

Internet Engineering Task Force (IETF)
Request for Comments: 6775
Updates: 4944
Category: Standards Track
ISSN: 2070-1721

Z. Shelby, Ed.
Sensinode
S. Chakrabarti
Ericsson
E. Nordmark
Cisco Systems
C. Bormann
Universitaet Bremen TZI
November 2012

Neighbor Discovery Optimization for IPv6 over Low-Power Wireless
Personal Area Networks (6LoWPANs)

Abstract

The IETF work in IPv6 over Low-power Wireless Personal Area Network (6LoWPAN) defines 6LoWPANs such as IEEE 802.15.4. This and other similar link technologies have limited or no usage of multicast signaling due to energy conservation. In addition, the wireless network may not strictly follow the traditional concept of IP subnets and IP links. IPv6 Neighbor Discovery was not designed for non-transitive wireless links, as its reliance on the traditional IPv6 link concept and its heavy use of multicast make it inefficient and sometimes impractical in a low-power and lossy network. This document describes simple optimizations to IPv6 Neighbor Discovery, its addressing mechanisms, and duplicate address detection for Low-power Wireless Personal Area Networks and similar networks. The document thus updates RFC 4944 to specify the use of the optimizations defined here.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6775>.

Copyright Notice

Copyright (c) 2012 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. The Shortcomings of IPv6 Neighbor Discovery	5
1.2. Applicability	6
1.3. Goals and Assumptions	7
1.4. Substitutable Features	8
2. Terminology	9
3. Protocol Overview	11
3.1. Extensions to RFC 4861	11
3.2. Address Assignment	12
3.3. Host-to-Router Interaction	13
3.4. Router-to-Router Interaction	14
3.5. Neighbor Cache Management	14
4. New Neighbor Discovery Options and Messages	15
4.1. Address Registration Option	15
4.2. 6LoWPAN Context Option	17
4.3. Authoritative Border Router Option	19
4.4. Duplicate Address Messages	20
5. Host Behavior	22
5.1. Forbidden Actions	22
5.2. Interface Initialization	22
5.3. Sending a Router Solicitation	23
5.4. Processing a Router Advertisement	23
5.4.1. Address Configuration	23
5.4.2. Storing Contexts	24
5.4.3. Maintaining Prefix and Context Information	24
5.5. Registration and Neighbor Unreachability Detection	25
5.5.1. Sending a Neighbor Solicitation	25
5.5.2. Processing a Neighbor Advertisement	25
5.5.3. Recovering from Failures	26
5.6. Next-Hop Determination	26
5.7. Address Resolution	27

5.8. Sleeping	27
5.8.1. Picking an Appropriate Registration Lifetime	27
5.8.2. Behavior on Wakeup	28
6. Router Behavior for 6LRs and 6LBRs	28
6.1. Forbidden Actions	28
6.2. Interface Initialization	29
6.3. Processing a Router Solicitation	29
6.4. Periodic Router Advertisements	30
6.5. Processing a Neighbor Solicitation	30
6.5.1. Checking for Duplicates	30
6.5.2. Returning Address Registration Errors	31
6.5.3. Updating the Neighbor Cache	31
6.5.4. Next-Hop Determination	32
6.5.5. Address Resolution between Routers	32
7. Border Router Behavior	32
7.1. Prefix Determination	33
7.2. Context Configuration and Management	33
8. Substitutable Feature Behavior	34
8.1. Multihop Prefix and Context Distribution	34
8.1.1. 6LBRs Sending Router Advertisements	35
8.1.2. Routers Sending Router Solicitations	35
8.1.3. Routers Processing Router Advertisements	35
8.1.4. Storing the Information	36
8.1.5. Sending Router Advertisements	36
8.2. Multihop Duplicate Address Detection	37
8.2.1. Message Validation for DAR and DAC	38
8.2.2. Conceptual Data Structures	39
8.2.3. 6LR Sending a Duplicate Address Request	39
8.2.4. 6LBR Receiving a Duplicate Address Request	39
8.2.5. Processing a Duplicate Address Confirmation	40
8.2.6. Recovering from Failures	40
9. Protocol Constants	41
10. Examples	42
10.1. Message Examples	42
10.2. Host Bootstrapping Example	43
10.2.1. Host Bootstrapping Messages	45
10.3. Router Interaction Example	46
10.3.1. Bootstrapping a Router	46
10.3.2. Updating the Neighbor Cache	47
11. Security Considerations	47
12. IANA Considerations	48
13. Interaction with Other Neighbor Discovery Extensions	49
14. Guidelines for New Features	49
15. Acknowledgments	52
16. References	52
16.1. Normative References	52
16.2. Informative References	53

1. Introduction

The IPv6-over-IEEE 802.15.4 [RFC4944] document specifies how IPv6 is carried over an IEEE 802.15.4 network with the help of an adaptation layer that sits between the Media Access Control (MAC) layer and the IP network layer. A link in a Low-power Wireless Personal Area Network (LoWPAN) is characterized as lossy, low-power, low-bit-rate, short-range; with many nodes saving energy with long sleep periods. Multicast as used in IPv6 Neighbor Discovery (ND) [RFC4861] is not desirable in such a wireless low-power and lossy network. Moreover, LoWPAN links are asymmetric and non-transitive in nature. A LoWPAN is potentially composed of a large number of overlapping radio ranges. Although a given radio range has broadcast capabilities, the aggregation of these is a complex Non-Broadcast Multiple Access (NBMA) [RFC2491] structure with generally no LoWPAN-wide multicast capabilities. Link-local scope is in reality defined by reachability and radio strength. Thus, we can consider a LoWPAN to be made up of links with undetermined connectivity properties as in [RFC5889], along with the corresponding address model assumptions defined therein.

This specification introduces the following optimizations to IPv6 Neighbor Discovery [RFC4861] specifically aimed at low-power and lossy networks such as LoWPANs:

- o Host-initiated interactions to allow for sleeping hosts.
- o Elimination of multicast-based address resolution for hosts.
- o A host address registration feature using a new option in unicast Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages.
- o A new Neighbor Discovery option to distribute 6LoWPAN header compression context to hosts.
- o Multihop distribution of prefix and 6LoWPAN header compression context.
- o Multihop Duplicate Address Detection (DAD), which uses two new ICMPv6 message types.

The two multihop items can be substituted by a routing protocol mechanism if that is desired; see Section 1.4.

The document defines three new ICMPv6 message options: the Address Registration Option (ARO), the Authoritative Border Router Option (ABRO), and the 6LoWPAN Context Option (6CO). It also defines two new ICMPv6 message types: the Duplicate Address Request (DAR) and the Duplicate Address Confirmation (DAC).

1.1. The Shortcomings of IPv6 Neighbor Discovery

IPv6 Neighbor Discovery [RFC4861] provides several important mechanisms used for router discovery, address resolution, Duplicate Address Detection, and Redirect messages, along with prefix and parameter discovery.

Following power-on and initialization of the network in IPv6 Ethernet networks, a node joins the solicited-node multicast address on the interface and then performs Duplicate Address Detection (DAD) for the acquired link-local address by sending a solicited-node multicast message to the link. After that, it sends multicast messages to the all-routers multicast address to solicit Router Advertisements (RAs). If the host receives a valid RA with the A (autonomous address configuration) flag, it autoconfigures the IPv6 address with the advertised prefix in the RA message. Besides this, the IPv6 routers usually send RAs periodically on the network. RAs are sent to the all-nodes multicast address. Nodes send Neighbor Solicitation/Neighbor Advertisement messages to resolve the IPv6 address of the destination on the link. The Neighbor Solicitation messages used for address resolution are multicast. The Duplicate Address Detection procedure and the use of periodic Router Advertisement messages assume that the nodes are powered on and reachable most of the time.

In Neighbor Discovery, the routers find the hosts by assuming that a subnet prefix maps to one broadcast domain, and then they multicast Neighbor Solicitation messages to find the host and its link-layer address. Furthermore, the DAD use of multicast assumes that all hosts that autoconfigure IPv6 addresses from the same prefix can be reached using link-local multicast messages.

Note that the L (on-link) bit in the Prefix Information Option (PIO) can be set to zero in Neighbor Discovery, which makes the host not use multicast Neighbor Solicitation (NS) messages for address resolution of other hosts, but routers still use multicast NS messages to find the hosts.

Due to the lossy nature of wireless communication and a changing radio environment, the IPv6-link node-set may change due to external physical factors. Thus, the link is often unstable, and the nodes appear to be moving without necessarily moving physically.

A LoWPAN can use two types of link-layer addresses: 16-bit short addresses and 64-bit unique addresses as defined in [RFC4944]. Moreover, the available link-layer payload size is on the order of less than 100 bytes; thus, header compression is very useful.

Considering the above characteristics in a LoWPAN, and the IPv6 Neighbor Discovery [RFC4861] protocol design, some optimizations and extensions to Neighbor Discovery are useful for the wide deployment of IPv6 over low-power and lossy networks (example: 6LoWPAN and other homogeneous low-power networks).

1.2. Applicability

In its Section 1, [RFC4861] foresees a document that covers operating IP over a particular link type and defines an exception to the otherwise general applicability of unmodified [RFC4861]. The present specification improves the usage of IPv6 Neighbor Discovery for LoWPANs in order to save energy and processing power of such nodes. This document thus updates [RFC4944] to specify the use of the optimizations defined here.

The applicability of this specification is limited to LoWPANs where all nodes on the subnet implement these optimizations in a homogeneous way. Although it is noted that some of these optimizations may be useful outside of 6LoWPANs, for example, in general IPv6 low-power and lossy networks and possibly even in combination with [RFC4861], the usage of such combinations is out of scope of this document.

In this document, we specify a set of behaviors between hosts and routers in LoWPANs. An implementation that adheres to this document MUST implement those behaviors. The document also specifies a set of behaviors (multihop prefix or context dissemination and, separately, multihop Duplicate Address Detection) that are needed in route-over configurations. An implementation of this specification MUST support those pieces, unless the implementation supports some alternative ("substitute") from some other specification.

The optimizations described in this document apply to different topologies. They are most useful for route-over and mesh-under configurations in Mesh topologies. However, Star topology configurations will also benefit from the optimizations due to reduced signaling, robust handling of the non-transitive link, and header compression context information.

1.3. Goals and Assumptions

The document has the following main goals and assumptions.

Goals:

- o Optimize Neighbor Discovery with a mechanism that is minimal yet sufficient for the operation in both mesh-under and route-over configurations.
- o Minimize signaling by avoiding the use of multicast flooding and reducing the use of link-scope multicast messages.
- o Optimize the interfaces between hosts and their default routers.
- o Provide support for sleeping hosts.
- o Disseminate context information to hosts as needed by 6LoWPAN header compression [RFC6282].
- o Disseminate context information and prefix information from the border to all routers in a LoWPAN.
- o Provide a multihop Duplicate Address Detection mechanism suitable for route-over LoWPANs.

Assumptions:

- o 64-bit Extended Unique Identifier (EUI-64) [EUI64] addresses are globally unique, and the LoWPAN is homogeneous.
- o All nodes in the network have an EUI-64 Interface ID in order to do address autoconfiguration and detect duplicate addresses.
- o The link-layer technology is assumed to be low-power and lossy, exhibiting undetermined connectivity, such as IEEE 802.15.4 [RFC4944]. However, the address registration mechanism might be useful for other link-layer technologies.
- o A 6LoWPAN is configured to share one or more global IPv6 address prefixes to enable hosts to move between routers in the LoWPAN without changing their IPv6 addresses.
- o When using the multihop DAD mechanism (Section 8.2), each 6LoWPAN Router (6LR) registers with all the 6LoWPAN Border Routers (6LBRs) available in the LoWPAN.

- o If IEEE 802.15.4 16-bit short addresses are used, then some technique is used to ensure the uniqueness of those link-layer addresses. That could be done using DHCPv6, Address Registration Option-based Duplicate Address Detection (specified in Section 8.2), or other techniques outside of the scope of this document.
- o In order to preserve the uniqueness of addresses (see Section 5.4 of [RFC4862]) not derived from an EUI-64, they must be either assigned or checked for duplicates in the same way throughout the LoWPAN. This can be done using DHCPv6 for assignment and/or using the Duplicate Address Detection mechanism specified in Section 8.2 (or any other protocols developed for that purpose).
- o In order for 6LoWPAN header compression [RFC6282] to operate correctly, the compression context must match for all the hosts, 6LRs, and 6LBRs that can send, receive, or forward a given packet. If Section 8.1 is used to distribute context information, this implies that all the 6LBRs must coordinate the context information they distribute within a single LoWPAN.
- o This specification describes the operation of ND within a single LoWPAN. The participation of a node in multiple LoWPANs simultaneously may be possible but is out of scope of this document.
- o Since the LoWPAN shares its prefix(es) throughout the network, mobility of nodes within the LoWPAN is transparent. Inter-LoWPAN mobility is out of scope of this document.

1.4. Substitutable Features

This document defines the optimization of Neighbor Discovery messages for the host-router interface and introduces two new mechanisms in a route-over topology.

Unless specified otherwise (in a document that defines a routing protocol that is used in a 6LoWPAN), this document applies to networks with any routing protocol. However, because the routing protocol may provide good alternate mechanisms, this document defines certain features as "substitutable", meaning they can be substituted by a routing protocol specification that provides mechanisms achieving the same overall effect.

The features that are substitutable (individually or in a group):

- o Multihop distribution of prefix and 6LoWPAN header compression context
- o Multihop Duplicate Address Detection

Thus, multihop prefix distribution (the ABRO) and the 6LoWPAN Context Option (6CO) for distributing header compression contexts go hand in hand. If substitution is intended for one of them, then both of them MUST be substituted.

Guidelines for feature implementation and deployment are provided in Section 14.

2. 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 [RFC2119].

This specification requires readers to be familiar with all the terms and concepts that are discussed in "Neighbor Discovery for IP version 6 (IPv6)" [RFC4861], "IPv6 Stateless Address Autoconfiguration" [RFC4862], "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [RFC4919], "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944], and "IP Addressing Model in Ad Hoc Networks" [RFC5889].

This specification makes extensive use of the same terminology defined in [RFC4861], unless otherwise defined below.

6LoWPAN link:

A wireless link determined by single IP hop reachability of neighboring nodes. These are considered links with undetermined connectivity properties as in [RFC5889].

6LoWPAN Node (6LN):

A 6LoWPAN node is any host or router participating in a LoWPAN. This term is used when referring to situations in which either a host or router can play the role described.

6LoWPAN Router (6LR):

An intermediate router in the LoWPAN that is able to send and receive Router Advertisements (RAs) and Router Solicitations (RSs) as well as forward and route IPv6 packets. 6LoWPAN routers are present only in route-over topologies.

6LoWPAN Border Router (6LBR):

A border router located at the junction of separate 6LoWPAN networks or between a 6LoWPAN network and another IP network. There may be one or more 6LBRs at the 6LoWPAN network boundary. A 6LBR is the responsible authority for IPv6 prefix propagation for the 6LoWPAN network it is serving. An isolated LoWPAN also contains a 6LBR in the network, which provides the prefix(es) for the isolated network.

Router:

Either a 6LR or a 6LBR. Note that nothing in this document precludes a node being a router on some interfaces and a host on other interfaces as allowed by [RFC2460].

Mesh-under:

A topology where nodes are connected to a 6LBR through a mesh using link-layer forwarding. Thus, in a mesh-under configuration, all IPv6 hosts in a LoWPAN are only one IP hop away from the 6LBR. This topology simulates the typical IP-subnet topology with one router with multiple nodes in the same subnet.

Route-over:

A topology where hosts are connected to the 6LBR through the use of intermediate layer-3 (IP) routing. Here, hosts are typically multiple IP hops away from a 6LBR. The route-over topology typically consists of a 6LBR, a set of 6LRs, and hosts.

Non-transitive link:

A link that exhibits asymmetric reachability as defined in Section 2.2 of [RFC4861].

IP-over-foo document:

A specification that covers operating IP over a particular link type, for example, [RFC4944] "Transmission of IPv6 Packets over IEEE 802.15.4 Networks".

Header compression context:

Address information shared across a LoWPAN and used by 6LoWPAN header compression [RFC6282] to enable the elision of information that would otherwise be sent repeatedly. In a "context", a (potentially partial) address is associated with a Context Identifier (CID), which is then used in header compression as a shortcut for (parts of) a source or destination address.

Registration:

The process during which a LoWPAN node sends a Neighbor Solicitation message with an Address Registration Option to a router creating a Neighbor Cache Entry (NCE) for the LoWPAN node with a specific timeout. Thus, for 6LoWPAN routers, the Neighbor Cache doesn't behave like a cache. Instead, it behaves as a registry of all the host addresses that are attached to the router.

3. Protocol Overview

These Neighbor Discovery optimizations are applicable to both mesh-under and route-over configurations. In a mesh-under configuration, only 6LoWPAN Border Routers and hosts exist; there are no 6LoWPAN routers in mesh-under topologies.

The most important part of the optimizations is the evolved host-to-router interaction that allows for sleeping nodes and avoids using multicast Neighbor Discovery messages except for the case of a host finding an initial set of default routers, and redoing such determination when that set of routers have become unreachable.

The protocol also provides for header compression [RFC6282] by carrying header compression information in a new option in Router Advertisement messages.

In addition, there are separate mechanisms that can be used between 6LRs and 6LBRs to perform multihop Duplicate Address Detection and distribution of the prefix and compression context information from the 6LBRs to all the 6LRs, which in turn use normal Neighbor Discovery mechanisms to convey this information to the hosts.

The protocol is designed so that the host-to-router interaction is not affected by the configuration of the 6LoWPAN; the host-to-router interaction is the same in a mesh-under and route-over configuration.

3.1. Extensions to RFC 4861

This document specifies the following optimizations and extensions to IPv6 Neighbor Discovery [RFC4861]:

- o Host-initiated refresh of Router Advertisement information. This removes the need for periodic or unsolicited Router Advertisements from routers to hosts.
- o No Duplicate Address Detection (DAD) is performed if EUI-64-based IPv6 addresses are used (as these addresses are assumed to be globally unique).

- o DAD is optional if DHCPv6 is used to assign addresses.
- o A new address registration mechanism using a new Address Registration Option between hosts and routers. This removes the need for routers to use multicast Neighbor Solicitations to find hosts and supports sleeping hosts. This also enables the same IPv6 address prefix(es) to be used across a route-over 6LoWPAN. It provides the host-to-router interface for Duplicate Address Detection.
- o A new Router Advertisement option, the 6LoWPAN Context Option, for context information used by 6LoWPAN header compression.
- o A new mechanism to perform Duplicate Address Detection across a route-over 6LoWPAN using the new Duplicate Address Request and Duplicate Address Confirmation messages.
- o New mechanisms to distribute prefixes and context information across a route-over network that uses a new Authoritative Border Router Option to control the flooding of configuration changes.
- o A few new default protocol constants are introduced, and some existing Neighbor Discovery protocol constants are tuned.

3.2. Address Assignment

Hosts in a 6LoWPAN configure their IPv6 addresses as specified in [RFC4861] and [RFC4862] based on the information received in Router Advertisement messages. The use of the M (managed address configuration) flag in this optimization is, however, more restrictive than in [RFC4861]. When the M flag is set, a host is assumed to use DHCPv6 to assign any non-EUI-64 addresses. When the M flag is not set, the nodes in the LoWPAN support Duplicate Address Detection; thus, a host can then safely use the address registration mechanism to check non-EUI-64 addresses for uniqueness.

6LRs MAY use the same mechanisms to configure their IPv6 addresses.

The 6LBRs are responsible for managing the prefix(es) assigned to the 6LoWPAN, using manual configuration, DHCPv6 Prefix Delegation [RFC3633], or other mechanisms. In an isolated LoWPAN, a Unique Local Address (ULA) [RFC4193] prefix SHOULD be generated by the 6LBR.

3.3. Host-to-Router Interaction

A host sends Router Solicitation messages at startup and also when the Neighbor Unreachability Detection (NUD) of one of its default routers fails.

Hosts receive Router Advertisement messages typically containing the Authoritative Border Router Option (ABRO) and may optionally contain one or more 6LoWPAN Context Options (6COs) in addition to the existing Prefix Information Options (PIOs) as described in [RFC4861].

When a host has configured a non-link-local IPv6 address, it registers that address with one or more of its default routers using the Address Registration Option (ARO) in an NS message. The host chooses a lifetime of the registration and repeats the ARO periodically (before the lifetime runs out) to maintain the registration. The lifetime should be chosen in such a way as to maintain the registration even while a host is sleeping. Likewise, mobile nodes that often change their point of attachment should use a suitably short lifetime. See Section 5.5 for registration details and Section 9 for protocol constants.

The registration fails when an ARO is returned to the host with a non-zero Status. One reason may be that the router determines that the IPv6 address is already used by another host, i.e., is used by a host with a different EUI-64. This can be used to support non-EUI-64-based addresses such as temporary IPv6 addresses [RFC4941] or addresses based on an Interface ID that is an IEEE 802.15.4 16-bit short address. Failure can also occur if the Neighbor Cache on that router is full.

The re-registration of an address can be combined with Neighbor Unreachability Detection (NUD) of the router, since both use unicast Neighbor Solicitation messages. This makes things efficient when a host wakes up to send a packet and needs to both perform NUD to check that the router is still reachable and refresh its registration with the router.

The response to an address registration might not be immediate, since in route-over configurations the 6LR might perform Duplicate Address Detection against the 6LBR. A host retransmits the Address Registration Option until it is acknowledged by the receipt of an Address Registration Option.

As part of the optimizations, address resolution is not performed by multicasting Neighbor Solicitation messages as in [RFC4861]. Instead, the routers maintain Neighbor Cache Entries for all registered IPv6 addresses. If the address is not in the Neighbor

Cache in the router, then the address either doesn't exist, is assigned to a host attached to some other router in the 6LoWPAN, or is external to the 6LoWPAN. In a route-over configuration, the routing protocol is used to route such packets toward the destination.

3.4. Router-to-Router Interaction

The new router-to-router interaction is only for the route-over configuration where 6LRs are present. See also Section 1.4.

6LRs MUST act like a host during system startup and prefix configuration by sending Router Solicitation messages and autoconfiguring their IPv6 addresses, unlike routers in [RFC4861].

When multihop prefix and context dissemination are used, then the 6LRs store the ABRO, 6CO, and prefix information received (directly or indirectly) from the 6LBRs and redistribute this information in the Router Advertisement they send to other 6LRs or send to hosts in response to a Router Solicitation. There is a Version Number field in the ABRO (see Section 4.3), which is used to limit the flooding of updated information between the 6LRs.

A 6LR can perform Duplicate Address Detection against one or more 6LBRs using the new Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages, which carry the information from the Address Registration Option. The DAR and DAC messages will be forwarded between the 6LR and 6LBRs; thus, the [RFC4861] rule for checking hop limit=255 does not apply to the DAR and DAC messages. Those multihop DAD messages MUST NOT modify any Neighbor Cache Entries on the routers, since we do not have the security benefits provided by the hop limit=255 check.

3.5. Neighbor Cache Management

The use of explicit registrations with lifetimes, plus the desire to not multicast Neighbor Solicitation messages for hosts, imply that we manage the Neighbor Cache Entries (NCEs) slightly differently than in [RFC4861]. This results in three different types of NCEs, and the types specify how those entries can be removed:

Garbage-collectible: Entries that are subject to the normal rules in [RFC4861] that allow for garbage collection when low on memory.

Registered: Entries that have an explicit registered lifetime and are kept until this lifetime expires or they are explicitly unregistered.

Tentative: Entries that are temporary with a short lifetime, which typically get converted to Registered entries.

Note that the type of the NCE is orthogonal to the states specified in [RFC4861].

When a host interacts with a router by sending Router Solicitations, this results in a Tentative NCE. Once a router has successfully had a node register with it, the result is a Registered NCE. When routers send RAs to hosts, and when routers receive RA messages or receive multicast NS messages from other routers, the result is Garbage-collectible NCEs. There can only be one kind of NCE for an IP address at a time.

Neighbor Cache Entries on routers can additionally be added or deleted by a routing protocol used in the 6LoWPAN. This is useful if the routing protocol carries the link-layer addresses of the neighboring routers. Depending on the details of such routing protocols, such NCEs could be either Registered or Garbage-collectible.

4. New Neighbor Discovery Options and Messages

This section defines new Neighbor Discovery message options used by this specification. The Address Registration Option is used by hosts, whereas the Authoritative Border Router Option and 6LoWPAN Context Option are used in the substitutable router-to-router interaction. This section also defines the new router-to-router Duplicate Address Request and Duplicate Address Confirmation messages.

4.1. Address Registration Option


The routers need to know the set of host IP addresses that are directly reachable and their corresponding link-layer addresses. This needs to be maintained as the radio reachability changes. For this purpose, an Address Registration Option (ARO) is introduced, which can be included in unicast NS messages sent by hosts. Thus, it can be included in the unicast NS messages that a host sends as part of NUD to determine that it can still reach a default router. The ARO is used by the receiving router to reliably maintain its Neighbor Cache. The same option is included in corresponding NA messages with a Status field indicating the success or failure of the registration. This option is always host initiated.

The information contained in the ARO is also included in the multihop DAR and DAC messages used between 6LRs and 6LBRs, but the option itself is not used in those messages.

The ARO is required for reliability and power saving. The lifetime field provides flexibility to the host to register an address that should be usable (continue to be advertised by the 6LR in the routing protocol, etc.) during its intended sleep schedule.

The sender of the NS also includes the EUI-64 [EUI64] of the interface from which it is registering an address. This is used as a unique ID for the detection of duplicate addresses. It is used to tell the difference between the same node re-registering its address and a different node (with a different EUI-64) registering an address that is already in use by someone else. The EUI-64 is also used to deliver an NA carrying an error Status code to the EUI-64-based link-local IPv6 address of the host (see Section 6.5.2).

When the ARO is used by hosts, an SLLAO (Source Link-Layer Address Option) [RFC4861] MUST be included, and the address that is to be registered MUST be the IPv6 source address of the NS message.



Fields:

- Type: 33
- Length: 8-bit unsigned integer. The length of the option in units of 8 bytes. Always 2.
- Status: 8-bit unsigned integer. Indicates the status of a registration in the NA response. MUST be set to 0 in NS messages. See below.
- Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Registration Lifetime: 16-bit unsigned integer. The amount of time in units of 60 seconds that the router should retain the NCE for the sender of the NS that includes this option.

EUI-64: 64 bits. This field is used to uniquely identify the interface of the Registered Address by including the EUI-64 identifier [EUI64] assigned to it unmodified.

The Status values used in NAs are:

Status	Description
0	Success
1	Duplicate Address
2	Neighbor Cache Full
3-255	Allocated using Standards Action [RFC5226]

Table 1

4.2. 6LoWPAN Context Option

The 6LoWPAN Context Option (6CO) carries prefix information for LoWPAN header compression and is similar to the PIO of [RFC4861]. However, the prefixes can be remote as well as local to the LoWPAN, since header compression potentially applies to all IPv6 addresses. This option allows for the dissemination of multiple contexts identified by a CID for use as specified in [RFC6282]. A context may be a prefix of any length or an address (/128), and up to 16 6COs may be carried in an RA message.




Figure 1: 6LoWPAN Context Option Format

Type: 34

Length: 8-bit unsigned integer. The length of the option (including the Type and Length fields) in units of 8 bytes. May be 2 or 3, depending on the length of the Context Prefix field.

Context Length: 8-bit unsigned integer. The number of leading bits in the Context Prefix field that are valid. The value ranges from 0 to 128. If it is more than 64, then the Length MUST be 3.

C: 1-bit context Compression flag. This flag indicates if the context is valid for use in compression. A context that is not valid MUST NOT be used for compression but SHOULD be used in decompression in case another compressor has not yet received the updated context information. This flag is used to manage the context life cycle based on the recommendations in Section 7.2.

CID: 4-bit Context Identifier for this prefix information. The CID is used by context-based header compression as specified in [RFC6282]. The list of CIDs for a LoWPAN is configured on the 6LBR that originates the context information for the 6LoWPAN.

Res, Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

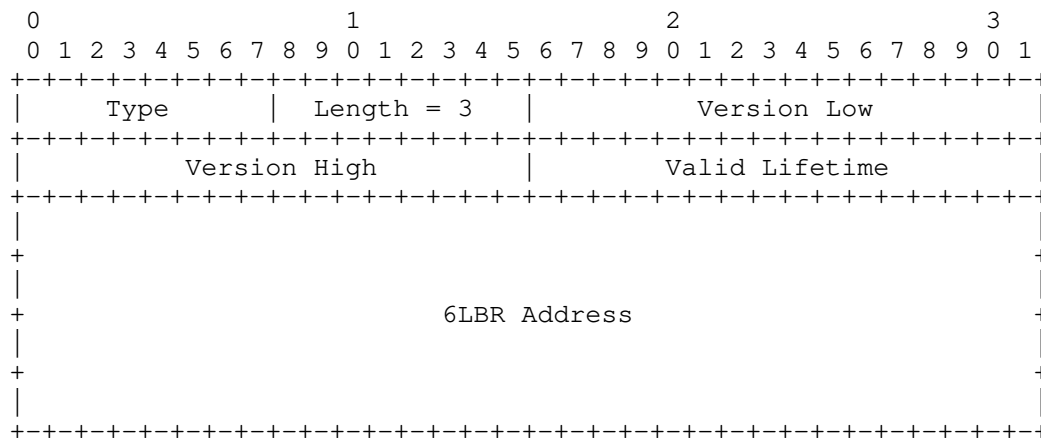
Valid Lifetime: 16-bit unsigned integer. The length of time in units of 60 seconds (relative to the time the packet is received) that the context is valid for the purpose of header compression or decompression. A value of all zero bits (0x0) indicates that this context entry MUST be removed immediately.

Context Prefix: The IPv6 prefix or address corresponding to the CID field. The valid length of this field is included in the Context Length field. This field is padded with zeros in order to make the option a multiple of 8 bytes.

4.3. Authoritative Border Router Option

The Authoritative Border Router Option (ABRO) is needed when RA messages are used to disseminate prefixes and context information across a route-over topology. In this case, 6LRs receive PIOs from other 6LRs. This implies that a 6LR can't just let the most recently received RA win. In order to be able to reliably add and remove prefixes from the 6LoWPAN, we need to carry information from the authoritative 6LBR. This is done by introducing a version number that the 6LBR sets and that 6LRs propagate as they propagate the prefix and context information with this ABRO. When there are multiple 6LBRs, they would have separate version number spaces. Thus, this option needs to carry the IP address of the 6LBR that originated that set of information.

The ABRO MUST be included in all RA messages in the case when RAs are used to propagate information between routers (as described in Section 8.2).



Fields:

Type: 35

Length: 8-bit unsigned integer. The length of the option in units of 8 bytes. Always 3.

Version Low, Version High: Together, Version Low and Version High constitute the Version Number field, a 32-bit unsigned integer where Version Low is the least significant 16 bits and Version High is the most significant

16 bits. The version number corresponding to this set of information contained in the RA message. The authoritative 6LBR originating the prefix increases this version number each time its set of prefix or context information changes.

Valid Lifetime: 16-bit unsigned integer. The length of time in units of 60 seconds (relative to the time the packet is received) that this set of border router information is valid. A value of all zero bits (0x0) assumes a default value of 10,000 (~one week).

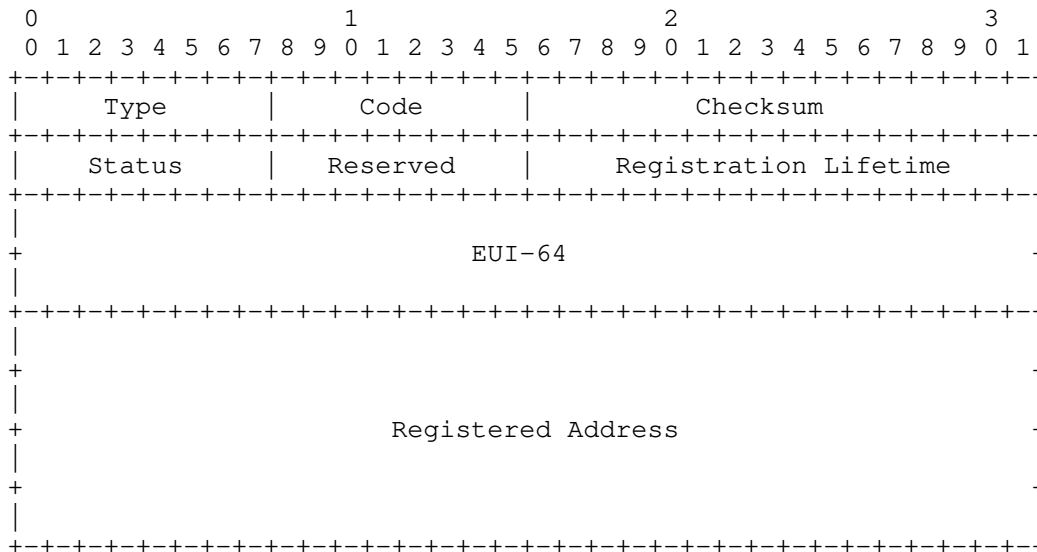
Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

6LBR Address: IPv6 address of the 6LBR that is the origin of the included version number.

4.4. Duplicate Address Messages

For the multihop DAD exchanges between a 6LR and 6LBR as specified in Section 8.2, there are two new ICMPv6 message types called the Duplicate Address Request (DAR) and the Duplicate Address Confirmation (DAC). We avoid reusing the NS and NA messages for this purpose, since these messages are not subject to the hop limit=255 check as they are forwarded by intermediate 6LRs. The information contained in the messages is otherwise the same as would be in an NS carrying an ARO, with the message format inlining the fields that are in the ARO.

The DAR and DAC use the same message format with different ICMPv6 type values, and the Status field is only meaningful in the DAC message.



IP fields:

- IPv6 Source: A non-link-local address of the sending router.
- IPv6 Destination: In a DAR, a non-link-local address of a 6LBR. In a DAC, this is just the source from the DAR.
- Hop Limit: Set to MULTI_HOP_HOPLIMIT on transmit. MUST be ignored on receipt.

ICMP Fields:

- Type: 157 for the DAR and 158 for the DAC.
- Code: Set to zero on transmit. MUST be ignored on receipt.
- Checksum: The ICMP checksum. See [RFC4443].
- Status: 8-bit unsigned integer. Indicates the status of a registration in the DAC. MUST be set to 0 in the DAR. See Table 1.
- Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Registration Lifetime: 16-bit unsigned integer. The amount of time in units of 60 seconds that the 6LBR should retain the DAD table entry (Section 8.2.2) for the Registered Address. A value of 0 indicates in a DAR that the DAD table entry should be removed.

EUI-64: 64 bits. This field is used to uniquely identify the interface of the Registered Address by including the EUI-64 identifier [EUI64] assigned to it unmodified.

Registered Address: 128-bit field. Carries the host address that was contained in the IPv6 Source field in the NS that contained the ARO sent by the host.

5. Host Behavior

Hosts in a LoWPAN use the ARO in the NS messages they send as a way to maintain the Neighbor Cache in the routers, thereby removing the need for multicast NSs to do address resolution. Unlike in [RFC4861], the hosts initiate updating the information they receive in RAs by sending RSs before the information expires. Finally, when NUD indicates that one or all default routers have become unreachable, then the host uses RSs to find a new set of default routers.

5.1. Forbidden Actions

A host **MUST NOT** multicast an NS message.

5.2. Interface Initialization

When the interface on a host is initialized, it follows the specification in [RFC4861]. A link-local address is formed based on the EUI-64 identifier [EUI64] assigned to the interface as per [RFC4944] or the appropriate IP-over-foo document for the link, and then the host sends RS messages as described in [RFC4861] Section 6.3.7.

There is no need to join the solicited-node multicast address, since nobody multicasts NSs in this type of network. A host **MUST** join the all-nodes multicast address.

5.3. Sending a Router Solicitation

The RS is formatted as specified in [RFC4861] and sent to the IPv6 all-routers multicast address (see [RFC4861] Section 6.3.7 for details). An SLLAO MUST be included to enable unicast RAs in response. An unspecified source address MUST NOT be used in RS messages.

If the link layer supports a way to send packets to some kind of all-routers anycast link-layer address, then that MAY be used to convey these packets to a router.

Since hosts do not depend on multicast RAs to discover routers, the hosts need to intelligently retransmit RSs whenever the default router list is empty, one of its default routers becomes unreachable, or the lifetime of the prefixes and contexts in the previous RA is about to expire. The RECOMMENDED rate of retransmissions is to initially send up to 3 (MAX_RTR_SOLICITATIONS) RS messages separated by at least 10 seconds (RTR_SOLICITATION_INTERVAL) as specified in [RFC4861], and then switch to slower retransmissions. After the initial retransmissions, the host SHOULD do truncated binary exponential backoff [ETHERNET] of the retransmission timer for each subsequent retransmission, truncating the increase of the retransmission timer at 60 seconds (MAX_RTR_SOLICITATION_INTERVAL). In all cases, the RS retransmissions are terminated when an RA is received. See Section 9 for protocol constants.

5.4. Processing a Router Advertisement

The processing of RAs is as in [RFC4861], with the addition of handling the 6CO and triggering address registration when a new address has been configured. Furthermore, the SLLAO MUST be included in the RA. Unlike in [RFC4861], the maximum value of the RA Router Lifetime field MAY be up to 0xFFFF (approximately 18 hours).

Should the host erroneously receive a PIO with the L (on-link) flag set, then that PIO MUST be ignored.

5.4.1. Address Configuration

Address configuration follows [RFC4862]. For an address not derived from an EUI-64, the M flag of the RA determines how the address can be configured. If the M flag is set in the RA, then DHCPv6 MUST be used to assign the address. If the M flag is not set, then the address can be configured by any other means (and duplicate detection is performed as part of the registration process).

Once an address has been configured, it will be registered by unicasting an NS with an ARO to one or more routers.

5.4.2. Storing Contexts

The host maintains a conceptual data structure for the context information it receives from the routers. This structure is called the context table. It includes the CID, the prefix (from the Context Prefix field in the 6CO), the Compression bit, and the Valid Lifetime. A context table entry that has the Compression bit clear is used for decompression when receiving packets but MUST NOT be used for compression when sending packets.

When a 6CO is received in an RA, it is used to add or update the information in the context table. If the CID field in the 6CO matches an existing context table entry, then that entry is updated with the information in the 6CO. If the Valid Lifetime field in the 6CO is zero, then the entry is immediately deleted.

If there is no matching entry in the context table, and the Valid Lifetime field is non-zero, then a new context is added to the context table. The 6CO is used to update the created entry.

When the 6LBR changes the context information, a host might not immediately notice. And in the worst case, a host might have stale context information. For this reason, 6LBRs use the recommendations in Section 7.2 for carefully managing the context life cycle. Nodes should be careful about using header compression in RA messages that include 6COs.

5.4.3. Maintaining Prefix and Context Information

The prefix information is timed out as specified in [RFC4861]. When the Valid Lifetime for a context table entry expires, the entry is placed in a receive-only mode, which is the equivalent of receiving a 6CO for that context with C=0. The entry is held in receive-only mode for a period of twice the default Router Lifetime, after which the entry is removed.

A host should inspect the various lifetimes to determine when it should next initiate sending an RS to ask for any updates to the information. The lifetimes that matter are the default Router Lifetime, the Valid Lifetime in the PIOs, and the Valid Lifetime in the 6CO. The host SHOULD unicast one or more RSs to the router well before the shortest of those lifetimes (across all the prefixes and all the contexts) expires and then switch to multicast RS messages if there is no response to the unicasts. The retransmission behavior for the RSs is specified in Section 5.3.

5.5. Registration and Neighbor Unreachability Detection

Hosts send unicast NS messages to register their IPv6 addresses, and also to do NUD to verify that their default routers are still reachable. The registration is performed by the host including an ARO in the NS it sends. Even if the host doesn't have data to send, but is expecting others to try to send packets to the host, the host needs to maintain its NCEs in the routers. This is done by sending NS messages with an ARO to the router well in advance of the Registration Lifetime expiring. NS messages are retransmitted up to MAX_UNICAST_SOLICIT times using a minimum timeout of RETRANS_TIMER until the host receives an NA message with an ARO.

Hosts that receive RA messages from multiple default routers SHOULD attempt to register with more than one of them in order to increase the robustness of the network.

Note that NUD probes can be suppressed by reachability confirmations from transport protocols or applications as specified in [RFC4861].

When a host knows it will no longer use a router it is registered to, it SHOULD de-register with the router by sending an NS with an ARO containing a lifetime of 0. To handle the case when a host loses connectivity with the default router involuntarily, the host SHOULD use a suitably low Registration Lifetime.

5.5.1. Sending a Neighbor Solicitation

The host triggers sending NS messages containing an ARO when a new address is configured, when it discovers a new default router, or well before the Registration Lifetime expires. Such an NS MUST include an SLLAO, since the router needs to record the link-layer address of the host. An unspecified source address MUST NOT be used in NS messages.

5.5.2. Processing a Neighbor Advertisement

A host handles NA messages as specified in [RFC4861], with added logic described in this section for handling the ARO.

In addition to the normal validation of an NA and its options, the ARO (if present) is verified as follows. If the Length field is not two, the option is silently ignored. If the EUI-64 field does not match the EUI-64 of the interface, the option is silently ignored.

If the Status field is zero, then the address registration was successful. The host saves the Registration Lifetime from the ARO for use to trigger a new NS well before the lifetime expires. If the Status field is not equal to zero, the address registration has failed.

5.5.3. Recovering from Failures

The procedure for maintaining reachability information about a neighbor is the same as in [RFC4861] Section 7.3, with the exception that address resolution is not performed.

The address registration procedure may fail for two reasons: no response to NSs is received (NUD failure), or an ARO with a failure Status (Status > 0) is received. In the case of NUD failure, the entry for that router will be removed; thus, address registration is no longer of importance. When an ARO with a non-zero Status field is received, this indicates that registration for that address has failed. A failure Status of one indicates that a duplicate address was detected, and the procedure described in [RFC4862] Section 5.4.5 is followed. The host MUST NOT use the address it tried to register. If the host has valid registrations with other routers, these MUST be removed by registering with each using a zero ARO lifetime.

A Status code of two indicates that the Neighbor Cache of that router is full. In this case, the host SHOULD remove this router from its default router list and attempt to register with another router. If the host's default router list is empty, it needs to revert to sending RSs as specified in Section 5.3.

Other failure codes may be defined in future documents.

5.6. Next-Hop Determination

The IP address of the next hop for a destination is determined as follows. Destinations to the link-local prefix (fe80::) are always sent on the link to that destination. It is assumed that link-local addresses are formed as specified in Section 5.2 from the EUI-64, and address resolution is not performed. Packets are sent to link-local destinations by reversing the procedure in Appendix A of [RFC4291].

Multicast addresses are considered to be on-link and are resolved as specified in [RFC4944] or the appropriate IP-over-foo document. Note that [RFC4944] only defines how to represent a multicast destination address in the LoWPAN header. Support for multicast scopes larger than link-local needs an appropriate multicast routing algorithm.

All other prefixes are assumed to be off-link [RFC5889]. Anycast addresses are always considered to be off-link. They are therefore sent to one of the routers in the default router list.

A LoWPAN node is not required to maintain a minimum of one buffer per neighbor as specified in [RFC4861], since packets are never queued while waiting for address resolution.

5.7. Address Resolution

The address registration mechanism and the SLLAO in RA messages provide sufficient a priori state in routers and hosts to resolve an IPv6 address to its associated link-layer address. As all prefixes except the link-local prefix and multicast addresses are always assumed to be off-link, multicast-based address resolution between neighbors is not needed.

Link-layer addresses for neighbors are stored in NCEs [RFC4861]. In order to achieve LoWPAN compression, most global addresses are formed using a link-layer address. Thus, a host can reduce memory usage by optimizing for this case and only storing link-layer address information if it differs from the link-layer address corresponding to the Interface ID of the IPv6 address (i.e., differs in more than the on-link/global bit being inverted).

5.8. Sleeping

It is often advantageous for battery-powered hosts in LoWPANs to keep a low duty cycle. The optimizations described in this document enable hosts to sleep, as further described in this section. Routers may want to cache traffic destined to a host that is sleeping, but such functionality is out of the scope of this document.

5.8.1. Picking an Appropriate Registration Lifetime

As all ND messages are initiated by the hosts, this allows a host to sleep or otherwise be unreachable between NS/NA message exchanges. The ARO attached to NS messages indicates to a router to keep the NCE for that address valid for the period in the Registration Lifetime field. A host should choose a sleep time appropriate for its energy characteristics and set a Registration Lifetime larger than the sleep time to ensure that the registration is renewed successfully (considering, for example, clock drift and additional time for potential retransmissions of the re-registration). External configuration of a host should also consider the stability of the network (how quickly the topology changes) when choosing its sleep

time (and thus Registration Lifetime). A dynamic network requires a shorter sleep time so that routers don't keep invalid NCEs for nodes longer than necessary.

5.8.2. Behavior on Wakeup

When a host wakes up from a sleep period, it SHOULD refresh its current address registrations that will time out before the next wakeup. This is done by sending NS messages with an ARO as described in Section 5.5.1. The host may also need to refresh its prefix and context information by sending a new unicast RS (the maximum Router Lifetime is about 18 hours, whereas the maximum Registration Lifetime is about 45.5 days). If after wakeup the host (using NUD) determines that some or all previous default routers have become unreachable, then the host will send multicast RSs to discover new default router(s) and restart the address registration process.

6. Router Behavior for 6LRs and 6LBRs

Both 6LRs and 6LBRs maintain the Neighbor Cache [RFC4861] based on the AROs they receive in NA messages from hosts, ND packets from other nodes, and, potentially, a routing protocol used in the 6LoWPAN as outlined in Section 3.5.

The routers SHOULD NOT garbage-collect Registered NCEs (see Section 3.4), since they need to retain them until the Registration Lifetime expires. Similarly, if NUD on the router determines that the host is UNREACHABLE (based on the logic in [RFC4861]), the NCE SHOULD NOT be deleted but rather retained until the Registration Lifetime expires. A renewed ARO should mark the cache entry as STALE. Thus, for 6LoWPAN routers, the Neighbor Cache doesn't behave like a cache. Instead, it behaves as a registry of all the host addresses that are attached to the router.

Routers MAY implement the Default Router Preference (Prf) extension [RFC4191] and use that to indicate to the host whether the router is a 6LBR or a 6LR. If this is implemented, then 6LRs with no route to a border router MUST set Prf to (11) for low preference, other 6LRs MUST set Prf to (00) for normal preference, and 6LBRs MUST set Prf to (01) for high preference.

6.1. Forbidden Actions

Even if a router in a route-over topology can reach both a host and another target, because of radio propagation it generally cannot know whether the host can directly reach the other target. Therefore, it cannot assume that Redirect will actually work from one host to another. Therefore, it SHOULD NOT send Redirect messages. The only

potential exception to this "SHOULD NOT" is when the deployment/implementation has a way to know how the host can reach the intended target. Hence, it is RECOMMENDED that the implementation by default does not send Redirect messages but can be configurable when the deployment calls for this. In contrast, for mesh-under topologies, the same considerations about Redirects apply as in [RFC4861].

A router MUST NOT set the L (on-link) flag in the PIOs, since that might trigger hosts to send multicast NSs.

6.2. Interface Initialization

The 6LBR router interface initialization behavior is the same as in [RFC4861]. However, in a dynamic configuration scenario (see Section 8.1), a 6LR comes up as a non-router and waits to receive the advertisement for configuring its own interface address first, before setting its interfaces to be advertising interfaces and turning into a router.

6.3. Processing a Router Solicitation

A router processes RS messages as specified in [RFC4861]. The differences relate to the inclusion of ABROs in the RA messages and the exclusive use of unicast RAs. If a 6LR has received an ABRO from a 6LBR, then it will include that option unmodified in the RA messages it sends. And, if the 6LR has received RAs -- whether with the same prefixes and context information or different -- from a different 6LBR, then it will need to keep those prefixes and that context information separately so that the RAs the 6LR sends will maintain the association between the ABRO and the prefixes and context information. The router can tell which 6LBR originated the prefixes and context information from the 6LBR Address field in the ABRO. When a router has information tied to multiple ABROs, a single RS will result in multiple RAs each containing a different ABRO.

When the ABRO Valid Lifetime associated with a 6LBR times out, all information related to that 6LBR MUST be removed. As an implementation note, it is recommended that RAs are sent sufficiently more frequently than the ABRO Valid Lifetime so that missing an RA does not result in removing all information related to a 6LBR.

An RS might be received from a host that has not yet registered its address with the router. Thus, the router MUST NOT modify an existing NCE based on the SLLAO from the RS. However, a router MAY create a Tentative NCE based on the SLLAO. Such a Tentative NCE SHOULD be timed out in TENTATIVE_NCE_LIFETIME seconds, unless a registration converts it into a Registered NCE.

A 6LR or 6LBR MUST include an SLLAO in the RAs it sends; this is required so that the hosts will know the link-layer address of the router. Unlike in [RFC4861], the maximum value of the RA Router Lifetime field MAY be up to 0xFFFF (approximately 18 hours).

Unlike [RFC4861], which suggests multicast RAs, this specification improves the exchange by always unicasting RAs in response to RSs. This is possible, since the RS always includes an SLLAO, which is used by the router to unicast the RA.

6.4. Periodic Router Advertisements

A router does not need to send any periodic RA messages, since the hosts will solicit updated information by sending RSs before the lifetimes expire.

However, if the routers use RAs to distribute prefix and/or context information across a route-over topology, that might require periodic RA messages. Such RAs are sent using the configurable MinRtrAdvInterval and MaxRtrAdvInterval as per [RFC4861].

6.5. Processing a Neighbor Solicitation

A router handles NS messages as specified in [RFC4861], with added logic described in this section for handling the ARO.

In addition to the normal validation of an NS and its options, the ARO is verified as follows (if present). If the Length field is not two, or if the Status field is not zero, then the NS is silently ignored.

If the source address of the NS is the unspecified address, or if no SLLAO is included, then any included ARO is ignored, that is, the NS is processed as if it did not contain an ARO.

6.5.1. Checking for Duplicates

If the NS contains a valid ARO, then the router inspects its Neighbor Cache on the arriving interface to see if it is a duplicate. It isn't a duplicate if (1) there is no NCE for the IPv6 source address of the NS or (2) there is such an NCE and the EUI-64 is the same. Otherwise, it is a duplicate address. Note that if multihop DAD (Section 8.2) is used, then the checks are slightly different, to take into account Tentative NCEs. In the case where it is a duplicate address, then the router responds with a unicast NA message with the ARO Status field set to one (to indicate that the address is a duplicate) as described in Section 6.5.2. In this case, there is no modification to the Neighbor Cache.

6.5.2. Returning Address Registration Errors

Address registration errors are not sent back to the source address of the NS due to a possible risk of L2 address collision. Instead, the NA is sent to the link-local IPv6 address with the Interface ID part derived from the EUI-64 field of the ARO as per [RFC4944]. In particular, this means that the universal/local bit needs to be inverted. The NA is formatted with a copy of the ARO from the NS, but with the Status field set to indicate the appropriate error.

The error is sent to the link-local address with the Interface ID derived from the EUI-64. Thus, if the ARO was from and for a short address, the L2 destination address for the NA with the ARO error will be the 64-bit unique address.

6.5.3. Updating the Neighbor Cache

If the ARO did not result in a duplicate address being detected as above, then if the Registration Lifetime is non-zero the router creates (if it didn't exist) or updates (otherwise) an NCE for the IPv6 source address of the NS. If the Neighbor Cache is full and a new entry needs to be created, then the router responds with a unicast NA with the ARO Status field set to two (to indicate that the router's Neighbor Cache is full) as described in Section 6.5.2.

The Registration Lifetime and the EUI-64 are recorded in the NCE. A unicast NA is then sent in response to the NS. This NA SHOULD include a copy of the ARO, with the Status field set to zero. A TLLAO (Target Link-Layer Address Option) [RFC4861] is not required in the NA, since the host already knows the router's link-layer address from RAs.

If the ARO contains a zero Registration Lifetime, then any existing NCE for the IPv6 source address of the NS MUST be deleted and an NA sent as above.

Should the Registration Lifetime in an NCE expire, then the router MUST delete the cache entry.

The addition and removal of Registered NCEs would result in notifying the routing protocol.

Note: If the substitutable multihop DAD (Section 8.2) is used, then the updating of the Neighbor Cache is slightly different due to Tentative NCEs.

6.5.4. Next-Hop Determination

In order to deliver a packet destined for a 6LN registered with a router, next-hop determination is slightly different for routers than for hosts (see Section 5.6). The routing table is checked to determine the next-hop IP address. A Registered NCE determines if the next-hop IP address is on-link. It is the responsibility of the routing protocol of the router to maintain on-link information about its registered neighbors. Tentative NCEs MUST NOT be used to determine on-link status of the registered nodes.

6.5.5. Address Resolution between Routers

There needs to be a mechanism somewhere for the routers to discover each other's link-layer addresses. If the routing protocol used between the routers provides this, then there is no need for the routers to use the ARO between each other. Otherwise, the routers SHOULD use the ARO. When routers use the ARO to register with each other and multihop DAD (Section 8.2) is in use, then care must be taken to ensure that there isn't a flood of ARO-carrying messages sent to the 6LBR as each router hears an ARO from their neighboring routers. The details for this scenario are out of scope of this document.

Routers MAY also use multicast NSs as in [RFC4861] to resolve each others link-layer addresses. Thus, routers MAY multicast NSs for other routers, for example, as a result of receiving some routing protocol update. Routers MUST respond to multicast NSs. This implies that routers MUST join the solicited-node multicast addresses as specified in [RFC4861].

7. Border Router Behavior

A 6LBR handles the sending of RAs and processing of NSs from hosts as specified above in Section 6. A 6LBR SHOULD always include an ABRO in the RAs it sends, listing itself as the 6LBR address. This requires that the 6LBR maintain the version number in stable storage and increase the version number when some information in its RAs changes. The information whose change affects the version is in the PIOs (the prefixes or their lifetimes) and in the 6CO (the prefixes, CIDs, or lifetimes).

In addition, a 6LBR is somehow configured with the prefix or prefixes that are assigned to the LoWPAN and advertises those in RAs as in [RFC4861]. In the case of route-over, those prefixes can be disseminated to all the 6LRs using the technique discussed in

Section 8.1. However, there might be mechanisms outside of the scope of this document that can be used as a substitute for prefix dissemination in the route-over topology (see Section 1.4).

If the 6LoWPAN uses header compression [RFC6282] with context, then the 6LBR needs to manage the CIDs and advertise those in RAs by including 6COs in its RAs so that directly attached hosts are informed about the CIDs. Below, we specify things to consider when the 6LBR needs to add, remove, or change the context information. In the case of route-over, the context information is disseminated to all the 6LRs using the technique discussed in Section 8, unless a different specification provides a substitute for this multihop distribution.

7.1. Prefix Determination

The prefix or prefixes used in a LoWPAN can be manually configured or can be acquired using DHCPv6 Prefix Delegation [RFC3633]. For a LoWPAN that is isolated from the network either permanently or occasionally, the 6LBR can assign a ULA prefix using [RFC4193]. The ULA prefix should be stored in stable storage so that the same prefix is used after a failure of the 6LBR. If the LoWPAN has multiple 6LBRs, then they should be configured with the same set of prefixes. The set of prefixes is included in the RA messages as specified in [RFC4861].

7.2. Context Configuration and Management

If the LoWPAN uses header compression [RFC6282] with context, then the 6LBR must be configured with context information and related CIDs. If the LoWPAN has multiple 6LBRs, then they MUST be configured with the same context information and CIDs. As noted in [RFC6282], maintaining consistency of context information is crucial for ensuring that packets will be decompressed correctly.

The context information carried in RA messages originates at 6LBRs and must be disseminated to all the routers and hosts within the LoWPAN. RAs include one 6CO for each context.

For the dissemination of context information using the 6CO, a strict life cycle SHOULD be used in order to ensure that the context information stays synchronized throughout the LoWPAN. New context information SHOULD be introduced into the LoWPAN with C=0, to ensure that it is known by all nodes that may have to perform header decompression based on this context information. Only when it is reasonable to assume that this information was successfully disseminated SHOULD an option with C=1 be sent, enabling the actual use of the context information for compression.

Conversely, to avoid the situation where nodes send packets that make use of previous values of contexts -- which would result in ambiguity when receiving a packet that uses a recently changed context -- old values of a context SHOULD be taken out of use for a while before new values are assigned to this specific context. That is, in preparation for a change of context information, its dissemination SHOULD continue for at least MIN_CONTEXT_CHANGE_DELAY with C=0. Only when it is reasonable to assume that the fact that the context is now invalid was successfully disseminated should the CID be taken out of dissemination or reused with a different Context Prefix field. In the latter case, dissemination of the new value again SHOULD start with C=0, as above.

8. Substitutable Feature Behavior

Normally, in a 6LoWPAN multihop network, the RA messages are used to disseminate prefixes and context information to all the 6LRs in a route-over topology. If all routers are configured to use a substitute mechanism for such information distribution, any remaining use of the 6LoWPAN-ND mechanisms is governed by the substitute specification.

There is also the option for a 6LR to perform multihop DAD (for IPv6 addresses not derived from an EUI-64) against a 6LBR in a route-over topology by using the DAR and DAC messages. This is substitutable because there might be other ways to either allocate a unique address, such as DHCPv6 [RFC3315], or use other future mechanisms for multihop DAD. Again, in this case, any remaining use of the 6LoWPAN-ND mechanisms is governed by the substitute specification.

To be clear: Implementations MUST support the features described in Sections 8.1 and 8.2, unless the implementation supports some alternative ("substitute") from some other specification.

8.1. Multihop Prefix and Context Distribution

The multihop distribution relies on RS messages and RA messages sent between routers, and using the ABRO version number to control the propagation of the information (prefixes and context information) that is being sent in the RAs.

This multihop distribution mechanism can handle arbitrary information from an arbitrary number of 6LBRs. However, the semantics of the context information requires that all the 6LRs use the same information whether they send, forward, or receive compressed packets. Thus, the manager of the 6LBRs needs to somehow ensure that the context information is in synchrony across the 6LBRs. This can be handled in different ways. One possible way to ensure it is to

treat the context and prefix information as originating from some logical or virtual source, which in essence means that it looks like the information is distributed from a single source.

If a set of 6LBRs behave as a single one (using mechanisms out of scope of this document) so that the prefixes and contexts and the ABRO version number will be the same from all the 6LBRs, then those 6LBRs can pick a single IP address to use in the ABRO.

8.1.1. 6LBRs Sending Router Advertisements

6LBRs supporting multihop prefix and context distribution MUST include an ABRO in each of their RAs. The ABRO Version Number field is used to keep prefix and context information consistent throughout the LoWPAN, along with the guidelines in Section 7.2. Each time any information in the set of PIOs or 6COs changes, the ABRO version is increased by one.

This requires that the 6LBR maintain the PIO, 6CO, and ABRO Version Number in stable storage, since an old version number will be silently ignored by the 6LRs.

8.1.2. Routers Sending Router Solicitations

In a 6LoWPAN, unless substituted, multihop distribution is done using RA messages. Thus, on interface initialization, a router (6LR) MUST send RS messages following the rules specified for hosts in [RFC4861]. This in turn will cause the routers to respond with RA messages that can then be used to initially seed the prefix and context information.

8.1.3. Routers Processing Router Advertisements

If multihop distribution is not done using RA messages, then the routers follow [RFC4861], which states that they merely do some consistency checks; in this case, nothing in Section 8.1 applies. Otherwise, the routers will check and record the prefix and context information from the received RAs, and use that information as follows.

If a received RA does not contain an ABRO, then the RA MUST be silently ignored.

The router uses the 6LBR Address field in the ABRO to check if it has previously received information from the 6LBR. If it finds no such information, then it just records the 6LBR address, Version, Valid Lifetime, and the associated prefixes and context information. If the 6LBR is previously known, then the Version Number field MUST be

compared against the recorded version number for that 6LBR. If the version number received in the packet is less than the stored version number, then the information in the RA is silently ignored. Otherwise, the recorded information and version number are updated.

8.1.4. Storing the Information

The router keeps state for each 6LBR that it sees with an ABRO. This includes the version number, the Valid Lifetime, and the complete set of PIOs and 6COs. The prefixes are timed out based on the Valid Lifetime in the PIO. The Context Prefix is timed out based on the Valid Lifetime in the 6CO.

While the prefixes and context information are stored in the router, their valid and preferred lifetimes are decremented as time passes. This ensures that when the router is in turn later advertising that information in the RAs it sends, the 'expiry time' doesn't accidentally move further into the future. For example, if a 6CO with a Valid Lifetime of 10 minutes is received at time T, and the router includes this in an RA it sends at time T+5 minutes, the Valid Lifetime in the 6CO it sends will be only 5 minutes.

8.1.5. Sending Router Advertisements

When multihop distribution is performed using RA messages, the routers MUST ensure that the ABRO always stays together with the prefixes and context information received with that ABRO. Thus, if the router has received prefix P1 with an ABRO saying it is from one 6LBR, and prefix P2 from another 6LBR, then the router MUST NOT include the two prefixes in the same RA message. Prefix P1 MUST be in an RA that includes an ABRO from the first 6LBR, etc. Note that multiple 6LBRs might advertise the same prefix and context information, but they still need to be associated with the 6LBRs that advertised them.

The routers periodically send RAs as in [RFC4861]. This is for the benefit of the other routers receiving the prefixes and context information. The routers also respond to RSs by unicasting RA messages. In both cases, the above constraint of keeping the ABRO together with 'its' prefixes and context information applies.

When a router receives new information from a 6LBR, that is, either it hears from a new 6LBR (a new 6LBR address in the ABRO) or the ABRO version number of an existing 6LBR has increased, then it is useful to send out a few triggered updates. The recommendation is to behave the same as when an interface has become an advertising interface as described in [RFC4861], that is, send up to three RA messages. This ensures rapid propagation of new information to all the 6LRs.

8.2. Multihop Duplicate Address Detection

The ARO can be used, in addition to registering an address in a 6LR, to have the 6LR verify that the address isn't used by some other host known to the 6LR. However, that isn't sufficient in a route-over topology (or in a LoWPAN with multiple 6LBRs), since some host attached to another 6LR could be using the same address. There might be different ways for the 6LRs to coordinate such duplicate address detection in the future, or addresses could be assigned using a DHCPv6 server that verifies uniqueness as part of the assignment.

This specification offers a substitutable simple technique for 6LRs and 6LBRs to perform DAD that reuses the information from the ARO in the DAR and DAC messages. This technique is not needed when the Interface ID in the address is based on an EUI-64, since those are assumed to be globally unique. The technique assumes that either the 6LRs register with all the 6LBRs or the network uses some out-of-scope mechanism to keep the DAD tables in the 6LBRs synchronized.

The multihop DAD mechanism is used synchronously the first time an address is registered with a particular 6LR. That is, the ARO is not returned to the host until multihop DAD has been completed against the 6LBRs. For existing registrations in the 6LR, multihop DAD needs to be repeated against the 6LBRs to ensure that the entry for the address in the 6LBRs does not time out, but that can be done asynchronously with the response to the hosts. One method to achieve this is to track how much is left of the lifetime the 6LR registered with the 6LBRs and to re-register with the 6LBR when this lifetime is about to run out.

For synchronous multihop DAD, the 6LR performs some additional checks to ensure that it has an NCE it can use to respond to the host when it receives a response from a 6LBR. This consists of checking for an already existing (Tentative or Registered) NCE for the Registered Address with a different EUI-64. If such a Registered NCE exists, then the 6LR SHOULD respond that the address is a duplicate. If such a Tentative NCE exists, then the 6LR SHOULD silently ignore the ARO, thereby relying on the host retransmitting the ARO. This is needed to handle the case when multiple hosts try to register the same IPv6 address at the same time. If no NCE exists, then the 6LR MUST create a Tentative NCE with the EUI-64 and the SLLAO. This entry will be used to send the response to the host when the 6LBR responds positively.

When a 6LR receives an NS containing an ARO with a non-zero Registration Lifetime and it has no existing Registered NCE, then with this mechanism the 6LR will invoke synchronous multihop DAD.

The 6LR will unicast a DAR message to one or more 6LBRs, where the DAR contains the host's address in the Registered Address field. The DAR will be forwarded by 6LRs until it reaches the 6LBR; hence, its IPv6 Hop Limit field will not be 255 when received by the 6LBR. The 6LBR will respond with a DAC message, which will have a hop limit less than 255 when it reaches the 6LR.

When the 6LR receives the DAC from the 6LBR, it will look for a matching (same IP address and EUI-64) (Tentative or Registered) NCE. If no such entry is found, then the DAC is silently ignored. If an entry is found and the DAC had Status=0, then the 6LR will mark the Tentative NCE as Registered. In all cases, when an entry is found, then the 6LR will respond to the host with an NA, copying the Status and EUI-64 fields from the DAC to an ARO in the NA. In case the status is an error, then the destination IP address of the NA is derived from the EUI-64 field of the DAC.

A Tentative NCE SHOULD be timed out `TENTATIVE_NCE_LIFETIME` seconds after it was created in order to allow for another host to attempt to register the IPv6 address.

8.2.1. Message Validation for DAR and DAC

A node MUST silently discard any received DAR and DAC messages for which at least one of the following validity checks is not satisfied:

- o If the message includes an IP Authentication Header, the message authenticates correctly.
- o ICMP Checksum is valid.
- o ICMP Code is 0.
- o ICMP Length (derived from the IP length) is 32 or more bytes.
- o The Registered Address is not a multicast address.
- o All included options have a length that is greater than zero.
- o The IP source address is not the unspecified address, nor is it a multicast address.

The contents of the Reserved field and of any unrecognized options MUST be ignored. Future backward-compatible changes to the protocol may specify the contents of the Reserved field or add new options; backward-incompatible changes may use different Code values.

Note that due to the forwarding of the DAR and DAC messages between the 6LR and 6LBR, there is no hop-limit check on receipt for these ICMPv6 message types.

8.2.2. Conceptual Data Structures

A 6LBR implementing multihop DAD needs to maintain some state separate from the Neighbor Cache. We call this conceptual data structure the DAD table. It is indexed by the IPv6 address -- the Registered Address in the DAR -- and contains the EUI-64 and the Registration Lifetime of the host that is using that address.

8.2.3. 6LR Sending a Duplicate Address Request

When a 6LR that implements multihop DAD receives an NS from a host, and subject to the above checks, the 6LR forms and sends a DAR to at least one 6LBR. The DAR contains the following information:

- o In the IPv6 source address, a global address of the 6LR.
- o In the IPv6 destination address, the address of the 6LBR.
- o In the IPv6 hop limit, MULTI_HOP_HOPLIMIT.
- o The Status field MUST be set to zero.
- o The EUI-64 and Registration Lifetime are copied from the ARO received from the host.
- o The Registered Address set to the IPv6 address of the host, that is, the sender of the triggering NS.

When a 6LR receives an NS from a host with a zero Registration Lifetime, then, in addition to removing the NCE for the host as specified in Section 6, a DAR is sent to the 6LBRs as above.

A router MUST NOT modify the Neighbor Cache as a result of receiving a DAR.

8.2.4. 6LBR Receiving a Duplicate Address Request

When a 6LBR that implements the substitutable multihop DAD receives a DAR from a 6LR, it performs the message validation specified in Section 8.2.1. If the DAR is valid, the 6LBR proceeds to look for the Registration Address in the DAD table. If an entry is found and the recorded EUI-64 is different than the EUI-64 in the DAR, then it

returns a DAC NA with the Status set to 1 ('Duplicate Address'). Otherwise, it returns a DAC with Status set to zero and updates the lifetime.

If no entry is found in the DAD table and the Registration Lifetime is non-zero, then an entry is created and the EUI-64 and Registered Address from the DAR are stored in that entry.

If an entry is found in the DAD table, the EUI-64 matches, and the Registration Lifetime is zero, then the entry is deleted from the table.

In both of the above cases, the 6LBR forms a DAC with the information copied from the DAR and the Status field is set to zero. The DAC is sent back to the 6LR, i.e., back to the source of the DAR. The IPv6 hop limit is set to MULTI_HOP_HOPLIMIT.

8.2.5. Processing a Duplicate Address Confirmation

When a 6LR implementing multihop DAD receives a DAC message, then it first validates the message per Section 8.2.1. For a valid DAC, if there is no Tentative NCE matching the Registered Address and EUI-64, then the DAC is silently ignored. Otherwise, the information in the DAC and in the Tentative NCE is used to form an NA to send to the host. The Status code is copied from the DAC to the ARO that is sent to the host. In the case where the DAC indicates an error (the Status is non-zero), the NA is returned to the host as described in Section 6.5.2, and the Tentative NCE for the Registered Address is removed. Otherwise, it is made into a Registered NCE.

A router MUST NOT modify the Neighbor Cache as a result of receiving a DAC, unless there is a Tentative NCE matching the IPv6 address and EUI-64.

8.2.6. Recovering from Failures

If there is no response from a 6LBR after RETRANS_TIMER [RFC4861], then the 6LR would retransmit the DAR to the 6LBR up to MAX_UNICAST_SOLICIT [RFC4861] times. After this, the 6LR SHOULD respond to the host with an ARO Status of zero.

9. Protocol Constants

This section defines the relevant protocol constants used in this document based on a subset of [RFC4861] constants. "*" indicates constants modified from [RFC4861], and "+" indicates new constants.

Additional protocol constants are defined in Section 4.

6LBR Constants:

MIN_CONTEXT_CHANGE_DELAY+ 300 seconds

6LR Constants:

MAX_RTR_ADVERTISEMENTS 3 transmissions

MIN_DELAY_BETWEEN_RAS* 10 seconds

MAX_RA_DELAY_TIME* 2 seconds

TENTATIVE_NCE_LIFETIME+ 20 seconds

Router Constants:

MULTIHOP_HOPLIMIT+ 64

Host Constants:

RTR_SOLICITATION_INTERVAL* 10 seconds

MAX_RTR_SOLICITATIONS 3 transmissions

MAX_RTR_SOLICITATION_INTERVAL+ 60 seconds

10. Examples

10.1. Message Examples




Figure 2: Basic Router Solicitation/Router Advertisement Exchange between a Node and a 6LR or 6LBR

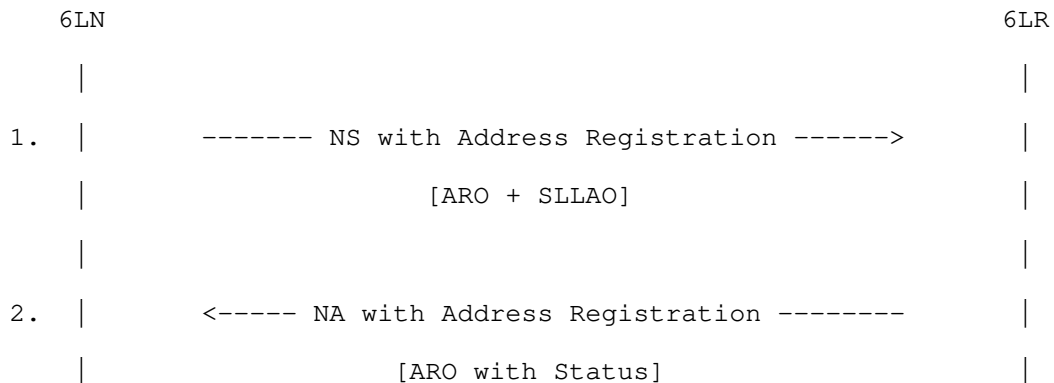


Figure 3: Neighbor Discovery Address Registration

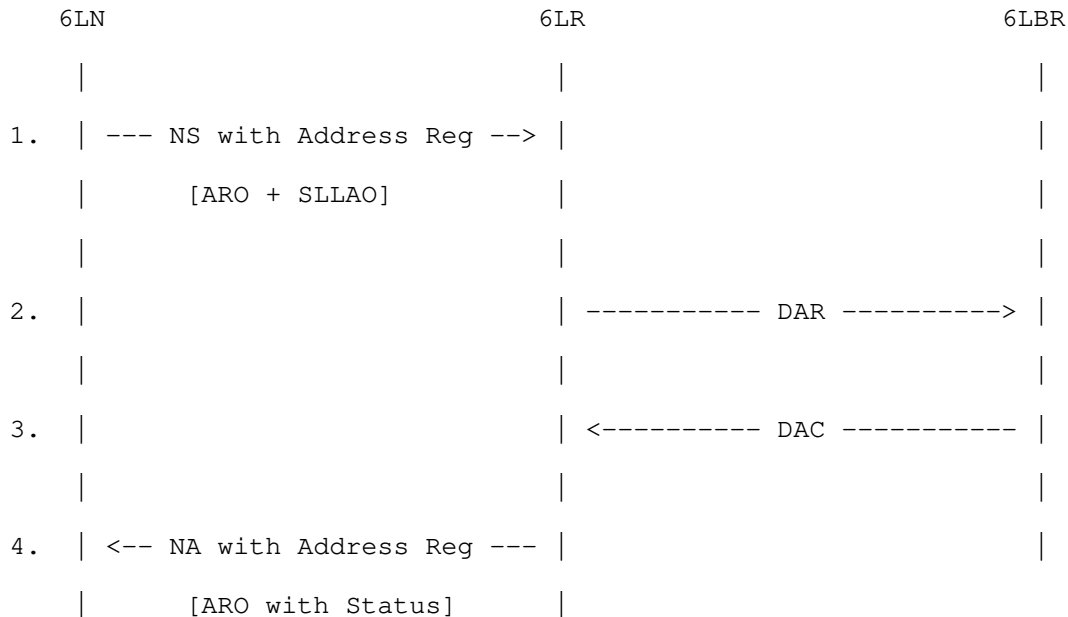


Figure 4: Neighbor Discovery Address Registration with Multihop DAD

10.2. Host Bootstrapping Example

The following example describes the address bootstrapping scenarios using the improved ND mechanisms specified in this document. It is assumed that the 6LN first performs a sequence of operations in order to get secure access at the link layer of the LoWPAN and obtain a key for link-layer security. The methods of how to establish link-layer security are out of scope of this document. In this example, an IEEE 802.15.4 6LN forms a 16-bit short IPv6 address without using DHCPv6 (i.e., the M flag is not set in the RAs).

1. After obtaining link-layer security, a 6LN assigns a link-local IPv6 address to itself. A link-local IPv6 address is configured based on the 6LN's EUI-64 link-layer address formed as per [RFC4944].
2. Next, the 6LN determines one or more default routers in the network by sending an RS to the all-routers multicast address with the SLLAO set to its EUI-64 link-local address. If the 6LN was able to obtain the link-layer address of a router through its link-layer operations, then the 6LN may form a link-local destination IPv6 address for the router and send it a unicast RS.

The 6LR responds with a unicast RA to the IP source address using the SLLAO from the RS (it may have created a Tentative NCE). See Figure 2.

3. In order to communicate more than one IP hop away, the 6LN configures a global IPv6 address. In order to save overhead, this 6LN wishes to configure its IPv6 address based on a 16-bit short address as per [RFC4944]. As the network is unmanaged (M flag not set in the RA), the 6LN randomly chooses a 16-bit link-layer address and forms a Tentative IPv6 address from it.
4. Next, the 6LN registers that address with one or more of its default routers by sending a unicast NS message with an ARO containing its Tentative global IPv6 address to register, the Registration Lifetime, and its EUI-64. An SLLAO is also included with the link-layer address corresponding to the address being registered. If a successful (Status 0) NA message is received, the address can then be used, and the 6LN assumes that it has been successfully checked for duplicates. If a duplicate address (Status 1) NA message is received, the 6LN then removes the temporary IPv6 address and 16-bit link-layer address and goes back to step 3. If a Neighbor Cache Full (Status 2) message is received, the 6LN attempts to register with another default router or, if none, goes back to step 2. See Figure 3. Note that an NA message returning an error would be sent back to the link-local EUI-64-based IPv6 address of the 6LN instead of the 16-bit (duplicate) address.
5. The 6LN now performs maintenance by sending a new NS address registration before the lifetime expires.

If multihop DAD and multihop prefix and context distribution are used, the effect of the 6LRs and hosts following the above bootstrapping process is a "wavefront" of 6LRs and hosts being configured, spreading outward from the 6LBRs: First, the hosts and 6LRs that can directly reach a 6LBR would receive one or more RAs and then configure and register their IPv6 addresses. Once that is done, they would enable the routing protocol and start sending out RAs. That would result in a new set of 6LRs and hosts to receive responses to their RSs, form and register their addresses, etc. That repeats until all of the 6LRs and hosts have been configured.

10.2.1. Host Bootstrapping Messages

This section provides specific message examples related to the bootstrapping process described above. When discussing messages, the following notation is used:

LL64: Link-local address based on the EUI-64, which is also the 802.15.4 long address.

GP16: Global address based on the 802.15.4 short address. This address may not be unique.

GP64: Global addresses derived from the EUI-64 address as specified in [RFC4944].

MAC64: EUI-64 address used as the link-layer address.

MAC16: IEEE 802.15.4 16-bit short address.

Note that some implementations may use LL64 and GP16 style addresses instead of LL64 and GP64. In the following, we will show an example message flow as to how a node uses LL64 to register a GP16 address for multihop DAD verification.

```
6LN-----RS----->6LR
  Src= LL64 (6LN)
  Dst= all-router-link-scope-multicast
  SLLAO= MAC64 (6LN)
```

```
6LR-----RA----->6LN
  Src= LL64 (6LR)
  Dst= LL64 (6LN)
```

Note: Source address of RA must be a link-local address (Section 4.2 of RFC 4861).

```
6LN-----NS Reg----->6LR
  Src= GP16 (6LN)
  Dst= LL64 (6LR)
  ARO
  SLLAO= MAC16 (6LN)
```

```
6LR-----DAR----->6LBR
  Src= GP64 or GP16 (6LR)
  Dst= GP64 or GP16 (6LBR)
  Registered Address= GP16 (6LN) and EUI-64 (6LN)
```

```

6LBR-----DAC----->6LR
Src= GP64 or GP16 (6LBR)
Dst= GP64 or GP16 (6LR)
Copy of information from DAR

```

If Status is a success:

```

6LR -----NA-Reg----->6LN
Src= LL64 (6LR)
Dst= GP16 (6LN)
ARO with Status = 0

```

If Status is not a success:

```

6LR -----NA-Reg----->6LN
Src= LL64 (6LR)
Dst= LL64 (6LN) --> Derived from the EUI-64 of ARO
ARO with Status > 0

```

Figure 5: Detailed Message Address Examples

10.3. Router Interaction Example

In the route-over topology, when a routing protocol is run across 6LRs, the bootstrapping and Neighbor Cache management are handled a little differently. The description in this paragraph provides only a guideline for an implementation.

At the initialization of a 6LR, it may choose to bootstrap as a host with the help of a parent 6LR if the substitutable multihop DAD is performed with the 6LBR. The Neighbor Cache management of a router and address resolution among the neighboring routers are described in Sections 6.5.3 and 6.5.5, respectively. In this example, we assume that the neighboring 6LoWPAN link is secure.

10.3.1. Bootstrapping a Router

In this scenario, the bootstrapping 6LR, 'R1', is multiple hops away from the 6LBR and surrounded by other 6LR neighbors. Initially, R1 behaves as a host. It sends a multicast RS and receives an RA from one or more neighboring 6LRs. R1 picks one 6LR as its temporary default router and performs address resolution via this default router. Note that if multihop DAD is not required (e.g., in a managed network or using EUI-64-based addresses), then it does not need to pick a temporary default router; however, it may still want to send the initial RS message if it wants to autoconfigure its address with the global prefix disseminated by the 6LBR.

Based on the information received in the RAs, R1 updates its cache with entries for all the neighboring 6LRs. Upon completion of the address registration, the bootstrapping router deletes the temporary entry of the default router, and the routing protocol is started.

Also note that R1 may refresh its multihop DAD registration directly with the 6LBR (using the next-hop neighboring 6LR determined by the routing protocol for reaching the 6LBR).

10.3.2. Updating the Neighbor Cache

In this example, there are three 6LRs: R1, R2, and R3. Initially, when R2 boots, it sees only R1, and accordingly R2 creates an NCE for R1. Now assume that R2 receives a valid routing update from router R3. R2 does not have any NCE for R3. If the implementation of R2 supports detecting link-layer addresses from the routing information packets, then it directly updates its Neighbor Cache using that link-layer information. If this is not possible, then R2 should perform multicast NS with the source set with its link-local or global address, depending on the scope of the source IP address received in the routing update packet. The target address of the NS message is the source IPv6 address of the received routing update packet. The format of the NS message is as described in Section 4.3 of [RFC4861].

More generally, any 6LR that receives a valid route update from a neighboring router for which it does not have any NCE is required to update its Neighbor Cache as described above.

The router (6LR and 6LBR) IP addresses learned via ND are not redistributed to the routing protocol.

11. Security Considerations

The security considerations of IPv6 ND [RFC4861] and address autoconfiguration [RFC4862] apply. Additional considerations can be found in [RFC3756].

There is a slight modification to those considerations, due to the fact that in this specification the M flag in the RAs disables the use of stateless address autoconfiguration for addresses not derived from EUI-64. Thus, a rogue router on the link can force the use of only DHCP for short addresses, whereas in [RFC4861] and [RFC4862] the rogue router could only cause the addition of DHCP and not disable stateless address autoconfiguration for short addresses.

This specification assumes that the link layer is sufficiently protected -- for instance, by using MAC-sublayer cryptography. Thus, its threat model is no different from that of IPv6 ND [RFC4861]. The first trust model listed in Section 3 of [RFC3756] applies here. However, any future 6LoWPAN security protocol that applies to ND for the 6LoWPAN protocol is out of scope of this document.

The multihop DAD mechanisms rely on DAR and DAC messages that are forwarded by 6LRs, and as a result the hop_limit=255 check on the receiver does not apply to those messages. This implies that any node on the Internet could successfully send such messages. We avoid any additional security issues due to this by requiring that the routers never modify the NCE due to such messages, and that they discard them unless they are received on an interface that has been explicitly configured to use these optimizations.

In some future deployments, one might want to use SEcure Neighbor Discovery (SEND) [RFC3971] [RFC3972]. This is possible with the ARO as sent between hosts and routers, since the address that is being registered is the IPv6 source address of the NS and SEND verifies the IPv6 source address of the packet. Applying SEND to the router-to-router communication in this document is out of scope.

12. IANA Considerations

This document registers three new ND option types under the subregistry "IPv6 Neighbor Discovery Option Formats":

- o Address Registration Option (33)
- o 6LoWPAN Context Option (34)
- o Authoritative Border Router Option (35)

The document registers two new ICMPv6 "type" numbers under the subregistry "ICMPv6 "type" Numbers":

- o Duplicate Address Request (157)
- o Duplicate Address Confirmation (158)

IANA has also created a new subregistry for the Status values of the Address Registration Option, under the ICMPv6 parameters registry.

Address Registration Option Status Values registry:

- o Possible values are 8-bit unsigned integers (0..255).
- o Registration procedure is "Standards Action" [RFC5226].
- o Initial allocation is as indicated in Table 2:

Status	Description
0	Success
1	Duplicate Address
2	Neighbor Cache Full
3-255	Allocated using Standards Action [RFC5226]

Table 2

13. Interaction with Other Neighbor Discovery Extensions

There are two classes of ND extensions that interact with this specification in different ways.

One class encompasses extensions to the DAD mechanisms in [RFC4861] and [RFC4862]. An example of this is Optimistic DAD [RFC4429]. Such extensions do not apply when this specification is being used, since it uses ARO for DAD (which is neither optimistic nor pessimistic -- always one round trip to the router to check DAD).

All other (non-DAD) ND extensions, be they path selection types like default router preferences [RFC4191], configuration types like DNS configuration [RFC6106], or other types like Detecting Network Attachment [RFC6059], are completely orthogonal to this specification and will work as is.

14. Guidelines for New Features

This section discusses guidelines of new protocol features defined in this document. It also sets some expectations for implementation and deployment of these features. This section is informative in nature: it does not override the detailed specifications of the previous sections but summarizes them and presents them in a compact form, to be used as checklists. The checklists act as guidelines to indicate the possible importance of a feature in terms of a deployment as per information available as of the writing of the document. Note that in some cases the deployment is 'SHOULD' where the implementation is

a 'MUST'. This is due to the presence of substitutable features; the deployment may use alternative methods for those. Therefore, implementing a configuration knob is recommended for the substitutable features. The lists emphasize conciseness over completeness.

Section	Description	Deploy	Implement
3.1	Host-initiated RA	MUST	MUST
3.2	EUI-64-based IPv6 address	MUST	MUST
	16-bit MAC-based address	MAY	SHOULD
	Other non-unique addresses	MAY	MAY
3.3	Host-initiated RS	MUST	MUST
	ABRO processing	SHOULD	MUST
4.1	Registration with ARO	MUST	MUST
4.2, 5.4	6CO	SHOULD	SHOULD
5.2	Joining solicited-node multicast	N/A	N/A
	Joining all-nodes multicast	MUST	MUST
	Using link-layer indication for NUD	MAY	MAY
5.5	6LoWPAN-ND NUD	MUST	MUST
5.8.2	Behavior on wakeup	SHOULD	SHOULD

Table 3: Guideline for 6LoWPAN-ND Features for Hosts

Section	Description	Deploy	Implement
3.1	Periodic RA	SHOULD NOT	SHOULD NOT
3.2	Address assignment during startup	SHOULD	MUST
3.3	Supporting EUI-64-based MAC hosts	MUST	MUST
	Supporting 16-bit MAC hosts	MAY	SHOULD
3.4, 4.3, 8.1.3, 8.1.4	ABRO processing/sending	SHOULD	MUST
8.1	Multihop prefix storing and redistribution	SHOULD	MUST
3.5	Tentative NCE	MUST	MUST
8.2	Multihop DAD	SHOULD	MUST
4.1, 6.5, 6.5.1 - 6.5.5	ARO support	MUST	MUST
4.2	6CO	SHOULD	SHOULD
6.3	Process RS/ABRO	MUST	MUST

Table 4: Guideline for 6LR Features in 6LoWPAN-ND

Section	Description	Deploy	Implement
3.1	Periodic RA	SHOULD NOT	SHOULD NOT
3.2	Address autoconf on router interface	MUST NOT	MUST NOT
3.3	EUI-64 MAC support on 6LoWPAN interface	MUST	MUST
8.1 - 8.1.1, 8.1.5	Multihop prefix distribution	SHOULD	MUST
8.2	Multihop DAD	SHOULD	MUST

Table 5: Guideline for 6LBR Features in 6LoWPAN-ND

15. Acknowledgments

The authors thank Pascal Thubert, Jonathan Hui, Richard Kelsey, Geoff Mulligan, Julien Abeille, Alexandru Petrescu, Peter Siklosi, Pieter De Mil, Fred Baker, Anthony Schoofs, Phil Roberts, Daniel Gavelle, Joseph Reddy, Robert Cragie, Mathilde Durvy, Colin O'Flynn, Dario Tedeschi, Esko Dijk, and Joakim Eriksson for useful discussions and comments that have helped shape and improve this document.

Additionally, the authors would like to recognize Pascal Thubert for contributing the original registration idea and for extensive contributions to earlier versions of the document, Jonathan Hui for original ideas on prefix/context distribution and extensive contributions to earlier versions of the document, Colin O'Flynn for useful "Error-to" suggestions (Section 6.5.2) and for contributions to the Examples section, Geoff Mulligan for suggesting the use of address registration as part of existing IPv6 ND messages, and Mathilde Durvy for helping to clarify router interaction.

16. References

16.1. Normative References

[ETHERNET]

"IEEE Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications", IEEE Std 802.3-2008, December 2008, <http://standards.ieee.org/getieee802/download/802.3-2008_section1.pdf>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2491] Armitage, G., Schuler, P., Jork, M., and G. Harter, "IPv6 over Non-Broadcast Multiple Access (NBMA) networks", RFC 2491, January 1999.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, September 2011.

16.2. Informative References

- [EUI64] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64(TM)) Registration Authority", <<http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>>.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, April 2006.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, August 2007.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.
- [RFC5889] Baccelli, E. and M. Townsley, "IP Addressing Model in Ad Hoc Networks", RFC 5889, September 2010.
- [RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, November 2010.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.

Authors' Addresses

Zach Shelby (editor)
Sensinode
Konekuja 2
Oulu 90620
Finland

Phone: +358407796297
EMail: zach@sensinode.com

Samita Chakrabarti
Ericsson

EMail: samita.chakrabarti@ericsson.com

Erik Nordmark
Cisco Systems

EMail: nordmark@cisco.com

Carsten Bormann
Universitaet Bremen TZI
Postfach 330440
Bremen D-28359
Germany

Phone: +49-421-218-63921
EMail: cabo@tzi.org