

6man Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 7, 2013

S. Krishnan
Ericsson
D. Anipko
D. Thaler
Microsoft
May 6, 2013

Packet loss resiliency for Router Solicitations
draft-ietf-6man-resilient-rs-01

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. Furthermore, on some links, unsolicited multicast Router Advertisements are never sent and the mechanism in this document is intended to work even in such scenarios.

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 7, 2013.

Copyright Notice

Copyright (c) 2013 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.	Conventions used in this document	3
2.	Proposed algorithm	4
2.1.	Stopping the retransmissions	4
3.	Configuring the use of retransmissions	5
4.	Known Limitations	5
5.	IANA Considerations	5
6.	Security Considerations	5
7.	Acknowledgements	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	6
	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 messages, each separated by at least RTR_SOLICITATION_INTERVAL (4) seconds. In certain scenarios, these router solicitations transmitted by the host might be lost.

The generic scenario is that the interface on the host comes up before it gets access to a router. Examples include:

- a. The host is connected to a bridged residential gateway over Ethernet or WiFi. 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.
- b. Access networks/links that turn off periodic RAs and only send RAs in response to RSs. In this case, if the link between the AP and the host comes up before the link between the AP and the Controller/Router, the host will never be able to connect.
- c. Links that are not multicast capable. In this case, sending an RA can only be triggered by an RS (as is the case, for instance, on ISATAP [\[RFC5214\]](#) links).

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. In general, the 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  4 seconds
MRT  3600 seconds
MRC  0
MRD  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 [\[SOLMAXRT\]](#). This is to ensure that the short term behavior of the RSs is similar to what is experienced in current networks, and longer term persistent retransmission behavior trends towards being similar to that of DHCPv6 [\[RFC3315\]](#) [\[SOLMAXRT\]](#).

2.1. Stopping the retransmissions

On multicast-capable links, the hosts following this specification MUST stop retransmitting the RSs when an RA 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 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.

3. Configuring the use of retransmissions

Implementations of this specification MAY provide a configuration option to enable or disable the use of such potentially infinite retransmissions. If the implementation provides such a configuration option, it MUST be able to enable/disable retransmissions on a per-interface basis.

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. IANA Considerations

This document does not require any IANA actions.

6. Security Considerations

This document does not present any additional security issues beyond those discussed in [[RFC4861](#)].

7. Acknowledgements

The author would like to thank Steve Baillargeon, and Erik Kline for their reviews and suggestions that made this document better.

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.
- [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.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.

[SOLMAXRT] Droms, R., "Modification to Default Value of SOL_MAX_RT", [draft-droms-dhc-dhcpv6-solmaxrt-update-02](#) (work in progress), January 2012.

8.2. Informative References

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", [RFC 5214](#), March 2008.

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
Microsoft
One Microsoft Way
Redmond, WA
USA

Phone: +1 425 703 7070
Email: danipko@microsoft.com

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA
USA

Email: dthaler@microsoft.com