6MAN WG
Internet-Draft
Expires: November 24, 2011

E. Nordmark
Cisco Systems, Inc.
May 23, 2011

Neighbor Unreachability Detection is too impatient
draft-nordmark-6man-impatient-nud-00.txt

Abstract

IPv6 Neighbor Discovery includes Neighbor Unreachability Detection. That function is very useful when a host has an alternative, for instance multiple default routers, since it allows the host to switch to the alternative in short time. This time is 3 seconds after the node starts probing. However, if there are no alternatives, this is far too impatient. This document proposes an approach where an implementation can choose the timeout behavior to be different based on whether or not there are alternatives.

Table of Contents

1. Introduction

IPv6 Neighbor Discovery [RFC4861] includes Neighbor Unreachability Detection, which detects when a neighbor is no longer reachable. The timeouts specified are very short (three transmissions spaced one second apart). That can be appropriate when there are alternative paths the packet can be sent. For example, if a host has multiple default routers in its Default Router List, or if the host has a Neigbor Cache Entry (NCE) created by a Redirect message. The effect of NUD reporting a failure in those cases is that the host will try the alternative; the next router in the Default Router List, or discard the NCE which will also send using a different router.

For that reason the timeouts where chosen to be short; this ensures that if a default router fails the host can use the next router in less than 45 seconds.

However, where there is no alternative there are several benefits in making NUD try probing for a longer time. One of those benefits is to be more robust against transient failures, such as spanning tree recovergence and other layer 2 issues that can take many seconds to resolve. Marking the NCE as unreachable in that case causes additional multicast on the network. Assuming there are IP packets to send, the lack of an NCE will result in multicast Neighbor Solicitations every second instead of the unicast Neighbor Solicitations that NUD sends.

As a result IPv6 is operationally more brittle than IPv4. For IPv4 there is no mandatory time limit on the retransmission behavior for ARP [RFC0826] which allows implementors to pick more robust schemes.

The following constant values in [RFC4861] seem to have been made part of IPv6 conformance testing: MAX_MULTICAST_SOLICIT, MAX_UNICAST_SOLICIT, RETRANS_TIMER. While such strict conformance testing seems consistent with the the specificiation, it means that we need to update the standard if we want to allow IPv6 Neighbor Discovery to be as operationally robust as ARP.

2. Proposed Remedy

We can clarify that the giving up after three packets spaced one second apart is only REQUIRED when there is an alternative, such as an additional default route or a redirect.

If implementations transmit more than MAX_*CAST_SOLICIT packets they MAY use binary exponential backoff of the retransmit timer. This is so that if we end up with implementations that try for a very long

time we don't end up with a steady background level of retransmissions.

3. Security Considerations

Relaxing the retransmission behavior for NUD has no impact on security. In particular, it doesn't impact applying Secure Neighbor Discovery [RFC3971].

4. IANA Considerations

This are no IANA considerations for this document.

5. References

5.1. Normative References

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

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

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

5.2. Informative References

[RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.

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

Author's Address

Erik Nordmark
Cisco Systems, Inc.
510 McCarthy Blvd.
Milpitas, CA, 95035
USA

Phone: +1 408 527 6625
Email: nordmark@cisco.com