6LoWPAN Working Group	Z. Shelby, Ed.
Internet-Draft	Sensinode
Updates: 4944 (if approved)	S. Chakrabarti
Intended status: Standards Track	Ericsson
Expires: April 26, 2012	E. Nordmark
	Cisco Systems
	October 24, 2011

Neighbor Discovery Optimization for Low Power and Lossy Networks (6LoWPAN)
draft-ietf-6lowpan-nd-18

Abstract

The IETF 6LoWPAN working group defines IPv6 over Low-power Wireless Personal Area Networks 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 traditional concept of IP subnets and IP links. IPv6 Neighbor Discovery was not designed for non-transitive wireless links. The traditional IPv6 link concept and heavy use of multicast make the protocol inefficient and sometimes impractical in a low power and lossy network. This document describes simple optimizations to IPv6 Neighbor Discovery, addressing mechanisms and duplicate address detection for 6LoWPAN and similar networks. The document, thus updates RFC 4944 to specify the use of the optimizations defined here.

Status of this Memo

Copyright Notice

Table of Contents

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 which sits between the MAC layer and the IP network layer. A link in a 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 [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 MultiAccess (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:

*Host-initiated interactions to allow for sleeping hosts.

*Elimination of multicast-based address resolution for hosts.

*A host address registration feature using a new option in unicast Neighbor Solicitation and Neighbor Advertisement messages.

*A new Neighbor Discovery option to distribute 6LoWPAN header compression context to hosts.

*Optional multihop distribution of prefix and 6LoWPAN header compression context.

*Optional multihop duplicate address detection which uses two new ICMPv6 message types.

The document defines three new ICMPv6 message options: the required Address Registration option and the optional Authoritative Border Router and 6LoWPAN Context options. It also defines two new ICMPv6 message types: the Duplicate Address Request and Duplicate Address Confirmation.

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, Redirect, 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-router address to solicit router advertisements. If the host receives a valid Router Advertisement with the "A" flag, it autoconfigures the IPv6 address with the advertised prefix in the Router Advertisement (RA) message. Besides this, the IPv6 routers usually send router advertisements periodically on the network. RAs are sent to the all-node 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 assumes 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 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 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.
In a LoWPAN, primarily two types of network topologies are found - star networks and mesh networks. A star network is similar to a regular IPv6 subnet with a router and a set of nodes connected to it via the same non-transitive link. But in Mesh networks, the nodes are capable of routing and forwarding packets. 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 center, some optimizations and extensions to Neighbor Discovery are useful for the wide deployment of IPv6 over low-powered and lossy networks such as 6LoWPANs.

1.2. Mesh-under and Route-over Concepts

In the 6LoWPAN context, often a link-layer mesh routing mechanism is referred to as "mesh-under" while routing/forwarding packets using IP-layer addresses is referred to as "route-over". The difference between mesh-under and route-over is similar to a bridged-network versus IP-routing using Ethernet. In a mesh-under network all nodes are on the same link which is served by one or more routers, which we call 6LoWPAN Border Routers (6LBR). In a route-over network, there are multiple links in the 6LoWPAN. Unlike fixed IP links, these link's members may be changing due to the nature of the low-power and lossy behavior of wireless technology. Thus a route-over network is made up of a flexible set of links interconnected by interior routers, which we call 6LoWPAN Routers (6LR).
This specification is applicable to both mesh-under and route-over networks. However, in route-over networks, we have two types of routers - 6LBRs and 6LRs. 6LoWPAN Border Routers sit at the boundary of the 6LoWPAN and the rest of the network while 6LoWPAN Routers are inside the LoWPAN. 6LoWPAN Routers are assumed to be running a routing protocol.
In a mesh-under configuration a 6LBR is acting as the IPv6 router where all the hosts in the LoWPAN are on the same link, thus they are only one IP hop away. No 6LoWPAN Routers exist in this topology as forwarding is handled by a link-layer mesh routing protocol.
In a route-over configuration, Neighbor Discovery operations take place between hosts and 6LRs or 6LBRs. The 6LR nodes are able to send and receive Router Advertisements, Router Solicitations as well as forward and route IPv6 packets. Here packet forwarding happens at the IP layer.
In both types of configurations, hosts do not take part in routing and forwarding packets and they act as simple IPv6 hosts.

1.3. 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 RFC 4861. The present specification optimizes the usage of IPv6 Neighbor Discovery for LoWPANs in order to save energy and processing power of such nodes. The document, thus updates RFC 4944 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 6LoWPAN, 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) which are OPTIONAL to use. An implementation of this specification SHOULD implement those optional to use pieces.
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 minimized signaling, robust handling of the non-transitive link, and header compression context information.

1.4. Goals and Assumptions

The document has the following main goals and assumptions.
Goals:

- *Optimize Neighbor Discovery with a mechanism that is minimal yet sufficient for the operation in both mesh-under and route-over configurations.

- *Minimize signaling by avoiding the use of multicast flooding and reducing the use of link-scope multicast messages.

- *Optimize the interfaces between hosts and their default routers.

- *Support for sleeping hosts.

- *Disseminate context information to hosts as needed by [I-D.ietf-6lowpan-hc].

- *Optionally disseminate context information and prefix information from the border to all routers in a LoWPAN.

- *Optional duplicate address detection mechanism suitable for route-over LoWPANs.

Assumptions:

- *EUI-64 addresses are globally unique.

- *All nodes in the network have an EUI-64 interface identifier in order to do address auto-configuration and detect duplicate addresses.

*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.

*A 6LoWPAN is configured to share one or more global IPv6 address prefixes to enable hosts to move between routers in the 6LoWPAN without changing their IPv6 addresses.

*When using the optional DAD mechanism of Section 8.2 it is assumed that 6LRs register with all the 6LBRs.

*If IEEE 802.15.4 16-bit short addresses are used, then some technique is used to ensure uniqueness of those link-layer addresses. That could be done using DHCPv6, the Address Registration Option based duplicate address detection (specified in Section 8.2) or other techniques outside of the scope of this document.

*In order to preserve the uniqueness of addresses 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).

*In order for [I-D.ietf-6lowpan-hc] 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 6LoWPAN.

*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.

*Since the 6LoWPAN shares one single prefix throughout the network, mobility of nodes within the LoWPAN is transparent. Inter-LoWPAN mobility is out-of-scope of this document.

1.5. Optional Features

This document defines the optimization of Neighbor Discovery messages host-router interfaces and introduces the communication in case of Route-over topology. The multihop prefix distribution by the 6LBR and multihop Duplicate Address Detection mechanisms, as well as 6LoWPAN context option are optional features for a 6LoWPAN deployment. A

guideline for feature implementation and deployment is provided at the end of the document.

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" *[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 who can communicate with other 6LoWPAN routers in the same LoWPAN. 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 hosts 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.

Registration:
The process during which a LoWPAN node sends an Neighbor Solicitation message with an Address Registration option to a Router creating a Neighbor Cache entry 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 [I-D.ietf-6lowpan-hc] by carrying header compression information in a new option in Router Advertisement messages.
In addition, there are optional and 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 RFC4861

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

*Host initiated refresh of Router Advertisement information. This removes the need for periodic or unsolicited Router Advertisements from routers to hosts.

*No Duplicate Address Detection (DAD) is performed if EUI-64 based IPv6 addresses are used (as these addresses are assumed to be globally unique).

*DAD is optional if DHCPv6 is used to assign addresses.

*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.

*A new optional Router Advertisement option for Context information used by 6LoWPAN header compression.

*A new optional mechanism to perform Duplicate Address Detection across a route-over 6LoWPAN using the new Duplicate Address Request and Confirmation messages.

*New optional mechanisms to distribute Prefixes and Context information across a route-over network which uses a new Authoritative Border Router option to control the flooding of configuration changes.

*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 address as specified in [RFC4861] and [RFC4862] based on the information received in Router Advertisement messages. The use of the M flag in this optimization is however more restrictive than in [RFC4861]. When the M flag is set a host is required to use DHCPv6 to assign any non-EUI-64 addresses. When the M flag is not set, the LoWPAN is required to 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 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 it suspects that one of its default routers has become unreachable (after Neighbor Unreachability Detection towards the router fails).

Hosts receive Router Advertisement messages typically containing the Authoritative Border Router option (ABRO) and may optionally contain one or more 6LoWPAN Context options (6CO) in addition to the existing Prefix Information options (PIO) 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 option 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 change their point of attachment often, should use a suitably short lifetime.

The registration can fail (an ARO option returned to the host with a non-zero Status) if the router determines that the IPv6 address is already used by another host, that is, 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 a IEEE 802.15.4 16-bit short addresses. 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 both need to 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 a 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, or 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 optional new router-to-router interaction is only for the route-over configuration where 6LRs are present. It is optional in this protocol since the functions it provides might be better provided by other protocol mechanisms, be it DHCPv6, link-layer mechanisms, the routing protocol, or something else. It is however assumed that all 6LRs in a network are configured to perform these functions homogeneously. Some mechanisms from this protocol might be used for router-to-router interaction, while others are provided by other

protocols. For instance, context information and/or prefix information might be disseminated using this protocol, while Duplicate Address Detection is done using some other protocol.
6LRs MAY 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 or context dissemination is 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 Solicitations. There is a version number field in the ABRO which is used to limit the flooding of updated information between the 6LRs.
Optionally the 6LRs can perform Duplicate Address Detection against one or more 6LBRs using the new Duplicate Address Request (DAR) and 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 (NCE) slightly differently than in [RFC4861]. This results in three different types of NCEs and the types specify how those entries can be removed: [RFC4861].

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 When a host interacts with a router by sending Router Solicitations this results in a Tentative NCE. Once a node successfully registers with a Router the result is a Registered NCE. When Routers send RAs to hosts, and when routers optionally 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 mandatory, whereas the Authoritative Border Router Option and 6LoWPAN Context Option are optional. This section also defines the optional and new Duplicate Address Request and 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 Neighbor Solicitation (NS) messages sent by hosts. Thus it can be included in the unicast NS messages that a host sends as part of Neighbor Unreachability Detection 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 Neighbor Advertisement (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 optional multihop DAR and DAC messages used between 6LRs to 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 which 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 it is registering an address from. 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 SLLA (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 Neighbor Solicitation message.

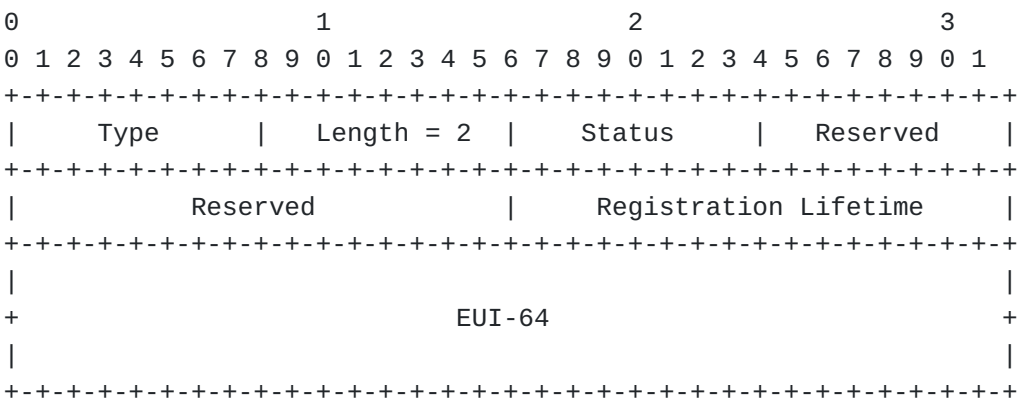

Fields:

Type: TBD1

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 a unit of 60 seconds that the router should retain the Neighbor Cache entry 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 Neighbor Advertisements are:

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

4.2. 6LoWPAN Context Option

The optional 6LoWPAN Context Option (6CO) carries prefix information for LoWPAN header compression, and is similar to the Prefix Information Option 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 Context Identifier (CID) for use as specified in [I-D.ietf-6lowpan-hc]. A context may be a prefix of any length or an address (/128), and up to 16 6LoWPAN Context options may be carried in an Router Advertisement message.

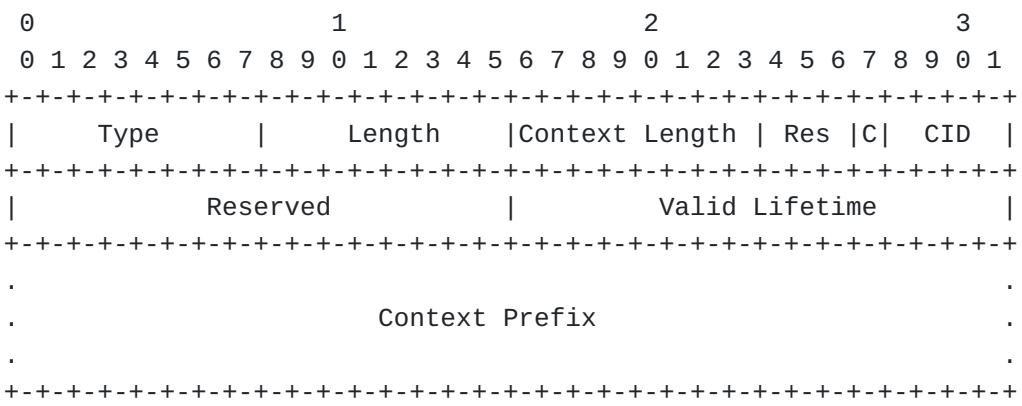

Type: TBD2

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 lifecycle based on the recommendations in Section 7.2.

CID: 4-bit Context Identifier for this prefix information. CID is used by context based header compression specified in [I-D.ietf-6lowpan-hc]. The list of CIDs for a LoWPAN is configured by 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 a unit 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 Context ID (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 optional Authoritative Border Router Option (ABRO) is needed when Router Advertisement (RA) messages are used to disseminate prefixes and context information across a route-over topology. In this case 6LRs receive Prefix Information options 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 which the 6LBR sets and 6LRs propagate as they propagate the prefix and context information with this Authoritative Border Router Option. 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 Authoritative Border Router option MUST be included in all Router Advertisement messages in the case when Router Advertisements are used to propagate information between routers (as described in Section 8.2).

Fields:

Type:

TBD3

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

Version Number: 16-bit unsigned integer. 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. This version number uses sequence number arithmetic as it may wrap around.

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 optional multihop DAD exchanges between 6LR and 6LBR specified in Section 8.2 there are two new ICMPv6 message types called the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC). We avoid reusing the Neighbor Solicitation and Neighbor Advertisement 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 are otherwise the same as would be in a Neighbor Solicitation carrying a Address Registration option, 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.

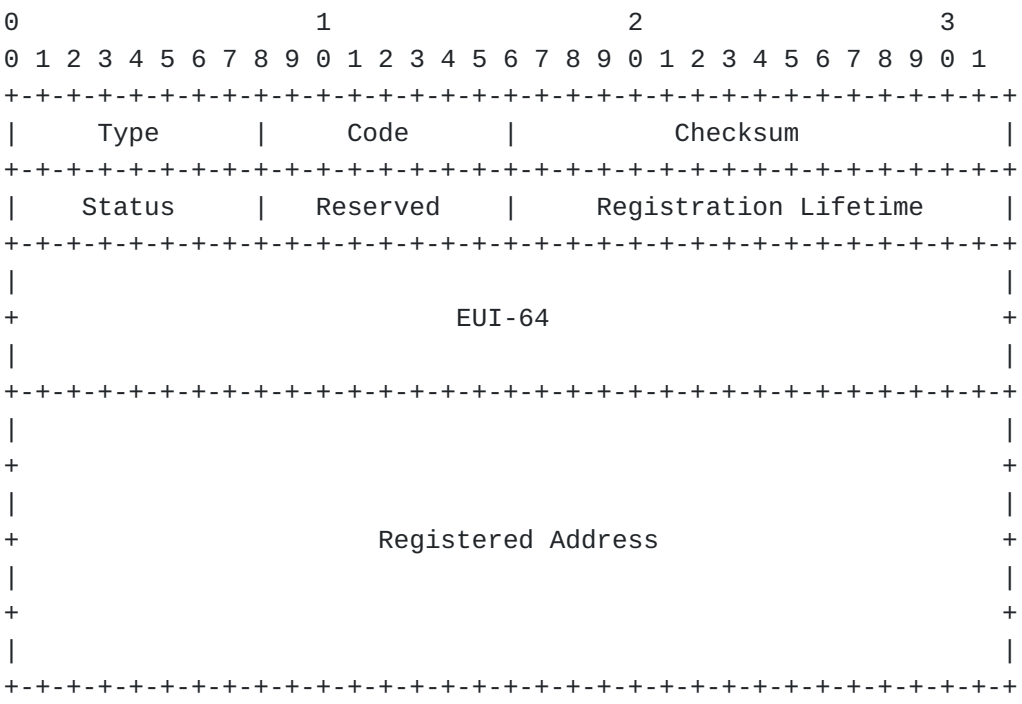

IP fields:

IPv6 source: A non link-local address of the sending router.

IPv6 destination: A non link-local address of the sending router. In a DAC this is just the source from the DAR.

Hop Limit: Set to MULTIHOP_HOPLIMIT on transmit. MUST be ignored on receipt.

ICMP Fields:

Type: TBD4 for DAR and TBD5 for 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 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 a unit of 60 seconds that the router should retain the Neighbor Cache entry for the sender of the NS that includes this option. A value of 0 indicates in an NS that the neighbor cache 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, which was contained in the IPv6 Source field in the NS that contained the ARO option sent by the host.

5. Host Behavior

Hosts in a LoWPAN use the Address Registration option in the Neighbor Solicitation messages they send as a way to maintain the Neighbor Cache in the routers thereby removing the need for multicast Neighbor Solicitations to do address resolution. Unlike in [RFC4861] the hosts initiate updating the information they receive in Router Advertisements by sending Router Solicitations before the information expires. Finally, when Neighbor Unreachability Detection indicates that one or all default routers have become unreachable, then the host uses Router Solicitations to find a new set of default routers.

5.1. Forbidden Actions

A host MUST NOT multicast a Neighbor Solicitation 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 Router Solicitation messages as described in [RFC4861] Section 6.3.7.
There is no need to join the Solicited-Node multicast address since nobody multicasts Neighbor Solicitations in this type of network. A host MUST join the all-nodes multicast address.

5.3. Sending a Router Solicitation

The Router Solicitation is formatted as specified in [RFC4861] and sent to the IPv6 All-Routers multicast address (see [RFC4861] Section 6.3.7 for details). An SLLA option MUST be included to enable unicast Router Advertisements 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 theses packets to a router.
Since hosts do not depend on multicast Router Advertisements to discover routers, the hosts need to intelligently retransmit Router Solicitations 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 are about to expire. The RECOMMENDED 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 binary exponential backoff of the retransmission timer for each subsequent retransmission. However, it is useful to have a maximum retransmission timer of 60 seconds (MAX_RTR_SOLICITATION_INTERVAL). In all cases the RS retransmissions are terminated when a RA is received.

5.4. Processing a Router Advertisement

The processing of Router Advertisements is as in [RFC4861] with the addition of handling the 6LoWPAN Context option and triggering address registration when a new address has been configured. Furthermore, the SLLA option 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 Prefix Information option with the 'L' (on-link) flag set, then that Prefix Information Option (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 a Neighbor Solicitation with the Address Registration option 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, which is called the Context Table. This includes the Context ID, 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 option is received in a Router Advertisement it is used to add or update the information in the Context Table. If the Context ID 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 lifecycle. 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 a Router Solicitation to ask for any updates to the information. The lifetimes that matter are the Default Router lifetime, the Valid Lifetime in the Prefix Information options, and the Valid Lifetime in the 6CO. The host SHOULD unicast one or more Router Solicitations to the router well before the minimum of those lifetimes (across all the prefixes and all the contexts) expire, and switch to multicast RS messages if there is no response to the unicasts. The retransmission behavior for the Router Solicitations is specified in Section 5.3.

5.5. Registration and Neighbor Unreachability Detection

Hosts send Unicast Neighbor Solicitation (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 Neighbor Solicitation 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 Neighbor Cache entries in the routers. This is done by sending NS messages with the 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 Neighbor Advertisement message with an ARO option.
Hosts that receive Router Advertisement 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 Neighbor Unreachability Detection 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 Neighbor Solicitation (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 a Source Link-Layer Address (SLLA) option, 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 Neighbor Advertisement messages as specified in [RFC4861], with added logic described in this section for handling the Address Registration option.
In addition to the normal validation of a Neighbor Advertisement and its options, the Address Registration option is verified as follows (if present). 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 Address Registration option 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 Neighbor Solicitations is received (NUD failure), or an Address Registration option 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 Address Registration option 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 has no more default routers it needs to revert to sending Router Solicitations 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.
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 SLLA option in Router Advertisement 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 Neighbor Cache entries [RFC4861]. In order to achieve LoWPAN compression, most global addresses are formed using a link-layer address. Thus a host can minimize 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 described further in this section. Routers may want to cache traffic destined to a host which is sleeping, but such functionality is out of the scope of this document.

5.8.1. Picking an Appropriate Registration Lifetime

As all Neighbor Discovery messages are initiated by the hosts, this allows a host to sleep or otherwise be unreachable between NS/NA message exchanges. The Address Registration option attached to NS messages indicates to a router to keep the Neighbor Cache entry 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 the registration is renewed successfully (considering e.g. clock drift and additional time for potential retransmissions of the re-registration). 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 neighbor cache entries for nodes longer than necessary.

5.8.2. Behavior on Wakeup

When a host wakes up from a sleep period it SHOULD maintain its current address registrations that will timeout before the next wakeup. This is done by sending Neighbor Solicitation messages with the Address Registration option as described in Section 5.5.1. The host may also need to refresh its prefix and context information by sending a new unicast Router Solicitation (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 Router Solicitations to discover new default router(s) and restart the address registration process.

6. Router Behavior for 6LR and 6LBR

Both 6LRs and 6LBRs maintain the Neighbor Cache [RFC4861] based on the Address Registration Options they receive in Neighbor Advertisement messages from hosts, Neighbor Discovery 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 Neighbor Cache entries (see Section 3.4) since they need to retain them until the Registration Lifetime expires. Similarly, if Neighbor Unreachability Detection on the router determines that the host is UNREACHABLE (based on the logic in [RFC4861]), the Neighbor Cache entry SHOULD NOT be deleted but be 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 Preferences [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

A router SHOULD NOT send Redirect messages in a route-over topology, but MAY send Redirect messages in a mesh-under topology. In route-over the link has non-transitive reachability and the router has no way to determine that the recipient of a Redirect message can reach the link-layer address.
A router MUST NOT set the 'L' (on-link) flag in the Prefix Information options, since that might trigger hosts to send multicast Neighbor Solicitations.

6.2. Interface Initialization

A router initializes its interface more or less as in [RFC4861]. However, a 6LR might want to wait to make its interfaces advertising (implicitly keeping the AdvSendAdvertisements flag clear) until it has received the prefix(es) and context information from its 6LBR. That is independent of whether prefixes and context information is disseminated using the methods specified in this document, or using some other method.

6.3. Processing a Router Solicitation

A router processes Router Solicitation messages as specified in [RFC4861]. The differences relate to the inclusion of Authoritative Border Router options in the Router Advertisement (RA) messages, and the exclusive use of unicast Router Advertisements. If a 6LR has received an ABRO from a 6LBR, then it will include that option unmodified in the Router Advertisement 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 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.
A Router Solicitation might be received from a host that has not yet registered its address with the router. Thus the router MUST NOT modify an existing Neighbor Cache entry based on the SLLA option from the Router Solicitation. However, a router MAY create a Tentative Neighbor Cache entry based on the SLLA option. Such a Tentative Neighbor Cache

entry SHOULD be timed out in TENTATIVE_NCE_LIFETIME seconds unless a registration converts it into a Registered NCE.
A 6LR or 6LBR MUST include a Source Link-layer address option in the Router Advertisements it sends. That 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 Router Advertisements, this specification optimizes the exchange by always unicasting RAs in response to RSs. This is possible since the RS always includes a SLLA option, which is used by the router to unicast the RA.

6.4. Periodic Router Advertisements

A router does not need to send any periodic Router Advertisement messages since the hosts will solicit updated information by sending Router Solicitations before the lifetimes expire.
However, if the routers use Router Advertisements to optionally distribute prefix and/or context information across a route-over topology, that might require periodic Router Advertisement messages. Such RAs are sent using the configurable MinRtrAdvInterval and MaxRtrAdvInterval as per [RFC4861].

6.5. Processing a Neighbor Solicitation

A router handles Neighbor Solicitation messages as specified in [RFC4861], with added logic described in this section for handling the Address Registration option.
In addition to the normal validation of a Neighbor Solicitation and its options, the Address Registration option is verified as follows (if present). If the Length field is not two, or if the Status field is not zero, then the Neighbor Solicitation is silently ignored.
If the source address of the NS is the unspecified address, or if no SLLA option 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. If there is no Neighbor Cache entry for the IPv6 source address of the NS, then it isn't a duplicate. If there is such a Neighbor Cache entry and the EUI-64 is the same, then it isn't a duplicate either. 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 Neighbor Cache entries. In the case it is a duplicate address then the router responds with a unicast Neighbor Advertisement (NA) message with the ARO Status field set to one (to indicate 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 IID 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.

6.5.3. Updating the Neighbor Cache

If 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) a Neighbor Cache entry 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 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 Neighbor Cache entry. A unicast Neighbor Advertisement (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 TLLA (Target Link-layer Address) option [RFC4861] is not required in the NA, since the host already knows the router's link-layer address from Router Advertisements.
If the ARO contains a zero Registration Lifetime then any existing Neighbor Cache entry for the IPv6 source address of the NS MUST be deleted, and a NA sent as above.
Should the Registration Lifetime in a Neighbor Cache entry expire, then the router MUST delete the cache entry.
The addition and removal of Registered Neighbor Cache entries would result in notifying the routing protocol.
Note: If the optional 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 hosts (see Section 5.6. The routing table is checked to determine the next hop IP address. A registered Neighbor Cache Entry (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 others' link-layer addresses. If the routing protocol used between the routers provides this, then there is no need for the routers to use the Address Registration option between each other. Otherwise, the routers MAY use the ARO. When routers use ARO to register with each other and the optional multihop DAD Section 8.2 is in use, then care should 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 is out of scope of this document.

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

7. Border Router Behavior

A 6LBR handles sending of Router Advertisements and processing of Neighbor Solicitations from hosts as specified above in section Section 6. A 6LBR SHOULD always include an Authoritative Border Router option in the Router Advertisements it sends, listing itself as the 6LBR Address. That requires that the 6LBR maintain the version number in stable storage, and increases the version number when some information in its Router Advertisements change. The information whose change affects the version are in the Prefix Information options (the prefixes or their lifetimes) and in the 6CO option (the prefixes, Context IDs, or lifetimes.)

In addition, a 6LBR is somehow configured with the prefix or prefixes that are assigned to the LoWPAN, and advertises those in Router Advertisements as in [RFC4861]. Optionally, in the case of route-over, those prefixes can be disseminated to all the 6LRs using the technique in Section 8.1. However, there might be mechanisms outside of the scope of this document that can be used instead for prefix dissemination with route-over.

If the 6LoWPAN uses Header Compression [I-D.ietf-6lowpan-hc] with context then the 6LBR needs to manage the context IDs, and advertise those in Router Advertisements by including 6CO options in its Router Advertisements so that directly attached hosts are informed about the context IDs. Below we specify things to consider when the 6LBR needs to add, remove, or change the context information. Optionally, in the case of route-over, the context information can be disseminated to all the 6LRs using the technique in Section 8. However, there might be mechanisms outside of the scope of this document that can be used instead for disseminating context information with route-over.

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 are included in the Router Advertisement messages as specified in [RFC4861].

7.2. Context Configuration and Management

If the LoWPAN uses Header Compression [I-D.ietf-6lowpan-hc] with context then the 6LBR may be configured with context information and related context IDs. If the LoWPAN has multiple 6LBRs, then they MUST be configured with the same context information and context IDs.
The context information carried in Router Advertisement (RA) messages originate 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 lifecycle SHOULD be used in order to ensure the context information stays synchronized throughout the LoWPAN. New context information SHOULD be introduced into the LoWPAN with C=0, to ensure it is known by all nodes that may have to decompress 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 that nodes send packets making use of previous values of contexts, resulting 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 context ID 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. Optional Behavior

Optionally the Router Advertisement messages can be used to disseminate prefixes and context information to all the 6LRs in a route-over topology. If all routers are configured to use another mechanism for such information distribution, this mechanism MAY stay unused.
There is also the option for a 6LR to perform multihop DAD (for non-EUI-64 derived IPv6 addresses) against a 6LBR in a route-over topology

by using the DAR and DAC messages. This is optional because there might be other ways to either allocate unique address, such as DHCPv6 [RFC3315], or other future mechanisms for multihop DAD.

8.1. Multihop Prefix and Context Distribution

The multihop distribution relies on Router Solicitation messages and Router Advertisement (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 6LNs use the same information, whether they send, forward, or receive compressed packets. Thus the manager of the 6LBRs need 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 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 option.

8.1.1. 6LBRs Sending Router Advertisements

6LBRs supporting multihop prefix and context distribution MUST include an ABRO in each of its 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 PIO or 6CO options change, 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

If multihop distribution is done using Router Advertisement (RA) messages, then on interface initialization a router SHOULD send some Router Solicitation messages similarly to how hosts do this in [RFC4861]. That will cause the routers to respond with RA messages which then can 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 and nothing in Section 8.1 applies. Otherwise the routers will check and record the prefix and context information from the receive RAs, and use that information as follows.
If a received RA does not contain a Authoritative Border Router option, 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 and Version 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. The comparison MUST be done the same way as TCP sequence number comparisons to handle the case when the version number wraps around. If the version number received in the packet is less than the stored version number (following [RFC1982] Section 3.2), then the information in the RA is silently ignored. Otherwise the recorded information and version number are updated.
By TCP sequence number comparison we mean that half of the version number space is "old" and half is "new". For example, if the current version number is 0x2, then anything between 0x80000003 (0x2-0x7fffffff) and 0x1 is old, and anything between 0x3 and 0x80000002 (0x2+0x8000000) is new.

8.1.4. Storing the Information

The router keeps state for each 6LBR that it sees with an ABRO. This includes the version number, and the complete set of Prefix Information options and 6LoWPAN Context options. The prefixes are timed out based on the Valid lifetime in the Prefix Information Option. The Context Prefix is timed out based on the Valid lifetime in the 6LoWPAN Context option.
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 Router Advertisements 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 a 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

If multihop distribution is performed using RA messages, then the routers MUST ensure that the ABRO always stay together with the prefixes and context information received with that ABRO. Thus if the router has received prefix P1 with 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 a RA that include a 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 Router Advertisements as in [RFC4861]. This is for the benefit of the other routers receiving the prefixes and context information. And the routers also respond to Router Solicitations by unicasting RA messages. In both cases the above constraint of keeping the ABRO together with 'its' prefixes and context information apply.
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 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 an optional and simple technique for 6LRs and 6LBRs to perform Duplicate Address Detection that reuses the information from Address Registration option 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 the 6LRs either register with all the 6LBRs, or that 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 option is not returned to the host until multihop DAD has been completed against the 6LBRs. For existing registrations in the 6LR the 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. For instance, by tracking how much is left of the lifetime the 6LR registered with the 6LBRs and re-registering with the 6LBR when this lifetime is about to run out.
For the synchronous multihop DAD the 6LR performs some additional checks to ensure that it has a Neighbor Cache entry 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) Neighbor Cache entry 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 Neighbor Cache entry exists, then the 6LR MUST create a Tentative Neighbor Cache entry with the EUI-64 and the SLLA option. This entry will be used to send the response to the host when the 6LBR responds positively.

When a 6LR receives a Neighbor Solicitation containing an Address Registration option with a non-zero Registration Lifetime and it has no existing Registered Neighbor Cache entry, then with this mechanism the 6LR will invoke synchronous multihop DAD.

The 6LR will unicast a Duplicate Address Request 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 Duplicate Address Confirmation 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) Neighbor Cache entry. 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 Neighbor Cache entry 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 option 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 Neighbor Cache entry 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 Duplicate Address Request and Confirmation messages that do not satisfy all of the following validity checks:

- *If the message includes an IP Authentication Header, the message authenticates correctly.
- *ICMP Checksum is valid.
- *ICMP Code is 0.
- *ICMP length (derived from the IP length) is 32 or more bytes.
- *The Registered Address is not a multicast address.

*All included options have a length that is greater than zero.

*The IP source address is not the unspecified address, nor 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 the optional 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 the optional 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:

*In the IPv6 source address, a global address of the 6LR.

*In the IPv6 destination address, the address of the 6LBR.

*In the IPv6 hop limit, MULTIHOP_HOPLIMIT.

*The Status field MUST be set to zero

*The EUI-64 and Registration lifetime are copied from the ARO received from the host.

*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 Neighbor Cache entry for the host as specified in Section 6, an DAR is sent to the 6LBRs as above.
A router MUST NOT modify the Neighbor Cache as a result of receiving a Duplicate Address Request.

8.2.4. 6LBR Receiving a Duplicate Address Request

When a 6LBR that implements the optional multihop DAD receives an 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 an 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 MULTIHOP_HOPLIMIT

8.2.5. Processing a Duplicate Address Confirmation

When a 6LR that implements the optional 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 Neighbor Cache entry matching the Registered address and EUI-64, then the DAC is silently ignored. Otherwise, the information in the DAC and in the Tentative Neighbor Cache entry 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 case of 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 Neighbor Cache entry for the Registered Address is removed. Otherwise it is made into a Registered Neighbor Cache entry.
A router MUST NOT modify the Neighbor Cache as a result of receiving a Duplicate Address Confirmation, unless there is a Tentative Neighbor Cache entry 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

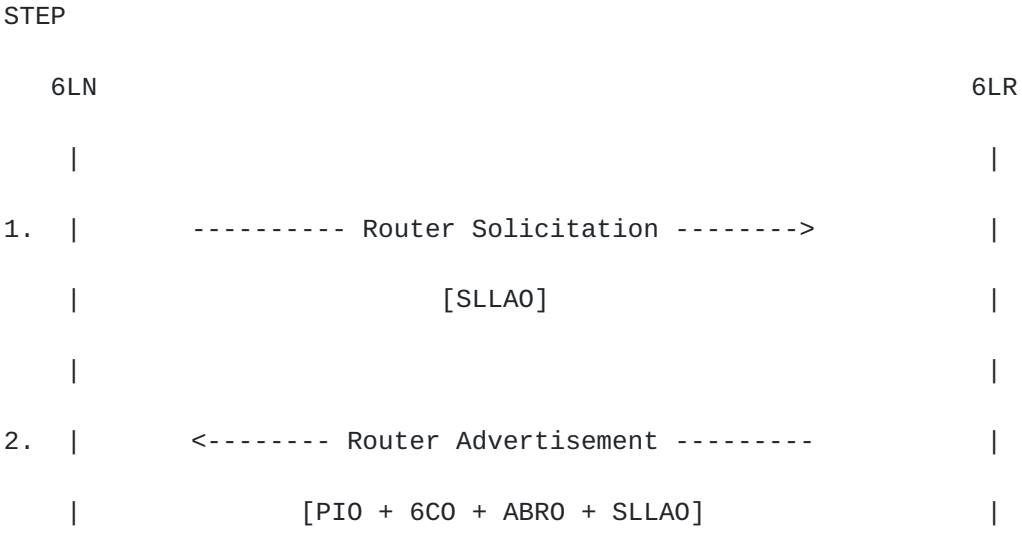

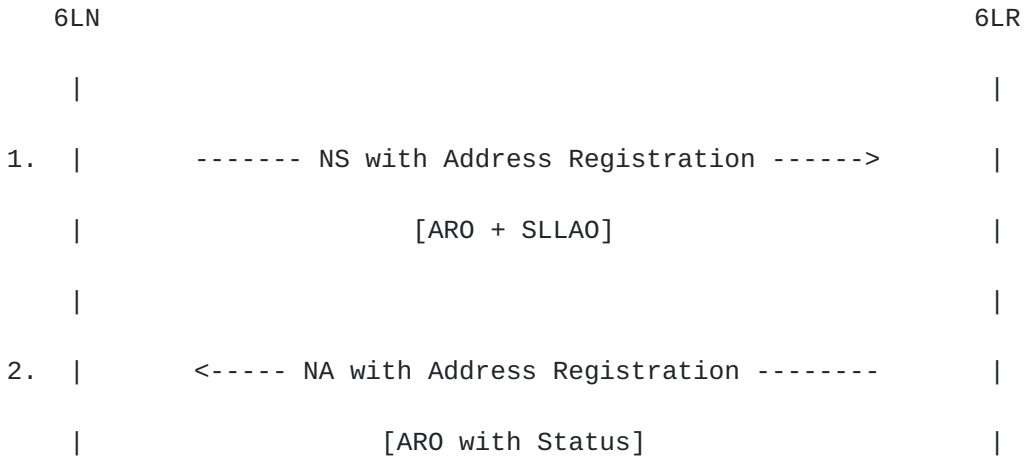

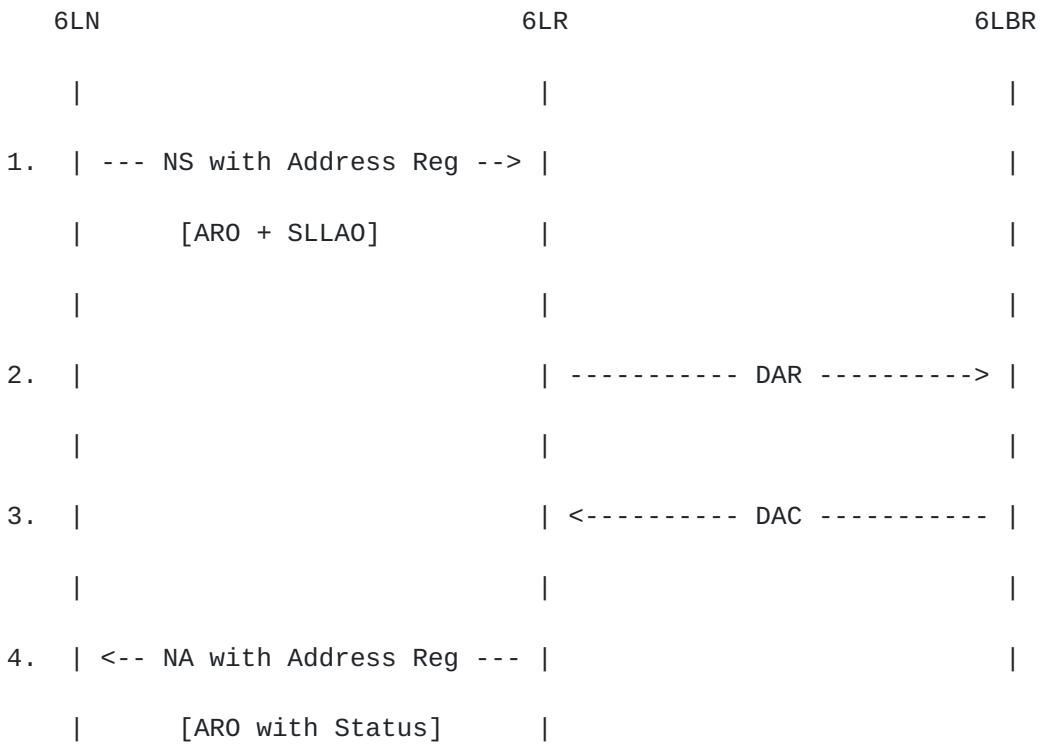

10.2. Host Bootstrapping Example

The following example describes the address bootstrapping scenarios using the optimized 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 the link-layer security is out of scope of this document. In this example an IEEE 802.15.4 6LN forms a 16-bit short-address based IPv6 addresses without using DHCPv6 (i.e., the M flag is not set in the Router Advertisements).

1. After obtaining link-level 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 SLLA Option 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 using the SLLA option from the RS (it may have created a tentative NCE). See Figure 5.
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 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 SLLA option 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 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 6. 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 is used, the effect of the 6LRs and hosts following the above bootstrapping is a "wavefront" of 6LRs and host being configured spreading from the 6LBRs. First the hosts and 6LRs that can directly reach a 6LBR would receive one or more RAs and configure and register their IPv6 addresses. Once that is done they would enable the routing protocol and start sending out Router Advertisements. That would result in a new set of 6LRs and hosts to receive responses to their Router Solicitations, 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 brings specific message examples to the previous bootstrapping process. 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 RFC 4944.
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, 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
```

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 optional multihop DAD is performed with the 6LBR. The neighbor cache management of a router and address resolution among the neighboring routers are described in Section 6.5.3 and Section 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 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, 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, R3. Initially when R2 boots it sees only R1, and accordingly R2 creates a neighbor cache entry for R1. Now assume R2 receives a valid routing update from router R3. R2 does not have any neighbor cache entry for R3. If the implementation of R2 supports detecting link-layer address from the routing information packets then it directly updates the its neighbor cache using that link-layer information. If this is not possible, then R2 should perform multicast NS with 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 neighbor cache entry is required to update its neighbor cache as described above.
The router (6LR and 6LBR) IP-addresses learned via Neighbor Discovery are not redistributed to the routing protocol.

11. Security Considerations

The security considerations of IPv6 Neighbor Discovery [RFC4861] apply. Additional considerations can be found in [RFC3756].
This specification expects that the link layer is sufficiently protected, for instance using MAC sublayer cryptography. In other words, model 1 from [RFC3756] applies. In particular, it is expected that the LoWPAN MAC provides secure unicast to/from Routers and secure broadcast from the Routers in a way that prevents tampering with or replaying the Router Advertisement messages. However, any future 6LoWPAN security protocol that applies to Neighbor Discovery for 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 Neighbor Cache entry due to such messages, and that they reject 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 [RFC3971] [RFC3972]. This is possible with the Address Registration option as sent between hosts and routers, since the address that is being registered is the IPv6 source address of the Neighbor Solicitation and SeND verifies the IPv6 source address of the packet. Applying SeND to the optional router-to-router communication in this document is out of scope.

12. IANA Considerations

The document requires three new Neighbor Discovery option types under the subregistry "IPv6 Neighbor Discovery Option Formats":

- *Address Registration Option (TBD1)
- *6LoWPAN Context Option (TBD2)
- *Authoritative Border Router Option (TBD3)

The document requires two new ICMPv6 types under the subregistry "ICMPv6 type Numbers":

- *Duplicate Address Request (TBD4)
- *Duplicate Address Confirmation (TBD5)

For the purpose of protocol interoperability testing of this specification, the following values are being used temporarily:

- *TBD1 = 31

*TBD2 = 32

*TBD3 = 33

*TBD4 = 155 XXX

*TBD3 = 156 XXX

This document also requests IANA to create a new registry for the Status values of the Address Registration Option.
[TO BE REMOVED: This registration should take place at the following location: http://www.iana.org/assignments/icmpv6-parameters]

13. Guideline for New Features

This section discusses a guideline of new features for implementation and deployment.

Section	Description	deploy	implement
3.1	Host initiated RA	MUST	MUST
3.2	EUI-64 based IPv6-address	MUST	MUST
	16bit-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	6lowpan Context Option	SHOULD	SHOULD
5.1	Re-direct Message Acceptance	MUST NOT	MUST NOT
	Joining Solicited Node Multicast	N/A	N/A
	Joining all-node Multicast	MUST	MUST
	Using link-layer indication for NUD	SHOULD	MAY
5.5	6lowpan-ND NUD	MUST	MUST
5.8.2	Behavior on wake-up	SHOULD	SHOULD

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

Section	Description	deploy	implement
3.4, 4.3, 8.1.3, 8.1.4	ABRO Processing/sending	MAY	SHOULD
8.1	Multihop Prefix storing and re-distribution	MAY	SHOULD
3.5	Tentative NCE	MUST	MUST
8.2	Multihop DAD	MAY	SHOULD
4.1, 6.5, 6.5.1 - 6.5.5	ARO Support	MUST	MUST
4.2	6LoWPAN Context Option	SHOULD	SHOULD
6.3	Process RS/ARO	MUST	MUST

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	MAY	SHOULD
8.2	Multihop DAD	MAY	SHOULD

Guideline for 6LBR features in 6LoWPAN-ND

14. Acknowledgments

The authors thank Pascal Thubert, Jonathan Hui, Carsten Bormann, 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 shaped and improve this document.
Additionally, the authors would like to recognize Carsten Bormann for the suggestions on the Context Prefix Option and contribution to earlier version of the draft, Pascal Thubert for contribution of the original registration idea and extensive contributions to earlier versions of the draft, Jonathan Hui for original ideas on prefix/ context distribution and extensive contributions to earlier versions of the draft, Colin O'Flynn for useful Error-to suggestions and contributions to the Examples section, Geoff Mulligan for suggesting the use of Address Registration as part of existing IPv6 Neighbor Discovery messages, and Mathilde Durvy for helping to clarify router interaction.

15. Changelog

Changes from -17 to -18:

*o Fixed nits related to IESG submission.

Changes from -16 to -17:

*o Removed unnecessary normative text from Assumptions.

*o Clarified the next-hop determination of multicast addresses.

*o Editorial improvements from WGLC review.

Changes from -15 to -16:

*o Added an applicability section (#133)

*o Updated document title to align with HC

*o Minor editing as result of WGLC review (#134)

Changes from -14 to -15:

*o Changed use of redirect to SHOULD NOT for route-over and MAY for mesh-under. (#130)

*o Changed the 16-bit lifetimes to a unit of 60 seconds (#131)

*o Added text to Section 5.4.2 adding a receive-only state to context entries that timeout. (#132)

Changes from -13 to -14:

*o Introduced the new DAR and DAC ICMPv6 message types for multihop DAD to avoid relying on the Length=4 checks for the ARO. This simplifies implementing the hop limit check.

*o Clarified the hop limit values for the multihop DAD messages by introducing the MULTIHOP_HOPLIMIT constant set to 64.

*o Clarified when a host should de-register from a router.

*o Added a section on next-hop determination for routers.

*o Removed the infinite lifetime from 6CO.

*o Increased MIN_CONTEXT_CHANGE_DELAY to 300 seconds.

Changes from -12 to -13:

*o Error-to solution added for returning NA messages carrying an error ARO option to the link-local EUI-64 based IPv6 address of the host (#126).

*o New examples added.

Changes from -11 to -12:

*o Version field of ABRO moved after Length for 32-bit alignment of the reserved space (#90).

*o Several clarifications were made on router interaction, including a new section with router interaction examples (#91).

*o Temporary Neighbor Cache Entry created upon host sending NS+ARO, and SLLAO removed from multihop DAD NS/NA messages (#87).

Changes from -10 to -11:

*o Reference to RFC1982 for version number comparison (#80)

*o RA Router Lifetime field use clarified (#81)

*o Make fields 16-bit rather than 32-bit where possible (#83)

*o Unicast RA clarification (#84)

*o Temporary ND option types (#85)

*o SLLA/TLLA clarification (#86)

*o GP16 as source address in initial NS clarification (#87)

Changes from -09 to -10:

*o Clarifications made to Section 8.2 (#66)

*o Explained behavior of Neighbor Cache (#67)

*o Clarified use of SLLAO in RS and NS messages (#68)

*o Added new term 6LN (#69)

*o Small clarification on 6CO flag (#70)

*o Defined host behavior on ARO failure better (#72)

*o Added bootstrapping example for a host (#73)

*o Added new Neighbor Cache Full ARO error (#74)

*o Added rule on the use of the M flag (#75)

Changes from -08 to -09:

*o Clean re-write of the draft (re-use of some introductory material)

*o Merged in draft-chakrabarti-6lowpan-ipv6-nd-simple-00

*o Changed address registration to an option piggybacked on NS/NA

*o New Authoritative Border Router option

*o New Address Registration Option

*o Separated Prefix Information and Content Information

*o Optional DAD to the edge

Changes from -07 to -08:

*o Removed Extended LoWPAN and Whiteboard related sections.

*o Included reference to the autoconf addressing model.

*o Added Optimistic Flag to 6AO.

*o Added guidelines on routers performing DAD.

*o Removed the NR/NC Advertising Interval.

*o Added assumption of uniform IID formation and DAD throughout a LoWPAN.

Changes from -06 to -07:

*o Updated addressing and address resolution (#60).

*o Changed the Address Option to 6LoWPAN Address Option, fixed S values (#61).

*o Added support for classic RFC4861 RA Prefix Information messages to be processed (#62).

*o Added a section on using 6LoWPAN-ND under a hard-wired RFC4861 stack (#63).

*o Updated the NR/NC message with a new Router flag, combined the Code and Status fields into one byte, and added the capability to carry 6IOs (#64).

*o Made co-existence with other ND mechanisms clear (#59).

*o Added a new Protocol Specification section with all mechanisms specified there (#59).

*o Removed dependencies and conflicts with RFC4861 wherever possible (#59).

*o Some editorial cleanup.

Changes from -05 to -06:

*o Fixed the Prf codes (#52).

*o Corrected the OIIO TID field to 8-bits. Changed the Nonce/OII order in both the OIIO and the NR/NC. (#53)

*o Corrected an error in Table 1 (#54).

*o Fixed asymmetric and a misplaced transient in the 6LoWPAN terminology section.

*o Added Updates RFC4861 to header

Changes from -04 to -05:

*o Meaning of the RA's M-bit changed to original [RFC4861] meaning (#46).

*o Terms "on-link" and "off-link" used in place of "on-link" and "off-link".

*o Next-hop determination text simplified (#49).

*o Neighbor cache and destination cache removed.

*o IID to link-layer address requirement relaxed.

*o NR/NC changes to enable on-link refresh with routers (#48).

*o Modified 6LoWPAN Information Option (#47).

*o Added a Protocol Constants section (#24)

*o Added the NR processing table (#51)

*o Considered the use of SeND on backbone NS/NA messages (#50)

Changes from -03 to -04:

*o Moved Ad-hoc LoWPAN operation to Section 7 and made ULA prefix generation a features useful also in Simple and Extended LoWPANs. (#41)

*o Added a 32-bit Owner Nonce to the NR/NC messages and the Whiteboard, removed the TID history. (#39)

*o Improved the duplicate OII detection algorithm using the Owner Nonce. (#39)

*o Clarified the use of Source and Target link-layer options in NR/NC. (#43)

*o Included text on the use of alternative methods to acquire addresses. (#38)

*o Removed S=2 from Address Option (not needed). (#36)

*o Added a section on router dissemination consistency. (#44)

*o Small improvements and extensive editing. (#42, #37, #35)

Changes from -02 to -03:

*o Updated terminology, with RFC4861 non-transitive link model.

*o 6LoWPAN and ND terminology separated.

*o Protocol overview explains RFC4861 diff in detail.

*o RR/RC is now Node Registration/Confirmation (NR/NC).

*o Added NR failure codes.

*o ER Metric now included in 6LoWPAN Summary Option for use in default router determination by hosts.

*o Examples of host data structures, and the Whiteboard given.

*o Whiteboard is supported by all Edge Routers for option simplicity.

*o Edge Router Specification chapter re-structured, clarifying optional Extended LoWPAN operation.

*o NS/NA now completely optional for nodes. No address resolution or NS/NA NUD required.

*o link-local operation now compatible with oDAD (was broken).

*o Exception to hop limit = 255 for NR/NC messages.

*o Security considerations improved.

*o ICMPv6 destination unreachable supported.

Changes from -01 to -02:

*o Fixed 16 != 0xff bug (ticket closed).

*o Specified use of ULAs in ad-hoc LoWPAN section 9 (ticket closed).

*o Terminology cleanup based on Alex's comments.

*o General editing improvements.

Changes from -00 to -01:

*o Specified the duplicate owner interface identifier procedures. A TID lollipop algorithm was sufficient (nonce unnecessary).

*o Defined fault tolerance using secondary bindings.

*o Defined ad-hoc network operation.

*o Removed the E flag from RA and the X flag from RR/RC.

*o Completed message examples.

*o Lots of improvements in text quality and consistency were made.

16. References

16.1. Normative References

[RFC1982]	Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, August 1996.
[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
[RFC5226]	Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
[RFC2460]	Deering, S.E. and R.M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
[RFC2491]	Armitage, G., Schulter, 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.
[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.

16.2. Informative References

[EUI64]	IEEE , "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", .
[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.
[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.

[I-D.ietf-6lowpan-hc]	Hui, J and P Thubert, "Compression Format for IPv6 Datagrams in Low Power and Lossy Networks (6LoWPAN)", Internet-Draft draft-ietf-6lowpan-hc-15, February 2011.

Authors' Addresses

Zach Shelby editor Shelby Sensinode Hallituskatu 13-17D Oulu, 90100 FINLAND Phone: +358407796297 EMail: zach@sensinode.com

Samita Chakrabarti Chakrabarti Ericsson EMail: samita.chakrabarti@ericsson.com

Erik Nordmark Nordmark Cisco Systems EMail: nordmark@cisco.com