RPL Option from the datagram and update the IPv6 Payload Length field accordingly.

6. Usage of the RPL Option

The RPL option is only for use within a RPL domain. RPL routers MUST process and include the RPL option when forwarding datagrams to other nodes within the RPL domain. Routers on the edge of a RPL domain MUST remove the RPL option when forwarding datagrams to nodes outside the RPL domain. The final destination of the datagram MAY ignore the RPL option.

7. Protocol Constants

RPL_OPTION_MAX_SIZE 128

8. Acknowledgements

The authors thank Vishwas Manral and Erik Nordmark for their comments and suggestions that helped shape this document.

9. IANA Considerations

The RPL option requires an IPv6 Option Number.

HEX	act	chg	rest
1	00	1	01011

The first two bits indicate that the IPv6 node may skip over this option and continue processing the header if it doesn't recognize the option type, and the third bit indicates that the Option Data may change en-route.

10. Security Considerations

This option may be used a several potential attacks since routers may be flooded by bogus datagram containing the RPL option. It is thus RECOMMENDED for routers to implement a rate limiter for datagrams using the RPL option.

11. References

11.1. Normative References

[I-D.ietf-roll-rpl]
Winter, T., Thubert, P., and R. Team, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-ietf-roll-rpl-08 (work in progress), May 2010.

[I-D.levis-roll-trickle]
Levis, P. and T. Clausen, "The Trickle Algorithm", draft-levis-roll-trickle-00 (work in progress), February 2010.

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

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

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.

11.2. Informative References

[I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-03 (work in progress), March 2010.

Authors' Addresses

Jonathan W. Hui
Arch Rock Corporation
501 2nd St. Ste. 410
San Francisco, California 94107
USA

Phone: +415 692 0828
Email: jhui@archrock.com

JP Vasseur
Cisco Systems, Inc
11, Rue Camille Desmoulins
Issy Les Moulineaux, 92782
France

Email: jpv@cisco.com