

Optimistic Duplicate Address Detection (DAD) for IPv6

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

Optimistic Duplicate Address Detection is an interoperable modification of the existing IPv6 Neighbor Discovery (RFC 2461) and Stateless Address Autoconfiguration (RFC 2462) processes. The intention is to minimize address configuration delays in the successful case, to reduce disruption as far as possible in the failure case, and to remain interoperable with unmodified hosts and routers.

Table of Contents

1. Introduction	3
1.1. Problem Statement	3
1.2. Definitions	4
1.3. Address Types	4
1.4. Abbreviations	5
2. Optimistic DAD Behaviors	6
2.1. Optimistic Addresses	6
2.2. Avoiding Disruption	6
2.3. Router Redirection	7
2.4. Contacting the Router	7
3. Modifications to RFC-Mandated Behavior	8
3.1. General	8
3.2. Modifications to RFC 2461 Neighbor Discovery	8
3.3. Modifications to RFC 2462 Stateless Address Autoconfiguration	9
4. Protocol Operation	10
4.1. Simple Case	10
4.2. Collision Case	10
4.3. Interoperation Cases	11
4.4. Pathological Cases	11
5. Security Considerations	12
Appendix A. Probability of Collision	13
A.1. The Birthday Paradox	13
A.2. Individual Moving Nodes	14
Normative References	15
Informative References	15
Acknowledgements	16

1. Introduction

Optimistic Duplicate Address Detection (DAD) is a modification of the existing IPv6 Neighbor Discovery (ND) [RFC2461] and Stateless Address Autoconfiguration (SLAAC) [RFC2462] processes. The intention is to minimize address configuration delays in the successful case, and to reduce disruption as far as possible in the failure case.

Optimistic DAD is a useful optimization because in most cases DAD is far more likely to succeed than fail. This is discussed further in Appendix A. Disruption is minimized by limiting nodes' participation in Neighbor Discovery while their addresses are still Optimistic.

It is not the intention of this memo to improve the security, reliability, or robustness of DAD beyond that of existing standards, but merely to provide a method to make it faster.

1.1. Problem Statement

The existing IPv6 address configuration mechanisms provide adequate collision detection mechanisms for the fixed hosts they were designed for. However, a growing population of nodes need to maintain continuous network access despite frequently changing their network attachment. Optimizations to the DAD process are required to provide these nodes with sufficiently fast address configuration.

An optimized DAD method needs to:

- * provide interoperability with nodes using the current standards.
- * remove the RetransTimer delay during address configuration.
- * ensure that the probability of address collision is not increased.
- * improve the resolution mechanisms for address collisions.
- * minimize disruption in the case of a collision.

It is not sufficient to merely reduce RetransTimer in order to reduce the handover delay, as values of RetransTimer long enough to guarantee detection of a collision are too long to avoid disruption of time-critical services.

1.2. Definitions

Definitions of requirements keywords ('MUST NOT', 'SHOULD NOT', 'MAY', 'SHOULD', 'MUST') are in accordance with the IETF Best Current Practice, RFC 2119 [RFC2119]

Address Resolution - Process defined by [RFC2461], section 7.2.

Neighbor Unreachability Detection (NUD) - Process defined by [RFC2461], section 7.3.

Standard Node - A Standard Node is one that is compliant with [RFC2461] and [RFC2462].

Optimistic Node (ON) - An Optimistic Node is one that is compliant with the rules specified in this memo.

Link - A communication facility or medium over which nodes can communicate at the link layer.

Neighbors - Nodes on the same link, which may therefore be competing for the same IP addresses.

1.3. Address Types

Tentative address (as per [RFC2462]) - an address whose uniqueness on a link is being verified, prior to its assignment to an interface. A Tentative address is not considered assigned to an interface in the usual sense. An interface discards received packets addressed to a Tentative address, but accepts Neighbor Discovery packets related to Duplicate Address Detection for the Tentative address.

Optimistic address - an address that is assigned to an interface and available for use, subject to restrictions, while its uniqueness on a link is being verified. This memo introduces the Optimistic state and defines its behaviors and restrictions.

Preferred address (as per [RFC2462]) - an address assigned to an interface whose use by upper-layer protocols is unrestricted. Preferred addresses may be used as the source (or destination) address of packets sent from (or to) the interface.

Deprecated address (as per [RFC2462]) - An address assigned to an interface whose use is discouraged, but not forbidden. A Deprecated address should no longer be used as a source address in new communications, but packets sent from or to Deprecated addresses are delivered as expected. A Deprecated address may continue to be used as a source address in communications where switching to a Preferred address causes hardship to a specific upper-layer activity (e.g., an existing TCP connection).

1.4. Abbreviations

DAD - Duplicate Address Detection. Technique used for SLAAC. See [RFC2462], section 5.4.

ICMP Redirect - See [RFC2461], section 4.5.

NA - Neighbor Advertisement. See [RFC2461], sections 4.4 and 7.

NC - Neighbor Cache. See [RFC2461], sections 5.1 and 7.3.

ND - Neighbor Discovery. The process described in [RFC2461].

NS - Neighbor Solicitation. See [RFC2461], sections 4.3 and 7.

RA - Router Advertisement. See [RFC2462], sections 4.2 and 6.

RS - Router Solicitation. See [RFC2461], sections 4.1 and 6.

SLAAC - StateLess Address AutoConfiguration. The process described in [RFC2462].

SLLAO - Source Link-Layer Address Option - an option to NS, RA, and RS messages, which gives the link-layer address of the source of the message. See [RFC2461], section 4.6.1.

TLLAO - Target Link-Layer Address Option - an option to ICMP Redirect messages and Neighbor Advertisements. See [RFC2461], sections 4.4, 4.5, and 4.6.1.

2. Optimistic DAD Behaviors

This non-normative section discusses Optimistic DAD behaviors.

2.1. Optimistic Addresses

[RFC2462] introduces the concept of Tentative (in 5.4) and Deprecated (in 5.5.4) addresses. Addresses that are neither are said to be Preferred. Tentative addresses may not be used for communication, and Deprecated addresses should not be used for new communications. These address states may also be used by other standards documents, for example, Default Address Selection [RFC3484].

This memo introduces a new address state, 'Optimistic', that is used to mark an address that is available for use but that has not completed DAD.

Unless noted otherwise, components of the IPv6 protocol stack should treat addresses in the Optimistic state equivalently to those in the Deprecated state, indicating that the address is available for use but should not be used if another suitable address is available. For example, Default Address Selection [RFC3484] uses the address state to decide which source address to use for an outgoing packet. Implementations should treat an address in state Optimistic as if it were in state Deprecated. If address states are recorded as individual flags, this can easily be achieved by also setting 'Deprecated' when 'Optimistic' is set.

It is important to note that the address lifetime rules of [RFC2462] still apply, and so an address may be Deprecated as well as Optimistic. When DAD completes without incident, the address becomes either a Preferred or a Deprecated address, as per [RFC2462].

2.2. Avoiding Disruption

In order to avoid interference, it is important that an Optimistic Node does not send any messages from an Optimistic Address that will override its neighbors' Neighbor Cache (NC) entries for the address it is trying to configure: doing so would disrupt the rightful owner of the address in the case of a collision.

This is achieved by:

- * Clearing the 'Override' flag in Neighbor Advertisements for Optimistic Addresses, which prevents neighbors from overriding their existing NC entries. The 'Override' flag is already defined [RFC2461] and used for Proxy Neighbor Advertisement.

- * Never sending Neighbor Solicitations from an Optimistic Address. NSes include a Source Link-Layer Address Option (SLLAO), which may cause Neighbor Cache disruption. NSes sent as part of DAD are sent from the unspecified address, without a SLLAO.
- * Never using an Optimistic Address as the source address of a Router Solicitation with a SLLAO. Another address, or the unspecified address, may be used, or the RS may be sent without a SLLAO.

An address collision with a router may cause a neighboring router's IsRouter flags for that address to be cleared. However, routers do not appear to use the IsRouter flag for anything, and the NA sent in response to the collision will reassert the IsRouter flag.

2.3. Router Redirection

Neighbor Solicitations cannot be sent from Optimistic Addresses, and so an ON cannot directly contact a neighbor that is not already in its Neighbor Cache. Instead, the ON forwards packets via its default router, relying on the router to forward the packets to their destination. In accordance with RFC 2461, the router should then provide the ON with an ICMP Redirect, which may include a Target Link-Layer Address Option (TLLAO). If it does, this will update the ON's NC, and direct communication can begin. If it does not, packets continue to be forwarded via the router until the ON has a non-Optimistic address from which to send an NS.

2.4. Contacting the Router

Generally, an RA will include a SLLAO, however this "MAY be omitted to facilitate in-bound load balancing over replicated interfaces" [RFC2461]. A node with only Optimistic Addresses is unable to determine the router's Link-Layer Address as it can neither send an RS to request a unicast RA, nor send an NS to request an NA. In this case, the ON will be unable to communicate with the router until at least one of its addresses is no longer Optimistic.

3. Modifications to RFC-Mandated Behavior

All normative text in this memo is contained in this section.

3.1. General

- * Optimistic DAD SHOULD only be used when the implementation is aware that the address is based on a most likely unique interface identifier (such as in [RFC2464]), generated randomly [RFC3041], or by a well-distributed hash function [RFC3972] or assigned by Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [RFC3315]. Optimistic DAD SHOULD NOT be used for manually entered addresses.

3.2. Modifications to RFC 2461 Neighbor Discovery

- * (modifies section 6.3.7) A node MUST NOT send a Router Solicitation with a SLLAO from an Optimistic Address. Router Solicitations SHOULD be sent from a non-Optimistic or the Unspecified Address; however, they MAY be sent from an Optimistic Address as long as the SLLAO is not included.
- * (modifies section 7.2.2) A node MUST NOT use an Optimistic Address as the source address of a Neighbor Solicitation.
- * If the ON isn't told the SLLAO of the router in an RA, and it cannot determine this information without breaching the rules above, it MUST leave the address Tentative until DAD completes despite being unable to send any packets to the router.
- * (modifies section 7.2.2) When a node has a unicast packet to send from an Optimistic Address to a neighbor, but does not know the neighbor's link-layer address, it MUST NOT perform Address Resolution. It SHOULD forward the packet to a default router on the link in the hope that the packet will be redirected. Otherwise, it SHOULD buffer the packet until DAD is complete.

3.3 Modifications to RFC 2462 Stateless Address Autoconfiguration

- * (modifies section 5.5) A host MAY choose to configure a new address as an Optimistic Address. A host that does not know the SLLAO of its router SHOULD NOT configure a new address as Optimistic. A router SHOULD NOT configure an Optimistic Address.
- * (modifies section 5.4.2) The host MUST join the all-nodes multicast address and the solicited-node multicast address of the Tentative address. The host SHOULD NOT delay before sending Neighbor Solicitation messages.
- * (modifies section 5.4) The Optimistic Address is configured and available for use on the interface immediately. The address MUST be flagged as 'Optimistic'.
- * When DAD completes for an Optimistic Address, the address is no longer Optimistic and it becomes Preferred or Deprecated according to the rules of RFC 2462.
- * (modifies section 5.4.3) The node MUST NOT reply to a Neighbor Solicitation for an Optimistic Address from the unspecified address. Receipt of such an NS indicates that the address is a duplicate, and it MUST be deconfigured as per the behaviour specified in RFC 2462 for Tentative addresses.
- * (modifies section 5.4.3) The node MUST reply to a Neighbor Solicitation for an Optimistic Address from a unicast address, but the reply MUST have the Override flag cleared (O=0).

4. Protocol Operation

This non-normative section provides clarification of the interactions between Optimistic Nodes, and between Optimistic Nodes and Standard Nodes.

The following cases all consider an Optimistic Node (ON) receiving a Router Advertisement containing a new prefix and deciding to autoconfigure a new Optimistic Address on that prefix.

The ON will immediately send out a Neighbor Solicitation to determine if its new Optimistic Address is already in use.

4.1. Simple Case

In the non-collision case, the Optimistic Address being configured by the new node is unused and not present in the Neighbor Caches of any of its neighbors.

There will be no response to its NS (sent from ::), and this NS will not modify the state of neighbors' Neighbor Caches.

The ON already has the link-layer address of the router (from the RA), and the router can determine the link-layer address of the ON through standard Address Resolution. Communications can begin as soon as the router and the ON have each other's link-layer addresses.

After the appropriate DAD delay has completed, the address is no longer Optimistic, and becomes either Preferred or Deprecated as per RFC 2462.

4.2. Collision Case

In the collision case, the Optimistic Address being configured by the new node is already in use by another node, and present in the Neighbor Caches (NCs) of neighbors that are communicating with this node.

The NS sent by the ON has the unspecified source address, ::, and no SLLAO. This NS will not cause changes to the NC entries of neighboring hosts.

The ON will hopefully already know all it needs to about the router from the initial RA. However, if it needs to it can still send an RS to ask for more information, but it may not include a SLLAO. This forces an all-nodes multicast response from the router, but will not disrupt other nodes' NCs.

In the course of establishing connections, the ON might have sent NAs in response to received NSes. Since NAs sent from Optimistic Addresses have O=0, they will not have overridden existing NC entries, although they may have resulted in a colliding entry being changed to state STALE. This change is recoverable through standard NUD.

When an NA is received from the collidee defending the address, the ON immediately stops using the address and deconfigures it.

Of course, in the meantime the ON may have sent packets that identify it as the owner of its new Optimistic Address (for example, Binding Updates in Mobile IPv6 [RFC3775]). This may incur some penalty to the ON, in the form of broken connections, and some penalty to the rightful owner of the address, since it will receive (and potentially reply to) the misdirected packets. It is for this reason that Optimistic DAD should be used only where the probability of collision is very low.

4.3. Interoperation Cases

Once the Optimistic Address has completed DAD, it acts exactly like a normal address, and so interoperation cases only arise while the address is Optimistic.

If an ON attempts to configure an address currently Tentatively assigned to a Standard Node, the Standard Node will see the Neighbor Solicitation and deconfigure the address.

If a node attempts to configure an ON's Optimistic Address, the ON will see the NS and deconfigure the address.

4.4. Pathological Cases

Optimistic DAD suffers from similar problems to Standard DAD; for example, duplicates are not guaranteed to be detected if packets are lost.

These problems exist, and are not gracefully recoverable, in Standard DAD. Their probability in both Optimistic and Standard DAD can be reduced by increasing the RFC 2462 DupAddrDetectTransmits variable to greater than 1.

This version of Optimistic DAD is dependent on the details of the router behavior, e.g., that the router includes SLLAOs in RAs and that the router is willing to redirect traffic for the ON. Where the router does not behave in this way, the behavior of Optimistic DAD inherently reverts to that of Standard DAD.

5. Security Considerations

There are existing security concerns with Neighbor Discovery and Stateless Address Autoconfiguration, and this memo does not purport to fix them. However, this memo does not significantly increase security concerns either.

Secure Neighbor Discovery (SEND) [RFC3971] provides protection against the threats to Neighbor Discovery described in [RFC3756]. Optimistic Duplicate Address Detection does not introduce any additional threats to Neighbor Discovery if SEND is used.

Optimistic DAD takes steps to ensure that if another node is already using an address, the proper link-layer address in existing Neighbor Cache entries is not replaced with the link-layer address of the Optimistic Node. However, there are still scenarios where incorrect entries may be created, if only temporarily. For example, if a router (while forwarding a packet) sends out a Neighbor Solicitation for an address, the Optimistic Node may respond first, and if the router has no pre-existing link-layer address for that IP address, it will accept the response and (incorrectly) forward any queued packets to the Optimistic Node. The Optimistic Node may then respond in an incorrect manner (e.g., sending a TCP RST in response to an unknown TCP connection). Such transient conditions should be short-lived, in most cases.

Likewise, an Optimistic Node can still inject IP packets into the Internet that will in effect be "spoofed" packets appearing to come from the legitimate node. In some cases, those packets may lead to errors or other operational problems, though one would expect that upper-layer protocols would generally treat such packets robustly, in the same way they must treat old and other duplicate packets.

Appendix A. Probability of Collision

In assessing the usefulness of Duplicate Address Detection, the probability of collision must be considered. Various mechanisms such as SLAAC [RFC2462] and DHCPv6 [RFC3315] attempt to guarantee the uniqueness of the address. The uniqueness of SLAAC depends on the reliability of the manufacturing process (so that duplicate L2 addresses are not assigned) and human factors if L2 addresses can be manually assigned. The uniqueness of DHCPv6-assigned addresses relies on the correctness of implementation to ensure that no two nodes can be given the same address.

"Privacy Extensions to SLAAC" [RFC3041] avoids these potential error cases by picking an Interface Identifier (IID) at random from 2^{62} possible 64-bit IIDs (allowing for the reserved U and G bits). No attempt is made to guarantee uniqueness, but the probability can be easily estimated, and as the following discussion shows, probability of collision is exceedingly small.

A.1. The Birthday Paradox

When considering collision probability, the Birthday Paradox is generally mentioned. When randomly selecting k values from n possibilities, the probability of two values being the same is:

$$Pb(n,k) = 1 - (n! / [(n-k)! \cdot n^k])$$

Calculating the probability of collision with this method is difficult, however, as one of the terms is $n!$, and $(2^{62})!$ is an unwieldy number. We can, however, calculate an upper bound for the probability of collision:

$$Pb(n,k) \leq 1 - ([(n-k+1)/n] ^ [k-1])$$

which lets us calculate that even for large networks the probability of any two nodes colliding is very small indeed:

$$\begin{aligned} Pb(2^{62}, 500) &\leq 5.4e-14 \\ Pb(2^{62}, 5000) &\leq 5.4e-12 \\ Pb(2^{62}, 50000) &\leq 5.4e-10 \\ Pb(2^{62}, 500000) &\leq 5.4e-08 \end{aligned}$$

The upper-bound formula used above was taken from "Random Generation of Interface Identifiers", by M. Bagnulo, I. Soto, A. Garcia-Martinez, and A. Azcorra, and is used with the kind permission of the authors.

A.2. Individual Nodes

When considering the effect of collisions on an individual node, we do not need to consider the Birthday Paradox. When a node moves into a network with K existing nodes, the probability that it will not collide with any of the distinct addresses in use is simply $1-K/N$. If it moves to such networks M times, the probability that it will not cause a collision on any of those moves is $(1-K/N)^M$; thus, the probability of it causing at least one collision is:

$$P_c(n,k,m) = 1 - [(1-k/n)^m]$$

Even considering a very large number of moves ($m = 600000$, slightly more than one move per minute for one year) and rather crowded networks ($k=50000$ nodes per network), the odds of collision for a given node are vanishingly small:

$$\begin{aligned} P_c(2^{62}, 5000, 600000) &= 6.66e-10 \\ P_c(2^{62}, 50000, 600000) &= 6.53e-09 \end{aligned}$$

Each such collision affects two nodes, so the probability of being affected by a collision is twice this. Even if the node moves into networks of 50000 nodes once per minute for 100 years, the probability of it causing or suffering a collision at any point are a little over 1 in a million.

$$P_c(2^{62}, 50000, 60000000) * 2 = 1.3e-06$$

Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

Informative References

- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

Acknowledgements

There is some precedent for this work in expired Internet-Drafts and in discussions in the MobileIP WG mailing list and at IETF-54. A similar concept occurs in the 'Optimistic' bit used by R. Koodli and C. Perkins in the now expired, "Fast Handovers in Mobile IPv6".

Thanks to Greg Daley, Richard Nelson, Brett Pentland and Ahmet Sekercioglu at Monash University CTIE for their feedback and encouragement. More information is available at:

<http://www.ctie.monash.edu.au/ipv6/fastho/>

Thanks to all the MobileIP and IPng/IPv6 WG members who have contributed to the debate, especially and alphabetically: Jari Arkko, Marcelo Bagnulo, JinHyeock Choi, Youn-Hee Han, James Kempf, Thomas Narten, Pekka Nikander, Erik Nordmark, Soohong 'Daniel' Park, Mohan Parthasarathy, Ed Remmel, Pekka Savola, Hesham Soliman, Ignatious Souvatzis, Jinmei Tatuya, Dave Thaler, Pascal Thubert, Christian Vogt, Vladislav Yasevich, and Alper Yegin.

This work has been supported by the Australian Telecommunications Cooperative Research Centre (ATcrc):

<http://www.telecommunications.crc.org.au/>

Author's Address

Nick 'Sharkey' Moore
Centre for Telecommunications and Information Engineering
Monash University 3800
Victoria, Australia

Comments should be sent to sharkey@zoic.org and/or the IPv6 Working Group mailing list. Please include 'RFC4429' in the Subject line.

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).