

Optimistic Duplicate Address Detection for IPv6
<[draft-ietf-ipv6-optimistic-dad-02.txt](#)>

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, or will be disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Optimistic Duplicate Address Detection is an interoperable modification of the existing IPv6 Neighbor Discovery ([RFC2461](#)) and Stateless Address Autoconfiguration ([RFC2462](#)) process. 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.

Nick 'Sharkey' Moore

Expires: 09 April 2005

[Page 1]

INTERNET-DRAFT

Optimistic DAD

09 September 2004

Table of Contents

Status of this Memo	1
Abstract	1
Table of Contents	2
1 . Introduction	3
1.1 Problem Statement	3
1.2 History	4
1.3 Definitions	4
1.4 Abbreviations	5
2 . Optimistic Behaviours	6
2.1 Probability of Collision	6
2.2 Optimistic Address Flag	6
2.3 Avoiding Disruption	7
2.4 Router Redirection	7
3 . Modifications to RFC-compliant behaviour	8
3.1 Modifications to RFC 2461 Neighbor Discovery	8
3.2 Modifications to RFC 2462 SAA	9
4 . Protocol Operation	10
4.1 Simple case	10
4.2 Collision case	11
4.3 Interoperation cases	11
4.4 Pathological cases	12
5 . Security Considerations	12
6 . IANA Considerations	12
Normative References	13
Informative References	13
Author's Address	14
Acknowledgments	14

Full Copyright Statement	14
Intellectual Property Statement	15
Disclaimer of Validity	15

Nick 'Sharkey' Moore

Expires: 09 April 2005

[Page 2]

INTERNET-DRAFT

Optimistic DAD

09 September 2004

[1](#). Introduction

Optimistic Duplicate Address Detection (DAD) is a modification of the existing IPv6 Neighbor Discovery (ND) [[RFC2461](#)] and Stateless Address Autoconfiguration (SAA) [[RFC2462](#)] process. 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 DAD is far more likely to succeed than fail for a well-distributed random address [[SOTO](#)]. 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, 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 static 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 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](#) History

There is some precedent for this work in previous Internet Drafts [[KOODLI](#)], and in discussions in the MobileIP WG mailing list and at IETF-54. This version of Optimistic DAD differs somewhat from previous versions in that it uses no additional flags or message types beyond those already defined, therefore allowing interoperation between Optimistic and Standard nodes.

Earlier versions of this work were presented by the author to the MobileIP WG at IETF-56, and to the IPv6 WG at IETF-59. An issues list was presented at IETF-60.

Working implementations of draft versions of this memo have been made by the author as a freely-available patch to Linux 2.4.18, and by Ed Remmel of Elmic Systems.

An implementation of this version by the author is in progress, and will be released as a freely-available patch to Linux 2.6.7.

[1.3](#) Definitions

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

Address Resolution - Process defined by [\[RFC2461\] section 7.2](#).

Neighbor Unreachability Detection - Process defined by [\[RFC2461\] section 7.3](#).

Tentative Address - an address for which a node has not yet completed DAD is regarded as Tentative: a single Neighbor Solicitation for this address or a single Neighbor Advertisement defending this address will cause the node to deconfigure the address and cease using it.

Deprecated Address - an address which should not be used if an alternative is available.

Optimistic Address - an address which is available for use despite DAD not being fully complete. This memo places restrictions on the use of Optimistic Addresses.

Preferred Address - an address which is neither Tentative, Deprecated or Optimistic.

Nick 'Sharkey' Moore

Expires: 09 April 2005

[Page 4]

INTERNET-DRAFT

Optimistic DAD

09 September 2004

Optimistic Node - An Optimistic Node is one which is compliant with the rules specified in this memo.

Standard Node - A Standard Node is one which is compliant with RFCs 2461 and 2462.

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 addresses.

[1.4](#) Abbreviations

DAD - Duplicate Address Detection. Technique used for SAA. 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 [\[RFC2461\] section 5.1](#) and 7.3.

ND - Neighbor Discovery. The process described in [\[RFC2461\]](#)

NS - Neighbor Solicitation. See [\[RFC2461\]](#) sections [4.3](#) and [7](#).

ON - Optimistic Node. A node which is behaving according to the rules of this memo.

RA - Router Advertisement. See [\[RFC2462\]](#) sections [4.2](#) and [6](#).

RS - Router Solicitation. See [\[RFC2461\]](#) sections [4.1](#) and [6](#).

SAA - 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 Behaviours

This section provides some discussion of Optimistic DAD Behaviours. [Section 3](#) provides more specific information on changes to RFC-mandated behaviours.

[2.1](#) Probability of Collision

Optimistic DAD is only a useful optimization when the probability of collision is very small. As such, the Optimistic algorithm should not be used for manually assigned addresses, where the collision probability is likely to be much higher than that for random addresses due to human error.

Modifications are required only to Optimistic nodes -- Optimistic nodes will interoperate with Standard nodes without significant advantage or incompatibility.

[2.2](#) Optimistic Addresses

[RFC2462] introduces the concept of Tentative (in 5.4) and Deprecated (in 5.5.4) Addresses. Addresses which 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', which is used to mark an address which is available for use but which has not completed DAD. Protocols which do not understand this state should treat it equivalently to 'Deprecated', to indicate that the address is available for use but should not be used if another suitable address is available. If address states are recorded as individual flags, this can easily be achieved by setting 'Deprecated' when 'Optimistic' is set. In any case, 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 a Preferred or Deprecated address, as per [[RFC2462](#)].

INTERNET-DRAFT

Optimistic DAD

09 September 2004

[2.3](#) Avoiding Disruption

In order to avoid interference, it is important that an Optimistic node does not send any messages from an Optimistic Address which 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. NSs include a Source Link Layer Address Option (SLLAO), which may cause Neighbor Cache disruption. NSs sent as part of DAD are sent from the unspecified address, without a SLLAO.
- * Never using a 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 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.

[2.4](#) Router Redirection

When the ON wants to contact another neighbor, but it cannot because the neighbor is not in its NC, it should instead forward the packet to the router, relying on the router to forward the packet. 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. Implementing this behaviour may be difficult and unnecessary, so it is left as an option to the implementor.

[3.](#) Modifications to RFC-mandated behaviour

[3.1](#) Modifications to [RFC 2461](#) Neighbor Discovery

- * (modifies 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 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 wait until DAD completes despite being unable to

send any packets to the router.

- * (modifies 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.

[3.2](#) Modifications to [RFC 2462](#) Stateless Address Autoconfiguration

- * (modifies 5.5) A host MAY choose to configure a new address as an Optimistic Address. A host which does not know the SLLAO of its router SHOULD NOT configure a new address as Optimistic. A router MUST NOT configure an Optimistic Address.

- * (modifies 5.4) As soon as the initial Neighbor Solicitation is sent, the Optimistic Address is configured on the interface and available for use immediately. The address MUST be flagged as 'Optimistic'. Protocols which do not understand this state SHOULD treat it equivalently to 'Deprecated'.
- * When the DAD completes for an Optimistic Address, the address is no longer Optimistic and it becomes Preferred or Deprecated according to the rules of [[RFC2462](#)].
- * (modifies 5.4.3) The node MUST NOT reply to a Neighbor Solicitation for an Optimistic Address from the unspecified address. This NS indicates that the address is a duplicate, and it MUST be deconfigured as per the behaviour specified in [RFC2462](#) for Tentative addresses.
- * (modifies 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

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

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

[4.1](#) Simple case

In the non-collision case, the 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 Optimistic Node 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 others' link-layer addresses.

After the appropriate DAD delay has completed, the address is no longer Optimistic, and becomes either Preferred or Deprecated as per [RFC2462](#).

[4.2](#) Collision case

In the collision case, the address being configured by the new node is already in use by another node, and present in the Neighbor Caches (NCs) of neighbors which 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 a broadcast response from the router, but will not disrupt other nodes' NCs.

In the course of establishing connections, the ON may send NAs either spontaneously or in response to received NSs. Since these NAs will have O=0, they will not override existing NC entries, although they may result in a colliding entry being changed to state STALE. This change is recoverable through standard NUD.

Of course, in the meantime the ON may have sent packets which identify it as the owner of its new Optimistic Address (for example, Binding Updates in [[MIPv6](#)]). 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 only be used 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 Optimistic Node 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 Optimistic Node's Optimistic Address, the Optimistic Node 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 [RFC2462](#) DupAddrDetectTransmits variable to greater than 1.

This version of Optimistic DAD is dependant on the details of the router behaviour, eg: 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 behaviour 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.

Further work will be required to integrate Optimistic DAD with Secure Neighbor Discovery [[SEND](#)].

[6.](#) IANA Considerations

This document has no actions for IANA.

Normative References

- [RFC2119] S. Bradner. "Key words for use in RFCs to Indicate Requirement Levels." Request for Comments (Best Current Practice) [2119](#) ([BCP 14](#)), Internet Engineering Task Force, March 1997.
- [RFC2461] T. Narten, E. Nordmark, W. Simpson. "Neighbor Discovery for IP Version 6 (IPv6)." Request for Comments (Draft Standard) [2461](#), Internet Engineering Task Force, December 1998.
- [RFC2462] S. Thomson, T. Narten. "IPv6 Stateless Address Autoconfiguration." Request for Comments (Draft Standard) [2462](#), Internet Engineering Task Force, December 1998.

Informative References

- [RFC3484] R. Draves. "Default Address Selection for Internet Protocol version 6 (IPv6)". Request for Comments (Proposed Standard) [3484](#), Internet Engineering Task Force, February 2003.
- [MIPV6] D. Johnson, C. Perkins, J. Arkko. Mobility Support in IPv6, revision 24 ([draft-ietf-mobileip-ipv6-24](#)). June 2003 ... Expired December 2003.
- [K00DLI] R. Koodli, C. Perkins. Fast Handovers in Mobile IPv6, revision 00 ([draft-koodli-mobileip-fastv6-00](#)). October 2000 ... Expired April 2001.
- [SOTO] M. Bagnulo, I. Soto, A. Garcia-Martinez, A. Azcorra. Random generation of interface identifiers, revision 00. ([draft-soto-mobileip-random-ids-00](#)). January 2002 ... Expired July 2002.
- [SEND] J. Arkko, J. Kempf, B. Sommerfeld, B. Zill, P. Nikander.

SEcure Neighbor Discovery (SEND), revision 03. ([draft-ietf-send-ndopt-03](#)). January 2004 ... Expires July 2004.

Nick 'Sharkey' Moore Expires: 09 April 2005

[Page 13]

INTERNET-DRAFT

Optimistic DAD

09 September 2004

Author's Address:

Nick 'Sharkey' Moore
<nick.moore@eng.monash.edu.au> or <sharkey@zoic.org>
Centre for Telecommunications and Information Engineering
Monash University 3800
Victoria, Australia

Comments should be sent to either of the above email addresses.

Acknowledgments

Thanks to Greg Daley, Brett Pentland, Richard Nelson and Ahmet Sekercioglu at Monash Uni 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, JinHyeock Choi, Youn-Hee Han, James Kempf, Thomas Narten, Richard Nelson, Pekka Nikander, Erik Nordmark, Soohong 'Daniel' Park, Ed Rummel, Pekka Savola, Hesham Soliman, Ignatious Souvatzis, Jinmei

Tatuya, Dave Thaler, Pascal Thubert, Vladislav Yasevich and Alper Yegin.

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

[<http://www.telecommunications.crc.org.au/>](http://www.telecommunications.crc.org.au/)

Funding for the RFC Editor function is currently provided by the Internet Society.

Full Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the

IETF's procedures with respect to rights in IETF Documents can be found in [BCP 78](#) and 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 which may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

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.