Network Working Group
Internet-Draft
Expires: April 19, 2004

S. Roy
A. Durand
J. Paugh
Sun Microsystems, Inc.
October 20, 2003

IPv6 Neighbor Discovery On-Link Assumption Considered Harmful
draft-ietf-v6ops-onlinkassumption-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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/1id-abstracts.txt.

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

This Internet-Draft will expire on April 19, 2004.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This document proposes a change to the IPv6 Neighbor Discovery conceptual host sending algorithm. According to the algorithm, when a host's default router list is empty, the host assumes that all destinations are on-link. This document describes how making this assumption causes problems, and describes how these problems outweigh the benefits of this part of the conceptual sending algorithm.

Table of Contents

1. Introduction

Neighbor Discovery for IPv6 [ND] defines a conceptual sending algorithm for hosts. This algorithm states that if a host's default router list is empty, then the host assumes that all destinations are on-link.

This assumption creates problems for destination address selection as defined in [ADDRSEL], and adds connection delays associated with unnecessary address resolution and neighbor unreachability detection. The behavior associated with the assumption is undefined in multihomed scenarios, and has some subtle security implications. All of these issues are discussed in this document.

2. Background

This part of Neighbor Discovery's [ND] conceptual sending algorithm was created to facilitate communication on a single link between systems manually configured with different global prefixes. For example, two systems that are manually configured with global addresses while on separate links are then plugged in back-to-back. They can still communicate with each other via their global addresses because they'll correctly assume that each is on-link.

Without the on-link assumption, the above scenario wouldn't work as seamlessly. One workaround would be to use link-local addresses for this communication. Another is to configure new global addresses using the same /64 prefix on these systems, either by manually configuring such addresses, or by placing a router on-link that advertises this prefix.

3. Problems

The on-link assumption causes the following problems.

3.1 First Rule of Destination Address Selection

Default Address Selection for IPv6 [ADDRSEL] defines a destination address selection algorithm that takes an unordered list of destination addresses as input, and produces a sorted list of destination addresses as output. The algorithm consists of destination address sorting rules, the first of which is "Avoid unusable destinations". The idea behind this rule is to place unreachable destinations at the end of the sorted list so that applications will be least likely to try to communicate with those addresses first.

The unreachability determination for a destination as it pertains to this rule is an implementation detail. One implementable method is to do a simple forwarding table lookup on the destination, and to deem the destination as reachable if the lookup succeeds. The Neighbor Discovery on-link assumption makes this method somewhat irrelevant, however, as an implementation of the assumption could simply be to insert an IPv6 default on-link route into the system's forwarding table when the default router list is empty. The side-effect is that the rule would always determine that all IPv6 destinations are reachable.

On a network where there is no IPv6 router (all off-link IPv6 destinations are unreachable) and there is off-link IPv4 connectivity, the on-link assumption causes the rule to not necessarily prefer reachable IPv4 destinations over unreachable IPv6 destinations. This results in unreachable destinations being placed at the front of the sorted list.

3.2 Delays Associated with Address Resolution

Users expect that applications quickly connect to a given destination regardless of the number of IP addresses assigned to that destination. If a destination name resolves to multiple addresses and the application attempts to communicate to each address until one succeeds, this process shouldn't take an unreasonable amount of time. It is therefore important that the system quickly determine if IPv6 destinations are unreachable so that the application can try other destinations when those IPv6 destinations are unreachable.

For an IPv6 enabled host deployed on a network that has no IPv6 routers, the result of the on-link assumption is that link-layer address resolution must be performed on all IPv6 addresses to which

the host sends packets. The Application will not receive acknowledgment of the unreachability of destinations that are not on-link until at least address resolution has failed, which is no less than three seconds (MAX_MULTICAST_SOLICIT * RETRANS_TIMER) (amplified by transport protocol delays). When the application has a large list of off-link unreachable IPv6 addresses followed by at least one reachable IPv4 address, the delay associated with NUD of each IPv6 addresses before successful communication with the IPv4 address is unacceptable.

3.3 Multi-homing Ambiguity

There is no defined way to implement this aspect of the sending algorithm on a multi-homed node. From an implementor's point of view, there are three ways to handle sending an IPv6 packet to a destination in the face of the on-link assumption on a multi-homed node:

1. Attempt to resolve the destination on a single link.

2. Attempt to resolve the destination on every link.

3. Drop the packet.

If the destination is indeed on-link, the first option may not succeed since the wrong link could be picked. The second option would always succeed in reaching the destination (assuming that it's reachable) but is more complex to implement. Dropping the packet is equivalent to not making the on-link assumption at all. In other words, if there is no route to the destination, don't attempt to send the packet.

4. Conclusion

This document suggests the following changes to the Neighbor Discovery [ND] specification:

> The last sentence of the second paragraph of section 5.2 ("Conceptual Sending Algorithm") should be removed. This sentence is currently, "If the Default Router List is empty, the sender assumes that the destination is on-link.
>
> Bullet item 3) in section 6.3.6 ("Default Router Selection") should be removed. The item currently reads, "If the Default Router List is empty, assume that all destinations are on-link as specified in Section 5.2."

The result of these changes is that destinations are considered unreachable when there is no routing information for that destination (through a default router or otherwise). Instead of attempting link-layer address resolution when sending to such a destination, a node should send an ICMPv6 Destination Unreachable message (code 0 - no route to destination) message up the stack.

5. Security Considerations

The on-link assumption discussed here introduces a security vulnerability to the Neighbor Discovery protocol described in section 4.2.2 of IPv6 Neighbor Discovery Trust Models and Threats [PSREQ] titled "Default router is 'killed'". There is a threat that a host's router can be maliciously killed in order to cause the host to start sending all packets on-link. The attacker can then spoof off-link nodes by sending packets on the same link as the host. The vulnerability is discussed in detail in [PSREQ].

Another security related side-effect of the on-link assumption has to do with VPN's. It has been observed that some commercially available VPN software solutions that don't support IPv6 send IPv6 packets to the local media in the clear (their security policy doesn't simply drop IPv6 packets). Consider a scenario where a system has a single Ethernet interface with VPN software that encrypts and encapsulates certain packets. The system attempts to send a packet to an IPv6 destination that it obtained by doing a DNS lookup, and the packet ends up going in the clear to the local media. A malicious second party could then spoof the destination on-link.

Normative References

[ADDRSEL] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[ND] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[PSREQ] Nikander, P., Kempf, J. and E. Nordmark, "IPv6 Neighbor Discovery trust models and threats", draft-ietf-send-psreq-04, October 2003.

Informative References

[AUTOCONF]
Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

Authors' Addresses

Sebastien Roy
Sun Microsystems, Inc.
1 Network Drive
UBUR02-212
Burlington, MA 01801

EMail: sebastien.roy@sun.com

Alain Durand
Sun Microsystems, Inc.
17 Network Circle
UMPK17-202
Menlo Park, CA 94025

EMail: alain.durand@sun.com

James Paugh
Sun Microsystems, Inc.
17 Network Circle
UMPK17-202
Menlo Park, CA 94025

EMail: james.paugh@sun.com

Appendix A. Acknowledgments

The authors gratefully acknowledge the contributions of Jim Bound, Mika Liljeberg, Erik Nordmark, Pekka Savola, and Ronald van der Pol.

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 standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

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 practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgment

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