Network Working Group
Request for Comments: 4943
Category: Informational

S. Roy
Sun Microsystems, Inc.
A. Durand
Comcast
J. Paugh
Nominum, Inc.

September 2007

IPv6 Neighbor Discovery On-Link Assumption Considered Harmful

Status of This Memo

Abstract

This document describes the historical and background information behind the removal of the "on-link assumption" from the conceptual host sending algorithm defined in Neighbor Discovery for IP Version 6 (IPv6). According to the algorithm as originally described, when a host's default router list is empty, the host assumes that all destinations are on-link. This is particularly problematic with IPv6-capable nodes that do not have off-link IPv6 connectivity (e.g., no default router). This document describes how making this assumption causes problems and how these problems outweigh the benefits of this part of the conceptual sending algorithm.

Table of Contents

1. Introduction

Neighbor Discovery for IPv6 [RFC4861] defines a conceptual sending algorithm for hosts. The version of the algorithm described in [RFC2461] states that if a host's default router list is empty, then the host assumes that all destinations are on-link. This memo documents the removal of this assumption in the updated Neighbor Discovery specification [RFC4861] and describes the reasons why this assumption was removed.

This assumption is problematic with IPv6-capable nodes that do not have off-link IPv6 connectivity. This is typical when systems that have IPv6 enabled on their network interfaces (either on by default or administratively configured that way) are attached to networks that have no IPv6 services such as off-link routing. Such systems will resolve DNS names to AAAA and A records, and they may attempt to connect to unreachable IPv6 off-link nodes.

The on-link assumption creates problems for destination address selection as defined in [RFC3484], and it adds connection delays associated with unnecessary address resolution and neighbor unreachability detection. The behavior associated with the assumption is undefined on multi-interface nodes and has some subtle security implications. All of these issues are discussed in this document.

2. Background on the On-link Assumption

This part of Neighbor Discovery's [RFC2461] conceptual sending algorithm was created to facilitate communication on a single link between systems configured with different global prefixes in the absence of an IPv6 router. For example, consider the case where two systems on separate links are manually configured with global addresses and 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, and the systems would need to be configured to share a common prefix such as the link-local prefix. On the other hand, the on-link assumption introduces several problems to various parts of the networking stack described in Section 3. As such, this document points out that the problems introduced by the on-link assumption outweigh the benefit that the assumption lends to this scenario. It is more beneficial to the end user to remove the on-link assumption from Neighbor Discovery and declare this scenario illegitimate (or a misconfiguration).

3. Problems

The on-link assumption causes the following problems.

3.1. First Rule of Destination Address Selection

Default Address Selection for IPv6 [RFC3484] 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 on-link assumption could potentially cause false positives when attempting unreachability determination for this rule. On a network where there is no IPv6 router (all off-link IPv6 destinations are unreachable), the on-link assumption states that destinations are assumed to be on-link. An implementation could interpret that as, if the default router list is empty, then all destinations are reachable on-link. This may cause the rule to prefer an unreachable IPv6 destination over a reachable IPv4 destination.

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 3 seconds (MAX_MULTICAST_SOLICIT * RETRANS_TIMER). This is greatly amplified by transport protocol delays. For example, [RFC1122] Section 4.2.3.5 requires that TCP retransmit for at least 3 minutes before aborting the connection attempt.

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 Neighbor Unreachability Detection (NUD) of each IPv6 address before successful communication with the IPv4 address is unacceptable.

3.3. Multi-interface Ambiguity

There is no defined way to implement this aspect of the sending algorithm on a node that is attached to multiple links. Specifically, a problem arises when a node is faced with sending a packet to an IPv6 destination address to which it has no route, and the sending node is attached to multiple links. With the on-link assumption, this node assumes that the destination is on-link, but on which link? 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-interface node:

1. Attempt to send the packet on a single link (either administratively pre-defined or using some algorithm).

2. Attempt to send the packet on every link.

3. Drop the packet.

If the destination is indeed on-link, the first option might not succeed since the wrong link could be picked. The second option might succeed in reaching the destination but is more complex to implement and isn't guaranteed to pick the correct destination. For example, there could be more than one node configured with the same address, each reachable through a different link. The address by itself does not disambiguate which destination the sender actually wanted to reach, so attempting to send the packet to every link is not guaranteed to reach the anticipated destination. The third option, 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. An implementation that behaves this way would require an administrator to configure routes to the destination in order to have reachability to the destination, thus eliminating the ambiguity.

3.4. Security-Related Issues

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 [RFC3756] 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 [RFC3756].

Another security-related side-effect of the on-link assumption has to do with virtual private networks (VPNs). 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 third party could then spoof the destination on-link.

4. Changes to RFC 2461

The following changes have been made to the Neighbor Discovery specification between [RFC2461] and [RFC4861]:

- The last sentence of the second paragraph of Section 5.2 ("Conceptual Sending Algorithm") was removed. This sentence was, "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") was removed. The item read, "If the Default Router List is empty, assume that all destinations are on-link as specified in Section 5.2."

- APPENDIX A was modified to remove on-link assumption related text in bullet item 1) under the discussion on what happens when a multihomed host fails to receive Router Advertisements.

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 removal of the on-link assumption from Neighbor Discovery addresses all of the security-related vulnerabilities of the protocol as described in Section 3.4.

6. Normative References

[RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.

[RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[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.

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

Appendix A. Acknowledgments

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

Authors' Addresses

Sebastien Roy
Sun Microsystems, Inc.
1 Network Drive
UBUR02-212
Burlington, MA 01803

EMail: sebastien.roy@sun.com

Alain Durand
Comcast
1500 Market Street
Philadelphia, PA 19102

EMail: alain_durand@cable.comcast.com

James Paugh
Nominum, Inc.
2385 Bay Road
Redwood City, CA 94063

EMail: jim.paugh@nominum.com

Full Copyright Statement

Intellectual Property