Internet Engineering Task Force
Internet-Draft
Updates: 4861 (if approved)
Intended status: Standards Track
Expires: April 10, 2013

M. Smith
IMOT

October 7, 2012

Mitigating IPv6 Router Neighbor Cache DoS Using Stateless Neighbor Discovery
draft-smith-6man-mitigate-nd-cache-dos-slnd-00

Abstract

The IPv6 neighbor discovery cache is vulernable to a Denial of Service attack that purposely exhausts the state used during the neighbor discovery address resolution process. This can be very disruptive when a router is successfully attacked.

This memo proposes a stateless form of neighbor discovery to be used by routers to eliminate the opportunity for this DoS attack. This method of stateless neighbor discovery would be used for unknown or untrusted packet sources, when the router's neighbor cache's state capacity reaches a medium to high threshold of use. Trusted packet sources would continue to be provided with traditional stateful neighbor discovery.

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 April 10, 2013.

Table of Contents

1. Introduction 3
1.1. Requirements Language 3
2. Terminology 3
3. Stateless Neighbor Discovery 3
3.1. SLND Variables 3
3.2. SLND Process 4
4. Consequences of Stateless Neighbor Discovery 5
4.1. Neighbor Advertisement Validation 6
4.2. Optimisation Functions 6
5. Trusted/Untrusted Source Prefix List 7
5.1. Configured Trusted and Untrusted Prefixes 7
5.2. Routing Information 8
5.3. Default to Untrusted 8
6. Acknowledgements 8
7. Security Considerations 9
8. Change Log [RFC Editor please remove] 9
9. References 9
9.1. Normative References 9
9.2. Informative References 9
Author's Address 9

1. Introduction

The IPv6 neighbor discovery cache [RFC4861] is vulernable to a Denial of Service attack that purposely exhausts the state used during the neighbor discovery address resolution process [RFC3756]. This can be very disruptive when a router is successfully attacked.

This memo proposes a stateless form of neighbor discovery to be used by routers to eliminate the opportunity for this DoS attack. This method of stateless neighbor discovery would be used for unknown or untrusted packet sources, when the router's neighbor cache's state capacity reaches a medium to high threshold of use. Trusted packet sources would continue to be provided with traditional stateful neighbor discovery.

1.1. Requirements Language

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. Terminology

Stateful Neighbor Discovery (SFND): Traditional neighbor discovery, as specified in [RFC4861]. This form of neighbor discovery maintains per packet destination state for all unresolved destinations during the neighbor discovery process. The neighbor cache's state capacity is intentionally exhausted to cause the neighbor cache Denial of Service attack.

Stateless Neighbor Discovery (SLND): The form of neighbor discovery described in this memo. This form of neighbor discovery does not maintain per packet destination state for unresolved destinations during the neighbor discovery process.

3. Stateless Neighbor Discovery

3.1. SLND Variables

To perform stateless neighbor discovery, four variables are maintained:

SLND Flag - This flag indicates whether or not the interface will perform SLND if necessary.

SLDN Activate Threshold - This variable specifies the threshold when

stateless neighbor discovery is activated. The threshold specifies a neighbor cache utilisation level. It is expressed as a percentage, with a default value of 80%. It may be either a per-interface or router global variable depending on whether the router implementation has per-interface neighbor caches or a global neighbor cache used by all interfaces.

SLND Active Flag - This flag indicates whether or not the interface is performing SLND for untrusted packet sources. It is maintained for each interface on the router.

Trusted/Untrusted Sources Prefix List ("TUSP List") - This variable specifies a list of trusted and/or untrusted packet source address prefixes. It is a per-interface variable, as different interfaces on the router may have different sets of trusted and/or untrusted packet sources. However, a router may maintain a single global TUSP List that is used by all interfaces that may perform SLND, which don't have an interface specific TUSP List.

SLND Neighbor Solicitation Rate Limit ("SLND NS Rate Limit") - This variable specifies a threshold for multicast Neighbor Solictiations when the interface is performing SLND, specified in packets per second. It is a per-interface attribute, as different interfaces may have different thresholds. The rate value should be appropriate to the multicast capabilities of the interface link technology, with a typical value being 10% of the multicast rate supported by the link. A router may maintain a global threshold that is applied to interfaces that do not have an interface specific rate limit.

3.2. SLND Process

The stateless neighbor discovery process may occur once a router has determined the outgoing interface for a packet, and that the packet's destination is on-link.

If the packet's destination address is present in the neighbor cache, and the link-layer address has been resolved, the packet is forwarded to it's destination.

If the packet's destination address is not present in the neighbor cache, and the SLND Flag is off, traditional stateful neighbor discovery is performed for the packet's destination.

If the packet's destination address is not present in the neighbor cache, and the SLND Flag is on, the packet's source address is compared to the TUSP List.

If the packet's source address is determined to be trusted,

traditional neighbor discovery is performed.

If the packet's source address is determined to be untrusted, stateless neighbor discovery is performed. The stateless neighbor discovery process is as follows:

1. The router determines if sending a multicast neighbor solicitation would exceed the SLND NS Rate Limit for the outgoing interface. If the SLND NS Rate Limit would be exceeded, drop the packet and do not proceed any further.

2. A multicast neighbor solicitation is sent by the router for the destination address in the packet. The packet is then dropped.

3. As some later point in time, the router is likely to receive a unicast neighbor advertisement, for a previously sent neighbor solicitation.

4. If the SLND Active Flag is off, the router ignores the neighbor advertisement.

5. If the SLND Active Flag is on, the router creates an entry in it's neighbor cache using the information received in the unicast neighbor advertisement. Stateless neighbor discovery is now complete.

The utilisation of the neighbor cache has to be measured to determine if it crosses the SLDN Activate Threshold. If the utilisation increases above the SLDN Activate Threshold, the SLND Active Flag is switched on, and if it decreases below the SLDN Activate Threshold, the SLND Active Flag is switched off. Neighbor cache utlisation should be measured and compared to the SLDN Activate Threshold when:

- entries are added to the neighbor cache, during either stateful or stateless neighbor discovery

- entries are removed from the neighbor cache when NUD discovers the neighbor has become unreachable

4. Consequences of Stateless Neighbor Discovery

During traditional stateful neighbor discovery, state is used to perform the following:

- ensure a received neighbor advertisement corresponds to a previously sent neighbor solicitation

- to retransmit a limited number of neighbor solicitations if previous solicitations remain unanswered

- to store a small number of packets that triggered the neighbor discovery process, so that they can be transmitted if neighbor discovery completes successfully

- to generate an ICMPv6 destination unreachable, address unreachble messages back to the traffic source, should the neighbor discovery process fail

Stateless neighbor discovery sacrifices these functions and the related state to mitigate the neighbor cache denial-of-service attack.

4.1. Neighbor Advertisement Validation

Ensuring received advertisements correspond to previously sent neighbor solicitations prevents on-link nodes from sending unsolicited neighbor advertisements, and then having them added to the router's neighbor cache without validation. Doing so would allow the on-link nodes to perform a neighbor cache denial of service attack, similar to the one this memo mitigates for off-link sources.

If neighbor advertisement validation is to occur, then the router is vulnerable to an off-link sourced neighbor cache DoS attack, but not vulnerable to an on-link sourced neighbor cache DoS attack. If this neighbor advertisement validation does not occur, then the reverse is the case.

Considering that on-link nodes will usually have a vested interest in the router continuing to operate, that there will be a much smaller set of on-link sources, and that they'll be far better known and possibly access controlled, the likelihood of an on-link sourced neighbor cache DoS is much lower than an off-link sourced neighbor cache DoS. It is therefore beneficial to sacrifice on-link neighbor cache DoS protection to gain off-link neighbor cache DoS protection. Also note that during the stateless neighbor discovery process proposed in this memo, neighbor advertisement validation is only sacrificed when an off-link sourced neighbor cache DoS appears to be taking place. Under normal circumstances on-link sourced neighbor advertisement validation will continue to occur.

4.2. Optimisation Functions

The nature of IPv6 is best effort, meaning that there is a possibility that packets may be lost as they transit the network, and that IPv6 will not make any attempt to recover lost packets. If an

application residing on an IPv6 node requires reliable packet delivery, it will need to utilise locally implemented reliable upper layer protocols such as TCP and SCTP, or implement it's own reliability mechanisms. These reliability mechanisms involve retransmitting packets.

The remaining uses of stateful neighbor discovery state are not assured of success. The limited number of neighbor solicitation retransmissions may not be enough, causing neighbor discovery to fail even though the target node exists. There may be more packets sent that trigger neighbor discovery than are stored for transmission when neighbor discovery completes successfully, causing them to be dropped. The ICMPv6 destination unreachable message may be dropped on the way back to the traffic originating node, perhaps intentionally by a network located firewall.

This means that these functions are useful but not essential optimisations. If necessary, they do not need to be performed, as the traffic source will retransmit it's packets, reinitiating the neighbor discovery process. This provides the opportunity to perform a stateless form of neighbor discovery if there is evidence that a neighbor cache DoS attack is occuring, mitigating the off-link sourced neighbor cache DoS attack.

5. Trusted/Untrusted Source Prefix List

The following information sources can be used to construct the trusted/untrusted source prefix list (TUSP List).

5.1. Configured Trusted and Untrusted Prefixes

The first TUSP List source is an operator configured list of prefixes and their lengths, each with a flag indicating whether traffic with source addresses that falls within the specified prefix is from a trusted or untrusted source.

How this list is evaluated would be implementation dependent, however it is likely to be either sequential from first to last entry, or using a longest match algorithm.

This list should have a default entry of the ULA prefix (fc00::/7) [RFC4193], flagged as a trusted source. An implementation must allow this entry to be removed.

5.2. Routing Information

The second TUSP List source is the network's routing information.

The network's routing information can be used to distinguish trusted and untrusted traffic sources. An advantage of using routing information for this purpose is that it will typically be dynamically and automatically distributed to all routers within the network, when dynamic routing protocols are used. This avoids routers in the network having to be manually reconfigured when subnets are added or removed from the network.

The contents of a stub network's route table is typically all the internal routes for the network, and then a default route used to reach the Internet. The list of internal routes can be used to distinguish between trusted and untrusted sources, with traffic sources matching internal routes being trusted, and all other traffic sources being untrusted.

In more complex routing environments, such as those using one or more IGPs and an EGP such as BGP, there may be other methods available to distinguish between trusted and untrusted sources. For example, routes carried in an IGP could be considered trusted, while routes carried in BGP are untrusted. For a network using BGP to carry all reachability information, except network transit and loopback interface routes, routes may be tagged with one or more BGP communities which indicate internal and therefore trusted prefixes.

A default route should never be used to define a trusted traffic source prefix. If a router's operator wishes to trust all traffic sources, they should configure ::/0 as a configured trusted prefix.

Implementations should provide convenient methods to use the network's routing information to distinguish between trusted and untrusted traffic source prefixes.

5.3. Default to Untrusted

Finally, should none of the previous trusted or untrusted source prefix information sources match the source address of traffic that would trigger neighbor discovery, the packet source should be considered untrusted.

6. Acknowledgements

7. Security Considerations

This memo proposes a security mitigation for an off-link sourced neighbor cache denial-of-service attack.

As discussed in Section 4.1, the method proposed creates an opportunity for an on-link sourced neighbor cache DoS attack, when mitigating the off-link sourced neighbor cache DoS. This is considered to be an acceptable security trade-off.

8. Change Log [RFC Editor please remove]

draft-smith-6man-mitigate-nd-cache-dos-slnd-00, initial version, 2012-09-04

9. References

9.1. Normative References

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

9.2. Informative References

[RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

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

Author's Address

Mark Smith
In My Own Time
PO BOX 521
HEIDELBERG, VIC 3084
AU

Email: markzzzsmith@yahoo.com.au