Network Working Group
Internet-Draft
Updates: 4862, 4861, 3971 (if approved)
Intended status: Standards Track
Expires: May 17, 2015

R. Asati
H. Singh
W. Beebee
C. Pignataro
Cisco Systems, Inc.
E. Dart
Lawrence Berkeley National Laboratory
W. George
Time Warner Cable

November 13, 2014

Enhanced Duplicate Address Detection
draft-ietf-6man-enhanced-dad-10

Abstract

IPv6 Loopback Suppression and Duplicate Address Detection (DAD) are discussed in Appendix A of RFC4862. That specification mentions a hardware-assisted mechanism to detect looped back DAD messages. If hardware cannot suppress looped back DAD messages, a software solution is required. Several service provider communities have expressed a need for automated detection of looped backed Neighbor Discovery (ND) messages used by DAD. This document includes mitigation techniques and outlines the Enhanced DAD algorithm to automate the detection of looped back IPv6 ND messages used by DAD. For network loopback tests, the Enhanced DAD algorithm allows IPv6 to self-heal after a loopback is placed and removed. Further, for certain access networks the document automates resolving a specific duplicate address conflict. This document updates RFC4861, RFC4862, and RFC3971.

Table of Contents

1. Introduction

IPv6 Loopback Suppression and Duplicate Address Detection (DAD) are discussed in Appendix A of [RFC4862]. That specification mentions a

hardware-assisted mechanism to detect looped back DAD messages. If hardware cannot suppress looped back DAD messages, a software solution is required. One specific DAD message is the Neighbor Solicitation (NS), specified in [RFC4861]. The NS is issued by the network interface of an IPv6 node for DAD. Another message involved

in DAD is the Neighbor Advertisement (NA). The Enhanced DAD algorithm specified in this document focuses on detecting an NS looped back to the transmitting interface during the DAD operation. Detecting a looped back NA does not solve the looped back DAD problem. Detection of any other looped back ND messages during the DAD operation is outside the scope of this document. This document also includes a section on Mitigation that discusses means already available to mitigate the DAD loopback problem. This document updates RFC4861, RFC4862, and RFC3971.

1.1. Two Deployment Problems

In each problem articulated below, the IPv6 link-local address DAD operation fails due to a looped back DAD probe. However, the looped back DAD probe exists for any IPv6 address type including global addresses.

Recently, service providers have reported a problem with DAD that is caused by looped back NS messages. The following is a description of the circumstances under which the problem arises. Loopback testing for troubleshooting purposes is underway on a circuit connected to an interface on a router. The interface on the router is enabled for IPv6. The interface issues a NS for the IPv6 link-local address DAD. The NS is reflected back to the router interface due to the loopback condition of the circuit, and the router interface enters a DAD-failed state. After the circuit troubleshooting has concluded and the loopback condition is removed, IPv4 will return to operation without further manual intervention. However, IPv6 will remain in DAD-failed state until manual intervention on the router restores IPv6 to operation.

There are other conditions which will also trigger similar problems with DAD Loopback. While the following example is not a common configuration, it has occurred in a large service provider network. It is necessary to address it in the proposed solution because the trigger scenario has the potential to cause significant IPv6 service

outages when it does occur. In this scenario, two broadband modems in the same home are served by the same service provider and both modems are served by one access concentrator and one layer-3 interface on the access concentrator. The two modems have the Ethernet ports of each modem connected to a network hub. The access concentrator serving the modems is the first-hop IPv6 router for the modems. The network interface of the access concentrator serving the two broadband modems is enabled for IPv6 and the interface issues a NS(DAD) message for the IPv6 link-local address. The NS message reaches one modem first and this modem sends the message to the network hub which sends the message to the second modem which forwards the message back to the access concentrator. The looped

back NS message causes the network interface on the access concentrator to be in a DAD-failed state. Such a network interface typically serves up to a hundred thousand broadband modems causing all the modems (and hosts behind the modems) to fail to get IPv6 online on the access network. Additionally, it may be tedious for the user of the access concentrator to find out which of the hundred thousand or more homes looped back the DAD message. Clearly there is a need for automated detection of looped back NS messages during DAD operations by a node.

2. Terminology

- DAD-failed state - Duplication Address Detection failure as specified in [RFC4862]. Note even Optimistic DAD as specified in [RFC4429] can fail due to a looped back DAD probe. This document covers looped back detection for Optimistic DAD as well.

- Looped back message - also referred to as a reflected message. The message sent by the sender is received by the sender due to the network or an Upper Layer Protocol on the sender looping the message back.

- Loopback - A function in which the router's layer-3 interface (or the circuit to which the router's interface is connected) is looped back or connected to itself. Loopback causes packets sent by the interface to be received by the interface and results in interface unavailability for regular data traffic forwarding. See more details in section 9.1 of [RFC2328]. The Loopback function is commonly used in an interface context to gain information on

the quality of the interface, by employing mechanisms such as ICMPv6 pings and bit-error tests. In a circuit context, this function is used in wide area environments including optical Dense Wave Division Multiplexing (DWDM) and SONET/SDH for fault isolation (e.g. by placing a loopback at different geographic locations along the path of a wide area circuit to help locate a circuit fault). The Loopback function may be employed locally or remotely.

- NS(DAD) - shorthand notation to denote an Neighbor Solicitation (NS) (as specified in [RFC4861]) with unspecified IPv6 source-address issued during DAD.

2.1. 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 [RFC2119].

3. Operational Mitigation Options

Two mitigation options are described below. The mechanisms do not require any change to existing implementations.

3.1. Disable DAD on an Interface

One can disable DAD on an interface so that there is no NS(DAD) issued to be looped back. DAD is disabled by setting the interface's DupAddrDetectTransmits variable to zero. While this mitigation may be the simplest, the mitigation has three drawbacks.

This mitigation would likely require careful analysis of configuration on such point-to-point interfaces, a one-time manual configuration on each of such interfaces, and more importantly, genuine duplicates in the link will not be detected.

A Service Provider router, such as an access concentrator, or network core router, SHOULD support this mitigation strategy.

3.2. Dynamic Disable/Enable of DAD Using Layer-2 Protocol

One or more layer-2 protocols MAY include provisions to detect the existence of a loopback on an interface circuit, usually by comparing protocol data sent and received. For example, the Point-to-Point Protocol (PPP) uses a magic number (section 6.4 of [RFC1661]) to detect a loopback on an interface.

When a layer-2 protocol detects that a loopback is present on an interface circuit, the device MUST temporarily disable DAD on the interface. When the protocol detects that a loopback is no longer present (or the interface state has changed), the device MUST (re-)enable DAD on that interface.

This solution requires no protocol changes. This solution SHOULD be enabled by default, and MUST be a configurable option if the layer-2 technology provides means for detecting loopback messages on an interface circuit.

This mitigation has several benefits. They are

1. It leverages layer-2 protocol's built-in loopback detection capability, if available.

2. It scales better since it relies on an event-driven model which requires no additional state or timer. This may be a significant scaling consideration on devices with hundreds or thousands of interfaces that may be in loopback for long periods of time (such

as while awaiting turn-up or during long-duration intrusive bit error rate tests).

3.3. Operational Considerations

The mitigation options discussed in the document do not require the devices on both ends of the circuit to support the mitigation functionality simultaneously, and do not propose any capability negotiation. The mitigation options discussed in this document are effective for unidirectional circuit or interface loopback (i.e. the the loopback is placed in one direction on the circuit, rendering the other direction non-operational).

The mitigation options may not be effective for the bidirectional loopback (i.e. the loopback is placed in both directions of the

circuit interface, so as to identify the faulty segment) if only one device followed a mitigation option specified in this document, since the other device would follow current behavior and disable IPv6 on that interface due to DAD until manual intervention restores it.

This is nothing different from what happens today (without the solutions proposed by this document) in case of unidirectional loopback. Hence, it is expected that an operator would resort to manual intervention for the devices not compliant with this document, as usual.

4. The Enhanced DAD Algorithm

The Enhanced DAD algorithm covers detection of a looped back NS(DAD) message. The document proposes use of the Nonce Option specified in the SEND document of [RFC3971]. The nonce is a random number as specified in [RFC3971]. If SEND is enabled on the router and the router also supports the Enhanced DAD algorithm (specified in this document), there is integration with the Enhanced DAD algorithm and SEND. (See more details in the Impact on SEND section 4.3.) Since a nonce is used only once, The NS(DAD) for each IPv6 address of an interface uses a different nonce. Additional details of the algorithm are included in section 4.2.

Six bytes of random nonce is sufficiently large to minimize collisions. However, if there is a collision because two nodes using the same Target Address in their NS(DAD) and generated the same random nonce, then the algorithm will incorrectly detect a looped back NS(DAD) when a genuine address collision has occurred. Since each looped back NS(DAD) event is logged to system management, the administrator of the network will have access to the information necessary to intervene manually. Also, because the nodes will have detected what appear to be looped back NS(DAD) messages, they will

continue to probe, and it is unlikely that they will choose the same nonce the second time (assuming quality random number generators).

The algorithm is capable of detecting any ND solicitation (NS and Router Solicitation) or advertisement (NA and Router Advertisement) that is looped back. However, for a sender to store a nonce and nonce related state for all ND messages has impact on memory and causes more complexity for the sender node. Therefore, this document

does not recommend using the algorithm outside of the DAD operation by an interface on a node.

4.1. Processing Rules for Senders

If a node has been configured to use the Enhanced DAD algorithm, when sending an NS(DAD) for a tentative or optimistic interface address the sender MUST generate a random nonce associated with the interface address, MUST store the nonce internally, and MUST include the nonce in the Nonce Option included in the NS(DAD). If the interface does not receive any DAD failure indications within RetransTimer milliseconds (see [RFC4861]) after having sent DupAddrDetectTransmits Neighbor Solicitations, the interface moves the Target Address to the assigned state.

If any probe is looped back within RetransTimer milliseconds after having sent DupAddrDetectTransmits NS(DAD) messages, the interface continues with another MAX_MULTICAST_SOLICIT number of NS(DAD) messages transmitted RetransTimer milliseconds apart. Section 2 of [RFC3971] defines a single-use nonce, so each Enhanced DAD probe uses a different nonce. If no probe is looped back within RetransTimer milliseconds after MAX_MULTICAST_SOLICIT NS(DAD) messages are sent, the probing stops. The probing MAY be stopped via manual intervention. When probing is stopped, the interface moves the Target Address to the assigned state.

4.2. Processing Rules for Receivers

If the node has been configured to use the Enhanced DAD algorithm and an interface on the node receives any NS(DAD) message where the Target Address matches the interface address (in tentative or optimistic state), the receiver compares the nonce included in the message, with any stored nonce on the receiving interface. If a match is found, the node SHOULD log a system management message, SHOULD update any statistics counter, and MUST drop the received message. If the received NS(DAD) message includes a nonce and no match is found with any stored nonce, the node SHOULD log a system management message for a DAD-failed state, and SHOULD update any statistics counter.

4.3. Impact on SEND

The SEND document uses the Nonce Option in the context of matching an NA with an NS. However, no text in SEND has an explicit mention of detecting looped back ND messages. As this document updates [RFC4862], SEND should be updated to integrate with the Enhanced DAD algorithm. A minor update to SEND would be to explicitly mention that the nonce in SEND is also used by SEND to detect looped back NS(DAD) messages during DAD operations by the node. In a mixed SEND environment with SEND and unsecured nodes, the lengths of the nonce used by SEND and unsecured nodes MUST be identical.

4.4. Changes to RFC 4862

The following text is added to the end of the Introduction section of [RFC4862].

A network interface of an IPv6 node should implement the Enhanced DAD algorithm. For example, if the interface on an IPv6 node is connected to a circuit that supports loopback testing, then the node should implement the Enhanced DAD algorithm that allows the IPv6 interface to self-heal after loopback testing is ended on the circuit. Another example is when the IPv6 interface resides on an access concentrator running DAD Proxy. The interface supports up to a hundred thousand IPv6 clients (broadband modems) connected to the interface. If the interface performs DAD for its IPv6 link-local address and the DAD probe is reflected back to the interface, the interface is stuck in DAD-failed state and IPv6 services to the hundred thousand clients is denied. Disabling DAD for such an IPv6 interface on an access concentrator is less desirable than implementing the algorithm because the network also needs to detect genuine duplicates in the interface downstream network. The Enhanced DAD algorithm also facilitates detecting a genuine duplicate for the interface on the access concentrator. (See the Actions to Perform on Detecting a Genuine Duplicate section of the Enhanced DAD document.)

The following text is added to the end of Appendix A of [RFC4862].

The Enhanced DAD algorithm from draft-ietf-6man-enhanced-dad is designed to detect looped back DAD probes. A network interface of an IPv6 node SHOULD implement the Enhanced DAD algorithm.

4.5. Changes to RFC 4861

The following text is appended to the RetransTimer variable description in section 6.3.2 of [RFC4861]:

The RetransTimer may be overridden by a link-specific document if a node supports the Enhanced DAD algorithm.

The following text is appended to the Source Address definition in section 4.3 of [RFC4861]:

If a node has been configured to use the Enhanced DAD algorithm, an NS with an unspecified source address adds the Nonce option to the message and implements the state machine of the Enhanced DAD algorithm.

4.6. Changes to RFC 3971

The following text is changed in section 5.3.2 of [RFC3971]:

The purpose of the Nonce option is to make sure that an advertisement is a fresh response to a solicitation sent earlier by the node.

The new text is included below:

The purpose of the Nonce option is to make sure that an advertisement is a fresh response to a solicitation sent earlier by the node. The nonce is also used to detect looped back NS messages when the network interface performs DAD [RFC4862]. Detecting looped back DAD messages is covered by the Enhanced DAD algorithm as specified in draft-ietf-6man-enhanced-dad. In a mixed SEND environment with SEND and unsecured nodes, the lengths of the nonce used by SEND and unsecured nodes MUST be identical.

5. Actions to Perform on Detecting a Genuine Duplicate

As described in the paragraphs above, the nonce can also serve to detect genuine duplicates even when the network has potential for looping back ND messages. When a genuine duplicate is detected, the node follows the manual intervention specified in section 5.4.5 of [RFC4862]. However, in certain networks such as an access network, if the genuine duplicate matches the tentative or optimistic IPv6 address of a network interface of the access concentrator, automated actions are recommended.

One example of a type of access network is cable broadband deployment where the access concentrator is the first-hop IPv6 router to several hundred thousand broadband modems. The router also supports proxying of DAD messages. The network interface on the access concentrator initiates DAD for an IPv6 address and detects a genuine duplicate due to receiving an NS(DAD) or an NA message. On detecting such a

duplicate, the access concentrator logs a system management message,

drops the received ND message, and blocks the modem on whose layer-2 service identifier the NS(DAD) or NA message was received on.

The network described above follows a trust model where a trusted router serves un-trusted IPv6 host nodes. Operators of such networks have a desire to take automated action if a network interface of the trusted router has a tentative or optimistic address duplicated by a host. Any other network that follows the same trust model MAY use the automated actions proposed in this section.

6. Security Considerations

This document does not improve nor reduce the security posture of [RFC4862]. The nonce can be exploited by a rogue deliberately changing the nonce to fail the looped back detection specified by the Enhanced DAD algorithm. SEND is recommended to circumvent this exploit. Additionally, the nonce does not protect against the DoS caused by a rogue node replying by a fake NA to all DAD probes. SEND is recommended to circumvent this exploit also. Disabling DAD has an obvious security issue before a remote node on the link can issue reflected NS(DAD) messages. Again, SEND is recommended for this exploit. Source Address Validation Improvement (SAVI) [RFC6620] also protects against various attacks by on-link rogues.

7. IANA Considerations

None.

8. Acknowledgements

Thanks (in alphabetical order by first name) to Bernie Volz, Dmitry Anipko, Eric Levy-Abegnoli, Eric Vyncke, Erik Nordmark, Fred Templin, Jouni Korhonen, Michael Sinatra, Ole Troan, Pascal Thubert, Ray Hunter, Suresh Krishnan, and Tassos Chatzithomaoglou for their guidance and review of the document. Thanks to Thomas Narten for encouraging this work. Thanks to Steinar Haug and Scott Beuker for describing the use cases.

9. Normative References

[RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.

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

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, April 2006.

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

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

[RFC6620] Nordmark, E., Bagnulo, M., and E. Levy-Abegnoli, "FCFS SAVI: First-Come, First-Served Source Address Validation Improvement for Locally Assigned IPv6 Addresses", RFC 6620, May 2012.

Authors' Addresses

Rajiv Asati
Cisco Systems, Inc.
7025 Kit Creek road
Research Triangle Park, NC 27709-4987
USA

Email: rajiva@cisco.com
URI: http://www.cisco.com/

Hemant Singh
Cisco Systems, Inc.

1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: +1 978 936 1622
Email: shemant@cisco.com
URI: http://www.cisco.com/

Wes Beebee
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
USA

Phone: +1 978 936 2030
Email: wbeebee@cisco.com
URI: http://www.cisco.com/

Carlos Pignataro
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
USA

Email: cpignata@cisco.com
URI: http://www.cisco.com/

Eli Dart
Lawrence Berkeley National Laboratory
1 Cyclotron Road, Berkeley, CA 94720
USA

Email: dart@es.net
URI: http://www.es.net/

Wes George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
USA

Email: wesley.george@twcable.com