

Internet Engineering Task Force (IETF)
Request for Comments: 7559
Updates: 4861
Category: Standards Track
ISSN: 2070-1721

S. Krishnan
Ericsson
D. Anipko
Unaffiliated
D. Thaler
Microsoft
May 2015

Packet-Loss Resiliency for Router Solicitations

Abstract

When an interface on a host is initialized, the host transmits Router Solicitations in order to minimize the amount of time it needs to wait until the next unsolicited multicast Router Advertisement is received. In certain scenarios, these Router Solicitations transmitted by the host might be lost. This document specifies a mechanism for hosts to cope with the loss of the initial Router Solicitations.

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/rfc7559>.

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. Conventions Used in This Document	3
2. Proposed Algorithm	3
2.1. Stopping the Retransmissions	3
3. Configuring the Use of Retransmissions	4
4. Known Limitations	4
5. Security Considerations	4
6. References	5
6.1. Normative References	5
6.2. Informative References	5
Acknowledgements	5
Authors' Addresses	6

1. Introduction

As specified in [RFC4861], when an interface on a host is initialized, in order to obtain Router Advertisements quickly, a host transmits up to MAX_RTR_SOLICITATIONS (3) Router Solicitation (RS) messages, each separated by at least RTR_SOLICITATION_INTERVAL (4) seconds. In certain scenarios, these Router Solicitations transmitted by the host might be lost. For example, the host is connected to a bridged residential gateway over Ethernet or Wi-Fi. LAN connectivity is achieved at interface initialization, but the upstream WAN connectivity is not active yet. In this case, the host just gives up after the initial RS retransmits.

Once the initial RSs are lost, the host gives up and assumes that there are no routers on the link as specified in Section 6.3.7 of [RFC4861]. The host will not have any form of Internet connectivity until the next unsolicited multicast Router Advertisement is received. These Router Advertisements are transmitted at most

MaxRtrAdvInterval seconds apart (maximum value 1800 seconds). Thus, in the worst-case scenario a host would be without any connectivity for 30 minutes. This delay may be unacceptable in some scenarios.

1.1. Conventions Used in This Document

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. Proposed Algorithm

To achieve resiliency to packet loss, the host needs to continue retransmitting the Router Solicitations until it receives a Router Advertisement, or until it is willing to accept that no router exists. If the host continues retransmitting the RSs at RTR_SOLICITATION_INTERVAL second intervals, it may cause excessive network traffic if a large number of such hosts exists. To achieve resiliency while keeping the aggregate network traffic low, the host can use some form of exponential backoff algorithm to retransmit the RSs.

Hosts complying to this specification MUST use the exponential backoff algorithm for retransmits that is described in Section 14 of [RFC3315] in order to continuously retransmit the Router Solicitations until a Router Advertisement is received. The hosts SHOULD use the following variables as input to the retransmission algorithm:

```
IRT (Initial Retransmission Time):    4 seconds
MRT (Maximum Retransmission Time):    3600 seconds
MRC (Maximum Retransmission Count):   0
MRD (Maximum Retransmission Duration): 0
```

The initial value IRT was chosen to be in line with the current retransmission interval (RTR_SOLICITATION_INTERVAL) that is specified by [RFC4861], and the maximum retransmission time MRT was chosen to be in line with the new value of SOL_MAX_RT as specified by [RFC7083]. This is to ensure that the short-term behavior of the RSs is similar to what is experienced in current networks, and that longer-term persistent retransmission behavior trends towards being similar to that of DHCPv6 [RFC3315] [RFC7083].

2.1. Stopping the Retransmissions

On multicast-capable links, the hosts following this specification SHOULD stop retransmitting the RSs when Router Discovery is successful (i.e., an RA with a non-zero Router Lifetime that results

in a default route is received). If an RA is received from a router and it does not result in a default route (i.e., Router Lifetime is zero), the host MUST continue retransmitting the RSs.

On non-multicast links, the hosts following this specification MUST continue retransmitting the RSs even after an RA that results in a default route is received. This is required because, in such links, sending an RA can only be triggered by an RS. Please note that such links have special mechanisms for sending RSs as well. For example, the mechanism specified in Section 8.3.4 of the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [RFC5214] unicasts the RSs to specific routers.

3. Configuring the Use of Retransmissions

Implementations of this specification are encouraged to provide a configuration option to enable or disable potentially infinite RS retransmissions. If a configuration option is provided, it MUST enable RS retransmissions by default. Providing an option to enable/disable retransmissions on a per-interface basis allows network operators to configure RS behavior in the most applicable way for each connected link.

4. Known Limitations

When an IPv6-capable host attaches to a network that does not have IPv6 enabled, it transmits 3 (MAX_RTR_SOLICITATIONS) Router Solicitations as specified in [RFC4861]. If it receives no Router Advertisements, it assumes that there are no routers present on the link and it ceases to send further RSs. With the mechanism specified in this document, the host will continue to retransmit RSs indefinitely at the rate of approximately 1 RS per hour. It is unclear how to differentiate between such a network with no IPv6 routers and a link where an IPv6 router is temporarily unreachable but could become reachable in the future.

5. Security Considerations

This document does not present any additional security issues beyond those discussed in [RFC4861] and those RFCs that update [RFC4861].

6. References

6.1. 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>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC7083] Droms, R., "Modification to Default Values of SOL_MAX_RT and INF_MAX_RT", RFC 7083, DOI 10.17487/RFC7083, November 2013, <<http://www.rfc-editor.org/info/rfc7083>>.

6.2. Informative References

- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, DOI 10.17487/RFC5214, March 2008, <<http://www.rfc-editor.org/info/rfc5214>>.

Acknowledgements

The authors would like to thank Steve Baillargeon, Erik Kline, Andrew Yourtchenko, Ole Troan, Erik Nordmark, Lorenzo Colitti, Thomas Narten, Ran Atkinson, Allison Mankin, Les Ginsberg, Brian Carpenter, Barry Leiba, Brian Haberman, Spencer Dawkins, Alia Atlas, Stephen Farrell, and Mehmet Ersue for their reviews and suggestions that made this document better.

Authors' Addresses

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Phone: +1 514 345 7900 x42871
EMail: suresh.krishnan@ericsson.com

Dmitry Anipko
Unaffiliated

Phone: +1 425 442 6356
EMail: dmitry.anipko@gmail.com

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA
United States

EMail: dthaler@microsoft.com