

6man WG
Internet-Draft
Updates: [4861](#) (if approved)
Intended status: Standards Track
Expires: April 30, 2015

E. Nordmark
Arista Networks
A. Yourtchenko
Cisco
S. Krishnan
Ericsson
October 27, 2014

IPv6 Neighbor Discovery Optional Unicast RS/RA Refresh
draft-nordmark-6man-rs-refresh-01

Abstract

IPv6 Neighbor Discovery relies on periodic multicast Router Advertisement messages to update timer values and to distribute new information (such as new prefixes) to hosts. On some links the use of periodic multicast messages to all host becomes expensive, and in some cases it results in hosts waking up frequently. Many implementations of [RFC 4861](#) also use multicast for solicited Router Advertisement messages, even though that behavior is optional.

This specification provides an optional mechanism for hosts and routers where instead of periodic multicast Router Advertisements the hosts are instructed (by the routers) to use unicast Router Solicitations to request refreshed Router Advertisements. This mechanism is enabled by configuring the router to include a new option in the Router Advertisement in order to allow the network administrator to choose host behavior based on whether periodic multicast are more efficient on their link or not. The routers can also tell whether the hosts are capable of the new behavior through a new flag in the Router Solicitations.

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."

Internet-Draft

Optional Unicast RS/RA Refresh

October 2014

This Internet-Draft will expire on April 30, 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.

Internet-Draft

Optional Unicast RS/RA Refresh

October 2014

Table of Contents

1.	Introduction	4
2.	Goals and Requirements	5
3.	Definition Of Terms	5
4.	Protocol Overview	5
5.	New Neighbor Discovery Flags and Options	6
5.1.	Introducing a Router Solicitation Flag	6
5.2.	Refresh Time option	6
6.	Conceptual Data Structures	7
7.	Host Behavior	7
7.1.	Sleep and Wakeup	8
7.2.	Movement	8
8.	Router Behavior	8
8.1.	Router and/or Interface Initialization	9
8.2.	Periodic Multicast RA for unmodified hosts	9
8.3.	Unsolicited RAs to share new information	9
9.	Router Advertisement Consistency	10
10.	Security Considerations	10
11.	IANA Considerations	10
12.	Acknowledgements	10
13.	Open Issues	10
14.	References	11
14.1.	Normative References	11
14.2.	Informative References	11
	Authors' Addresses	12

1. Introduction

IPv6 Neighbor Discovery [[RFC4861](#)] was defined at a time when local area networks had different properties than today. A common link was the yellow-coax shared wire Ethernet, where a link-layer multicast and unicast worked the same - send the packet on the wire and the interested receivers will pick it up. Thus the network cost (ignoring any processing cost on the receivers that might not completely filter out Ethernet multicast addresses that they did not want) and the reliability of sending a link-layer unicast and multicast was the same. Furthermore, the hosts at the time was always on and connected. Powering on and off the workstation/PC hosts at the time was slow and disruptive process.

Under the above assumptions it was quite efficient to maintain the shared state of the link such as the prefixes and their lifetimes using periodic multicast Router Advertisement messages. It was also efficient to use multicast Neighbor Solicitations for address resolution as a slight improvement over the broadcast use in ARP. And finally, checking for a potential duplicate IPv6 address using multicast was efficient and natural.

There are still links, such a satellite links, where periodic multicast advertisements is the most efficient and reliable approach to keep the hosts up to date. However other links have different performance and reliability for multicast than for unicast (see for instance [[I-D.vyncke-6man-mcast-not-efficient](#)] which discusses WiFi links). Cellular networks which employ paging and support sleeping hosts have different issues (see e.g.,

[[I-D.garneij-6man-nd-m2m-issues](#)] that would benefit from having the hosts wake up and request information from the routers instead of the routers periodically multicasting the information.

Since different links types and deployments have different needs, this specification provides mechanism by which the routers can determine whether all the hosts support the RS refresh, and the hosts only employ the RS refresh when instructed by the routers using an option in the Router Advertisement.

The operator retains the option to use unsolicited multicast Router Advertisement to announce new or removed information. That can be useful for uncommon cases while allowing using a higher refresh time for normal network operations.

The specification does not assume that all hosts on the link implement the new capability. As soon as there are router(s) on a link which supports these optimizations, then the updated hosts on the link can sleep better, while co-existing on the same link with

unmodified hosts.

[2.](#) Goals and Requirements

The key goal is to allow the operator to choose whether unicast RS refresh is more efficient than periodic multicast RAs, while preserving the timely and scalable reconfiguration capabilities that a periodic RA model provides.

In addition, an operator might want to be notified whether the link includes hosts that do not support the new mechanism. Potential router implementations can react dynamically to that information, or can log events to system management when hosts appear which do not implement this new capability.

The assumption is that host which implement this specification also implement [[I-D.ietf-6man-resilient-rs](#)] as that ensures resiliency to packet loss at host initialization.

[3.](#) Definition Of Terms

The keywords "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](#)].

[4.](#) Protocol Overview

The hosts include a new flag in the Router Solicitation message, which allows the routers to report to system management whether there are hosts that do not support the RS refresh on the link.

If the network administrator has configured the routers to send the new Refresh Timer option, then the option will be included in all the Router Advertisements. This option includes the time interval when the hosts should unicast Router Solicitations.

The host maintains the value of the Refresh Timer option (RTO) by recording it in the default router list. A value of zero can be used to indicate that a router did not include a Refresh Timer option.

The host calculates a timeout after it has received a RTO - either per router or per link. If it is maintained per link then the host SHOULD use the minimum Refresh Timer it has received from the routers on the link. The timeout is a random value uniformly distributed between 0.5 and 1.5 times the Refresh Timer value (in order to avoid


synchronization of the timers across hosts. [TBD: Add SYNC reference from [RFC 4861](#).] When this timer fires the host sends one unicast Router Solicitation to the router (if maintained per router) or to all the routers in the default router list for link (if maintained per link.)

[5.](#) New Neighbor Discovery Flags and Options

This specification introduces a option used in the RAs which both indicates that the router can handle RS refresh using unicast RA, and a flag for the RS that indicates to the router that the host will do RS refresh if the router so wishes.

[5.1.](#) Introducing a Router Solicitation Flag

A node which implements this specification sets the R flag in all the Router Solicitation messages it sends. That allows the router to determine whether there are legacy hosts on the link.




New fields:

R-flag:	When set indicates that the sending node is capable of doing unicast RS refresh.
---------	--

Reserved: Field is reduced from 32 bits to 31 bits. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

5.2. Refresh Time option

A router which implements this specification can be configured to operate without periodic multicast Router Advertisements. When the operator configures this mode of operation, then the router **MUST** include this new option in the RA.



Fields:

Type: TBD ND option code value (IANA)

Length: 8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 bytes. The value 0 is invalid. Value is 1 for this option.

Refresh Time: 16-bit unsigned integer. Units is seconds. The all-ones value (65535) means infinite.

Reserved: 32 bits. This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

6. Conceptual Data Structures

In addition to the Conceptual Data structures in [[RFC4861](#)] a host records the received RT0 value in the default router list. It also maintains a timeout - either per link or per default router.

7. Host Behavior

See Protocol Overview section above.

A host implementing this specification SHOULD also implement [[I-D.ietf-6man-resilient-rs](#)]. That ensures that a router that has been configured to not send periodic RA messages will always receive an RS from the host as the host initializes.

If there is no RT0 in the received Router Advertisements, then the host behavior does not change. However, if RT0s start appearing in RAs after the initial RAs, the host SHOULD start performing RS refresh. As the last router that included RT0 options time out from the default router list, the host SHOULD stop sending RS refresh messages.

The host MUST join the all-nodes multicast address as in [[RFC4861](#)]

since the routers MAY send multicast RAs for important changes.

Some links might have routers with different configuration where some router includes RTO in the RA and others do not. Hosts MAY make the simplifying assumption that if any router on the link includes RTO then the host can use RS refresh to all the routers in the default router list. Also, the routers might advertise different refresh time, and hosts MAY use the minimum of the time received from any router that remains in the default router list. Note that [Section 9](#) says that routers SHOULD report such inconsistencies to system management.

[7.1.](#) Sleep and Wakeup

The protocol allows the sleepy nodes to complete its sleep schedule without waking up due to multicast Router Advertisement messages and the host is not required to wake up solely for the purposes of performing RS refresh. This assumes that sleepy nodes perform a RS refresh when they wake up. If hosts do wake up due to multicast RAs, then the host only needs to perform a refresh on wakeup if the Refresh timeout has expired while the host was sleeping.

[7.2.](#) Movement

When a host wakes up it can combine movement detecting (DNA), NUD, and refreshing its prefixes etc by sending a unicast RS to each of its existing default router(s). If it receives unicast RA from a router, then it can mark the router as REACHABLE.

Note that DNA [[RFC6059](#)] specifies using NS messages since many IPv6 routers delay (and multicast) solicited RAs and DNA wants to avoid that delay. Routers which implement this specification SHOULD unicast solicited RAs, hence if a router included the RTO then the host can use RS for DNA. For non-RTO routers the host MAY choose to use NS for DNA as in [[RFC6059](#)].

[8.](#) Router Behavior

See Protocol Overview section.

A router implementing this specification (and including RTO in the RAs) SHOULD also respond to unicast RS messages (that do not have an unspecified source address) with unicast RAs. If a RS message has an unspecified source address then the host MAY respond with a RA unicast at layer 2 (sent to the link-layer source address of the RS), or it MAY follow the rate-limited multicast RA procedure in

[[RFC4861](#)].

The RECOMMENDED default configuration for routers is to have RTO disabled.

[8.1.](#) Router and/or Interface Initialization

This specification does not change the initialization procedure. Thus a router multicasts some initial Router Advertisements (MAX_INITIAL_RTR_ADVERTISEMENTS) at system startup or interface initialization as specified in [[RFC4861](#)] and its updates.

[8.2.](#) Periodic Multicast RA for unmodified hosts

By default a router MUST send periodic multicast RAs as specified in [[RFC4861](#)]. A router can be configured to omit those, which can be used in particular deployments. If they are omitted, then there MUST be a mechanism to prevent or detect the existence of unmodified hosts on the link. That be be performed at deployment time (e.g., only hosts which are known to support RTO are configured with the layer 2 security keys), or the routers detect any RSs which do not include the R-flag and report this to system management, or dynamically enable periodic multicast RAs when observing at least one RS without the R-flag.

Note that such dynamic detection is not bullet proof. If a host does not implement RS refresh nor implements resilient RS [[I-D.ietf-6man-resilient-rs](#)], then the host might receive a multicast RA (from router initialization or the periodic multicast RAs) without the router ever receiving a RS from the host. Such a host would function as long as the routers are sending periodic multicast RAs.

[8.3.](#) Unsolicited RAs to share new information

When a router has new information to share (new prefixes, prefixes that should be immediately deprecated, etc) it MAY multicast up to MAX_INITIAL_RTR_ADVERTISEMENTS number of Router Advertisements.

On links where multicast is expensive the router MAY instead unicast up to MAX_INITIAL_RTR_ADVERTISEMENTS number of Router Advertisements to the hosts in its neighbor cache.

. Note that such new information is not likely to reach sleeping hosts until those hosts refresh by sending a RS.

[9.](#) Router Advertisement Consistency

The routers follows [section 6.2.7 in \[RFC4861\]](#) by receiving RAs from other routers on the link. In addition to the checks in that section, the routers SHOULD verify that the RT0 have the same Refresh Time, and report to system management if they differ. While the host will pick the lowest time and operate correctly, it is not useful to use different Refresh Times for different routers.

[10.](#) Security Considerations

These optimizations are not known to introduce any new threats against Neighbor Discovery beyond what is already documented for IPv6 [\[RFC3756\]](#).

[Section 11.2 of \[RFC4861\]](#) applies to this document as well.

The mechanisms in this document work with SeND [\[RFC3971\]](#).

[11.](#) IANA Considerations

A new flag (R-flag) in the Router Solicitation message has been introduced by carving out a bit from the Reserved field. There is currently no IANA registry for RS flags. Perhaps one should be created?

This document needs a new Neighbor Discovery option type for the RT0.

[12.](#) Acknowledgements

The original idea came up in a discussion with Suresh Krishnan. Comments from Erik Kline and Samita Chakrabarti have helped improve the document.

This document has been discussed in the efficient-nd design team.

13. Open Issues

Should we make the Refresh Time 32 bits instead of 16? 16 bits implies maximum of 18 hours and in some deployments a refresh time measured in days might be desirable.

Should we update the DNA procedures [[RFC6059](#)]? We can use a unicast RS with this approach since that will result in an

Nordmark, et al.

Expires April 30, 2015

[Page 10]

Internet-Draft

Optional Unicast RS/RA Refresh

October 2014

immediate unicast RA which would include any updated prefixes. Note that a RS can not have an unspecified source and a SLLA0, hence some care would be needed in the interaction with DAD.

Would it be worth-while to try to remove unchanged information from the refreshed RAs? If so it could be done by including some epoch number in the RS and RA, and if the RS contains the current epoch then the RA would not need to include any options except the epoch number indicating that none of the options are the same as before.

14. References

14.1. Normative References

[I-D.ietf-6man-resilient-rs]

Krishnan, S., Anipko, D., and D. Thaler, "Packet loss resiliency for Router Solicitations", [draft-ietf-6man-resilient-rs-04](#) (work in progress), October 2014.

[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.

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

14.2. Informative References

[I-D.garneij-6man-nd-m2m-issues]

Garneij, F., Chakrabarti, S., and S. Krishnan, "Impact of IPv6 Neighbor Discovery on Cellular M2M Networks", [draft-garneij-6man-nd-m2m-issues-00](#) (work in progress), July 2014.

[I-D.vyncke-6man-mcast-not-efficient]

Vyncke, E., Thubert, P., Levy-Abegnoli, E., and A. Yourtchenko, "Why Network-Layer Multicast is Not Always Efficient At Datalink Layer", [draft-vyncke-6man-mcast-not-efficient-01](#) (work in progress), February 2014.

[RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor

Nordmark, et al.

Expires April 30, 2015

[Page 11]

Internet-Draft

Optional Unicast RS/RA Refresh

October 2014

Discovery (ND) Trust Models and Threats", [RFC 3756](#), May 2004.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.

[RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", [RFC 6059](#), November 2010.

Authors' Addresses

Erik Nordmark
Arista Networks
Santa Clara, CA
USA

Email: nordmark@acm.org

Andrew Yourtchenko
Cisco
7a de Kleetlaan
Diegem, 1831

Belgium

Phone: +32 2 704 5494

Email: ayourtch@cisco.com

Suresh Krishnan

Ericsson

8400 Decarie Blvd.

Town of Mount Royal, QC

Canada

Phone: +1 514 345 7900 x42871

Email: suresh.krishnan@ericsson.com