IPv6 maintenance Working Group (6man)
Internet-Draft
Updates: 4861 (if approved)
Intended status: Standards Track
Expires: August 18, 2014

F. Gont
SI6 Networks / UTN-FRH
R. Bonica
Juniper Networks
W. Liu
Huawei Technologies
February 14, 2014

Validation of Neighbor Discovery Source Link-Layer Address (SLLA) and Target Link-layer Address (TLLA) options draft-gont-6man-lla-opt-validation-00

Abstract

This memo documents two scenarios in which an on-link attacker emits a crafted IPv6 Neighbor Discovery (ND) packet that poisons its victim's neighbor cache. In the first scenario, the attacker causes a victim to map a local IPv6 address to a local router's own link-layer address. In the second scenario, the attacker causes the victim to map a unicast IP address to a link layer broadcast address. In both scenarios, the attacker can exploit the poisoned neighbor cache to perform a subsequent forwording-loop attack, thus potentially causing a Denial of Service.

Finally, this memo specifies simple validations that the recipient of an ND message can execute in order to protect itself against the above-mentioned threats.

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 August 18, 2014.

Table of Contents

1. Introduction

IPv6 [RFC2460] nodes use a Neighbor Discovery (ND) [RFC4861] mechanism to discover on-link neighbors and learn their link layer addresses. Having discovered an on-link neighbor and learned its link layer address, an IPv6 node stores that information in a local data structure, called the "neighbor cache".

ND defines the following ICMPv6 [RFC4443] messages:

- Router Solicitation (RS)

- Router Advertisement (RA)

- Neighbor Solicitation (NS)
- Neighbor Advertisement (NA)
- Redirect

ND also defines a Source Link-Layer Address (SLLA) option and a Target Link-Layer Address (TLLA) option. The RS, RA, and NS messages all typically contain the SLLA option, that contains the link layer address of the node sending the message. The NA and Redirect messages contain the TLLA option, that maps a target IPv6 address that is contained by the NA or Redirect message to a link layer address.

This memo documents two scenarios in which an on-link attacker emits a crafted ND packet that poisons its victim's neighbor cache. In the first scenario, the attacker causes a victim to map an IPv6 address to a the victim router's own link-layer address. In the second scenario, the attacker causes the victim to map a unicast IP address to the link layer broadcast or multicast address. In both scenarios, the attacker can subsequently exploit the poisoned neighbor cache to perform a forwording-loop attack, thus potentially causing a Denial of Service (DoS).

Finally, this memo specifies simple validations that the recipient of an ND message can execute in order to protect itself against the above-mentioned threats.

2. Terminology

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].

3. ND-based Forwarding-Loop Attacks

3.1. Mapping an IPv6 Address to a Local Router's Own Link-layer Address

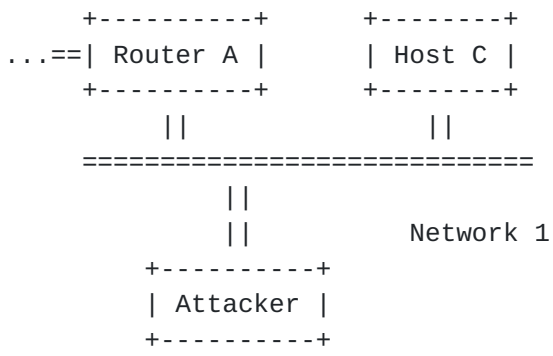

Figure 1: Unicast Forwarding Loop

In Figure 1, the Attacker sends Router A a crafted ND message. The aforementioned ND message contains the Target Addres set to Host C's IPv6 address, and a TLLA option set to Router A's link-layer address. The ND message causes Router A to map Host C's IPv6 address to the link layer address of Router A's interface to Network 1. This sets up the scenario for a subsequent attack.

A packet is sent to Router A with the IPv6 Destination Address set to that of Host C. Router A forwards the packet on Network 1, specifying its own Network 1 interface as the link layer destination. Because Router A specified itself as the link layer destination, Router A receives the packet and forwards it again. This process repeats until the IPv6 Hop Limit is decremented to 0 (and hence the packet is discarded). In this scenario, the amplification factor is equal to the Hop Limit minus one.

An attacker can realize this attack by sending either of the following:

- An ND message whose SLLA maps an IPv6 address to the link layer address of the victim router's (Router A's in our case) interface to the local network (Network 1 in our case)

- An ND message whose TLLA maps an IPv6 address to the link layer address of the victim router's (Router A's in our case) interface to the local network (Network 1 in our case)

3.2. Mapping a Unicast IPv6 Address to A Broadcast Link-Layer Address

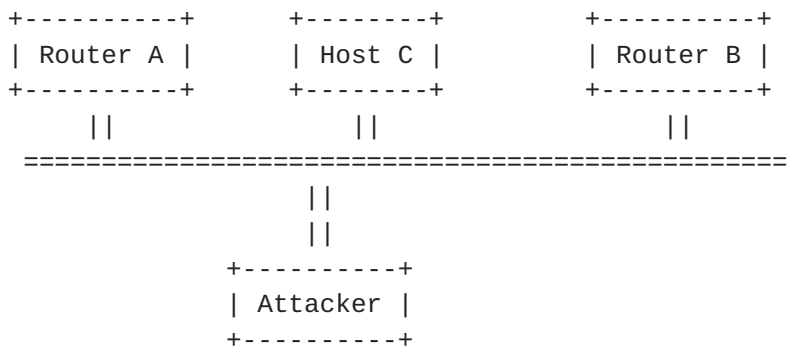

Figure 2: Broadcast Forwarding Loop

In Figure 2, the Attacker sends one crafted ND message to Router A, and one crafted ND message to Router B. Each crafted ND message contains the Target Addres set to Host C's IPv6 address, and a TLLA option set to the Ethernet broadcast address (ff:ff:ff:ff:ff:ff). These ND messages causes each router to map Host C's IPv6 address to the Ethernet broadcast address. This sets up the scenario for a subsequent attack.

Subsequently, the Attacker sends a packet to the Ethernet broadcast address (ff:ff:ff:ff:ff:ff), with an IPv6 Destination Address equal to the IPv6 address of Host C. Upon receipt, both Router A and Router C decrement the Hop Limit of the packet, and resend it to the Ethernet broadcast address. As a result, both Router A and Router B receive two copies of the same packet (one sent by Router A, and another sent by Router B). This would result in a "chain reaction" that would only disappear once the Hop Limit of each of the packets is decremented to 0. The equation in Figure 3 describes the amplification factor for this scenario :

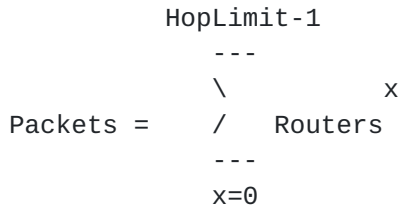

Figure 3: Maximum amplification factor

This equation does not take into account ICMPv6 Redirect messages that each of the Routers could send, nor the possible ICMPv6 "time exceeded in transit" error messages that each of the routers could send to the Source Address of the packet when each of the "copies" of the original packet is discarded as a result of their Hop Limit being decremented to 0.

An attacker can realize this attack by sending either of the following:

- An ND message whose SLLA maps an IPv6 address not belonging to the victim routers to the broadcast link-layer address

- An ND message whose TLLA maps an IPv6 address not belonging to the victim routers to the broadcast link-layer address

  NOTE: the IPv6 Destination Address of the attack packet should not belong to any of the victim routers, such that they forward the packet rather than "consume" it.

An additional mitigation would be for routers to not forward IPv6 packets on the same interface if the link-layer destination address of the received packet was a broadcast or multicast address.

4. Implications of Malicious Link-layer Address Options

If SLLA or TLLA options are allowed to contain broadcast (e.g., the IEEE 802 "ff:ff:ff:ff:ff:ff") or multicast (e.g., the IEEE 802 "33:33:00:00:00:01") addresses, traffic directed to the corresponding IPv6 address would be sent to the broadcast or multicast address specified in the SLLA or TLLA option. This could have multiple implications:

- It would have a negative impact on the performance of the nodes attached to the network and on the network itself, as packets sent to these addresses would need to be delivered to multiple nodes (and processed by them) unnecessarily.

- An attacker could easily capture traffic on a switched network, without the need to forward packets to their intended destinations, as the corresponding packets would be delivered to all (in the case of broadcast) or multiple (in the case of multicast) nodes.

- Packets could result in forwarding loops at routers, as a router forwarding a packet to the corresponding address would receive itself a copy of the forwarded packet. The loop would end only when the Hop Limit is eventually decremented to 0. The problem would be exacerbated if multiple routers are present on the same link. Section 3 of this document contains further analysis of this vulnerability.

Additionally, if SLLA or TLLA options are allowed to contain the receiving router's own link-layer address, the victim router would

receive a copy of the very same packets it means to forward to other destinations. This could have the following implications:

- o It would have a negative impact on the performance of the victim router and of the network itself, as a single packet would be sent multiple times (up to 255) on the local network, thus serving as an amplification vector.

5. Validation Checks for the Source Link-Layer Address Option

The Source link-layer address option contains the link-layer address of the sender of the packet. It is used by Neighbor Solicitation, Router Solicitation, and Router Advertisement messages.

The following figure illustrates the syntax of the source link-layer address:

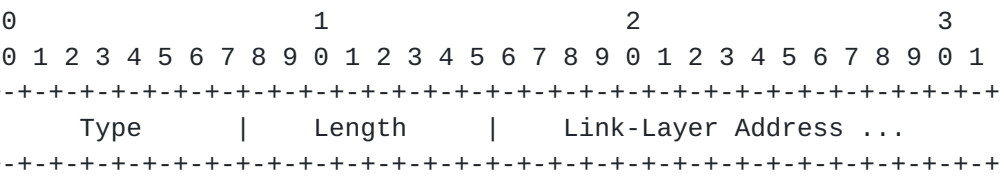

Figure 4: ND Source link-layer address option

The Type field is set to 1. The Length field specifies the length of the option (including the Type and Length octets) in units of 8 octets. A node that receives an ICMPv6 message with this option MUST verify that the Length field is valid for the underlying link layer. For example, for IEEE 802 addresses the Length field MUST be 1 [RFC2464]. If the packet does not pass this check, it MUST be silently dropped.

> NOTE: The Link-Layer Address field contains the link-layer address. The length, contents, and format of this field varies from one link layer to another, and is specified in specific documents that describes how IPv6 operates over different link layers.

Additionally, the SLLA option MUST NOT not contain a broadcast or multicast address. If the option does not pass this check, the Neighbor Discovery message carrying the option MUST be discarded. Finally, nodes MUST NOT allow the SLLA option to contain one of the receiving node's link-layer addresses. If the option does not pass this check, the Neighbor Discovery message carrying the option MUST discarded.

6. Validation Checks for the Target Link-Layer Address Option

The Target link-layer address option contains the link-layer address of the Target of the packet. It is used by Neighbor Advertisement and Redirect messages.

The following figure illustrates the syntax of the Target link-layer address:

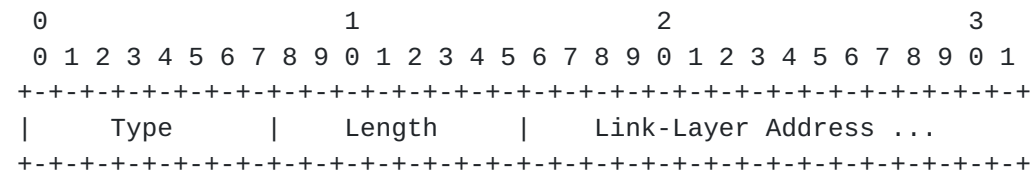

Figure 5: ND Target link-layer address option format

The Type field is set to 2. The Length field specifies the length of the option (including the Type and Length octets) in units of 8 octets. A node that receives a ND message with this option MUST verify that the Length field is valid for the underlying link-layer. For example, for IEEE 802 addresses the Length field MUST be 1 [RFC2464]. If the packet does not pass this check, it MUST be silently dropped.

A node that receives a ND message with this option MUST verify that the Length field is valid for the underlying link layer. For example, for IEEE 802 addresses the Length field MUST be 1 [RFC2464]. If the packet does not pass this check, it MUST be silently dropped.

The TLLA option MUST NOT not contain a broadcast or multicast address. If the option does not pass this check, the Neighbor Discovery message carrying the option MUST be discarded. Finally, nodes MUST NOT allow the source link-layer address to contain one of the receiving node's link-layer addresses. If the option does not pass this check, the Neighbor Discovery message carrying the option MUST discarded.

7. IANA Considerations

There are no IANA registries within this document. The RFC-Editor can remove this section before publication of this document as an RFC.

8. Security Considerations

This document discusses how the Neighbor DIscovery SLLA and TLLA options can be leveraged to perform a number of attacks, and specifies sanity checks to be enforced by Neighbor Discovery implementations, such that these vulnerabilities are eliminated.

9. Acknowledgements

This document is based on the technical report "Security Assessment of the Internet Protocol version 6 (IPv6)" [CPNI-IPv6] authored by Fernando Gont on behalf of the UK Centre for the Protection of National Infrastructure (CPNI).

10. References

10.1. Normative References

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

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.

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

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

10.2. 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).

Authors' Addresses

Fernando Gont
SI6 Networks / UTN-FRH
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: http://www.si6networks.com

Ronald P. Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
US

Phone: 571 250 5819
Email: rbonica@juniper.net

Will (Shucheng) Liu
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: liushucheng@huawei.com