

Internet Engineering Task Force (IETF)
Request for Comments: 7113
Updates: 6105
Category: Informational
ISSN: 2070-1721

F. Gont
Huawei Technologies
February 2014

Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)

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 formally updates RFC 6105, such that the aforementioned RA-Guard evasion vectors are eliminated.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc7113>.

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.

Table of Contents

1. Introduction	2
2. Evasion Techniques for Some RA-Guard Implementations	3
2.1. Attack Vector Based on IPv6 Extension Headers	3
2.2. Attack Vector Based on IPv6 Fragmentation	4
3. RA-Guard Implementation Advice	6
4. Other Implications	9
5. Security Considerations	9
6. Acknowledgements	10
7. References	10
7.1. Normative References	10
7.2. Informative References	11
Appendix A. Assessment Tools	12

1. Introduction

IPv6 Router Advertisement Guard (RA-Guard) is a mitigation technique for attack vectors based on ICMPv6 Router Advertisement [RFC4861] messages. [RFC6104] describes the problem statement of "Rogue IPv6 Router Advertisements", and [RFC6105] specifies the "IPv6 Router Advertisement Guard" functionality.

The concept behind RA-Guard is that a Layer-2 (L2) 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 L2 device unless they are received on a specified port of the L2 device. Clearly, the effectiveness of RA-Guard relies on the ability of the L2 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 previously mentioned 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 [RFC3315] attacks.


2. Evasion Techniques for Some RA-Guard Implementations

The following subsections describe two different vectors that have been found to be effective for the evasion of popular implementations of RA-Guard. Section 2.1 describes an attack vector based on the use of IPv6 Extension Headers with 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 against all existing implementations of RA-Guard.

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 [RFC4861] 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 L2 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, which results in two fragments. The following figure illustrates the "original" attack packet, prior to fragmentation, and the two resulting fragments that 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 L2 device were to implement IPv6 fragment reassembly, it would be impossible for the device to identify the ICMPv6 type of the message.


An L2 device could, however, at least detect that an ICMPv6 message (of 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 L2 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 that 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 to "60" (Destination Options header); 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 ports that face interfaces that are not allowed to send ICMPv6 Router Advertisement messages, such that the vulnerabilities discussed in this document are eliminated:

1. If the IPv6 Source Address of the packet is not a link-local address (fe80::/10), RA-Guard must pass the packet.

RATIONALE: This prevents RA-Guard from dedicating processing cycles to filtering packets that originate off-net and that, if they are RA's, would not be accepted by the host. Section 6.1.2 of [RFC4861] requires nodes to discard Router Advertisement messages if their IPv6 Source Address is not a link-local address.

2. If the Hop Limit is not 255, RA-Guard must pass the packet.

RATIONALE: This prevents RA-Guard from dedicating processing cycles to filtering packets that originate off-net and that, if they are RA's, would not be accepted by the destination host. Section 6.1.2 of [RFC4861] requires nodes to discard Router Advertisement messages if their Hop Limit is not 255.

3. RA-Guard must parse the entire IPv6 header chain present in the packet, to identify whether the packet is a Router Advertisement message.

NOTE: RA-Guard implementations must not enforce a limit on the number of bytes they can inspect (starting from the beginning of the IPv6 packet), since this could introduce false positives: legitimate packets could be dropped simply because the RA-Guard device does not parse the entire IPv6 header chain present in the packet. An implementation that has such an implementation-specific limit must not claim compliance with this specification, and must pass the packet when such implementation-specific limit is reached.

4. When parsing the IPv6 header chain, if the packet is a First Fragment (i.e., a packet containing a Fragment Header with the Fragment Offset set to 0) and it fails to contain the entire IPv6 header chain (i.e., all the headers starting from the IPv6 header up to, and including, the upper-layer header), RA-Guard must drop the packet and should log the packet drop event in an implementation-specific manner as a security fault.

RATIONALE: [RFC7112] specifies that the First Fragment (i.e., the fragment with the Fragment Offset set to 0) must contain the entire IPv6 header chain, and allows intermediate systems such as routers to drop those packets that fail to comply with this requirement.

NOTE: This rule should only be applied to 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).

5. When parsing the IPv6 header chain, if the packet is identified to be an ICMPv6 Router Advertisement message or the packet contains an unrecognized Next Header value [IANA-IP-PROTO], RA-Guard must drop the packet, and should log the packet drop event in an implementation-specific manner as a security fault. RA-Guard must provide a configuration knob that controls whether packets with unrecognized Next Header values are dropped; this configuration knob must default to "drop".

RATIONALE: By definition, Router Advertisement messages are required to originate on-link, have a link-local IPv6 Source Address, and have a Hop Limit value of 255 [RFC4861]. [RFC7045] requires that nodes be configurable with respect to

whether packets with unrecognized headers are forwarded, and allows the default behavior to be that such packets be dropped.

6. In all other cases, RA-Guard must pass the packet as usual.

NOTE: For the purpose of enforcing the RA-Guard filtering policy, an Encapsulating Security Payload (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 a packet.

If a packet is dropped due to this filtering policy, then the packet drop event should be logged in an implementation-specific manner as a security fault. The logging mechanism should include a drop counter dedicated to RA-Guard packet drops.

In order to protect current end-node IPv6 implementations, Rule #4 has been defined as a default rule to drop packets that cannot be positively identified as not being Router Advertisement (RA) messages (because the packet is a fragment that fails to include 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 #1 and Rule #2 require that packets that would not pass the required validation checks for RA messages (Section 6.1.2 of [RFC4861]) be passed without further inspection. In any case, as noted in [RFC7112], IPv6 packets that fail to include the entire IPv6 header chain are virtually impossible to police with state-less filters and firewalls and, hence, are unlikely to survive in real networks. [RFC7112] requires that hosts employing fragmentation include the entire IPv6 header chain in the First Fragment (the fragment with the Fragment Offset set to 0), thus eliminating the aforementioned false positives.

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 that they discard those packets otherwise.

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 dropped.

Finally, we note that IPv6 implementations that allow overlapping fragments (i.e., that do not comply with [RFC5722]) might still be subject of RA-based attacks. However, a recent assessment of IPv6 implementations [SI6-FRAG] with respect to their fragment reassembly policy seems to indicate that most current implementations comply with [RFC5722].

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 countermeasures for such evasion attack are analogous to those described in Section 3 of this document.

[DHCPv6-Shield] specifies a mechanism to protect against rogue DHCPv6 servers, while taking into consideration the evasion techniques discussed in 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 implementers such that those evasion vulnerabilities are eliminated.

As noted in Section 3, IPv6 implementations that allow overlapping fragments (i.e., that do not comply with [RFC5722]) might still be subject of RA-based attacks. However, most current implementations seem to comply with [RFC5722].

We note that if an attacker sends a fragmented ICMPv6 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 IPv6 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 datagrams in which at

least one of the fragments is missing. [CPNI-IPv6] discusses some resource management strategies that could be implemented for the IPv6 fragment reassembly buffer.

We note that the most effective and efficient mitigation for these attacks would rely on the prohibiting the use of IPv6 fragmentation with Router Advertisement messages (as specified by [RFC6980]), 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.

Finally, we note that RA-Guard only mitigates attack vectors based on ICMPv6 Router advertisement messages. Protection against similar attacks based on other messages (such as DHCPv6) is considered out of the scope of this document and is left for other documents (e.g., [DHCPv6-Shield]).

6. Acknowledgements

The author would like to thank Ran Atkinson, who provided very detailed comments and suggested text that was incorporated into this document.

The author would like to thank Ran Atkinson, Karl Auer, Robert Downie, Washam Fan, David Farmer, Mike Heard, Marc Heuse, Nick Hilliard, Ray Hunter, Joel Jaeggli, Simon Perreault, Arturo Servin, Gunter van de Velde, James Woodyatt, and Bjoern A. Zeeb, 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).

7. References

7.1. Normative References

- [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.
- [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.
- [RFC5722] Krishnan, S., "Handling of Overlapping IPv6 Fragments", RFC 5722, December 2009.
- [RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, February 2011.
- [RFC6980] Gont, F., "Security Implications of IPv6 Fragmentation with IPv6 Neighbor Discovery", RFC 6980, August 2013.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, December 2013.
- [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", RFC 7112, January 2014.

7.2. Informative References

- [CPNI-IPv6] Gont, F., "Security Assessment of the Internet Protocol version 6 (IPv6)", UK Centre for the Protection of National Infrastructure, (available on request).
- [DHCPv6-Shield] Gont, F., Liu, W., and G. Van de Velde, "DHCPv6-Shield: Protecting Against Rogue DHCPv6 Servers", Work in Progress, October 2013.
- [IANA-IP-PROTO] IANA, "Assigned Internet Protocol Numbers", <<http://www.iana.org/assignments/protocol-numbers/>>.
- [NDPMon] "NDPMon - IPv6 Neighbor Discovery Protocol Monitor", <<http://ndpmon.sourceforge.net/>>.
- [RFC6104] Chown, T. and S. Venaas, "Rogue IPv6 Router Advertisement Problem Statement", RFC 6104, February 2011.
- [SI6-FRAG] SI6 Networks, "IPv6 NIDS evasion and improvements in IPv6 fragmentation/reassembly", 2012, <<http://blog.si6networks.com/2012/02/ipv6-nids-evasion-and-improvements-in.html>>.

- [SI6-IPv6] "SI6 Networks' IPv6 toolkit",
<<http://www.si6networks.com/tools/ipv6toolkit>>.
- [THC-IPV6] "The Hacker's Choice IPv6 Attack Toolkit",
<<http://www.thc.org/thc-ipv6/>>.
- [rafixd] "rafixd", <<http://www.kame.net/dev/cvsweb2.cgi/kame/kame/kame/rafixd/>>.
- [ramond] "ramond", <<http://ramond.sourceforge.net/>>.

Appendix A. Assessment Tools

[SI6-IPv6] is a publicly available set of tools (for Linux, *BSD, and Mac OS) that implements the techniques described in this document.

[THC-IPv6] is a publicly available set of tools (for Linux) that implements some of the techniques described in this document.

Author's Address

Fernando Gont
Huawei Technologies
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
EMail: fgont@si6networks.com

