IPv6 Operations Working Group (v6ops)
Internet-Draft
Intended status: Informational
Expires: December 10, 2011

F. Gont
UK CPNI

June 8, 2011

IPv6 Router Advertisement Guard (RA-Guard) Evasion
draft-gont-v6ops-ra-guard-evasion-01

Abstract

The IPv6 Router Advertisement Guard (RA-Guard) mechanism is commonly employed to mitigate attack vectors based on forged ICMPv6 Router Advertisement messages. Many existing IPv6 deployments rely on RA-Guard as the first line of defense against the aforementioned attack vectors. This document describes possible ways in which current RA-Guard implementations can be circumvented, and discusses possible mitigations.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

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 10, 2011.

Table of Contents

1. Introduction

IPv6 Router Advertisement Guard (RA-Guard) is a mitigation technique for attack vectors based on ICMPv6 Router Advertisement messages. [RFC6104] describes the problem statement of "Rogue IPv6 Router Advertisements", and [RFC6105] specifies the "IPv6 Router Advertisement Guard" functionality.

The basic concept behind RA-Guard is that a layer-2 device filters ICMPv6 Router Advertisement messages, according to a number of different criteria. The most basic filtering criterion is that Router Advertisement messages are discarded by the layer-2 device unless they are received on a specified port of the layer-2 device. Clearly, the effectiveness of the RA Guard mitigation relies on the ability of the layer-2 device to identify ICMPv6 Router Advertisement messages.

As part of the project "Security Assessment of the Internet Protocol version 6 (IPv6)" [CPNI-IPv6], we have devised two techniques for circumventing the RA-Guard protection, which are described in the following sections of this document. These techniques, and the corresponding tools to assess their effectiveness, had so far been made available only to vendors, in the hopes that they could implement counter-measures before they were publicly disclosed. However, since there has been some public discussion about these issues, it was deemed as appropriate to publish the present document.

It should be noted that the aforementioned techniques could also be exploited to evade network monitoring tools such as NDPMon [NDPMon], ramond [ramond], and rafixd [rafixd], and could probably be exploited to perform stealth DHCPv6 attacks.

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. Router Advertisement Guard (RA Guard) Evasion Vulnerability

The following subsections describe two different vectors for evading the RA-Guard protection. Section 2.1 describes an attack vector based on the use of IPv6 Extension Headers with the ICMPv6 Router Advertisement messages, which may be used to circumvent the RA-Guard protection of those implementations that fail to process an entire IPv6 header chain when trying to identify the ICMPv6 Router Advertisement messages. Section 2.2 describes an attack method based on the use of IPv6 fragmentation, possibly in conjunction with the use of IPv6 Extension Headers. This later vector is expected to be effective with all existing implementations of the RA-Guard functionality.

2.1. Attack Vector based on IPv6 Extension Headers

While there is currently no legitimate use for IPv6 Extension Headers in ICMPv6 Router Advertisement messages, Neighbor Discovery implementations allow the use of Extension Headers with these messages, by simply ignoring the received options. We believe that some implementations may simply try to identify ICMPv6 Router Advertisement messages by looking at the "Next Header" field of the fixed IPv6 header, rather than following the entire header chain. As a result, these implementations would fail to identify any ICMPv6 Router Advertisement messages that include any Extension Headers (for example, Hop by Hop Options header, Destination Options Header, etc.).

The following figure illustrates the structure of ICMPv6 Router Advertisement messages that implement this RA-Guard evasion technique:

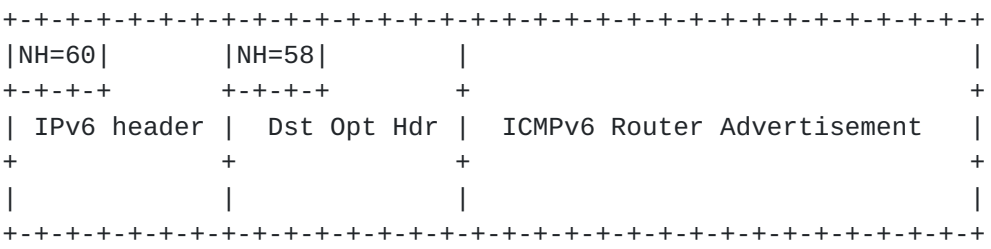

2.2. Attack vector based on IPv6 fragmentation

While the attack vector described in Section 2.1 may be effective with implementations that fail to process the entire header chain, it can easily be mitigated by an RA-Guard implementation, since all the information needed to identify ICMPv6 Router Advertisement messages is present in the attack packets.

This section presents a different attack vector, which aims at making it virtually impossible for a layer-2 device to identify ICMPv6 Router Advertisements by leveraging the IPv6 Fragment Header. The basic idea behind this attack vector is that if the forged ICMPv6 Router Advertisement is fragmented into at least two fragments, the layer-2 device implementing "RA-Guard" would be unable to identify the attack packet, and would thus fail to block it.

A first variant of this attack vector would be an original ICMPv6 Router Advertisement message preceded with a Destination Options Header, that results in two fragments. The following figure illustrates the "original" attack packet, prior to fragmentation, and the two resulting fragments which are actually sent as part of the attack.

Original packet:

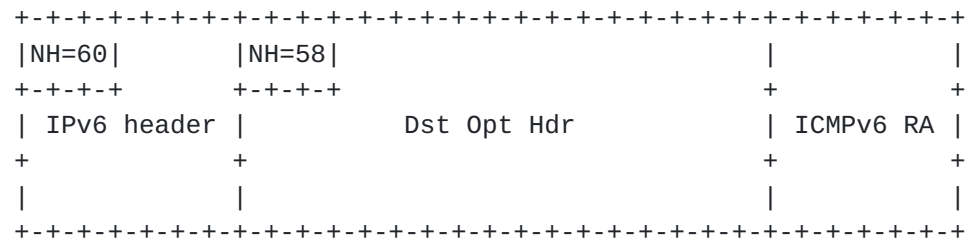

First fragment:

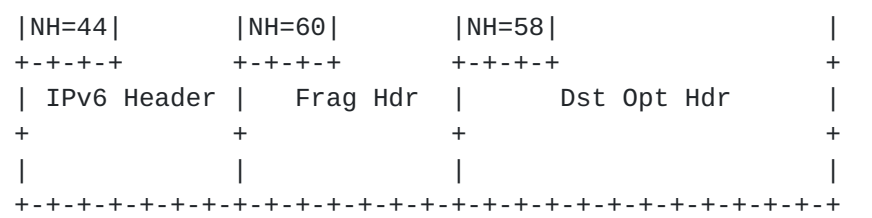

Second fragment:

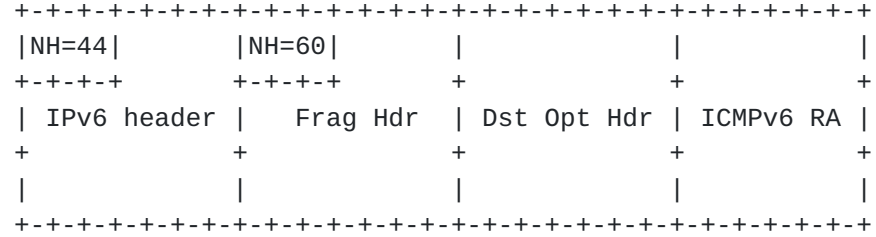

It should be noted that the "Hdr Ext Len" field of the Destination

Options Header is present in the first fragment (rather than the second). Therefore, it would be impossible for a device processing only the second fragment to locate the ICMPv6 header contained in that fragment, since it is unknown how many bytes should be "skipped" to get to the next header following the Destination Options Header.

Thus, by leveraging the use of the Fragment Header together with the use of the Destination Options header, the attacker is able to conceal the type and contents of the ICMPv6 message he is sending (an ICMPv6 Router Advertisement in this example). Unless the layer-2 device were to implement IPv6 fragment reassembly, it would be impossible for the device to identify the ICMPv6 type of the message.

> A layer-2 device could, however, at least detect that that an ICMPv6 message (or some type) is being sent, since the "Next Header" field of the Destination Options header contained in the first fragment is set to "58" (ICMPv6).

It is possible to take this idea further, such that it is also impossible for the layer-2 device to detect that the attacker is sending an ICMPv6 message in the first place. This can be achieved with an original ICMPv6 Router Advertisement message preceded with two Destination Options Headers, that results in two fragments. The following figure illustrates the "original" attack packet, prior to fragmentation, and the two resulting packets which are actually sent as part of the attack.

Original packet:

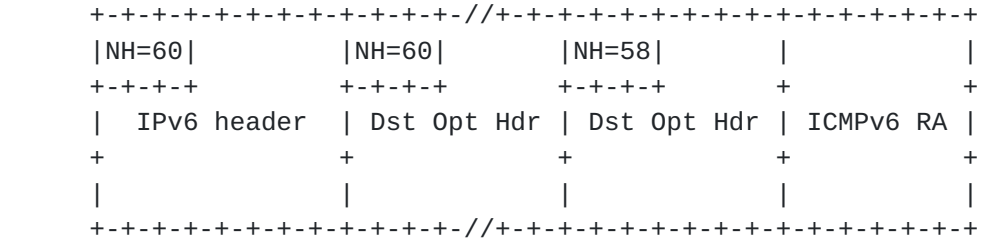

First fragment:

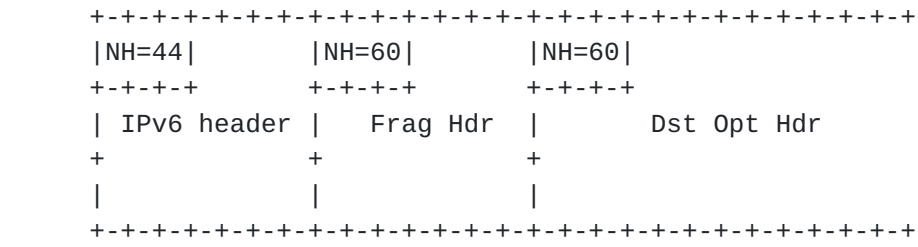

Second fragment:

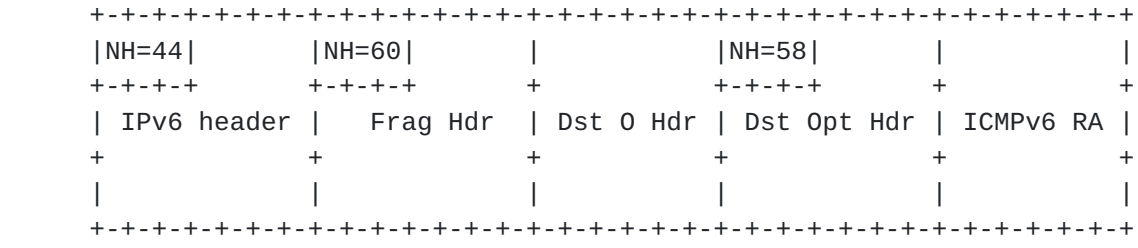

In this variant, the "Next Header" field of the Destination Options header contained in the first fragment is set "60" (Destination Options header), and thus it is impossible for a device processing only the first fragment to detect that an ICMPv6 message is being sent in the first place.

The second fragment presents the same challenges as the second fragment of the previous variant. That is, it would be impossible for a device processing only the second fragment to locate the second Destination Options header (and hence the ICMPv6 header), since the "Hdr Ext Len" field of the first Destination Options header is present in the first fragment (rather than the second).

3. Mitigations

The most effective and efficient mitigation for the RA-Guard evasion vulnerability discussed in this document would be to prohibit the use of IPv6 Extension Headers in Neighbor Discovery messages, as proposed in [I-D.gont-6man-nd-extension-headers].

Nevertheless, an administrator might want to mitigate these vulnerabilities by deploying more advanced filtering. The following filtering rules could be implemented as part of an "RA-Guard" implementation, such that the vulnerabilities discussed in this document can be mitigated:

- When trying to identify an ICMPv6 Router Advertisement message, follow the IPv6 header chain, enforcing a limit on the maximum number of Extension Headers that is allowed for each packet. If such limit is exceeded, block the packet.

- If the layer-2 device is unable to identify whether the packet is an ICMPv6 Router Advertisement message or not (i.e., the packet is a fragment, and the necessary information is missing), and the IPv6 Source Address of the packet is a link-local address or the unspecified address (::), block the packet.

- In all other cases, pass the packet as usual.

This filtering policy assumes that host implementations require that the IPv6 Source Address of ICMPv6 Router Advertisement messages be a link-local address, and that they discard the packet if this check fails, as required by the current IETF specifications [RFC4861]. Unfortunately, it should be noted that the aforementioned filtering policy might be inefficient to implement (if at all possible), and might also result (at least in theory) in false positives.

4. Other Implications

A similar concept to that of "RA-Guard" has been implemented for protecting against forged DHCPv6 messages. Such protection can be circumvented with the same techniques discussed in this document, and the counter-measures for such evasion attack are analogous to those described in Section 3 of this document.

5. Security Considerations

This document describes a number of techniques to circumvent a mechanism known as "RA-Guard", which many organizations deploy as a "first line of defense" against attacks based on forged Router Advertisements.

The most effective and efficient mitigation for these attacks would be to prohibit the use of IPv6 extension headers (as proposed by [I-D.gont-6man-nd-extension-headers]), such that the RA-Guard protection cannot be easily circumvented. However, since this mitigation requires an update to existing implementations, in the short term some network administrators might want to mitigate these issues by implemented the more advanced filtering policy described in Section 3.

6. Acknowledgements

The author would like to thank Karl Auer, Robert Downie, David Farmer, Marc Heuse, and Arturo Servin, for providing valuable comments on earlier versions of this document.

This document resulted from the project "Security Assessment of the Internet Protocol version 6 (IPv6)" [CPNI-IPv6], carried out by Fernando Gont on behalf of the UK Centre for the Protection of National Infrastructure (CPNI). The author would like to thank the UK CPNI, for their continued support.

7. References

7.1. Normative References

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

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

7.2. Informative References

[RFC6104] Chown, T. and S. Venaas, "Rogue IPv6 Router Advertisement Problem Statement", RFC 6104, February 2011.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, February 2011.

[I-D.gont-6man-nd-extension-headers] Gont, F. and U. CPNI, "Security Implications of the Use of IPv6 Extension Headers with IPv6 Neighbor Discovery", draft-gont-6man-nd-extension-headers-00 (work in progress), May 2011.

[CPNI-IPv6] Gont, F., "Security Assessment of the Internet Protocol version 6 (IPv6)", UK Centre for the Protection of National Infrastructure, (to be published).

[NDPMon] "NDPMon - IPv6 Neighbor Discovery Protocol Monitor", <http://ndpmon.sourceforge.net/>.

[rafixd] "rafixd", <http://www.kame.net/dev/cvsweb2.cgi/kame/kame/kame/rafixd/>.

[ramond] "ramond", <http://ramond.sourceforge.net/>.

[THC-IPV6] "THC-IPV6", <http://www.thc.org/thc-ipv6/>.

Appendix A. Changes from previous versions of the draft (to be removed by the RFC Editor before publication of this document as a RFC

A.1. Changes from draft-gont-v6ops-ra-guard-evasion-00

- Minor editorial changes
- The discussion of the challenge represented by a combination of fragmentation and Destination Options headers was improved/ clarified.
- In Section 2.2, in the illustration of the second variant of the attack (fragmentation combined with two Destination Optios headers), the figure corresponding to the "first fragment" was corrected.
- Clarified the filtering rules in Section 3.

Appendix B. Assessment tools

CPNI has produced assessment tools, which have not yet been made publicly available. If you think that you would benefit from these tools to assess the security of your network or of your RA-Guard implementation, we might be able to provide a copy of the tools (please contact Fernando Gont at fernando@gont.com.ar).

[THC-IPV6] is a publicly-available set of tools that implements some of the techniques described in this document.

Appendix C. Advice and guidance to vendors

Vendors are urged to contact CSIRTUK (csirt@cpni.gsi.gov.uk) if they think they may be affected by the issues described in this document. As the lead coordination centre for these issues, CPNI is well placed to give advice and guidance as required.

CPNI works extensively with government departments and agencies, commercial organisations and the academic community to research vulnerabilities and potential threats to IT systems especially where they may have an impact on Critical National Infrastructure's (CNI).

Other ways to contact CPNI, plus CPNI's PGP public key, are available at http://www.cpni.gov.uk.

Author's Address

Fernando Gont
Centre for the Protection of National Infrastructure

Email: fernando@gont.com.ar
URI: http://www.gont.com.ar