IPv6 Operations Working Group (v6ops)
Internet-Draft
Intended status: BCP
Expires: September 9, 2012

F. Gont
UK CPNI

March 8, 2012

Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)
draft-ietf-v6ops-ra-guard-implementation-02

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. However, some implementations of RA-Guard have been found to be prone to circumvention by employing IPv6 Extension Headers. This document describes the evasion techniques that affect the aforementioned implementations, and provides advice on the implementation of RA-Guard, such that the RA-Guard evasion vectors are eliminated.

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 September 9, 2012.

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.

Some popular RA-Guard implementations have been found to be easy to circumvent by employing IPv6 extension headers [CPNI-IPv6]. This document describes such evasion techniques, and provides advice to RA-Guard implementers such that the aforementioned evasion vectors can be eliminated.

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. Evasion techniques for some Router Advertisement Guard (RA Guard) implementations

The following subsections describe two different vectors that have been found to be effective for the evasion of popular implementations of 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 has been found to be effective with all existing implementations of the RA-Guard mechanism.

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. Some RA-Guard implementations try to identify ICMPv6 Router Advertisement messages by simply looking at the "Next Header" field of the fixed IPv6 header, rather than following the entire header chain. As a result, such implementations fail to identify any ICMPv6 Router Advertisement messages that include any Extension Headers (for example, a Hop by Hop Options header, a Destination Options Header, etc.), and can be easily circumvented.

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

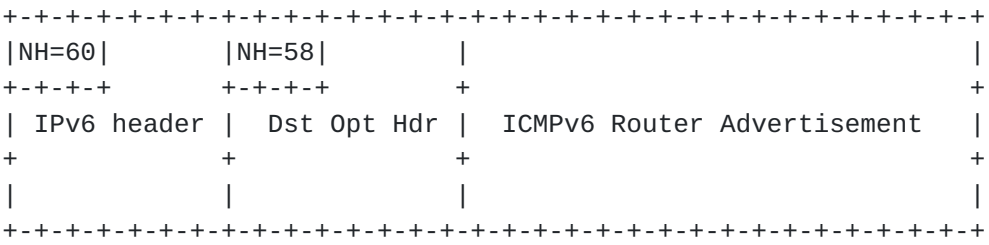

2.2. Attack vector based on IPv6 fragmentation

This section presents a different attack vector, which has been found to be effective against all implementations of RA-Guard. 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:

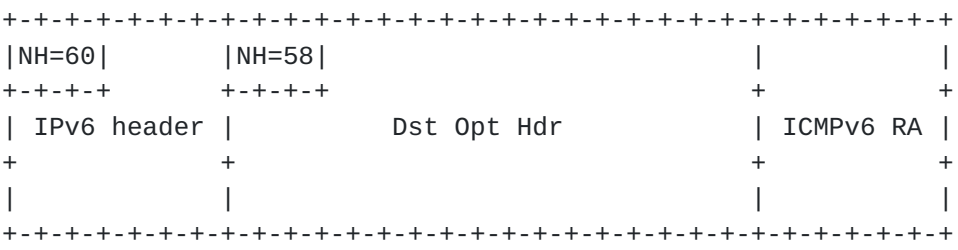

First fragment:

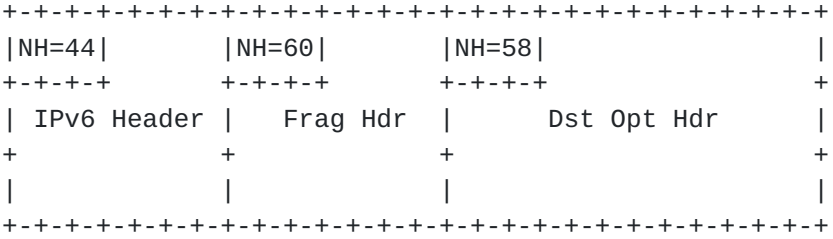

Second fragment:

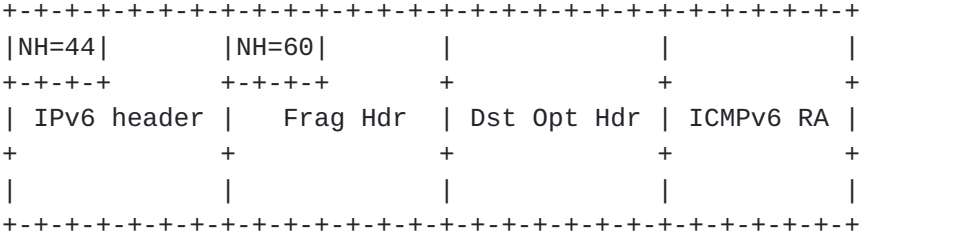

It should be noted that the "Hdr Ext Len" field of the Destination Options Header is present in the first fragment (rather than the second). Therefore, it is impossible for a device processing only the second fragment to locate the ICMPv6 header contained in that fragment, since it is unknown how many 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).

This idea can be taken 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:

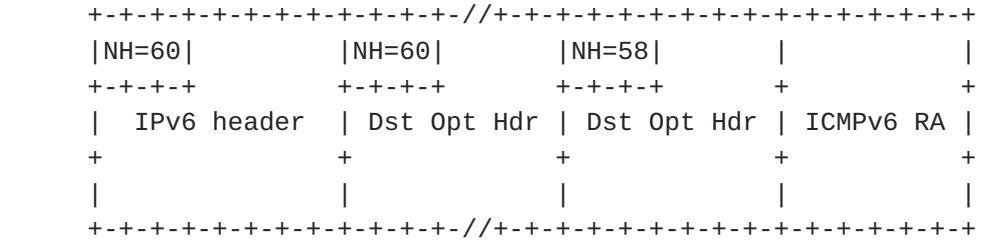

First fragment:

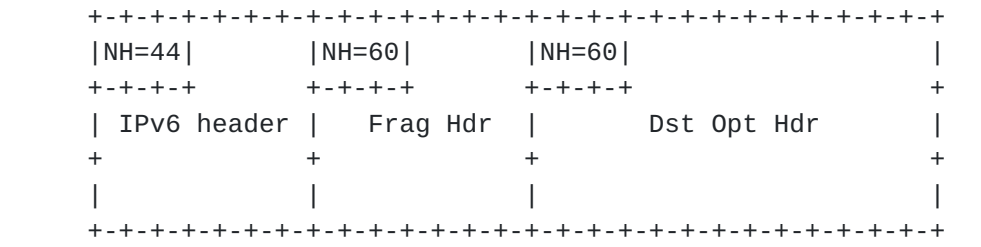

Second fragment:

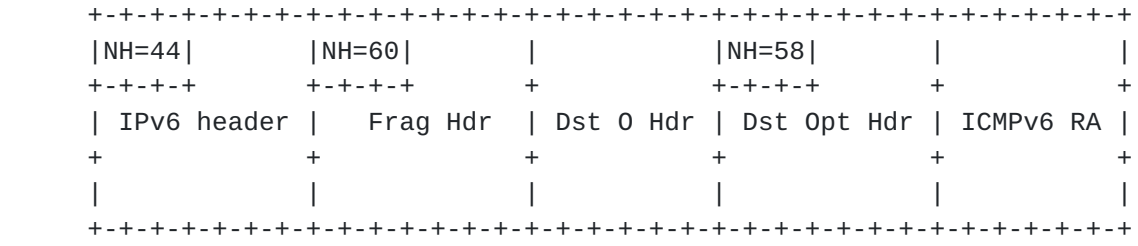

In this variant, the "Next Header" field of the Destination Options header contained in the first fragment is set "60" (Destination Options header), and thus it is impossible for a device processing only the first fragment to detect that an ICMPv6 message is being sent in the first place.

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. RA-Guard implementation advice

The following filtering rules MUST be implemented as part of an "RA-Guard" implementation on those ports that are not allowed to send ICMPv6 Router Advertisement messages, such that the vulnerabilities discussed in this document are eliminated:

1. 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 hit before the upper-layer protocol is identified, silently drop the packet.

2. If the packet is identified to be an ICMPv6 Router Advertisement message, silently drop the packet.

3. 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 first-fragment, and the necessary information is missing), the IPv6 Source Address of the packet is a link-local address or the unspecified address (::), and the Hop Limit is 255, silently drop the packet.

   Note: This rule should only be applied to non-fragmented IPv6 datagrams and IPv6 fragments with a Fragment Offset of 0 (non-first fragments can be safely passed, since they will never reassemble into a complete datagram if they are part of a Router Advertisement received on a port where such packets are not allowed).

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

Note: For the purpose of enforcing the RA-Guard filtering policy, an ESP header [RFC4303] should be considered to be an "upper-layer protocol" (that is, it should be considered the last header in the IPv6 header chain). This means that packets employing ESP would be passed by the RA-Guard device to the intended destination. If the destination host does not have a security association with the sender of the aforementioned IPv6 packet, the packet would be dropped. Otherwise, if the packet is considered valid by the IPsec implementation at the receiving host and encapsulates a Router Advertisement message, it is up to the receiving host what to do with such packet.

In order to protect current end-node IPv6 implementations, Rule #3 has been defined as a default rule to drop packets that cannot be positively identified as RA packets or not (perhaps due to the fact that it contains fragments that do not contain the entire IPv6 header

chain). This means that, at least in theory, RA-Guard could result in false-positive blocking of some legitimate non-RA packets that could not be positively identified as being non-RA. In order to reduce the likelihood of false positives, Rule #3 also requires that an RA-Guard implementation check, before dropping an unidentifiable packet, that it has an IPv6 Source Address that is a link-local address or the unspecified address (::), and that the Hop Limit is 255. In any case, as noted in [I-D.gont-6man-oversized-header-chain], IPv6 packets that fail to include the entire IPv6 header chain are anyway unlikely to survive in real networks. Whilst currently legitimate from a specifications standpoint, they are virtually impossible to police with state-less filters and firewalls, and are hence likely to be blocked by such filters and firewalls.

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]. Additionally, it assumes that hosts require the Hop Limit of Neighbor Discovery messages to be 255, and discard those packets otherwise.

Finally, note that the aforementioned filtering rules implicitly handle the case of fragmented packets: if the RA-Guard device fails to identify the upper-layer protocol as a result of the use of fragmentation, the corresponding packets would be silently dropped.

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 that have been found to be effective to circumvent popular RA-Guard implementations, and provides advice to RA-Guard implementations such that those evasion vulnerabilities are eliminated.

We note that if an attacker sends a fragmented Router Advertisement message on a port not allowed to send such packets, the first-fragment would be dropped, and the rest of the fragments would be passed. This means that the victim node would tie memory buffers for the aforementioned fragments, which would never reassemble into a complete datagram. If a large number of such packets were sent by an attacker, and the victim node failed to implement proper resource management for the fragment reassembly buffer, this could lead to a Denial of Service (DoS). However, this does not really introduce a new attack vector, since an attacker could always perform the same attack by sending forged fragmented datagram in which at least one of the fragments is missing. [CPNI-IPv6] discusses some resource management strategies that could be implemented for the fragment reassembly buffer.

Finally, we note that most effective and efficient mitigation for these attacks would be to prohibit the use of IPv6 fragmentation with Router Advertisement messages (as proposed by [I-D.gont-6man-nd-extension-headers]), such that the RA-Guard functionality is easier to implement. However, since such mitigation would require an update to existing implementations, it cannot be relied upon in the short or near term.

6. Acknowledgements

The author would like to thank Ran Atkinson, Karl Auer, Robert Downie, Washam Fan, David Farmer, Marc Heuse, Ray Hunter, Simon Perreault, Arturo Servin, and Gunter van de Velde, for providing valuable comments on earlier versions of this document.

The author would like to thank Arturo Servin, who presented this document at IETF 81.

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.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.

[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-oversized-header-chain] Gont, F. and V. Manral, "Security and Interoperability Implications of Oversized IPv6 Header Chains", draft-gont-6man-oversized-header-chain-00 (work in progress), February 2012.

[I-D.gont-6man-nd-extension-headers] Gont, F., "Security Implications of the Use of IPv6 Extension Headers with IPv6 Neighbor Discovery", draft-gont-6man-nd-extension-headers-02 (work in progress), January 2012.

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

[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-ietf-v6ops-ra-guard-implementation-00

- The filtering rules in Section 3 have been further clarified.

A.2. Changes from draft-gont-v6ops-ra-guard-implementation-01

- Document resubmitted as draft-ietf to reflect wg adoption.

A.3. Changes from draft-gont-v6ops-ra-guard-implementation-00

- Miscellaneous (minor) editorial changes.
- The filtering rules in Section 3 have been polished.

A.4. Changes from draft-gont-v6ops-ra-guard-evasion-01

- The contents were updated to reflect that the evasion vulnerabilities are based on implementation flaws, rather than on the RA-Guard "concept" itself.
- The I-D now focuses on providing advice to RA-Guard implementers.

Appendix B. Assessment tools

CPNI has produced assessment tools (which have not yet been made publicly available) to assess RA-Guard implementations with respect to the issues described in this document. If you think that you would benefit from these tools, 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: fgont@si6networks.com
URI: http://www.cpni.gov.uk