

## Goals of Detecting Network Attachment in IPv6

### Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2005).

### Abstract

When a host establishes a new link-layer connection, it may or may not have a valid IP configuration for Internet connectivity. The host may check for link change (i.e., determine whether a link change has occurred), and then, based on the result, it can automatically decide whether its IP configuration is still valid. During link identity detection, the host may also collect necessary information to initiate a new IP configuration if the IP subnet has changed. In this memo, this procedure is called Detecting Network Attachment (DNA). DNA schemes should be precise, sufficiently fast, secure, and of limited signaling.

### Table of Contents

1. Introduction .....	2
2. Problems in Detecting Network Attachment .....	3
2.1. Wireless Link Properties .....	3
2.2. Link Identity Detection with a Single RA .....	3
2.3. Delays .....	4
3. Goals for Detecting Network Attachment .....	5
3.1. Goals List .....	6
4. Security Considerations .....	6
5. Acknowledgements .....	7
6. References .....	8
6.1. Normative References .....	8
6.2. Informative References .....	8

## 1. Introduction

When a host has established a new link-layer connection, it can send and receive some IPv6 packets on the link, including those used for configuration. On the other hand, the host has Internet connectivity only when it is able to exchange packets with off-link destinations.

When a link-layer connection is established or re-established, the host may not know whether its existing IP configuration is still valid for Internet connectivity. A subnet change might have occurred when the host changed its point of attachment.

In practice, the host doesn't know which of its addresses are valid on the newly attached link. It also doesn't know whether its existing default router is on this link or whether its neighbor cache entries are valid. Correct configuration of each of these components is necessary in order to send packets on and off the link.

To examine the status of the existing configuration, a host may check whether a 'link change' has occurred. In this document, the term 'link' is as defined in RFC 2461 [1]. The notion 'link' is not identical with the notion 'subnet', as defined in RFC 3753 [2]. For example, there may be more than one subnet on a link, and a host connected to a link may be part of one or more of the subnets on the link.

Today, a link change necessitates an IP configuration change. Whenever a host detects that it has remained at the same link, it can usually assume its IP configuration is still valid. Otherwise, the existing one is no longer valid, and a new configuration must be acquired. Therefore, to examine the validity of an IP configuration, all that is required is that the host checks for link change.

In the process of checking for link change, a host may collect some of the necessary information for a new IP configuration, such as on-link prefixes. So, when an IP subnet change has occurred, the host can immediately initiate the process of getting a new IP configuration. This may reduce handoff delay and minimize signaling.

Rapid attachment detection is required for a device that changes subnet while having on-going sessions. This may be the case if a host is connected intermittently, is a mobile node, or has urgent data to transmit upon attachment to a link.

Detecting Network Attachment (DNA) is the process by which a host collects the appropriate information and detects the identity of its currently attached link to ascertain the validity of its IP configuration.

DNA schemes are typically run per interface. When a host has multiple interfaces, the host separately checks for link changes on each interface.

It is important to note that DNA process does not include the actual IP configuration procedure. For example, with respect to DHCP, the DNA process may determine that the host needs to get some configuration information from a DHCP server. However, the process of actually retrieving the information from a DHCP server falls beyond the scope of DNA.

This document considers the DNA procedure only from the IPv6 point of view, unless explicitly mentioned otherwise. Thus, the term "IP" is to be understood to denote IPv6, by default. For the IPv4 case, refer to [7].

## 2. Problems in Detecting Network Attachment

A number of issues make DNA complicated. First, wireless connectivity is not as clear-cut as wired connectivity. Second, it's difficult for a single Router Advertisement (RA) message to indicate a link change. Third, the current Router Discovery specification specifies that routers wait a random delay of 0-.5 seconds prior to responding with a solicited RA. This delay can be significant and may result in service disruption.

### 2.1. Wireless Link Properties

Unlike in wired environments, what constitutes a wireless link is variable both in time and space. Wireless links do not have clear boundaries. This may be illustrated by the fact that a host may be within the coverage area of multiple (802.11) access points at the same time. Moreover, connectivity on a wireless link can be very volatile, which may make link identity detection hard. For example, it takes time for a host to check for link change. If the host ping-pongs between two links and doesn't stay long enough at a given link, it can't complete the DNA procedure.

### 2.2. Link Identity Detection with a Single RA

Usually, a host gets the information necessary for IP configuration from RA messages. Based on the current definition [1], it's difficult for a host to check for link change upon receipt of a single RA.

To detect link identity, a host may compare the information in an RA, such as router address or prefixes, with the locally stored information.

The host may use received router addresses to check for link change. The router address in the source address field of an RA is of link-local scope, however, so its uniqueness is not guaranteed outside a link. If it happens that two different router interfaces on different links have the same link-local address, the host can't detect that it has moved from one link to another by checking the router address in RA messages.

The set of all global prefixes assigned to a link can represent link identity. The host may compare the prefixes in an incoming RA with the currently stored ones. An unsolicited RA message, however, can omit some prefixes for convenience [1], and it's not easy for a host to attain and retain all the prefixes on a link with certainty. Therefore, neither the absence of a previously known prefix nor the presence of a previously unknown prefix in the RA guarantees that a link change has occurred.

### 2.3. Delays

The following issues cause DNA delay that may result in communication disruption.

#### 1) Delay for receiving a hint

A hint is an indication that a link change might have occurred. This hint itself doesn't confirm a link change, but initiates appropriate DNA procedures to detect the identity of the currently attached link.

Hints come in various forms and differ in how they indicate a possible link change. They can be link-layer event notifications [6], the lack of RA from the default router, or the receipt of a new RA. The time taken to receive a hint also varies.

As soon as a new link-layer connection has been made, the link layer may send a link-up notification to the IP layer. A host may interpret the new link-layer connection as a hint for a possible link change. With link-layer support, a host can receive such a hint almost instantly.

Mobile IPv6 [4] defines the use of RA Interval Timer expiry for a hint. A host keeps monitoring for periodic RAs and interprets the lack of them as a hint. It may implement its own policy to determine the number of missing RAs needed to interpret that as a hint. Thus, the delay depends on the Router Advertisement interval.

Without schemes such as those above, a host must receive a new RA from a new router to detect a possible link change. The detection time then also depends on the Router Advertisement frequency.

Periodic RA beaconing transmits packets within an interval varying randomly between MinRtrAdvInterval to MaxRtrAdvInterval seconds. Because a network attachment is unrelated to the advertisement time on the new link, hosts are expected to arrive, on average, halfway through the interval. This is approximately 1.75 seconds with Neighbor Discovery [1] advertisement rates.

## 2) Random delay execution for RS/RA exchange

Router Solicitation and Router Advertisement messages are used for Router Discovery. According to [1], it is sometimes necessary for a host to wait a random amount of time before it may send an RS, and for a router to wait before it may reply with an RA.

According to RFC 2461 [1], the following apply:

- Before a host sends an initial solicitation, it SHOULD delay the transmission for a random amount of time between 0 and MAX\_RTR\_SOLICITATION\_DELAY (1 second).
- Furthermore, any RA sent in response to a Router Solicitation MUST be delayed by a random time between 0 and MAX\_RA\_DELAY\_TIME (0.5 seconds).

## 3. Goals for Detecting Network Attachment

The DNA working group has been chartered to define an improved scheme for detecting IPv6 network attachment. In this section, we define the goals that any such solution should aim to fulfill.

DNA solutions should correctly determine whether a link change has occurred. Additionally, they should be sufficiently fast so that there would be no or at most minimal service disruption. They should neither flood the link with related signaling nor introduce new security holes.

When defining new solutions, it is necessary to investigate the usage of available tools, Neighbor Solicitation/Neighbor Advertisement messages, RS/RA messages, link-layer event notifications [6], and other features. This will allow precise description of procedures for efficient DNA Schemes.

### 3.1. Goals List

- G1 DNA schemes should detect the identity of the currently attached link to ascertain the validity of the existing IP configuration. They should recognize and determine whether a link change has occurred and initiate the process of acquiring a new configuration if necessary.
- G2 DNA schemes should detect the identity of an attached link with minimal latency lest there should be service disruption.
- G3 If a host has not changed a link, DNA schemes should not falsely assume a link change, and an IP configuration change should not occur.
- G4 DNA schemes should not cause undue signaling on a link.
- G5 DNA schemes should make use of existing signaling mechanisms where available.
- G6 DNA schemes should make use of signaling within the link (particularly link-local scope messages), because communication off-link may not be achievable in the case of a link change.
- G7 DNA schemes should be compatible with security schemes such as Secure Neighbor Discovery [3].
- G8 DNA schemes should not introduce new security vulnerabilities. The node supporting DNA schemes should not expose itself or other nodes on a link to additional man-in-the-middle, identity-revealing, or denial-of-service attacks.
- G9 Nodes (such as routers or hosts) that support DNA schemes should work appropriately with unmodified nodes that do not.
- G10 Hosts, especially in wireless environments, may perceive routers reachable on different links. DNA schemes should take into consideration the case where a host is attached to more than one link at the same time.

### 4. Security Considerations

The DNA process is intimately related to the Neighbor Discovery protocol [1] and its trust model and threats have much in common with those presented in RFC 3756 [5]. Nodes connected over wireless interfaces may be particularly susceptible to jamming, monitoring, and packet-insertion attacks.

With unsecured DNA schemes, it is inadvisable for a host to adjust its security based on which network it believes it is attached to. For example, it would be inappropriate for a host to disable its personal firewall because it believed that it had connected to a home network.

Even in the case where authoritative information (routing and prefix state) are advertised, wireless network attackers may still prevent soliciting nodes from receiving packets. This may cause unnecessary IP configuration change in some devices. Such attacks may be used to make a host preferentially select a particular configuration or network access.

Devices receiving confirmations of reachability (for example, from upper-layer protocols) should be aware that unless these indications are sufficiently authenticated, reachability may falsely be asserted by an attacker. Similarly, even if such reachability tests are known to originate from a trusted source, they should be ignored for reachability confirmation if the packets are not fresh or have been replayed. This may reduce the effective window for attackers replaying otherwise authentic data.

It may be dangerous to receive link-change notifications from the link layer and network layer, if they are received from devices that are insufficiently authenticated. In particular, notifications that authentication has completed at the link layer may not imply that a security relationship is available at the network layer. Additional authentication may be required at the network layer to justify modification of IP configuration.

## 5. Acknowledgements

Erik Nordmark has contributed significantly to work predating this document. Also Ed Remmell's comments on the inconsistency of RA information were most illuminating. The authors wish to express our appreciation to Pekka Nikander for valuable feedback. We gratefully acknowledge the generous assistance we received from Shubhranshu Singh for clarifying the structure of the arguments. Thanks to Brett Pentland, Nick Moore, Youn-Hee Han, JaeHoon Kim, Alper Yegin, Jim Bound, and Jari Arkko for their contributions to this document.

## 6. References

### 6.1. Normative References

- [1] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [2] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.
- [3] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

### 6.2. Informative References

- [4] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [5] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.
- [6] Yegin, A., "Link-layer Event Notifications for Detecting Network Attachments", work in progress, July 2005.
- [7] Aboba, B., "Detecting Network Attachment (DNA) in IPv4", work in progress, June 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

Greg Daley  
CTIE Monash University  
Centre for Telecommunications and Information Engineering  
Monash University  
Clayton 3800 Victoria  
Australia

Phone: +61 3 9905 4655  
EMail: greg.daley@eng.monash.edu.au



## Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgement

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