Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 30, 2016

M. Smith

February 27, 2016

Further Mitigating Router ND Cache Exhaustion DoS Attacks Using Solicited-Node Group Membership
draft-smith-v6ops-mitigate-rtr-dos-mld-slctd-node-02

Abstract

For each of their IPv6 unicast or anycast addresses, nodes join a Solicited-Node multicast group, formed using the lower 24 bits of the address. This Solicited-Node group membership could be used by routers to further mitigate a Neighbor Discovery cache Denial of Service attack.

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 30, 2016.

Copyright Notice

Table of Contents

1. Introduction

When an IPv6 unicast or anycast address is added to or removed from an interface, a node is also required to join or leave the Solicited-Node multicast group that corresponds to the address [RFC4291][RFC6434], using the Multicast Listener Discovery (MLD) protocol [RFC2710][RFC3810]. The Solicited-Node multicast group the node joins or leaves is determined by appending the lower 24 bits of the unicast or anycast address, usually part of the Interface Identifier (IID), to the IPv6 multicast prefix FF02:0:0:0:0:1:FF00::/104 [RFC4291].

The current use of Solicited-Node multicast groups is to avoid having to link layer broadcast Neighbor Discovery (ND) Neighbor Solicitations to all nodes on the link (ARP's [RFC0826] behaviour for most ARP Requests). Instead, Neighbor Solicitations are sent to the Solicited-Node multicast group formed from the target address of the Neighbor Solicitation.

The use of Solicited-Node multicast groups for Neighbor Solicitations allows nodes to possibly filter Neighbor Solicitations they aren't interested in in their link layer network interface, avoiding interrupting the node's general purpose CPU (see sections 7.4 and 7.5 of [RFC1112] for further discussion), and possibly for the link layer forwarding device(s) to avoid sending Neighbor Solicitations to nodes that do not have the target address [RFC4541]. Facilitating link layer network interface multicast filtering and reducing the flooding scope of multicasts on a link helps increase the number of nodes that can be attached to a link.

As the addition or removal of unicast or anycast addresses triggers Solicited-Node multicast group joins or leaves, this mechanism is in effect a low resolution address range presence registration protocol, registering portions of the on-link address range for which there are unicast or anycast addresses present. The presence of a Solicited-Node multicast group on a link indicates that at least one unicast or anycast address that maps to the Solicited-Node multicast group is present. Conversely, the absence of a Solicited-Node multicast group on a link indicates that no unicast or anycast addresses are present that would map to the corresponding Solicited-Node multicast group.

MLDv2 joins for Solicited-Node multicast groups could also be used as a link-local address registration method for at least one of each nodes' link-local addresses, as link-local unicast addresses are used as MLDv2 source addresses, excepting MLDv2 joins for Solicited-Node multicast groups when a link-local address is not available [RFC3590]. It would not be possible to do this reliably with MLDv1 Solicited-Node multicast group joins as MLDv1 listeners will suppress joins for their own groups if they hear a join for the same groups from another listener.

This presence or absence of Solicited-Node multicast groups could be used by a router to determine if it needs to send Neighbor Solicitations for unresolved addresses on to the link. If the to-be-resolved address maps to a non-existent Solicited-Node multicast group, the router could discard the packet, rather than sending a Neighbor Solicitation to the corresponding Solicited-Node multicast group for the packet's destination and possibly queuing the trigger packet while neighbor discovery occurs. Discarding trigger packets that map to absent Solicited-Node multicast groups could be a further Neighbor Discovery cache Denial of Service (DoS) attack [RFC3756] mitigation technique.

For links with prefixes with lengths shorter than or equal to /104, such as the common /64 [RFC7421], the total number of Solicited-Node multicast groups possible on a link is 2^24, or 16 777 216 groups. The number of Solicited-Node multicast groups present on a link is equal to the number of IPv6 unicast or anycast addresses present on the link which have unique lower 24 bits, used to form the Solicited-Node multicast group address.

For most links the number of present Solicited-Node multicast groups present will be in the order of 10s, 100s or perhaps on rarer occasions in the low 1000s. This means that Neighbor Solicitations do not have to be sent for very large numbers of unresolved unicast or anycast addresses for which the corresponding Solicited-Node multicast group is not present. This would significantly reduce the attack surface for the ND cache exhaustion denial of service attack.

For example, if a link has 1000 present Solicited-Node multicast groups, then Neighbor Solicitations do not have to be sent for addresses that would map to the absent 16 776 216 Solicited-Node multicast groups, more than 99.99% of the possible on-link Solicited-Node multicast groups.

This memo describes how a router could collect Solicited-Node multicast group membership and how it could use this information as part of its neighbor presence discovery procedure, for the purposes of further mitigating the ND cache exhaustion attack.

Note that this method has been independently suggested by Greg Daley and perhaps others.

2. Method

2.1. Tracking Solicited-Node Multicast Group Presence

To track Solicited-Node multicast group presence on a link, a router uses the multicast listener discovery procedures specified in [RFC2710] or [RFC3810], without modification.

Note that the procedures specified in [RFC2710] and [RFC3810] do not require that a router performing them is to forward multicast packets, or is to be participating in a multicast routing protocol with other multicast routers. The ND cache DoS mitigation method described in this memo can be used regardless of whether the other routers in the network, including other on-link routers, are performing multicast forwarding.

If a router using this ND cache DoS mitigation method is not performing multicast forwarding, it may choose to only track Solicited-Node multicast group presence, ignoring the presence information it receives for other multicast groups. This may usefully reduce the router’s resources consumption. If a router using this optimisation becomes a multicast forwarding router, it will need to collect presence information for all on-link multicast groups, using the Querier Election procedure [RFC2710][RFC3810], as though it had just been attached to the link, and had no knowledge of the presence any of the multicast groups.

A router with two or more interfaces attached to the same link only needs to operate MLD on one of those interfaces [RFC3810]; the list of on-link Solicited-Node multicast groups would be used across all of these interfaces when mitigating ND cache DoSes.

2.2. Neighbor Presence Discovery

When a router receives a packet for a destination for which it does not have a neighbor cache entry, it uses the [RFC4291] specified method to form a Solicited-Node multicast group address from the packet's destination address.

The router then compares the resulting Solicited-Node multicast group address with its list of present Solicited-Node multicast groups on the link.

If the Solicited-Node multicast group is present, the router then performs the address resolution procedure for the packet's destination IPv6 address as specified in [RFC4861], starting with sending a Neighbor Solicitation towards the Solicited-Node multicast group that corresponds to the address.

Alternatively, when the Solicited-Node multicast group is not present, the router operates in one of two mitigation modes.

2.2.1. Strict Mitigation Mode

When operating in Strict Mitigation Mode, the router discards all packets whose destination address Solicited-Node multicast groups do not match any of the Solicited-Node multicast groups present on the link.

Strict Mitigation Mode makes the decision to perform Neighbor Discovery dependent on the successful discovery of the Solicited-Node multicast groups on the link by MLD. This means that if the router is assembling a list of present Solicited-Node multicast groups from scratch, such as after the router has been intialised, or when an interface comes online, there will be a period where Neighbor Discovery for existing nodes will not occur, while the full set of present Solicited-Node multicast groups are discovered. To off-link hosts sending traffic to the possible on-link hosts, this will appear to be a period of packet loss. These hosts are expected to have implemented methods to recover from transient failures of transmission, such as packet retransmission, if necessary [RFC1958].

This mode of operation is appropriate when it is known that all attached nodes announce their Solicited-Node multicast group membership for their addresses, and MLD operation on the link is known to be reliable. An example scenario would be a large Internet content provider's environment, where the content network routers and content servers are being operated by the same organisation.

2.2.2. Relaxed Mitigation Mode

When operating in Relaxed Mitigation Mode, under normal non-DoS circumstances the router will also perform the address resolution procedure for packets whose destination address Solicited-Node multicast group does not match any of the Solicited-Node multicast groups present on the link.

However, when there is an indication that a neighbor cache Denial of Service attack might be occurring, the router treats packets whose destination address Solicited-Node multicast group does not match a link present Solicited-Node multicast group with lower importance to those packets whose do.

Indicators that a neighbor cache Denial of Service attack might be occurring could be many failed address resolution attempts over a short period of time, rapid and unexpected consumption of neighbor cache resources (rapid consumption for a short period of time after the link or router has come on-line could be expected), or some other pattern of neighbor cache Denial of Service attack specific behaviour.

If a neighbor cache Denial of Service attack appears to be occurring, an implementation could immediately start discarding packets whose destination address Solicited-Node multicast group does not match those present on the link. A less harsh alternative would be to start discarding some of these packets, increasing the discard rate as neighbor cache resources are increasingly consumed by the Denial of Service attack.

This mode of operation would be appropriate when it is not known if all nodes will announce their Solicited Node multicast group membership, possibly due to some nodes being pre-[RFC2710] implementations or if MLD operation is not known to be reliable. Example scenarios would be residential or public Internet access networks, where the support for or reliability of MLD joins for Solicited Node multicast groups is not known. Specific to the residential network case, where the technical ability of the router operator is not known and likely to be low, Relaxed Mitigation Mode would be the safest default.

3. MLD Reliability

MLD is currently being used for two purposes:

- to join and leave multicast groups so that multicast applications will receive routed multicast traffic they are interested in receiving [RFC2710][RFC3810], and

- o to advise link layer devices of node multicast group membership to allow the link layer devices to limit to which devices multicast traffic is sent, instead of flooding multicast traffic to all attached devices [RFC4541]. Specific to this memo's topic, nodes using MLD to join Solicited-Node multicast groups for their addresses allows link layer devices to limit to which nodes multicast Neighbor Solicitations are sent.

For the first purpose, partial or complete failure of MLD to successfully join the intended multicast group(s) will likely cause the respective multicast application(s) to not function adequately or completely. While likely to be unacceptable to the application(s') user(s), the effects of the failure are limited to the impacted application(s); some multicast applications may function, and other unicast-only applications will not be impacted.

For the second purpose, partial or complete failure of MLD operation means the link layer device will not forward multicast traffic to devices for groups for which MLD joins failed. As with the first MLD purpose, application operation is likely to be impacted. MLD join failures for Solicited-Node multicast groups would mean that Duplicate Address Detection [RFC4861] and Neighbor Discovery [RFC4861] for the node's addresses will fail. IPv6 unicast connectivity for the effected node could be severely impacted, and possibly fail completely.

For this memo's method, when operating in Strict Mitigation Mode, partial or complete failure of MLD for Solicited-Node multicast group joins will cause Neighbor Discovery to fail for routers implementing this neighbor cache Denial of Service attack mitigation. The effected nodes will be unreachable for traffic sources beyond the impacted router.

With this memo's method, when operating in Relaxed Mitigation Mode, partial or complete failure of MLD for Solicited-Node multicast group joins will cause the router to consider neighbor discovery for the effected node's addresses to be of lower importance. Under normal, non-neighbor cache Denial of Service circumstances, these nodes will receive equal service to those who've successfully joined the Solicited-Node multicast groups via MLD. If a neighbor cache Denial of Service occurs, these MLD failed nodes will either have less success at or complete failure of being discovered by the router performing neighbor discovery. In this situation, some rather than all of the nodes will have been impacted by the Denial of Service attack, which is an improvement over the attack impacting all nodes.

It is important to note that failure of neighbor discovery during a neighbor cache Denial of Service attack will only impact nodes that

have not been previously discovered by the router. If a node has been previously discovered, its neighbor information will already reside in the router's neighbor cache, and its currency will be maintained by Neighbor Unreachability Detection [RFC4861].

Due to the number of significant consequences of MLD failure, including those introduced by this memo's method, MLD should be configured to operate reliably if the default MLD reliability related parameter values are not adequate [RFC2710][RFC3810]. Although [RFC6636] provides advice for tuning MLD operation for mobile and wireless networks, some of the advice and considerations might be more generally applied.

4. Security Considerations

The method described in this memo further mitigates the ND cache exhaustion DoS attack. It does not prevent it.

Using this method, neighbor presence discovery will occur for any of the unicast or anycast addresses that map to the present Solicited-Node multicast groups. As a Solicited-Node multicast group can map to up to 2^40 unicast or anycast addresses (for a /64 prefix, 2^(64 - 24)), the ND implementation is likely to continue to be vulnerable to a ND cache exhaustion denial of service for addresses covered by the present Solicited-Node multicast groups. While the number of non-existent addresses that can be targetted remains very large, it is very significantly smaller than the targettable non-existent addresses possible in the on-link prefixes without this measure.

The severity of this threat depends on two factors:

- the number of Solicited-Node multicast groups present on the link, and
- the ability of the off-link attacker to stumble upon or discover non-existent addresses that map to present Solicited-Node multicast groups.

The severity of the threat is lower with lesser numbers of Solicited-Node multicast groups, and less predictable and sparsely distributed Solicited-Node multicast group addresses.

[RFC7217] specifies the use of stable yet random and unpredictable IIDs, on a per-prefix basis. This will increase the number of present Solicited-Node multicast groups, by up to the number of prefixes multiplied by the number of hosts implementing [RFC7217]. This will reduce the effectiveness of the measure proposed in this memo. However, it will also conversely increase the effectiveness of

this measure, as the IIDs and therefore the Solicited-Node multicast groups become less predictable and more sparsely distributed.

To protect against ND cache DoS attacks for non-existent addresses that map to present Solicited-Node multicast groups, other ND cache protection measures, such as those described in [RFC6583] should also be implemented.

When a packet is sent to a destination that is unresolved and is not covered by a present Solicited-Node multicast group, a copy could be sent to an [RFC6018] greynet collector for further analysis. For example, packet sent to destinations falling outside the present Solicited-Node multicast groups could be an indication of an attempt to discover nodes via address probing.

5. Acknowledgements

Review and comments were provided by (in alphabetical order) Fred Baker, Lorenzo Colitti and Ray Hunter. Lorenzo expressed a concern about MLD reliability and its consequences, which prompted the creation of the two modes of mitigation and the MLD reliability discussion.

This memo was prepared using the xml2rfc tool.

6. Change Log [RFC Editor please remove]

draft-smith-v6ops-mitigate-rtr-dos-mld-slctd-node-00, initial version, 2014-04-08

draft-smith-v6ops-mitigate-rtr-dos-mld-slctd-node-01, 2016-02-08

- two modes of response to possible neighbor cache DoS
- MLD reliability discussion
- Expansion on the purpose of using solicited-node multicast groups

draft-smith-v6ops-mitigate-rtr-dos-mld-slctd-node-02, 2016-02-28

- "discard" rather than "drop" to show intentional dropping
- text about Strict Mode MLD discovery period

7. Informative References

[RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982, <http://www.rfc-editor.org/info/rfc826>.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <http://www.rfc-editor.org/info/rfc1112>.

[RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <http://www.rfc-editor.org/info/rfc1958>.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <http://www.rfc-editor.org/info/rfc2710>.

[RFC3590] Haberman, B., "Source Address Selection for the Multicast Listener Discovery (MLD) Protocol", RFC 3590, DOI 10.17487/RFC3590, September 2003, <http://www.rfc-editor.org/info/rfc3590>.

[RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <http://www.rfc-editor.org/info/rfc3756>.

[RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <http://www.rfc-editor.org/info/rfc3810>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://www.rfc-editor.org/info/rfc4291>.

[RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <http://www.rfc-editor.org/info/rfc4541>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <http://www.rfc-editor.org/info/rfc4861>.

[RFC6018] Baker, F., Harrop, W., and G. Armitage, "IPv4 and IPv6 Greynets", RFC 6018, DOI 10.17487/RFC6018, September 2010, <http://www.rfc-editor.org/info/rfc6018>.

[RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <http://www.rfc-editor.org/info/rfc6434>.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, DOI 10.17487/RFC6583, March 2012, <http://www.rfc-editor.org/info/rfc6583>.

[RFC6636] Asaeda, H., Liu, H., and Q. Wu, "Tuning the Behavior of the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) for Routers in Mobile and Wireless Networks", RFC 6636, DOI 10.17487/RFC6636, May 2012, <http://www.rfc-editor.org/info/rfc6636>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <http://www.rfc-editor.org/info/rfc6724>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <http://www.rfc-editor.org/info/rfc7217>.

[RFC7421] Carpenter, B., Ed., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", RFC 7421, DOI 10.17487/RFC7421, January 2015, <http://www.rfc-editor.org/info/rfc7421>.

Author’s Address

Mark Smith
PO BOX 521
HEIDELBERG, VIC 3084
AU

Email: markzzzsmith@gmail.com