

Internet Engineering Task Force (IETF)
Request for Comments: 6554
Category: Standards Track
ISSN: 2070-1721

J. Hui
JP. Vasseur
Cisco Systems
D. Culler
UC Berkeley
V. Manral
Hewlett Packard Co.
March 2012

An IPv6 Routing Header for Source Routes with
the Routing Protocol for Low-Power and Lossy Networks (RPL)

Abstract

In Low-Power and Lossy Networks (LLNs), memory constraints on routers may limit them to maintaining, at most, a few routes. In some configurations, it is necessary to use these memory-constrained routers to deliver datagrams to nodes within the LLN. The Routing Protocol for Low-Power and Lossy Networks (RPL) can be used in some deployments to store most, if not all, routes on one (e.g., the Directed Acyclic Graph (DAG) root) or a few routers and forward the IPv6 datagram using a source routing technique to avoid large routing tables on memory-constrained routers. This document specifies a new IPv6 Routing header type for delivering datagrams within a RPL routing domain.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6554>.

Copyright Notice

Copyright (c) 2012 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. Introduction2
1.1. Requirements Language3
2. Overview3
3. Format of the RPL Routing Header6
4. RPL Router Behavior8
4.1. Generating Source Routing Headers8
4.2. Processing Source Routing Headers9
5. Security Considerations11
5.1. Source Routing Attacks11
5.2. ICMPv6 Attacks12
6. IANA Considerations12
7. Acknowledgements12
8. References12
8.1. Normative References12
8.2. Informative References13

1. Introduction

The Routing Protocol for Low-Power and Lossy Networks (RPL) is a distance vector IPv6 routing protocol designed for Low-Power and Lossy Networks (LLNs) [RFC6550]. Such networks are typically constrained in resources (limited communication data rate, processing power, energy capacity, memory). In particular, some LLN configurations may utilize LLN routers where memory constraints limit nodes to maintaining only a small number of default routes and no other destinations. However, it may be necessary to utilize such memory-constrained routers to forward datagrams and maintain reachability to destinations within the LLN.

To utilize paths that include memory-constrained routers, RPL relies on source routing. In one deployment model of RPL, more-capable routers collect routing information and form paths to arbitrary destinations within a RPL routing domain. However, a source routing mechanism supported by IPv6 is needed to deliver datagrams.

This document specifies the Source Routing Header (SRH) for use strictly between RPL routers in the same RPL routing domain. A RPL routing domain is a collection of RPL routers under the control of a single administration. The boundaries of routing domains are defined by network management by setting some links to be exterior, or inter-domain, links.

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].

2. Overview


The format of the SRH draws from that of the Type 0 Routing header (RH0) [RFC2460]. However, the SRH introduces mechanisms to compact the source route entries when all entries share the same prefix with the IPv6 Destination Address of a packet carrying an SRH, a typical scenario in LLNs using source routing. The compaction mechanism reduces consumption of scarce resources such as channel capacity.

The SRH also differs from RH0 in the processing rules to alleviate security concerns that led to the deprecation of RH0 [RFC5095]. First, RPL routers implement a strict source route policy where each and every IPv6 hop between the source and destination of the source route is specified within the SRH. Note that the source route may be a subset of the path between the actual source and destination and is discussed further below. Second, an SRH is only used between RPL routers within a RPL routing domain. RPL Border Routers, responsible for connecting other RPL routing domains and IP domains that use other routing protocols, do not allow datagrams already carrying an SRH header to enter or exit a RPL routing domain. Third, a RPL router drops datagrams that include multiple addresses assigned to any interfaces on that router to avoid forwarding loops.


There are two cases that determine how to include an SRH when a RPL router requires the use of an SRH to deliver a datagram to its destination.

1. If the SRH specifies the complete path from source to destination, the router places the SRH directly in the datagram itself.
2. If the SRH only specifies a subset of the path from source to destination, the router uses IPv6-in-IPv6 tunneling [RFC2473] and places the SRH in the outer IPv6 header. Use of tunneling ensures that the datagram is delivered unmodified and that ICMP errors return to the source of the SRH rather than the source of the original datagram.

In a RPL network, Case 1 occurs when both source and destination are within a RPL routing domain and a single SRH is used to specify the entire path from source to destination, as shown in the following figure:



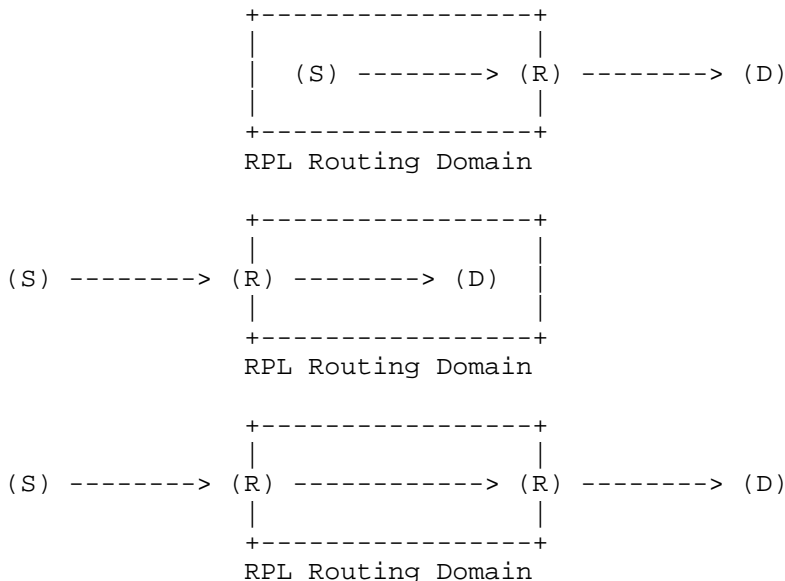
In the above scenario, datagrams traveling from source, S, to destination, D, have the following packet structure:




S's address is carried in the IPv6 header's Source Address field.

D's address is carried in the last entry of the SRH for all but the last hop, when D's address is carried in the IPv6 header's Destination Address field of the packet carrying the SRH.

In a RPL network, Case 2 occurs for all datagrams that have a source and/or destination outside the RPL routing domain, as shown in the following diagram:




In the scenarios above, R may indicate a RPL Border Router (when connecting to other routing domains) or a RPL Router (when connecting to hosts). The datagrams have the following structure when traveling within the RPL routing domain:



Note that the outer header (including the SRH) is added and removed by the RPL router.

Case 2 also occurs whenever a RPL router needs to insert a source route when forwarding a datagram. One such use case with RPL is to have all RPL traffic flow through a Border Router and have the Border Router use source routes to deliver datagrams to their final destination. When including the SRH using tunneled mode, the Border Router would encapsulate the received datagram unmodified using IPv6-in-IPv6 and include an SRH in the outer IPv6 header.



In the above scenario, datagrams travel from S to D through the Low-Power and Lossy Network Border Router (LBR). Between S and the LBR, the datagrams are routed using the DAG built by the RPL and do not contain an SRH. The LBR encapsulates received datagrams unmodified using IPv6-in-IPv6 and the SRH is included in the outer IPv6 header.

3. Format of the RPL Routing Header

The Source Routing Header has the following format:

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
+-----+-----+-----+-----+			
Next Header		Hdr Ext Len	
CmprI		CmprE	
Pad		Reserved	
+-----+-----+-----+-----+			
:			
: Addresses[1..n] :			
:			
+-----+-----+-----+-----+			

- Next Header: 8-bit selector. Identifies the type of header immediately following the Routing header. Uses the same values as the IPv6 Next Header field [RFC2460].
- Hdr Ext Len: 8-bit unsigned integer. Length of the Routing header in 8-octet units, not including the first 8 octets. Note that when Addresses[1..n] are compressed (i.e., value of CmprI or CmprE is not 0), Hdr Ext Len does not equal twice the number of Addresses.

Routing Type	8-bit selector. Identifies the particular Routing header variant. An SRH should set the Routing Type to 3.
Segments Left	8-bit unsigned integer. Number of route segments remaining, i.e., number of explicitly listed intermediate nodes still to be visited before reaching the final destination. The originator of an SRH sets this field to n, the number of addresses contained in Addresses[1..n].
CmprI	4-bit unsigned integer. Number of prefix octets from each segment, except than the last segment, (i.e., segments 1 through n-1) that are elided. For example, an SRH carrying full IPv6 addresses in Addresses[1..n-1] sets CmprI to 0.
CmprE	4-bit unsigned integer. Number of prefix octets from the last segment (i.e., segment n) that are elided. For example, an SRH carrying a full IPv6 address in Addresses[n] sets CmprE to 0.
Pad	4-bit unsigned integer. Number of octets that are used for padding after Address[n] at the end of the SRH.
Reserved	This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Address[1..n]	Vector of addresses, numbered 1 to n. Each vector element in [1..n-1] has size (16 - CmprI) and element [n] has size (16-CmprE). The originator of an SRH places the next (first) hop's IPv6 address in the IPv6 header's IPv6 Destination Address and the second hop's IPv6 address as the first address in Address[1..n] (i.e., Address[1]).

The SRH shares the same basic format as the Type 0 Routing header [RFC2460]. When carrying full IPv6 addresses, the CmprI, CmprE, and Pad fields are set to 0 and the only difference between the SRH and Type 0 encodings is the value of the Routing Type field.

A common network configuration for a RPL routing domain is that all routers within a RPL routing domain share a common prefix. The SRH introduces the CmprI, CmprE, and Pad fields to allow compaction of the Address[1..n] vector when all entries share the same prefix as

the IPv6 Destination Address field of the packet carrying the SRH. The CmprI and CmprE fields indicate the number of prefix octets that are shared with the IPv6 Destination Address of the packet carrying the SRH. The shared prefix octets are not carried within the Routing header and each entry in Address[1..n-1] has size (16 - CmprI) octets and Address[n] has size (16 - CmprE) octets. When CmprI or CmprE is non-zero, there may exist unused octets between the last entry, Address[n], and the end of the Routing header. The Pad field indicates the number of unused octets that are used for padding. Note that when CmprI and CmprE are both 0, Pad MUST carry a value of 0.

The SRH MUST NOT specify a path that visits a node more than once. When generating an SRH, the source may not know the mapping between IPv6 addresses and nodes. Minimally, the source MUST ensure that IPv6 addresses do not appear more than once and the IPv6 Source and Destination addresses of the encapsulating datagram do not appear in the SRH.

Multicast addresses MUST NOT appear in an SRH or in the IPv6 Destination Address field of a datagram carrying an SRH.

4. RPL Router Behavior

4.1. Generating Source Routing Headers

To deliver an IPv6 datagram to its destination, a router may need to generate a new SRH and specify a strict source route. When the router is the source of the original packet and the destination is known to be within the same RPL routing domain, the router SHOULD include the SRH directly within the original packet. Otherwise, the router MUST use IPv6-in-IPv6 tunneling [RFC2473] and place the SRH in the tunnel header. Using IPv6-in-IPv6 tunneling ensures that the delivered datagram remains unmodified and that ICMPv6 errors generated by an SRH are sent back to the router that generated the SRH.

When using IPv6-in-IPv6 tunneling, in order to respect the IPv6 Hop Limit value of the original datagram, a RPL router generating an SRH MUST set the Segments Left to less than the original datagram's IPv6 Hop Limit value upon forwarding. In the case that the source route is longer than the original datagram's IPv6 Hop Limit, only the initial hops (determined by the original datagram's IPv6 Hop Limit) should be included in the SRH. If the RPL router is not the source of the original datagram, the original datagram's IPv6 Hop Limit field is decremented before generating the SRH. After generating the SRH, the RPL router decrements the original datagram's IPv6 Hop Limit value by the SRH Segments Left value. Processing the SRH Segments

Left and original datagram's IPv6 Hop Limit fields in this way ensures that ICMPv6 Time Exceeded errors occur as would be expected on more traditional IPv6 networks that forward datagrams without tunneling.

To avoid fragmentation, it is desirable to employ MTU sizes that allow for the header expansion (i.e., at least 1280 + 40 (outer IP header) + SRH_MAX_SIZE), where SRH_MAX_SIZE is the maximum path length for a given RPL network. To take advantage of this, however, the communicating endpoints need to be aware of the MTU along the path (i.e., through Path MTU Discovery). Unfortunately, the larger MTU size may not be available on all links (e.g., 1280 octets on IPv6 Low-Power Wireless Personal Area Network (6LoWPAN) links). However, it is expected that much of the traffic on these types of networks consists of much smaller messages than the MTU, so performance degradation through fragmentation would be limited.

4.2. Processing Source Routing Headers

As specified in [RFC2460], a routing header is not examined or processed until it reaches the node identified in the Destination Address field of the IPv6 header. In that node, dispatching on the Next Header field of the immediately preceding header causes the Routing header module to be invoked.

The function of the SRH is intended to be very similar to the Type 0 Routing header defined in [RFC2460]. After the routing header has been processed and the IPv6 datagram resubmitted to the IPv6 module for processing, the IPv6 Destination Address contains the next hop's address. When forwarding an IPv6 datagram that contains an SRH with a non-zero Segments Left value, if the IPv6 Destination Address is not on-link, a router MUST drop the datagram and SHOULD send an ICMP Destination Unreachable (ICMPv6 Type 1) message with ICMPv6 Code set to 7 to the packet's Source Address. This ICMPv6 Code indicates that the IPv6 Destination Address is not on-link and the router cannot satisfy the strict source route requirement. When generating ICMPv6 error messages, the rules in Section 2.4 of [RFC4443] MUST be observed.

To detect loops in the SRH, a router MUST determine if the SRH includes multiple addresses assigned to any interface on that router. If such addresses appear more than once and are separated by at least one address not assigned to that router, the router MUST drop the packet and SHOULD send an ICMP Parameter Problem, Code 0, to the Source Address. While this loop check does add significant per-packet processing overhead, it is required to mitigate bandwidth exhaustion attacks that led to the deprecation of RH0 [RFC5095].

The following describes the algorithm performed when processing an SRH:

```
if Segments Left = 0 {
  proceed to process the next header in the packet, whose type is
  identified by the Next Header field in the Routing header
}
else {
  compute n, the number of addresses in the Routing header, by
  n = (((Hdr Ext Len * 8) - Pad - (16 - CmprE)) / (16 - CmprI)) + 1

  if Segments Left is greater than n {
    send an ICMP Parameter Problem, Code 0, message to the Source
    Address, pointing to the Segments Left field, and discard the
    packet
  }
  else {
    decrement Segments Left by 1

    compute i, the index of the next address to be visited in
    the address vector, by subtracting Segments Left from n

    if Address[i] or the IPv6 Destination Address is multicast {
      discard the packet
    }
    else if 2 or more entries in Address[1..n] are assigned to
      local interface and are separated by at least one
      address not assigned to local interface {
      send an ICMP Parameter Problem (Code 0) and discard the
      packet
    }
    else {
      swap the IPv6 Destination Address and Address[i]

      if the IPv6 Hop Limit is less than or equal to 1 {
        send an ICMP Time Exceeded -- Hop Limit Exceeded in
        Transit message to the Source Address and discard the
        packet
      }
      else {
        decrement the Hop Limit by 1

        resubmit the packet to the IPv6 module for transmission
        to the new destination
      }
    }
  }
}
```

RPL routers are responsible for ensuring that an SRH is only used between RPL routers:

1. For datagrams destined to a RPL router, the router processes the packet in the usual way. For instance, if the SRH was included using tunneled mode and the RPL router serves as the tunnel endpoint, the router removes the outer IPv6 header, at the same time removing the SRH as well.
2. Datagrams destined elsewhere within the same RPL routing domain are forwarded to the correct interface.
3. Datagrams destined to nodes outside the RPL routing domain are dropped if the outermost IPv6 header contains an SRH not generated by the RPL router forwarding the datagram.

5. Security Considerations

5.1. Source Routing Attacks

The RPL message security mechanisms defined in [RFC6550] do not apply to the RPL Source Route Header. This specification does not provide any confidentiality, integrity, or authenticity mechanisms to protect the SRH.

[RFC5095] deprecates the Type 0 Routing header due to a number of significant attacks that are referenced in that document. Such attacks include bypassing filtering devices, reaching otherwise unreachable Internet systems, network topology discovery, bandwidth exhaustion, and defeating anycast.

Because this document specifies that the SRH is only for use within a RPL routing domain, such attacks cannot be mounted from outside a RPL routing domain. As specified in this document, RPL routers MUST drop datagrams entering or exiting a RPL routing domain that contain an SRH in the IPv6 Extension headers.

Such attacks, however, can be mounted from within a RPL routing domain. To mitigate bandwidth exhaustion attacks, this specification requires RPL routers to check for loops in the SRH and drop datagrams that contain such loops. Attacks that include bypassing filtering devices and reaching otherwise unreachable Internet systems are not as relevant in mesh networks since the topologies are, by their very nature, highly dynamic. The RPL routing protocol is designed to provide reachability to all devices within a RPL routing domain and may utilize routes that traverse any number of devices in any order.

Even so, these attacks and others (e.g., defeating anycast and routing topology discovery) can occur within a RPL routing domain when using this specification.

5.2. ICMPv6 Attacks

The generation of ICMPv6 error messages may be used to attempt denial-of-service attacks by sending an error-causing SRH in back-to-back datagrams. An implementation that correctly follows Section 2.4 of [RFC4443] would be protected by the ICMPv6 rate-limiting mechanism.

6. IANA Considerations

This document defines a new IPv6 Routing Type, the "RPL Source Route Header", and has been assigned number 3 by IANA.

This document defines a new ICMPv6 Destination Unreachable Code, "Error in Source Routing Header", and has been assigned number 7 by IANA.

7. Acknowledgements

The authors thank Jari Arkko, Ralph Droms, Adrian Farrel, Stephen Farrell, Richard Kelsey, Suresh Krishnan, Erik Nordmark, Pascal Thubert, Sean Turner, and Tim Winter for their comments and suggestions that helped shape this document.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", RFC 5095, December 2007.

8.2. Informative References

[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, March 2012.

Authors' Addresses

Jonathan W. Hui
Cisco Systems
170 West Tasman Drive
San Jose, California 95134
USA

Phone: +408 424 1547
EMail: jonhui@cisco.com

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

EMail: jpv@cisco.com

David E. Culler
UC Berkeley
465 Soda Hall
Berkeley, California 94720
USA

Phone: +510 643 7572
EMail: culler@cs.berkeley.edu

Vishwas Manral
Hewlett Packard Co.
19111 Pruneridge Ave.
Cupertino, California 95014
USA

EMail: vishwas.manral@hp.com