

DNA Working Group
Internet-Draft
Expires: December 25, 2005

S. Narayanan
Panasonic
G. Daley
Monash University CTIE
N. Montavont
LSIIT - ULP
June 23, 2005

Detecting Network Attachment in IPv6 - Best Current Practices for hosts.
[draft-ietf-dna-hosts-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on December 25, 2005.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

Hosts experiencing rapid link-layer changes may require efficient IP configuration change detection procedures than traditional fixed hosts. DNA is defined as the process by which a host collects appropriate information and detects the identity of its currently attached link to ascertain the validity of its IP configuration.

This document details best current practice for Detecting Network Attachment in IPv6 hosts using existing Neighbor Discovery procedures. Since there is no explicit link identification mechanism in the existing Neighbor Discovery for IP Version 6, the document describes implicit mechanisms for identifying the current link.

Table of Contents

1.	Introduction	5
1.1	Structure of this Document	5
2.	Terms and Abbreviations	5
3.	Background & Motivation for DNA	6
3.1	Issues	7
4.	Detecting Network Attachment Steps	7
4.1	Making use of Prior Information	7
4.2	Duplicate Address Detection	8
4.3	Link identification	9
4.3.1	Same link	9
4.3.2	Link change	10
4.4	Multicast Listener State	10
4.5	Reachability detection	10
5.	Initiation of DNA Procedures	11
5.1	Actions Upon Hint Reception	12
5.2	Hints Due to Network Layer Messages	12
5.3	Handling Hints from Other Layers	13
5.4	Timer and Loss Based Hints	13
5.5	Simultaneous Hints	14
5.6	Hint Validity and Hysteresis	14
5.7	Hint Management for Inactive Hosts	15
6.	IP Hosts Configuration	15
6.1	Router and Prefix list	15
6.2	IPv6 Addresses	16
6.2.1	Autoconfiguration	16
6.2.2	Dynamic Host Configuration	16
6.3	Neighbor cache	17
6.4	Mobility Management	17
7.	Complications to Detecting Network Attachment	18
7.1	Packet Loss	18
7.2	Router Configurations	18
7.3	Overlapping Coverage	18
7.4	Multicast Snooping	19
7.5	Link Partition	19
8.	Security Considerations	19
8.1	Authorization and Detecting Network Attachment	20
8.2	Addressing	20
9.	Constants	21

10.	Acknowledgments	21
11.	References	21
11.1	Normative References	21
11.2	Informative References	22
	Authors' Addresses	23
A.	Example State Transition Diagram	24
	Intellectual Property and Copyright Statements	25

1. Introduction

When operating in changing environments, IPv6 hosts may experience variations in reachability or configuration state over time. For hosts accessing the Internet over wireless media, such changes may be caused by changes in wireless propagation or host motion.

Detecting Network Attachment (DNA) in IPv6 is the task of checking for changes in the validity of a host's IP configuration [[15](#)]. Changes may occur on establishment or disconnection of a link-layer. For newly connected interfaces, they may be on a link different from the existing configuration of the node.

In these and other cases, IP addressing and default routing configuration of the node may become invalid preventing packet transfer. DNA uses IPv6 Neighbour Discovery to provide information about the reachability and identity of the link.

DNA focuses on determining whether the current configuration is valid, leaving the actual practice of re-configuration to other subsystems, if the current configuration is invalid.

This document presents the best current practices for IPv6 hosts to address the task of Detecting Network Attachment in changing and wireless environments.

1.1 Structure of this Document

[Section 3](#) of this document provides background and motivation for Detecting Network Attachment.

Elaboration of specific practices for hosts in detecting network attachment continues in [Section 4](#), while [Section 5](#) discuss the initiation of DNA procedures.

[Section 6](#) describes interactions with other protocols, particularly upon link-change, while [Section 7](#) describes environmental challenges to detection of network attachment.

[Section 8](#) provides security considerations, and details a number of issues which arise due to wireless connectivity and the changeable nature of DNA hosts' Internet connections.

2. Terms and Abbreviations

Access network: A network where hosts are present. Especially, a network used for the support of visiting wireless hosts.

Attachment: The process of entering a new cell. Attachment (and detachment) may cause a link-change.

Cell: A system constituted by the propagation range of a wireless base station and its serviced hosts. Dependent on topology, one or many cells may be part of the same link.

Hint: An indication from other subsystems or protocol layers that link-change may have occurred.

Link: A link is the range across which communications can pass without being forwarded through a router [1].

Link-Change: Link-Change occurs when a host moves from a point-of-attachment on a link, to another point-of-attachment where it is unable to reach devices belonging to a link, without being forwarded through a router.

Point-of-Attachment: A link-layer base-station, VLAN or port through which a device attempts to reach the network. Changes to a host's point-of-attachment may cause link-change.

Reachability Detection: Determination that a device (such as a router) is currently reachable, over both a wireless medium, and any attached fixed network. This is typically achieved using Neighbor Unreachability Detection procedure [1].

Wireless Medium: A physical layer which incorporates free space electromagnetic or optical propagation. Such media are susceptible to mobility and interference effects, potentially resulting in high packet loss probabilities.

3. Background & Motivation for DNA

Hosts on the Internet may be connected by various media. It has become common that hosts have access through wireless media and are mobile, and have a variety of interfaces through which internet connectivity is provided. The frequency of configuration change for wireless and nomadic devices are high due to the vagaries of wireless propagation or the motion of the hosts themselves. Detecting Network Attachment is a strategy to assist such rapid configuration changes by determining whether they are required.

Due to these frequent link-layer changes, an IP configuration change

detection mechanism for DNA needs to be efficient and rapid to avoid unnecessary configuration delays upon link-change.

In an wireless environment, there will typically be a trade-off between configuration delays and the channel bandwidth utilized or host's energy used to transmit packets. This trade-off affects choices as to whether hosts probe for configuration information, or wait for network information. DNA seeks to assist hosts by providing information about network state, which may allow hosts to appropriately make decisions regarding such trade-offs.

Even though DNA is restricted to determining whether change is needed, in some circumstances the process of obtaining information for the new configuration may occur simultaneously with the detection to improve the efficiency of these two steps.

3.1 Issues

The following features of [RFC 2461](#) make the detection of link identity difficult:

Routers are not required to include all the prefixes they support in a single router advertisement message [\[1\]](#).

The default router address is link-local address and hence may only be unique within one link [\[1\]](#).

While neighbor cache entries are valid only on a single link, link-local addresses may be duplicated across many links, and only global addressing can be used to identify if there is a link change.

4. Detecting Network Attachment Steps

4.1 Making use of Prior Information

A device that has recently been attached to a particular wireless base station may still have state regarding the IP configuration valid for use on that link. This allows a host to begin any configuration procedures before checking the ongoing validity and security of the parameters.

The experimental protocols FMIPv6 [\[19\]](#) and CARD [\[20\]](#) each provide ways to generate such information using network-layer signaling, before arrival on a link. Additionally, the issue is the same when a host disconnects from one cell and returns to it immediately, or movement occurs between a pair of cells (the ping-pong effect).

A IP host MAY store L2 to L3 mapping information for the links for a period of time in order to use the information in the future. When a host attaches itself to a point-of-attachment for which it has an L2 to L3 mapping, if the stored record doesn't contain the prefix the host is using, the host SHOULD conclude that it has changed link and initiate a new configuration procedure.

If the host finds the prefix it is using in the stored record, a host MAY conclude that it is on the same link, but SHOULD undertake reachability testing as described in [Section 4.5](#). In this case, the host MUST undertake Duplicate Address Detection [3][8] to confirm that there are no duplicate addresses on the link.

The host MUST age this cached information based on the possibility that the link's configuration has changed and MUST NOT store information beyond either the remaining router or address lifetime or (at the outside) MAC_CACHE_TIME time-units.

If the assumptions attached to the stored configuration are incorrect the configuration cost may be increased, or even cause disruption of services to other devices. Hosts MUST NOT cause any disruption of the IP connectivity to other devices due to the invalidity or staleness of their configuration.

[4.2](#) Duplicate Address Detection

When a host connects to a new link, it needs to create a link-local address. But to ensure that the link-local address is unique on a link, Duplication Address Detection (DAD) MUST be performed [3] by sending NS targeted at the link-local address undergoing validation.

Optimistic Duplicate Address Detection allows addresses to be used while they are being checked, without marking addresses as tentative. Procedures defined in optimistic DAD [8] ensure that persistent invalid neighbour cache entries are not created. This may allow faster DNA procedures, by avoiding use of unspecified source addresses in RS's and consequently allowing unicast Router Advertisements responses [8]. It is RECOMMENDED that hosts follow the recommendations of optimistic DAD [8] to reduce the DAD delay.

Link-local addresses SHOULD be treated as either optimistic or tentative, and globally unique addresses SHOULD NOT be used in a way which creates neighbor cache state on their peers, while DNA procedures are underway.

While hosts performing DNA do not know if they have arrived on a new link, they SHOULD treat their addresses as if they were. The different treatment of IP addressing comes from the fact that on the

global addresses cannot have an address conflict if they move to a topologically incorrect network where link-local addresses may. Even though global addresses will not collide, the incorrect creation of neighbor cache entries on legacy peers may cause them some harm.

In the case that the host has not changed link and if the time elapsed since the hint is less than the DAD completion time (minus a packet transmission and propagation time), the host MAY reclaim the address by sending Neighbor Advertisement messages as if another host had attempted DAD while the host was away. This will prevent DAD completion by another node upon NA reception.

If a host has not been present on a link to defend its address, and has been absent for a full DAD delay (minus transmission time) the host MUST undertake the full DAD procedure for each address from that link it wishes to configure [3][8].

4.3 Link identification

4.3.1 Same link

An IP host MUST conclude that it is on the same link if any of the following events happen.

- Reception of a RA with the prefix known to be on the link from one of its default router address, even if it is the link-local address of the router.

- Reception of a RA from a known router's global address, present in a Prefix Information Option containing the R-"Router Address" flag [5].

A host SHOULD conclude that it is on the same link if any of the following events happen.

- Reception of a RA with a prefix that is known to be on the current link.

- Reception of data packets addressed to its current global address if the message was sent from or through a device which is known to be fixed (such as a router).

- Confirmation of a global address entry with the Router 'R' flag set in its neighbor cache[1].

[4.3.2](#) Link change

A host makes use of Router Discovery messages to determine that it has moved to a new link. Since the content of an existing received RA is not sufficient to identify the absence of any other prefix, additional inference is required for fast and accurate link-change detection.

Complete Prefix Lists provide a robust mechanism for link-change detection with even unmodified non-DNA routers if link-layer indications are available [\[24\]](#)[\[18\]](#). These procedures provide mechanisms to build confidence that a host knows all of a link's prefixes and so may rapidly identify a newly received RA as being from a different link.

A host SHOULD maintain a complete prefix list as recommended by [\[24\]](#) to identify if the link has changed.

[4.4](#) Multicast Listener State

Multicast routers on a link are aware of which groups are in use within a link. This information is used to undertake initiation of multicast routing for off-link multicast sources to the link [\[9\]](#) [\[11\]](#).

When a node arrives on a link, it may need to send Multicast Listener Discovery (MLD) reports, if the multicast stream is not already being received on the link. If it is an MLDv2 node it SHOULD send state change reports upon arrival on a link [\[11\]](#).

Since the identity of the link is tied to the presence and identity of routers, initiation of these procedures may be undertaken when DNA procedures have been completed. In the absence of received data packets from a multicast stream, it is unlikely that a host will be able to determine if the multicast stream is being received on a new link, and will have to send state change reports, in addition to responses to periodic multicast queries [\[9\]](#) [\[11\]](#).

For link scoped multicast, reporting needs to be done to ensure that packet reception in the link is available due to multicast snoopers. Some interaction is possible when sending messages for the purpose of DNA on a network where multicast snooping is in use. This issue is described in [Section 7.4](#).

[4.5](#) Reachability detection

If an IP node needs to confirm bi-directional reachability to its

default router either a NS-NA or an RS-RA message exchange can be used to conduct reachability testing. It is notable that the choice of whether the messages are addressed to multicast or unicast address will have different reachability implications. The reachability implications from the hosts' perspective for the four different message exchanges defined by [RFC 2461](#) [1] are presented in the table below. The host can confirm bi-directional reachability from the neighbor discovery or router discovery message exchanges except when a multicast RA is received at the host for its RS message. In this case, using IPv6 Neighbour Discovery procedures, the host cannot know whether the multicast RA is in response to its solicitation message or whether it is a periodic un-solicited transmission from the router [1].

Exchanges:	Upstream	Downstream
multicast NS/NA	Y	Y
unicast NS/NA	Y	Y
RS/multicast RA	N	Y
RS/unicast RA	Y	Y

Successful exchange of the messages listed in the table indicate the corresponding links to be operational. Lack of reception of response from the router may indicate that reachability is broken for one or both of the transmission directions or it may indicate an ordinary packet loss event in either direction.

Link-change detection incorporates message reception which may itself create neighbour reachability state. When this is the case, a host SHOULD rely upon existing Neighbor Discovery procedures in order to provide and maintain reachability detection [1].

5. Initiation of DNA Procedures

Link change detection procedures are initiated when information is received either directly from the network or from other protocol layers within the host. This information indicates that network reachability or configuration is suspect and is called a hint.

Hints MAY be used to update a wireless host's timers or probing behavior in such a way as to assist detection of network attachment. Hints SHOULD NOT be considered to be authoritative for detecting IP configuration change by themselves.

In some cases, hints will carry significant information (for example a hint indicating PPP IPv6CP open state [4]), although details of the configuration parameters may be available only after other IP configuration procedures. Implementers are encouraged to treat hints as though they may be incorrect, and require confirmation.

Hosts MUST ensure that untrusted hints do not cause unnecessary configuration changes, or significant consumption of host resources or bandwidth. In order to achieve this aim, a host MAY implement hysteresis mechanisms such as token buckets, hint weighting or holddown timers in order to limit the effect of excessive hints (see [Section 5.6](#)).

5.1 Actions Upon Hint Reception

Upon reception of a hint that link change detection may have occurred, a host SHOULD send Router Solicitation messages to determine the routers and prefixes which exist on a link. Hosts SHOULD apply rate limiting and/or hysteresis to this behaviour as appropriate to the link technology subject to the reliability of the hints.

Router Advertisements received as a result of such solicitation have a role in determining if existing configuration is valid, and may be used to construct prefix lists for a new link [24].

5.2 Hints Due to Network Layer Messages

Hint reception may be due to network-layer messages such as unexpected Router Advertisements, multicast listener queries or ICMPv6 error messages [1][9][6]. In these cases, the authenticity of the messages MUST be verified before expending resources to initiate DNA procedure.

When a host arrives on a new link, hints received due to external IP packets will typically be due to multicast messages. Actions based on multicast reception from untrusted sources are dangerous due to the threat of multicast response flooding. This issue is discussed further in [Section 8](#).

State changes within the network-layer itself may initiate link-change detection procedures. Existing subsystems for router and neighbor discovery, address leasing and multicast reception maintain their own timers and state variables. Changes to the state of one or more of these mechanisms may hint that link change has occurred, and initiate detection of network attachment.

5.3 Handling Hints from Other Layers

Events at other protocol layers may provide hints of link change to network attachment detection systems. Two examples of such events are TCP retransmission timeout and completion of link-layer access procedures [[18](#)].

While hints from other protocol layers originate from within the host's own stack, the network layer SHOULD NOT treat hints from other protocol layers as authoritative indications of link change.

This is because state changes within other protocol layers may be triggered by untrusted messages, come from compromised sources (for example 802.11 WEP Encryption [[21](#)]) or be inaccurate with regard to network-layer state.

While these hints come from the host's own stack, such hints may actually be due to packet reception or non-reception events at the originating layers. As such, it may be possible for other devices to instigate hint delivery on a host or multiple hosts deliberately, in order to prompt packet transmission, or configuration modification.

Therefore, hosts SHOULD NOT change their configuration state based on hints from other protocol layers. A host MAY adopt an appropriate link change detection strategy based upon hints received from other layers, with suitable caution and hysteresis, as described in [Section 5.6](#).

5.4 Timer and Loss Based Hints

Other hints may be received due to timer expiry, particularly In some cases the expiry of these timers may be a good hint that DNA procedures are necessary.

Since DNA is likely to be used when communicating with devices over wireless links, suitable resilience to packet loss SHOULD be incorporated into the DNA initiation system. Notably, non-reception of data associated with end-to-end communication over the Internet may be caused by reception errors at either end or because of network congestion. Hosts SHOULD NOT act immediately on packet loss indications, delaying until it is clear that the packet losses aren't caused by transient events.

Use of the Advertisement Interval Option specified in [[5](#)] follows these principles. Routers sending this option indicate the maximum interval between successive router advertisements. Hosts receiving this option monitor for multiple successive packet losses and initiate change discovery.

5.5 Simultaneous Hints

Some events which generate hints may affect a number of devices simultaneously.

For example, if a wireless base station goes down, all the hosts on that base station are likely to initiate link-layer configuration strategies after losing track of the last beacon or pilot signal from the base station.

As described in [\[1\]](#)[\[6\]](#), a host SHOULD delay randomly before acting on a widely receivable advertisement, in order to avoid response implosion.

Where a host considers it may be on a new link and learns this from a hint generated by a multicast message, the host SHOULD defer 0-1000ms in accordance with [\[1\]](#)[\[3\]](#). Please note though that a single desynchronization is required for any number of transmissions subsequent to a hint, regardless of which messages need to be sent.

In link-layers where sufficient serialization occurs after an event experienced by multiple hosts, each host MAY avoid the random delays before sending solicitations specified in [\[1\]](#).

5.6 Hint Validity and Hysteresis

In some cases, hints can be generated by lower-layer protocols at an elevated rate, which do not reflect actual changes in IP configuration. In other cases, hints may also be received prior to the availability of the medium for network-layer packets.

Additionally, since packet reception at the network and other layers are a source for hints, it is possible for traffic patterns on the link to create hints, through chance or malicious intent. Therefore, it may be necessary to classify hint sources and types for their relevance and recent behavior.

When experiencing a large number of hints, a host SHOULD employ hysteresis techniques to prevent excessive use of network resources. The host MAY change the weight of particular hints, to devalue them if their accuracy has been poor, they suggest invalid configurations, or are suspicious (refer to [Section 8](#)).

It is notable, that such hysteresis may cause sub-optimal change detection performance, and may themselves be used to block legitimate hint reception.

5.7 Hint Management for Inactive Hosts

If a host does not expect to send or receive packets soon, it MAY choose to defer detection of network attachment. This may preserve resources on latent hosts, by removing any need for packet transmission when a hint is received.

These hosts SHOULD delay 0-1000ms before sending a solicitation, and MAY choose to wait up to twice the advertised Router Advertisement Interval (plus the random delay) before sending a solicitation [5].

One benefit of inactive hosts' deferral of DNA procedures is that herd-like host configuration testing is mitigated when base-station failure or simultaneous motion occur. When latent hosts defer DNA tests, the number of devices actively probing for data simultaneously is reduced to those hosts which currently support active data sessions.

When a device begins sending packets, it will be necessary to test bidirectional reachability with the router (whose current Neighbor Cache state is STALE). As described in [1], the host will delay before probing to allow for the probability that upper layer packet reception confirms reachability.

6. IP Hosts Configuration

Various protocols within IPv6 provide their own configuration processes. A host will have collected various configuration information using these protocols during its presence on a link. Following is the list of steps the host needs to take if a link-change has occurred.

- Invalidation of router and prefix list

- Invalidation of IPv6 addresses

- Removing neighbor cache entries

- Completion of DAD for Link-Local Addresses.

- Initiating mobility signaling

The following sub-sections elaborate on these steps.

6.1 Router and Prefix list

Router Discovery is designed to provide hosts with a set of locally

configurable prefixes and default routers. These may then be configured by hosts for access to the Internet [1].

It allows hosts to discover routers and manage lists of eligible next hop gateways, and is based on IPv6 Neighbor Discovery. When one of the routers in the router list is determined to be no longer reachable, its destination cache entry is removed, and new router is selected from the list. If a currently configured router is unreachable, it is quite likely that other devices on the same link are no longer reachable.

On determining that link-change has occurred, the default router list SHOULD have entries removed which are related to unreachable routers, and consequently these routers' destination cache entries SHOULD be removed [1]. If no eligible default routers are in the default router list, Router Solicitations MAY be sent, in order to discover new routers.

6.2 IPv6 Addresses

6.2.1 Autoconfiguration

Unicast source addresses are required to send all packets on the Internet, except a restricted subset of local signaling such as router and neighbor solicitations.

In dynamic environments, hosts need to undertake automatic configuration of addresses, and select which addresses to use based on prefix information advertised in Router Advertisements. Such configurations may be based on either Stateless Address Autoconfiguration [3] or DHCPv6 [13].

For any configured address, Duplicate Address Detection (DAD) MUST be performed [3]. DAD defines that an address is treated tentatively until either series of timeouts expire after probe transmissions or an address owner defends its address. Tentative addresses cannot modify peers' neighbor cache entries, nor can they receive packets.

As described in [Section 4.2](#), messages used in DNA signaling should be treated as unconfirmed, due to the chance of link change. Optimistic DAD is designed to allow use of addressing while they are being checked for validity. Careful use of these addresses may contribute to faster DNA operation [8].

6.2.2 Dynamic Host Configuration

Dynamic Host Configuration Procedures for IPv6 define their own detection procedures [13]. In order to check if the current set of

configuration is valid, a host can send a 'Confirm' message with a sample of its current configuration, which is able to be responded to by any DHCP relay on a link.

If the replying relay knows it is not on the same link, it may respond, indicating that the host's configuration is invalid. Current use of this technique is hampered by the lack of wide scale deployment of DHCPv6 and hence the detection mechanism doesn't work when the host moves to a link which doesn't contain DHCP relays or servers.

Upon link change, any configuration learned from DHCP which is link or administrative domain specific may have become invalid. Subsequent operation of DHCP on the new link may therefore be necessary.

6.3 Neighbor cache

Neighbor caches allow for delivery of packets to peers on the same link. Neighbor cache entries are created by router or neighbor discovery signaling, and may be updated either by upper-layer reachability confirmations or explicit neighbor discovery exchanges.

In order to determine which link-layer address a peer is at, nodes send solicitations to the link-local solicited-node multicast address of their peer. If hosts are reachable on this address, then they will respond to the solicitation with a unicast response. Information from these responses are stored in neighbour cache entries.

When link change occurs, the reachability of all existing neighbor cache entries is likely to be invalidated, if link change prevents packet reception from an old link. For these links, the neighbor cache entries SHOULD be removed when a host moves to a new link (although it MAY be possible to keep keying and authorization information for such hosts cached on a least-recently-used basis [7]).

Reachability of a single node may support the likelihood of reaching the rest of the link, for example if a particular access technology relays such messages through wireless base stations.

6.4 Mobility Management

Mobile IPv6 provides global mobility signaling for hosts wishing to preserve sessions when their configured address becomes topologically incorrect [5]. This system relies upon signaling messages and tunnel movement to provide reachability at a constant address, while

directing packets to its visited network.

The Mobile IPv6 [RFC3775](#) [5] defines 'movement detection' procedures, which themselves rely upon Neighbor Discovery, to initiate mobility signaling. These procedures allow for some modification of Neighbor Discovery to enable faster change or movement detection. When a host identifies that it is on a new link, if it is Mobile-IPv6 enabled host, it MAY initiate mobility signaling with its home agent and correspondent node.

7. Complications to Detecting Network Attachment

Detection of network attachment procedures can be delayed or may be incorrect due to different factors. This section gives some examples where complications may interfere with DNA processing.

7.1 Packet Loss

Generally, packet loss introduces significant delays in validation of current configuration or discovery of new configuration. Because most of the protocols rely on timeout to consider that a peer is not reachable anymore, packet loss may lead to erroneous conclusions.

Additionally, packet loss rates for particular transmission modes (multicast or unicast) may differ, meaning that particular classes of DNA tests have higher chance of failure due to loss. Hosts SHOULD attempt to verify tests through retransmission where packet loss is prevalent.

7.2 Router Configurations

Each router can have its own configuration with respect to sending RA, and the treatment of router and neighbor solicitations. Different timers and constants might be used by different routers, such as the delay between Router Advertisements or delay before replying to an RS. If a host is changing its IPv6 link, the new router on that link may have a different configuration and may introduce more delay than the previous default router of the host. The time needed to discover the new link can then be longer than expected by the host.

7.3 Overlapping Coverage

If a host can be attached to different links at the same time with the same interface, the host will probably listen to different routers, at least one on each link. To be simultaneously attached to several links may be very valuable for a MN when it moves from one

access network to another. If the node can still be reachable through its old link while configuring the interface for its new link, packet loss can be minimized.

Such a situation may happen in a wireless environment if the link layer technology allows the MN to be simultaneously attached to several points of attachment and if the coverage area of access points are overlapping.

For the purposes of DNA, it is necessary to treat each of these points-of-attachment separately, otherwise incorrect conclusions of link-change may be made even if another of the link-layer connections is valid.

7.4 Multicast Snooping

When a host is participating in DNA on a link where multicast snooping is in use, multicast packets may not be delivered to the LAN-segment to which the host is attached until MLD signaling has been performed [9][11] [17]. Where DNA relies upon multicast packet delivery (for example, if a router needs to send a Neighbor Solicitation to the host), its function will be degraded until after an MLD report is sent.

Where it is possible that multicast snooping is in operation, hosts MUST send MLD group joins (MLD reports) for solicited nodes' addresses swiftly after starting DNA procedures.

7.5 Link Partition

Link partitioning occurs when a link's internal switching or relaying hardware fails, or if the internal communications within a link are prevented due to topology changes or wireless propagation.

When a host is on a link which partitions, only a subset of the addresses or devices it is communicating with may still be available. Where link partitioning is rare (for example, with wired communication between routers on the link), existing router and neighbor discovery procedures may be sufficient for detecting change.

8. Security Considerations

Detecting Network Attachment is a mechanism which allows network messages to change the node's belief about its IPv6 configuration state. As such, it is important that the need for rapid testing of link change does not lead to situations where configuration is invalidated by malicious third parties, nor that information passed to configuration processes exposes the host to other attacks.

Since DNA relies heavily upon IPv6 Neighbor Discovery, the threats which are applicable to those procedures also affect Detecting Network Attachment. These threats are described in [12].

Some additional threats are outlined below.

8.1 Authorization and Detecting Network Attachment

Hosts connecting to the Internet over wireless media may be exposed to a variety of network configurations with differing robustness, controls and security.

When a host is determining if link change has occurred, it may receive messages from devices with no advertised security mechanisms purporting to be routers, nodes sending signed router advertisements but with unknown delegation, or routers whose credentials need to be checked [12]. Where a host wishes to configure an unsecured router, it SHOULD confirm bidirectional reachability with the device, and it MUST mark the device as unsecured as described in [7].

In any case, a secured router SHOULD be preferred over an unsecured one, except where other factors (unreachability) make the router unsuitable. Since secured routers' advertisement services may be subject to attack, alternative (secured) reachability mechanisms from upper layers, or secured reachability of other devices known to be on the same link may be used to check reachability in the first instance.

8.2 Addressing

While a DNA host is checking for link-change, and observing DAD, it may receive a DAD defense NA from an unsecured source.

SEND says that DAD defenses MAY be accepted even from non SEND nodes for the first configured address [7].

While this is a valid action in the case where a host collides with another address owner after arrival on a new link, In the case that the host returns immediately to the same link, such a DAD defense NA message can only be a denial-of-service attempt.

If a non-SEND node forges a DAD defense for an address which is still in peers' neighbor cache entries, a host may send a SEND protected unicast neighbor solicitation without a source link-layer address option to one of its peers (which also uses SEND). If this peer is reachable, it will not have registered the non-SEND DAD defense NA in its neighbor cache, and will send a direct NA back to the soliciting host. Such an NA reception disproves the DAD defense NA's validity.

Therefore, a SEND host performing DNA which receives a DAD defense from a non-SEND node SHOULD send a unicast Neighbor Solicitation to a STALE or REACHABLE secure neighbor cache entry, omitting source link-layer address options. In this case, the host should pick an entry which is likely to have a corresponding entry on the peer. If the host responds within a RetransTimer interval, then the DAD defense was an attack, and the host SHOULD inform its systems administrator without disabling the address.

9. Constants

MAC_CACHE_TIME: 30 minutes

10. Acknowledgments

Thanks to JinHyeock Choi and Erik Nordmark for their significant contributions. Bernard Aboba's work on DNA for IPv4 strongly influenced this document.

11. References

11.1 Normative References

- [1] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [3] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [4] Haskin, D. and E. Allen, "IP Version 6 over PPP", [RFC 2472](#), December 1998.
- [5] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [6] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC2463](#) 2463, December 1998.
- [7] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.
- [8] Moore, N., "Optimistic Duplicate Address Detection for IPv6", [draft-ietf-ipv6-optimistic-dad-02](#) (work in progress), September 2004.

11.2 Informative References

- [9] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), October 1999.
- [10] Haberman, B., "Source Address Selection for the Multicast Listener Discovery (MLD) Protocol", [RFC 3590](#), September 2003.
- [11] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [12] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", [RFC 3756](#), May 2004.
- [13] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [14] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", [RFC 3513](#), April 2003.
- [15] Choi, J., "Detecting Network Attachment in IPv6 Goals", [draft-ietf-dna-goals-04](#) (work in progress), December 2004.
- [16] Fikouras, N A., K'onsgen, A J., and C. G'org, "Accelerating Mobile IP Hand-offs through Link-layer Information", Proceedings of the International Multiconference on Measurement, Modelling, and Evaluation of Computer-Communication Systems (MMB) 2001, September 2001.
- [17] Christensen, M., Kimball, K., and F. Solensky, "Considerations for IGMP and MLD Snooping Switches", [draft-ietf-magma-snoop-12](#) (work in progress), February 2005.
- [18] Yegin, A., "Link-layer Event Notifications for Detecting Network Attachments", [draft-ietf-dna-link-information-00](#) (work in progress), September 2004.
- [19] Koodli, R., "Fast Handovers for Mobile IPv6", [draft-ietf-mipshop-fast-mipv6-03](#) (work in progress), October 2004.
- [20] Liebsch, M., "Candidate Access Router Discovery", [draft-ietf-seamoby-card-protocol-08](#) (work in progress), September 2004.
- [21] O'Hara, B. and G. Ennis, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", ANSI/IEEE

Std 802.11, 1999.

- [22] Yamamoto, S., "Detecting Network Attachment Terminology",
[draft-yamamoto-dna-term-00](#) (work in progress), February 2004.
- [23] Manner, J. and M. Kojo, "Mobility Related Terminology",
[draft-ietf-seamoby-mobility-terminology-06](#) (work in progress),
February 2004.
- [24] Choi, J. and E. Nordmark, "DNA with unmodified routers: Prefix
list based approach", [draft-j-dna-cpl-00](#) (work in progress),
October 2005.

Authors' Addresses

Sathya Narayanan
Panasonic Digital Networking Lab
Two Research Way, 3rd Floor
Princeton, NJ 08536
USA

Phone: 609 734 7599
Email: sathya@Research.Panasonic.COM
URI:

Greg Daley
Centre for Telecommunications and Information Engineering
Department of Electrical and Computer Systems Engineering
Monash University
Clayton, Victoria 3800
Australia

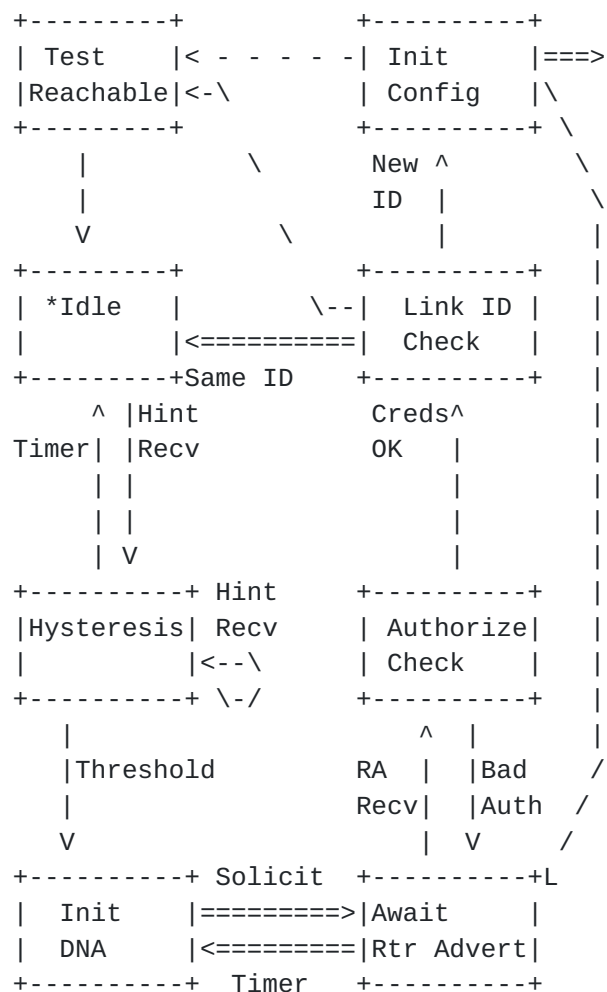
Phone: +61 3 9905 4655
Email: greg.daley@eng.monash.edu.au

Nicolas Montavont
 LSIIT - Univerity Louis Pasteur
 Pole API, bureau C444
 Boulevard Sebastien Brant
 Illkirch 67400
 FRANCE

Phone: (33) 3 90 24 45 87
 Email: montavont@dpt-info.u-strasbg.fr
 URI: <http://www-r2.u-strasbg.fr/~montavont/>

Appendix A. Example State Transition Diagram

Below is an example state diagram which indicates relationships between the practices in this document.



Intellectual Property Statement

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.

Disclaimer of Validity

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.

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.

Acknowledgment

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

