

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 29, 2010

G. Daley
NetStar Networks
E. Nordmark
Sun Microsystems
N. Moore
October 26, 2009

**Tentative Options for Link-Layer Addresses in IPv6 Neighbour Discovery
draft-ietf-dna-tentative-04.txt**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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 April 29, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

The proposed IPv6 Duplicate Address Detection (DAD) Optimization "Optimistic DAD" defines a set of recoverable procedures which allow a node to make use of an address before DAD completes. Essentially, Optimistic DAD forbids usage of certain Neighbour Discovery options which could pollute active neighbour cache entries, while an address is tentative.

This document defines a new option and procedures to replace cache polluting options, in a way which is useful to tentative nodes. These procedures are designed to be backward compatible with existing devices which support IPv6 Neighbour Discovery.

Table of Contents

1.	Terminology	4
2.	Introduction	4
2.1.	Tentative Option format	5
2.2.	Tentative Option semantics	5
3.	Sending solicitations containing Tentative Options	6
3.1.	Sending Neighbour Solicitations with Tentative Options	6
3.2.	Sending Router Solicitations with Tentative Options	6
4.	Receiving Tentative Options	6
4.1.	Handling Tentative Options	7
4.2.	Receiving Neighbour Solicitations containing Tentative Options	7
4.3.	Receiving a Router Solicitation containing a Tentative Option	8
5.	IANA Considerations	8
6.	Security Considerations	8
7.	Acknowledgments	10
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	10
Appendix A.	Constraints imposed by IPv6 Neighbour Discovery	11
A.1.	Constraints on Neighbour Solicitations	11
A.2.	Constraints on Router Solicitations	11
Appendix B.	Interactions with legacy nodes	11
B.1.	Legacy Neighbour Solicitation processing	12
B.2.	Legacy Router Solicitation processing	12
	Authors' Addresses	13

1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

2. Introduction

Source Link-Layer Address Options (SLLAOs) are sent in Neighbour discovery messages in order to notify neighbours of a mapping between a specific IPv6 Network layer address and a link-layer (or MAC) address. Upon reception of a Neighbour Discovery message containing such an option, nodes update their neighbour cache entries with the IP to link-layer address mapping in accordance with procedures defined in IPv6 Neighbour Discovery [[RFC4861](#)].

Optimistic DAD [[RFC4429](#)] prevents usage of these options in Router and Neighbour Solicitation messages from a tentative address (while Duplicate Address Detection is occurring). This is because receiving a Neighbour Solicitation (NS) or Router Solicitation (RS) containing an SLLAO would otherwise overwrite an existing cache entry, even if the cache entry contained the legitimate address owner, and the solicitor was a duplicate address.

Neighbour Advertisement (NA) messages don't have such an issue, since the Advertisement message contains a flag which explicitly disallows overriding of existing cache entries, by the target link-layer address option carried within.

The effect of preventing SLLAOs for tentative addresses is that communications with these addresses are sub-optimal for the tentative period. Sending solicitations without these options causes an additional round-trip for Neighbour Discovery if the advertiser does not have an existing neighbour cache entry for the solicitor. In some cases, multicast advertisements will be scheduled, where Neighbour Discovery is not possible on the advertiser.

This document proposes Tentative Options which designed to replace the existing Source Link-Layer Address Options available in IPv6

Neighbour Discovery, when a device is performing Optimistic DAD.

Operations in this document are safe with respect to existing nodes that implement IPv6 Neighbor Discovery [[RFC4861](#)] and Stateless Address Autoconfiguration [[RFC4862](#)]. Expected behaviours of legacy devices are summarized in [Appendix B](#).

2.1. Tentative Option format

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      Type      |      Length      |      Link-Layer Address ...
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Fields:

Type TBD (Requires IANA Allocation)

Length The length of the option (including the type and length fields) in units of 8 octets.

Link-Layer Address
 The variable length link-layer address.

Description
 The Tentative option contains the link-layer address of the sender of the packet. It is used in the Neighbour Solicitation and Router Solicitation packets.

2.2. Tentative Option semantics

The Tentative Option (TO) functions in the same role as the Source Link-Layer Address option defined for [[RFC4861](#)], but it MUST NOT override an existing neighbour cache entry.

The differing neighbour cache entry MUST NOT be affected by the reception of the Tentative Option. This ensures that tentative addresses are unable to modify legitimate neighbour cache entries.

In the case where an entry is unable to be added to the neighbour cache, a node MAY send responses direct to the link-layer address specified in the TO.

For these messages, no Neighbour Cache entry may be created, although response messages may be directed to a particular unicast address.

These procedures are discussed further in [Section 4.3](#).

[3.](#) Sending solicitations containing Tentative Options

Tentative Options may be sent in Router and Neighbour Solicitations, as described below.

In a case where it is safe to send a Source Link-Layer Address Option, a host SHOULD NOT send a TO, since the message may be misinterpreted by legacy nodes. Importantly, a node MUST NOT send a Tentative Option in the same message where a Source Link-Layer Address Option is sent.

Particularly, Tentative Options are permitted to be sent when addressing state precludes sending the SLLAO, but a neighbour cache entry will be created on a peer node [[RFC4429](#)].

[3.1.](#) Sending Neighbour Solicitations with Tentative Options

Neighbour Solicitations sent to unicast addresses MAY contain a Tentative Option.

Since delivery of a packet to a unicast destination requires prior knowledge of the destination's hardware address, unicast Neighbour Solicitation packets may only be sent to destinations for which a neighbour cache entry already exists.

For example, if checking bidirectional reachability to a router, it may be possible to send a Neighbour Solicitation with Tentative Option to the router's advertised address.

As discussed in [[RFC4861](#)], the peer device may not have a cache entry even if the soliciting host does, in which case reception of the Tentative Option may create a neighbour cache entry, without the need for Neighbour Discovering the original solicitor.

[3.2.](#) Sending Router Solicitations with Tentative Options

Any Router Solicitation from a Preferred, Deprecated or Optimistic address MAY be sent with a Tentative Option [[RFC4429](#)].

[4.](#) Receiving Tentative Options

Receiving a Tentative Option allows nodes to unicast responses to solicitations without performing Neighbour Discovery.

It does this by allowing the solicitation to create STALE neighbour cache entries if one doesn't exist, but only update an entry if the link-layer address in the option matches the entry.

4.1. Handling Tentative Options

Use of Tentative Options is only defined for Neighbour and Router Solicitation messages.

In any other received message, the presence of the option is silently ignored, that is, the packet is processed as if the option was not present.

It is REQUIRED that the same validation algorithms for Neighbour and Router Solicitations received with TO as in the IPv6 Neighbour Discovery specification [[RFC4861](#)], are used.

In the case that a solicitation containing a Tentative Option is received, The only processing differences occur in checking and updating the neighbour cache entry. Particularly, there is no reason to believe that the host will remain tentative after receiving a responding advertisement.

As defined in [Section 2.1](#), Tentative Options do not overwrite existing neighbour cache entries where the link-layer addresses of the option and entry differ.

If a solicitation from a unicast source address is received where no difference exists between the TO and an existing neighbour cache entry, the option MUST be treated as if it were an SLLAO after message validation, and processed accordingly.

In the case that a cache entry is unable to be created or updated due to existence of a conflicting neighbour cache entry, it MUST NOT update the neighbour cache entry.

4.2. Receiving Neighbour Solicitations containing Tentative Options

The Tentative Option is only allowed in Neighbour Solicitations with specified source addresses for which SLLAO is not required.

A Neighbour Solicitation message received with a TO and an unspecified source address MUST be silently discarded.

Upon reception of a Tentative Option in a Neighbour Solicitation for which the receiver has the Target Address configured, a node checks to see if there is a neighbour cache entry with conflicting link-layer address.

If no such entry exists, the neighbour cache of the receiver SHOULD be updated, as if the Tentative Option was a SLLAO.

Sending of the solicited Neighbour Advertisement then proceeds normally, as defined in [section 7.2.4 of \[RFC4861\]](#).

If there is a conflicting neighbour cache entry, the node processes the solicitation as defined in [Section 7.2.4 of \[RFC4861\]](#), except that the Neighbour Cache entry MUST NOT be modified.

4.3. Receiving a Router Solicitation containing a Tentative Option

In IPv6 Neighbour Discovery [[RFC4861](#)], responses to Router Solicitations are either sent to the all-nodes multicast address, or may be sent to the solicitation's source address if it is a unicast address.

Including a Tentative Option in the solicitation allows a router to choose to send a packet directly to the link-layer address even in situations where this would not normally be possible.

For Router Solicitations with unicast source addresses, neighbour caches SHOULD be updated with the link-layer address from a Tentative Option if there is no differing neighbour cache entry. In this case, Router Advertisement continues as in [Section 6.2.6 of \[RFC4861\]](#).

For received solicitations with a differing link-layer address to that stored in the neighbour cache, the node processes the solicitation as defined in [Section 6.2.6 of \[RFC4861\]](#), except that the Neighbour Cache entry MUST NOT be modified.

5. IANA Considerations

IANA action of options for IPv6 Neighbor Discovery require RFC Approval [[RFC2780](#)].

This document asks the IANA to allocate the Tentative Option for link-layer addressing (Section [Section 2.1](#)) from the IPv6 Neighbour Discovery options for IPv6.

6. Security Considerations

The use of the Tentative Option in Neighbour and Router Solicitation messages acts in a similar manner to SLLAO, updating neighbour cache entries, in a way which affects the transmission of subsequent packets.

An attacker may cause messages may be sent to another node by an advertising node (a reflector), without creating any ongoing state on the reflector.

This is attack requires one solicitation for each advertisement and the advertisement has to go to a unicast MAC destination. That said, the size of the advertisement may be significantly larger than the solicitation, or the attacker and reflector may be on a medium with greater available bandwidth than the victim.

For link-layers where it isn't possible to spoof the link-layer source address this allows a slightly increased risk of reflection attacks from nodes which are on-link.

Additionally, since a SEND host must always advertise using SEND options and signatures, a non-SEND attacker may cause excess computation on both a victim node and a router by causing SEND advertisement messages to be transmitted to a particular MAC address and the lall-nodes multicast. SEND specifies guidelines to hosts receiving unsolicited advertisements in order to mitigate such attacks [[RFC3971](#)].

While this is the same effect as experienced when accepting SLLAO from non-SEND nodes, the lack of created neighbour cache entries on the advertiser may make such attacks more difficult to trace.

Modification of Neighbour Discovery messages on the network is possible, unless SEND is used. [[RFC3971](#)] provides a protocol specification in which soliciting nodes sign ND messages with a private key and use addresses generated from this key.

Even if SEND is used, the lifetime of a neighbour cache entry may be extended by continually replaying a solicitation message to a particular router or hosts. Since this may be achieved for any Neighbour or Router Solicitation message, corresponding advertisements to the original transmitters of these solicitation messages may occur.

SEND defines use of Timestamp values to protect a device from attack through replay of previously sent messages. Although this applies to Neighbour and Router Solicitation messages, granularity of the timestamp allows the messages to be used for up to five minutes [[RFC3971](#)].

All Router and Neighbour Solicitations using SEND contain a Nonce option, containing a random identifier octet string. Since SEND messages are digitally signed, and may not be easily modified, replay attacks will contain the same Nonce option, as was used in the

original solicitation.

While the Nonce Option included in a transmission to another node may not vary within one short solicitation period (the host may itself replay solicitations in the case of packet loss), the presence of the timestamp option ensures that for later solicitations, a different Timestamp and Nonce will be used.

Therefore, a receiver seeing a solicitation with the same Timestamp and Nonce (and signature) for more than either of MAX_RTR_SOLICITATIONS (for Router Solicitations), MAX_UNICAST_SOLICIT or MAX_MULTICAST_SOLICIT (for Neighbour Solicitations), SHOULD ignore further solicitations with this (Nonce, Timestamp, Source) triple, ensuring that no modification is made to neighbour cache entries. This applies to any solicitation packet capable of carrying a SEND payload, whether they use a Tentative Option or SLLAO.

Stations noticing such an attack SHOULD notify their administrator of the attempt at Denial-of-service.

7. Acknowledgments

Erik Nordmark coined a proposal for Tentative version of the SLLAO during a conversation with JinHyeock Choi and Greg Daley.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection for IPv6", RFC [RFC4429](#), April 2006.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.

8.2. Informative References

- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers",

[BCP 37](#), [RFC 2780](#), March 2000.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.

[Appendix A](#). Constraints imposed by IPv6 Neighbour Discovery

Hosts which send and receive Tentative Options may be interacting with legacy nodes which support IPv6 Neighbour Discovery procedures, but do not understand the new option.

For these nodes, the presence of the option is silently ignored, that is, the packet is processed as if the option was not present. Therefore all messages sent with Tentative Options MUST be compliant with the existing requirements for options and addressing specified in the IPv6 Neighbour Discovery RFC [[RFC4861](#)].

[A.1](#). Constraints on Neighbour Solicitations

As described in [Section 7.2.2 of \[RFC4861\]](#), packets sent to solicited nodes' multicast addresses MUST contain Source Link-Layer Address options.

Neighbour solicitations to multicast addresses MUST NOT contain Tentative Options

Neighbour Solicitations to unicast addresses SHOULD include a link-layer address (if the sender has one) as a Source Link-Layer Address option.

Unicast neighbour solicitations without Source Link-Layer Address Options MAY contain Tentative Options, if the solicitor has a Link-Layer address.

[A.2](#). Constraints on Router Solicitations

As described in [Section 6.3.7 of \[RFC4861\]](#), Router Solicitations SHOULD contain Source Link-Layer Address Options.

Router Solicitations without Source Link-Layer Address options MAY contain a Tentative Option.

[Appendix B](#). Interactions with legacy nodes

Devices which do not implement Tentative Options will act as if no option was placed within the Neighbour Discovery message. The

following sections summarize how legacy hosts will interact with messages containing Tentative Options.

B.1. Legacy Neighbour Solicitation processing

A node can include the Tentative Option in a unicast NS (and no SLLAO option) when the transmitter's address is either preferred, tentative or optimistic.

An [RFC 2461](#) host receiving such a packet will "see" a packet without an SLLAO option, which is allowed in [RFC4861](#).

If the recipient host has an existing neighbour cache entry for the transmitter, it can then send a Neighbour Advertisement.

Where no neighbour cache entry exists, the recipient will send a multicast NS (containing its own SLLAO) in order for the original transmitter to respond with an NA. Upon reception of the original transmitter's NA, an NA is sent back to the origin.

The Tentative Option MUST NOT be included in an NS message which has no source address.

An [RFC 2461](#) host sees an NS without a source address as a Duplicate Address Detection message.

Reception of duplicate address detection messages may cause side-effects on other hosts, which may cause them to treat addresses as invalid.

B.2. Legacy Router Solicitation processing

A node can include the Tentative Option in an RS with a unicast source address (and no SLLAO option) when the transmitter's address is either tentative or optimistic.

An [RFC 2461](#) router receiving such a packet will "see" a packet without an SLLAO option, which is allowed in [RFC4861](#).

If the router has an existing neighbour cache entry for this host, it may send a Unicast RA in response, but may send a multicast in preference.

If no neighbour cache entry exists, some routers will not be able to provide a unicast response. These routers will schedule a multicast response.

Other routers may attempt to perform neighbour discovery (by sending a multicast NS), and unicast a response when a neighbour cache entry has been created.

Authors' Addresses

Greg Daley
NetStar Australia Pty Ltd
Lvl 9/636 St Kilda Rd
Melbourne, Victoria 3004
Australia

Phone: +61 401 772 770
Email: gdaley@netstarnetworks.com

Erik Nordmark
Sun Microsystems, Inc.
17 Network Circle
Mountain View, CA
USA

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

Nick "Sharkey" Moore

Email: sharkey@zoic.org

