Universität Karlsruhe (TH)
Institut für Telematik

TELEMATICS TECHNICAL REPORTS

Analysis of IPv6 Relocation Delays

Christian Vogt, Roland Bless, Mark Doll, Gregory Daley
{chvogt,bless,doll}@tm.uka.de, greg.daley@eng.monash.edu.au

July 18, 2005

TM-2005-4 (revised)

ISSN 1613-849X

http://doc.tm.uka.de/tr/

Institute of Telematics, University of Karlsruhe
Zirkel 2, D-76128 Karlsruhe, Germany

Abstract

As mobile nodes move to new links, they need to resume their communications in a timely fashion. To communicate, IPv6 nodes need to complete router discovery, address auto-configuration, and other tasks. Unfortunately, this depends on prior completion of the Multicast Listener Discovery (MLD) protocol, introducing a mandatory delay of up to one second. This document discusses where this can be problematic and proposes solutions that can alleviate the problem.

1 Introduction

Mobile nodes require fast router discovery and auto-configuration of global IPv6 addresses. New mechanisms defined in [1], [2], and [3] attend to this.

IPv6 Neighbor Discovery [1] accelerates router discovery in that it allows mobile nodes to forego the recommended backoff for the first Router Solicitation (RS) after interface (re-) initialization. Optimistic Duplicate Address Detection (ODAD) [2] allows early use of new IPv6 addresses once an initial Neighbor Solicitation (NS) has been sent. Nodes who support the Tentative Source Link-Layer Address Option (TSLLAO) [3] can save additional address-resolution round trips during router discovery and address auto-configuration.

Still, mobile nodes must send a MLD Report [4][5] at some stage during router discovery or address auto-configuration. MLD Reports are subject to a random backoff between 0 and MAX_RTR_SOLICITATION_DELAY (1 second) time RFC 2462bis [6]. There is currently no acceleration for this, defeating the benefits of the aforementioned optimizations in many situations.

This document identifies such situations and suggests two approaches to improve them: relaxing the delay for MLD Reports under certain conditions or permitting use of optimistic addresses prior to the initial NS. The document is considered a basis for further mailing list discussion. It is supposed to aid the revision process of existing DNA Working Group documents rather than to evolve itself towards a separate document.

2 History

This problem was first raised and discussed on the IPv6 mailing list [7]. The outcome was to not incorporate a delay relaxation for the MLD Report into RFC 2461bis and RFC 2462bis, as it was unclear then whether this could negatively impact other on-link nodes (both mobile and stationary). Instead, the consensus was to discuss this in greater detail in the DNA Working Group.

The problem appeared again in a similar context on the DNA mailing list [8].

3 Problem Statement

This section identifies five situations, A through E, in which IPv6 relocation is substantially delayed in spite of optimizations defined by RFC 2461bis, ODAD, and TSLLAO.

In all situations, after changing links, the mobile node avoids using its configured global unicast addresses during router discovery, since it does not know before reception of a Router Advertisement (RA) whether or not it has changed IPv6 attachment. Also, in order to avoid Neighbor Cache pollution of on-link neighbors, the mobile node must handle its configured link-local unicast addresses as if those where tentative.

It is also assumed that the mobile node implements ODAD. Nonetheless, the identified issues are essentially the same when the mobile node uses non-optimized RFC 2462bis.

3.1 Situation A

Routers on the new IP link support the TSLLAO for RSs. The mobile node solicits a unicast RA by sending an RS with an TSLLAO from the unspecified address.

The TSLLAO allows the router to unicast the RA without performing address resolution. What follows is the message exchange from the mobile node's perspective:

1. MN sends an RS from the unspecified address with an TSLLAO.
2. MN receives an IPv6-multicast RA by link-layer unicast.

3. MN sends an MLD Report for an optimistic address.

4. MN sends an NS for the optimistic address, initiating ODAD.

The RS (step 1) may be sent immediately, even though this is the first message transmitted after (re-) enabling the interface [1].

From a standard's perspective, it is debatable whether or not the MLD Report must be delayed. RFC 2462bis says in section 5.4.2:

> Even if the Neighbor Solicitation is not going to be the first message to be sent, the node SHOULD delay joining the solicited-node multicast address by a random delay between 0 and MAX_RTR_SOLICITATION_DELAY if the address being checked is configured by a router advertisement message sent to a multicast address.

The RA (step 2) is an IPv6 multicast, yet a link-layer unicast. Since there is only a single recipient in this case, omitting the delay for the MLD Report would be feasible. But the mobile node must inspect the link-layer destination address in order to determine whether the RA was a multicast or a unicast. This may not always be possible.

3.2 Situation B

Routers on the new IP link support the TSLLAO for RSs. The mobile node solicits a unicast RA by sending an RS with an TSLLAO from an optimistic address. Overall, it sends and receives the following messages in sequence:

1. MN sends an MLD Report for an optimistic address.

2. MN sends an NS for the optimistic address, initiating ODAD.

3. MN sends an RS with an TSLLAO from the optimistic address.

4. MN receives a unicast RA.

The MLD Report (step 1) is delayed for up to 1 second (MAX_RTR_SOLICITATION_DELAY) because it is the first message transmitted after (re-) enabling the interface [6].

3.3 Situation C

Routers on the new IP link do not support the TSLLAO for RSs. The mobile node solicits an RA. Routers unicast solicited RAs.

Soliciting a unicast RA requires the mobile node to send an RS from an optimistic address (without a SLLAO). This is the complete message exchange:

1. MN sends an MLD Report for an optimistic address.

2. MN sends an NS for the optimistic address, initiating ODAD.

3. MN sends an RS from the optimistic address.

4. MN receives an NS for the optimistic address from the router.

5. MN sends a Neighbor Advertisement (NA) for the optimistic address to the router.

6. MN receives a unicast RA.

The MLD Report (step 1) is delayed for up to 1 second (MAX_RTR_SOLICITATION_DELAY) because it is the first message transmitted after (re-) enabling the interface [6].

3.4 Situation D

Routers on the new IP link do not support the TSLLAO for RSs, but do send unsolicited multicast RAs every 30ms to 70ms according to rules in [9]. The mobile node sends and receives the following messages in sequence:

1. MN receives a multicast RA.
2. MN sends an MLD Report for an optimistic address.
3. MN sends an initial NS for the optimistic address (ODAD).

The MLD Report (step 2) is delayed for two reasons. First, it is the first message transmitted after (re-) enabling the interface. This calls for a delay of up to 1 second (MAX_RTR_SOLICITATION_DELAY) [6]. Second, the MLD Report is sent in response to a multicast RA. This also calls for a delay of up to 1 seconds (MAX_RTR_SOLICITATION_DELAY) [6].

3.5 Situation E

Routers on the new IP link do not support the TSLLAO for RSs. The mobile node solicits a multicast RA.

The mobile node can send the RS from the unspecified address, eliminating the requirement to initiate ODAD at this point. Overall, it sends an receives the following messages:

1. MN sends an RS from the unspecified address.
2. MN receives a multicast RA.
3. MN sends an MLD Report for an optimistic address.
4. MN sends an NS for the optimistic address, initiating ODAD.

The RS (step 1) may be sent immediately, even though this is the first message transmitted after (re-) enabling the interface [1].

The multicast RA (step 2) is delayed for up to 3 seconds (maximum of MAX_RA_DELAY_TIME and MIN_DELAY_BETWEEN_RAS) [1].

In addition, the MLD Report (step 3) is delayed for up to 1 second (MAX_RTR_SOLICITATION_DELAY) because it is sent in response to a multicast RA [6].

4 Possible Solutions

The following approaches can be used to eliminate the high IPv6 relocation delays identified in section 3.

4.1 Eliminating the Delay for MLD Reports

The random backoff for MLD Reports is ”RECOMMENDED” in RFC 2461bis and RFC 2462bis to resolve contention amongst multiple neighbors when booting up simultaneously. It has little benefit when applied by a mobile node after IPv6 relocation.

It hence seems feasible to eliminate the backoff for MLD Reports after IPv6 relocation. In particular, both of the following two conditions would have to be met:

- The mobile node has received a trigger from its local link layer indicating that an interface, which was previously operational, has gone down and come up again.
- The MLD Report is either the first message transmitted on the new link or it is sent in response to a unicast RA indicating IPv6 relocation.

4.2 Using Optimistic Addresses Before the Initial NS

ODAD allows (limited) use of optimistic addresses after an initial NS has been sent. This NS must be preceeded by a MLD Report for the corresponding solicited-nodes multicast address. The random backoff for the MLD Report foils the benefits of ODAD.

Fortunately, it appears feasible to allow (limited) use of an optimistic address even before the initial NS is sent. The delay for the MLD Report can so be moved to a non-critical phase.

Note that mobile nodes would still have to send a MLD Report prior to sending a RS from an optimistic address in situation C: Since the RS does not include a SLLAO, the router will have to invoke address resolution for the optimistic address. Sending the MLD Report ensures that the mobile node can receive the NS in the presence of snooping switches. Note that RFC 2461bis currently does not require transmission of a MLD Report prior to the RS in case of interface (re-) initialization.

4.3 Increasing Robustness

The optimum desynchonization delays for signaling messages such as the MLD Report are highly application- and environment-specific. RFC 2461bis and RFC 2462bis are tailored towards stationary nodes, so the delays of up to MAX_RTR_SOLICITATION_DELAY (1 second) are acceptable and make sense. However, for mobile nodes, such delays will badly impact real-time applications like VoIP. It may be possible in rare deployment scenarios to hard-code application- and environment-specific behavior into the nodes. But this approach is infeasible in general, simply because the applications and environments are unknown a priori.

One solution to this problem is to differentiate between messages for which loss is recoverable and messages for which it is not. E.g., loss of an RS can be detected by the mobile node by lack of a corresponding RA. In this (unlikely) case, the mobile node loses some time, but will eventually retransmit the RS. On the other hand, loss of a MLD Report or NS sent during ODAD cannot be detected, because there is no expected response. This might lead to an address duplicate, which is currently not recoverable.

In this sense, a mobile node should retransmit an MLD Report and RS after an appropriate time period if it fails to receive a corresponding RA. For the purpose of ODAD, the mobile node should retransmit MLD Reports and NSs in background mode two, three, or more times. An optimistic address would then be considered unique only after none of these transmissions resulted in a response. The mobile node may use plenty of desynchronization delay without negatively affecting applications because all transmission occur in background mode.

Specifically, mobile nodes SHOULD retransmit multiple NSs during ODAD. Each NS SHOULD be preceded by a retransmitted MLD Report. Thus, robustness to loss of (undelayed) MLD Report can be increased.

5 Brief Analysis

The conditions defined in section 4.1 should be conservative enough to limit the proposed optimization to non-critical situations. In particular, the random backoff for MLD Reports should be retained for responses to multicast RAs as well as when the node boots up. This assumes that a reliable indication for these conditions can be implemented.

Generally, collisions of MLD and IPv6 Neighbor Discovery messages are more likely to occur in a mobile environment than in a stationary because mobile nodes use these messages for movement detection and IPv6 relocation. Delaying the MLD Report does not mitigate this, however.

Retransmitted MLD Reports and RSs, as described in section 4.3, may repair router-discovery failure in situation C due to a collided (undelayed) MLD Report. Similarly, retransmitting MLD Reports and NSs, during ODAD, is likely to resolve collisions and ensure correct operation of ODAD in situations A through E.

ODAD may still fail when neighbors located at different sides of a snooping switch simultaneously attempt to auto-configure the same IPv6 address. If one of the nodes' MLD Report collides, but its NS does not, that node will not receive the competing node's NS. The result is that the former "wins" this race condition and the latter "loses" it.

More serious is a situation in which both MLD Reports get lost. The MLD state will eventually be repaired by the retransmitted MLD Report of both the mobile node and the neighbor. But there will be two nodes on the link using the same address. The IPv6 protocol suite currently does not recover from configured duplicate addresses.

Transmission of MLD Reports can lead to undesired signaling overhead. Furthermore, with respect to router discovery and address auto-configuration, MLD Reports have a benefit only in the presence of snooping switches [10]. The crux is that a mobile node usually does not know a new link's topology at attachment time, so it seems that omission of MLD Reports would become feasible only with appropriate support from the link layer.

6 Security Considerations

The optimizations proposed in this document affect desynchronization delays and retransmission policies applied by mobile nodes. Malicious nodes can always choose their own delays and policies, of course. As a consequence, the optimizations proposed in this document do not imply security concerns in addition to those which already exist in the standard IPv6 protocol suite [1][6], in ODAD [2], or in TSLLAO [3].

7 Acknowledgments

Credits go to Gregory Daley for a valuable discussion on interactions between MLD and router discovery plus address auto-configuration. Thanks also to Jari Arkko, who thoroughly reviewed version 00 of this draft. Last, but not least, the authors appreciate the contributions of folks to the initial thread on the IPv6 mailing list [7].

References

[1] T. Narten, "Neighbor Discovery for IP version 6 (IPv6)," IETF Internet Draft draft-ietf-ipv6-2461bis-03.txt, work in progress, May 2005.

[2] N. Moore, "Optimistic Duplicate Address Detection for IPv6," IETF Internet Draft draft-ietf-ipv6-optimistic-dad-05.txt, work in progress, February 2005.

[3] G. Daley, "Tentative Source Link-Layer Address Options for IPv6 Neighbour Discovery," IETF Internet Draft draft-daley-ipv6-tsllao-01.txt, work in progress, February 2005.

[4] S. E. Deering, W. C. Fenner, and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6," RFC 2710, October 1999.

[5] R. Vida, L. H. M. K. Costa, S. Fdida, S. Deering, B. Fenner, I. Kouvelas, and B. Haberman, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6," RFC 3810, June 2004.

[6] S. Thomson, "IPv6 Stateless Address Autoconfiguration," IETF Internet Draft draft-ietf-ipv6-rfc2462bis-08.txt, work in progress, May 2005.

[7] "Comment on RFCs 2461bis and 2462bis," Discussion on the IPv6 mailing list, May 27, 2005. [Online]. Available: http://www1.ietf.org/mail-archive/web/ipv6/current/msg05084.html

[8] "DAD and MLD Interaction," Discussion on the DNA mailing list, June 13, 2005. [Online]. Available: http://webcamserver.eng.monash.edu.au/~warchive/dna/2005-06/msg00098.html

[9] D. Johnson, C. E. Perkins, and J. Arkko, "Mobility Support in IPv6," RFC 3775, June 2004.

[10] K. Kimball and F. Solensky, "Considerations for IGMP and MLD Snooping Switches," IETF Internet Draft draft-ietf-magma-snoop-12.txt, work in progress, February 2005.