DNA WG
Internet-Draft
Expires: October 23, 2005

JinHyeock. Choi
Samsung AIT
Erik. Nordmark
Sun
April 21, 2005

DNA solution framework
draft-ietf-dna-soln-frame-00.txt

Status of this Memo

Copyright Notice

Abstract

This draft captures the authors' opinions and is intended to serve as input to the discussion for the solution in the DNA WG. It presents a few assumptions for DNA solution. The draft proposes the solution to be based on 1) link identity, 2) RS/RA exchange formed so that a host can determine whether it has moved from a single RA, and 3) Quick delivery of an RA. The draft sketches what rough shape DNA solution could take, including the necessary interaction with

Duplicate Address Detection and the Multicast Listener Discovery protocol. It also enumerate a few tasks to be worked on and issues to be resolved for efficient DNA solution.

Table of Contents

1. DNA Overview

Upon establishing a new link-layer connection, a host implementing the DNA solution should detect the identity of the currently attached link to ascertain whether it is attached to the same link as before, or attached to a different link. If the host is attached to a different link, it also needs to acquire the IP configuration for the new link. The DNA solution needs to be fast, precise, secure and have little signaling overhead.[4]

2. Basic Assumptions

We propose to design DNA solution based upon the following assumptions.

2.1 DNA solution based on link identity detection

When a host establishes a new link-layer connection, in order to check whether its IP configuration is still valid, the host checks for link change, i.e. determine whether it still remains attached to the same link or not. The term 'link' used in this document is as defined in [1]. NOTE that that definition is completely different than the definition of the term 'link' in IEEE 802 standards.

If a link change has occurred, a host assumes that its IP configuration is no longer valid. Thus it needs at least a new default router and a new IP address. If it has remained attached to the same link, the host assumes its IP configuration is still valid, and performs no further DNA operations.

2.2 Using a RS/RA exchange to determine the link identity

A Router Advertisement message is necessary when the host has attached to a different link, since the RA contains the new configuration information. This means that the number of messages needed for DNA can be minimized if the Router Advertisement message also contains all the information needed to determine whether or not there was a link change. In the general case, the host needs to send a Router Solicitation message so that it can quickly receive a Router Advertisement. Hence we end up with a suggested approach based on using a single RS/RA exchange to both determine the link identity, and to provide the host with the configuration information for a new link.

See [5] for the DTs different approaches to handle this.

In order for a single RA message to be useful both to detect a link change, and, should the host have attached to a different link, useful to initiate the new IP configuration, the RA message needs to include at least:

1) The information to indicate link change

2) Router address (to select new default router, in case of a link change)

3) Prefix(es) (to form a new IP address, in case of a link change)

4) Link-layer address of a router (to immediately send a packet)

We need to investigate whether the above is sufficient.

2.3 Quick delivery of an RA

To quickly check for link change, a host has to receive an RA with minimum latency. This is difficult due to the random delays for RAs in response to RSs and rate limiting of multicast RAs in Neighbor Discovery [1]. For fast DNA solution, we need to find a way to quickly deliver an RA to a host upon a new link-layer connection.

3. DNA Solution Sketch

3.1 Solution components

For efficient DNA solution, we may need the following components.

1. RA message optimized for DNA, which 1) properly indicate link change and 2) carries necessary information for a new IP configuration.

2. A way to quickly deliver an RA to a host upon a new link-layer connection.

3. Optimistic DAD [17] and TSLLAO [18] that is being specified in the IPv6 WG.

4. A procedure to apply DAD during the DNA procedure that is both efficient and safe should there be a duplicate.

5. A procedure for MLD so that the multicast groups are reported on a new link.

6. A procedure that handles DHCPv6 address (and other) configuration for those cases when stateless address autoconfiguration is not used.

3.2 Solution procedure

With the above, the DNA procedure might be as follows:

Step [0] Network attachment

A host has established a new link-layer connection.

Step [1] Hint

The host receives a hint that a link change might have occurred. This triggers the host to initiate DNA procedure. For instance, the hint might consist of the link-layer (device driver) providing a link Up event notification to the IP layer of the host.

Since the host doesn't know whether it is still attached to the same link, it needs to take the conservative approach and assumes it might have moved. Thus it switches to operating in optimistic DAD mode [17] at this point in time (but since it might still be on the same link, it would be overkill to immediately send a DAD probe). Since there might be MLD snooping switches in the network, the host must

use MLD to join, at least the solicited node multicast addresses that correspond to its IP addresses, at this point in time, so that it can receive Neighbor Solicitation messages that might indicate that an address is a duplicate.

Step [2] RA acquisition

The host acquires an RA optimized for DNA with minimum latency. This procedure may be initiated either by the host or network.

Either an AP (which implements [13]) immediately sends an RA to the host, or the host sends an RS to all-routers and one or more routers on the link responds to the RS with an RA. The first RA from the router(s) should not be delayed. Several approaches to accomplish this have been considered by the design team - see [5].

The RS should have the link-local address of the host as the source address. The RS message needs to contain the TSLLA option [18] to allow for optimal delivery of the RA in the case when the router supports TSLLAO, but should not contain a SLLAO option, since the link-local address might be a duplicate and DAD has not yet been completed.

If the router implements TSLLAO, the RA would be unicast to the host; otherwise the RA would either be multicast to all-nodes, or the router would perform an NS/NA exchange with the host before unicasting the RA to the host.

Step [3] Link identity detection

Using the mechanism which will be selected by the WG (see [5] for ones that are being considered), the host determines whether this is the same link as before.

If it is the same, the host assumes that its IP configuration is still valid. No further DNA operation is performed and all the host needs to do is turn off the optimistic DAD mode. (Since the host didn't move to a different link, we can rely on the DAD which was performed when the host was first attached to this link. However, there has been some discussion whether or not DAD should be redone if a host, independently of DNA, has been disconnected from the link for some time.)

If the host determines that it is attached to a new link, it immediately initiates a new IP configuration. The RA contains enough information to discard old default routers and prefixes, and configure new ones. (Should there be no "addrconf" prefixes in the RA, the host would presumably use DHCPv6 for address assignment which

would take one or two more roundtrips.)

At this point in time it also makes sense for the host to perform DAD by sending a DAD probe for each configured IP address. When the DAD probing has completed the host can turn off the optimistic DAD mode.

If neither the old link, nor the potentially new link, use the new DNA solution for identifying the link, then the host needs to use prefix based link determination [11] which might require multiple RAs and even multiple RSs being sent before it can determine whether or not it is attached to a new link.

Step [4] Multicast group reporting

Should the host have moved to a new link, it needs to send MLD reports for all the multicast groups it belongs to, in order to quickly re-establish reception for the multicast groups. (Note that the solicited-node multicast groups must be handled earlier - as part of the DAD procedure.) There might be cases when multicast reception is critical, where it would be beneficial to send the MLD reports earlier (during step 1) so that, in the case the host has moved to a new link, any interruption in multicast reception is minimized, even if this results in unneeded MLD report packets in the case when the host did not move.

3.3 Work items

It's our opinion that DNA WG (or IPv6 WG in the case of Optimistic DAD and TSLLAO) needs to work on the following subparts of the DNA problem:

1. A link identification mechanism from the ones identified in [5]

2. Complete Prefix List for the case when the above new mechanism is not available[11]

3. Immediate RA responses to RS from the ones identified in [5].

4. Optionally RAs that are sent by APs [13]

5. Optimistic DAD [17] and TSLLAO [18].

6. A procedure to apply DAD during the DNA procedure that is both efficient and safe should there be a duplicate.

7. A procedure for MLD so that the multicast groups are reported on a new link.

8. A procedure that triggers DHCPv6 address (and other) configuration for those cases when stateless address autoconfiguration is not used.

4. Issues

In this section, we enumerate the issues to be resolved for efficient DNA solution. We don't claim that the list is exhaustive.

4.1 Checking for link change with Link Identifier

[This discussion on this section pre-dates the design team discussion, thus it might no longer be relevant.]

Usually a host receives configuration information in one or more RA (Router Advertisement) messages. But it's difficult for a host to correctly check for link change using a single RA message. No information in RA can properly indicate whether a link change has occurred or not. Neither router address nor prefix can do.

It may be better to design a new way to represent a link, 'Link Identifier'. Each link has its unique and explicit Link Identifier and all routers attached to the link advertise the same Link Identifier in their RAs. With an explicit Link Identifier, an RA can represent a link identity and hosts can check for link change immediately without resorting to approximate knowledge.

When a host receives an RA with the same Link Identifier, it still remains at the same link. If it receives an RA with a different Link Identifier, a link change has occurred and the host is attached to a different link.

We need to investigate an efficient method to design an explicit Link Identifier. We may define a new option for Link Identifier. In [9] Erik Nordmark proposed a new option, Location Indication Option, which can server as Link Identifier. Also Brett Pentland and all submitted a draft on Link Identifier [10].

Other approaches can also be used, and many have been discussed in the Design Team. What is important is that the host can tell whether it remains on the same link, or has moved to a different link, from the first Router Advertisement message it receives.

See [5] for different approaches being discussed in the Design Team on how to handle this.

4.2 RA optimized for DNA

To design RA message optimized for DNA, we need to consider what kinds of information it needs to contain. We already presented 4 necessary information in Section 2.2 and also it may be useful for an RA to include the following:

1') A global IP address of a router

2') All prefixes that the router advertises

4.3 Quick delivery of an RA upon link-layer connection

Upon a new link-layer connection, it may take too long to receive a RA. A host may passively wait until it receives a periodic RA or, with link-layer hint, actively send an RS message and receive a solicited RA in response.

For the first case, the time to receive a RA depends on RA advertisement interval and it may take many minutes. Even in the second case there is a delay. A router MUST wait random amount of time between 0 and 0.5 sec before replying an RA [1]. And if the router multicasts RAs in response to RSs, the MIN_DELAY_BETWEEN_RAS in [1] is also a potential problem which must be looked into. Otherwise this would add the worst-case delay of 3.5 seconds until an RA is received.

To remedy this, currently two methods are proposed, FRD [13] and FastRA [14], [15]. In FRD, an AP caches a suitable RA and sends it immediately upon a new link-layer association. FastRA [14] allows a router to send an RA without delay with some safety mechanism. [15] defines a secured mechanism that allows routers to make decisions about which router responds fastest, and additionally allows other routers to avoid random delays.

Also see [5] for different approaches to handle this that have been discussed in the Design Team.

4.4 Links without Link Identification support

A host may visit a link that doesn't support the new DNA solution for link identification. There are a few cases to consider.

1 Moving from one link using DNA link identification to another link using DNA link identification (and the link are identified as being different).

2 Moving from a link using DNA link identification to a link which is not using it.

3 Moving from a link not using DNA link identification to a link which is using it.

4 Moving from a link not using DNA link identification to a link which is also not using it.

In all those cases, the host needs to be able to perform efficient DNA.

If the host can always tell from a single RA whether or not the link is using DNA link identification, then the second and third case above are easy, because the host must have moved to a different link.

This approach requires that 1) all the routers on a link are configured uniformly to either use DNA link identification or not, and 2) all the RAs contain at least a bit to indicate when DNA link identification is used.

5. IANA Considerations

No new message formats or services are defined in this document.

6. Security Considerations

Because DNA schemes are based on Neighbor Discovery, its trust models and threats are similar to the ones presented in [8]. Nodes connected over wireless interfaces may be particularly susceptible to jamming, monitoring and packet insertion attacks. Use of SEND [7] to secure Neighbor Discovery are important in achieving reliable detection of network attachment. DNA schemes SHOULD incorporate the solutions developed in IETF SEND WG if available, where assessment indicates such procedures are required.

7. Acknowledgment

The authors wish to express our appreciation to Greg Daley, Thomas Narten, Pekka Nikander and Alper Yegin for their valuable feedback. Also Thanks to Samita Chakrabarti, Youn-Hee Han, Gabriel Montenegro, Nick Moore, Brett Pentland, Ed Remmell and Margaret Wasserman for their contributions to this draft.

8. References

8.1 Normative References

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

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

[3] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

[4] Choi, J., "Detecting Network Attachment in IPv6 Goals", draft-ietf-dna-goals-04 (work in progress), December 2004.

8.2 Informative References

[5] Pentland, B., "An Overview of Approaches to Detecting Network Attachment in IPv6", draft-dnadt-dna-discussion-00 (work in progress), February 2005.

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

[7] Arkko, J., Kempf, J., Sommerfeld, B., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", draft-ietf-send-ndopt-06 (work in progress), July 2004.

[8] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.

[9] Nordmark, E., "MIPv6: from hindsight to foresight?", draft-nordmark-mobileip-mipv6-hindsight-00 (work in progress), November 2001.

[10] Pentland, B., "Router Advertisement Link Identification for Mobile IPv6 Movement Detection", draft-pentland-mobileip-linkid-03 (work in progress), October 2004.

[11] Nordmark, E. and J. Choi, "DNA with unmodified routers: Prefix list based approach", draft-ietf-dna-cpl-00 (work in progress), April 2005.

[12] Choi, J., "Router Advertisement Issues for Movement Detection/ Detection of Network Attachment", draft-jinchoi-ipv6-cra-00 (work in progress), October 2003.

[13] Choi, J., "Fast Router Discovery with RA Caching", draft-jinchoi-dna-frd-00 (work in progress), July 2004.

[14] Kempf, J., Khalil, M., and B. Pentland, "IPv6 Fast Router Advertisement", draft-mkhalil-ipv6-fastra-05 (work in progress), July 2004.

[15] Daley, G., "Deterministic Fast Router Advertisement Configuration", draft-daley-dna-det-fastra-01 (work in progress), October 2004.

[16] Narayanan, S., "Recommendations to achieve efficient Router Reachability Detection in IPv6 networks", draft-narayanan-dna-rrd-01 (work in progress), February 2005.

[17] Moore, N., "Optimistic Duplicate Address Detection for IPv6", draft-ietf-ipv6-optimistic-dad-05 (work in progress), February 2005.

[18] Daley, G., "Tentative Source Link-Layer Address Options for IPv6 Neighbour Discovery", draft-daley-ipv6-tsllao-01 (work in progress), February 2005.

[19] Yegin, A., "Link-layer Event Notifications for Detecting Network Attachments", draft-ietf-dna-link-information-01 (work in progress), February 2005.

[20] Aboba, B., "Detection of Network Attachment (DNA) in IPv4", draft-ietf-dhc-dna-ipv4-11 (work in progress), April 2005.

Authors' Addresses

JinHyeock Choi
Samsung AIT
Communication & N/W Lab
P.O.Box 111 Suwon 440-600
KOREA

Phone: +82 31 280 9233
Email: jinchoe@samsung.com

Erik Nordmark
Sun Microsystems
17 Network Circle
Menlo Park, CA 94025
USA

Phone: +1 650 786 2921
Email: erik.nordmark@sun.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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