

6MAN  
Internet-Draft  
Intended status: Standards Track  
Expires: April 18, 2014

P. Thubert  
E. Levy-Abegnoli  
Cisco  
October 17, 2013

Wireless Neighbor Discovery Stateful Address Identification and Location  
exchange  
[draft-thubert-6man-wind-sail-00](#)

## Abstract

This draft proposes an extension to IPv6 Neighbor Discovery to exchange Stateful Address Identification and Location between State Maintaining Entities located over a backbone link about attached nodes that are attached to the backbone via a Wireless Link, in order to maintain all the entities up-to-date and maintain reachability as the attached nodes move.

## Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on April 18, 2014.

## Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
--------------------	------------------------	-------------------

Internet-Draft

WiND-SAIL

October 2013

<a href="#">2.</a>	Terminology . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Overview . . . . .	<a href="#">6</a>
<a href="#">4.</a>	General Context . . . . .	<a href="#">8</a>
<a href="#">4.1.</a>	Efficient ND . . . . .	<a href="#">8</a>
<a href="#">4.2.</a>	Proxying classical ND . . . . .	<a href="#">10</a>
<a href="#">4.3.</a>	Federating Large LLNs . . . . .	<a href="#">11</a>
<a href="#">5.</a>	New types and formats . . . . .	<a href="#">12</a>
<a href="#">6.</a>	Validation Interface Operations . . . . .	<a href="#">14</a>
<a href="#">6.1.</a>	Child to Parent Operations . . . . .	<a href="#">15</a>
<a href="#">6.2.</a>	Parent to Child Operations . . . . .	<a href="#">16</a>
<a href="#">6.2.1.</a>	Address validation and registration . . . . .	<a href="#">16</a>
<a href="#">6.2.2.</a>	Registration update . . . . .	<a href="#">17</a>
<a href="#">6.3.</a>	Registration deletion . . . . .	<a href="#">18</a>
<a href="#">7.</a>	Security Considerations . . . . .	<a href="#">18</a>
<a href="#">8.</a>	IANA Considerations . . . . .	<a href="#">18</a>
<a href="#">9.</a>	Acknowledgments . . . . .	<a href="#">19</a>
<a href="#">10.</a>	References . . . . .	<a href="#">19</a>
<a href="#">10.1.</a>	Normative References . . . . .	<a href="#">19</a>
<a href="#">10.2.</a>	Informative References . . . . .	<a href="#">20</a>
	Authors' Addresses . . . . .	<a href="#">21</a>

## [1.](#) Introduction

"Neighbor Discovery for IP version 6" [[RFC4861](#)] (IPv6 ND) relies heavily on multicast signaling messages on the local Link. Conceptually, multicast is supposed to avoid broadcast messages, but, in most practical cases, its operation at the link level is that of a broadcast. This did not matter much at the time ND was originally designed, when an Ethernet network was more or less a single shared wire, but since then, large scale switched fabrics, low-power sleeping devices, mobile wireless devices and virtual machines have changed the landscape dramatically.

The overhead of multicast in IPv6 ND has become significant and is now a major annoyance in multiple scenarios, in particular for wireless nodes. With WIFI, a multicast message will consume the wireless link on all Access Points around a switched fabric and will be transmitted at the lowest speed possible in order to ensure the maximum reception by all other wireless nodes. This means that in an environment where bandwidth is scarce, a single multicast packet may consume the bandwidth for hundreds of unicast packets. Sadly, IPv6 ND is a major source of multicast messages in wireless devices, since

such messages are triggered each time a wireless device changes its point of attachment.

A similar situation can be seen in a datacenter, where Virtual Machine (VM) mobility also triggers floods of multicast messages,

which become a major hassle as the number of VMs grows to the tens of thousands and above. At the IETF, a Working Group was created to discuss Address Resolution in Massive Datacenters (ARMD), but the work did not go to completion. The problem with IPv6 ND multicast is still present, only getting worse as the scale and degree of mobility augments with the massive introduction of new mobile devices such as virtualized appliances, IoT and BYOD.

At the same time, the need to better control the ownership, utilization and location of IP addresses has become predominant in managed networks. The Source-Address Validation Improvements (SAVI) Working Group has proposed methods to locate, validate the ownership, and police the utilization of IPv6 addresses by snooping IPv6 ND and DHCP operations. But snooping requires being on the path of the protocols and is limited in particular in and for unicast responses.

Mobile nodes such as BYOD may change their point of attachment in the network but an eventual renumbering can be disruptive to existing connections. Virtual devices - typically VMs in a datacenter - also move though in a different fashion, from a physical device to the next. In any case, the need to maintain a same IPv6 address across movements implies the creation of very large, eventually multi-link, subnets. In such a large subnet, it might be difficult with the existing protocols to differentiate duplication from a rapid sequence of movements. And if it is indeed a sequence of movements, then it might be difficult to select the freshest information, and additional signaling is required to obtain the actual location of an address in a deterministic fashion.

In a modern managed switched fabric, a number of devices host IPv6

State Maintaining Entities (6SMEs) that hold Stateful Address Identification and Location (SAIL) information about the entity that owns an IPv6 address. A 6SME needs to reascertain periodically the state that it maintains and eliminate stale information. It is of common interest between all 6SMEs to share their information and help one another learn new state, update existing state and remove stale state rapidly. A Binding Table maintained by a secured registration protocol is certainly a more robust basis for 6SME activity than a classical IPv6 NDP [[RFC4861](#)] Neighbor Cache management coupled with protocol snooping as currently found with SAVI [[RFC6620](#)].

Mobile IPv6 [[RFC6275](#)] introduced such a registration protocol to maintain a tunnel and enable an IPv6 ND proxy operation over a Home Network. Applied to IPv6 Neighbor Discovery, the registration model balances the benefits of distributed Stateless Address AutoConfiguration (SLAAC) [[RFC4862](#)] for scalability and autonomic behaviours with the capability to reject or recuse an autoconfigured address on an exception basis - based for instance on administrative

policies -, which is a desired feature for managed networks that classically are operated with DHCPv6 [[RFC3115](#)]. In that sense, the ND registration allows a scalable hybrid of managed and non-managed networks while minimizing the total number of multicast messages between hosts, as well as between hosts and routers.

An IPv6 ND registration mechanism was standardized as Neighbor Discovery Optimization for Low-power and Lossy Networks [[RFC6775](#)]. The host to SME router operation is generalized by wireless ND [I-D.chakrabarti-nordmark-6man-efficient-nd] for devices that are not necessarily attached to a LLN but may still benefit from registration. [[RFC6775](#)] also introduces a protocol between SMEs based on new ICMP messages. This draft extends that model in order to allow for a distributed, eventually hierarchical set of SMEs to share and maintain SAIL states.

## [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](#)].

Readers are expected 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](#)], Neighbor Discovery Optimization for Low-power and Lossy Networks [[RFC6775](#)] and "Multi-link Subnet Support in IPv6" [I-D.ietf-ipv6-multilink-subnets].

Readers may benefit from reading the "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks" [[RFC6550](#)] specification; "Multi-Link Subnet Issues" [[RFC4903](#)]; "Mobility Support in IPv6" [[RFC6275](#)]; "Neighbor Discovery Proxies (ND Proxy)" [[RFC4389](#)]; "FCFS SAVI: First-Come, First-Served Source Address Validation Improvement for Locally Assigned IPv6 Addresses" [[RFC6620](#)]; and "Optimistic Duplicate Address Detection" [[RFC4429](#)] prior to this specification for a clear understanding of the art in ND-proxying and binding.

Additionally, this document uses terminology from [I-D.ietf-roll-terminology], and reuses or introduces the following terminology:

6LoWPAN Router (6LR): Please refer to [[RFC6775](#)].

6LoWPAN Border Router (6LBR): Please refer to [[RFC6775](#)].

DAR and DAC messages: Please refer to [[RFC6775](#)].

Multi-link subnet: Please refer to [I-D.ietf-ipv6-multilink-subnets].

**Backbone:** A link that forms the core of a multi-link subnet. All ARs and legacy devices are connected to the Backbone and classical ND operation ensure connectivity over that link.

**attached link:** An abstract link in a multi-link subnet other than the backbone. That link is classically not implemented as a fixed wire, and may only provide non-continuous connectivity, in particular with support for mobility. An attached link may be for instance a classical WiFi (IEEE802.11) link, a link in a

wireless mesh network, or an overlay tunnel.

attached node: A device in a multi-link subnet that is not directly connected to the Backbone but reachable via an attached link.

Backbone Router (BBR): A BBR is an IPv6 router that connects attached links to a Backbone link and enables the connectivity of an attached node by proxying IPv6 NDP over the Backbone for that node, either with the node MAC address, or its own. The BBR is a 6SME that can obtain attachment states from attached nodes by different methods, for instance by snooping IPv6 NDP or DHCPv6, by learning host routes acting as a RPL root, by accepting ND registrations acting as an AR or an IR.

Stateful Address Identification and Location (SAIL): As opposed to a cache entry, a SAIL state is Stateful in that it is obtained and maintained through a (secured) registration mechanism. A SAIL state may include for instance a secured identification of the owner of the address (e.g. a trusted token, a public key or a certificate), the position of the IPv6 address in the network (e.g. VLAN, Access Switch or Access Point), or the mapping of the IPv6 address with a MAC address. Some of this information may be stable, for instance a owner Identification, while other may be transient, for instance the Access Point identifier in a mobility scenario or the MAC address mapping in the case of NDP proxy operations.

State Maintaining Entity (SME): An entity that hold SAIL information. SMEs are implemented in devices such as security appliances such as Network Access Controllers (NACs), SAVI switches that protect the ownership of an IPv6 address and control the ingress of the network, Wireless LAN Controllers (WLCs) that terminate a CAPWAP tunnel and must rapidly re-enable reachability for a mobile device both at layer 2 and layer 3, as well as overlay terminators such as used for network virtualization (NV03). Overlay termination may operate both at layer 2 or layer 3, and may be found in data centers and enterprise networks to support mobility or extend the layer 2 fabric over a Layer 3 infrastructure, as well as in Service Provider networks to support IPv6 mobility.

**Binding:** The association of an IPv6 address with some SAIL state. A registrar maintains a binding table to store and query such associations.

**Registering Node (RN):** An IPv6 node that obtains and retains ownership of an IPv6 address through the process of IPv6 ND registration.

**Authoritative Registrar (AR):** A 6SME that stores authoritative information about a registration. An AR is the reference for address and SAIL state binding within its domain of authority, e.g. a specific subset of addresses within a subnet. There can be multiple ARs in a subnet and domains may overlap for redundancy and balancing.

**Intermediate Registrar (IR):** A 6SME that stores information about a registration as part of the registration flow. IRs form a directed acyclic graph (DAG) that is directed towards ARs. A registration from an RN will be addressed to an IR and will follow the IR DAG till it reaches a node that can grant the ownership, typically a AR.

**Registration Interface (RIF):** The interface between an RN and an IR. The RIF is typically implemented using Wireless ND, but can also be implicitly implemented by snooping IPv6 ND, e.g. as suggested by SAVI.

**Validation Interface (VIF):** The interface between a child IR and a parent IR or AR. The VIF is typically implemented using this specification which extends the DAR and DAC messages as defined in [[RFC6775](#)].

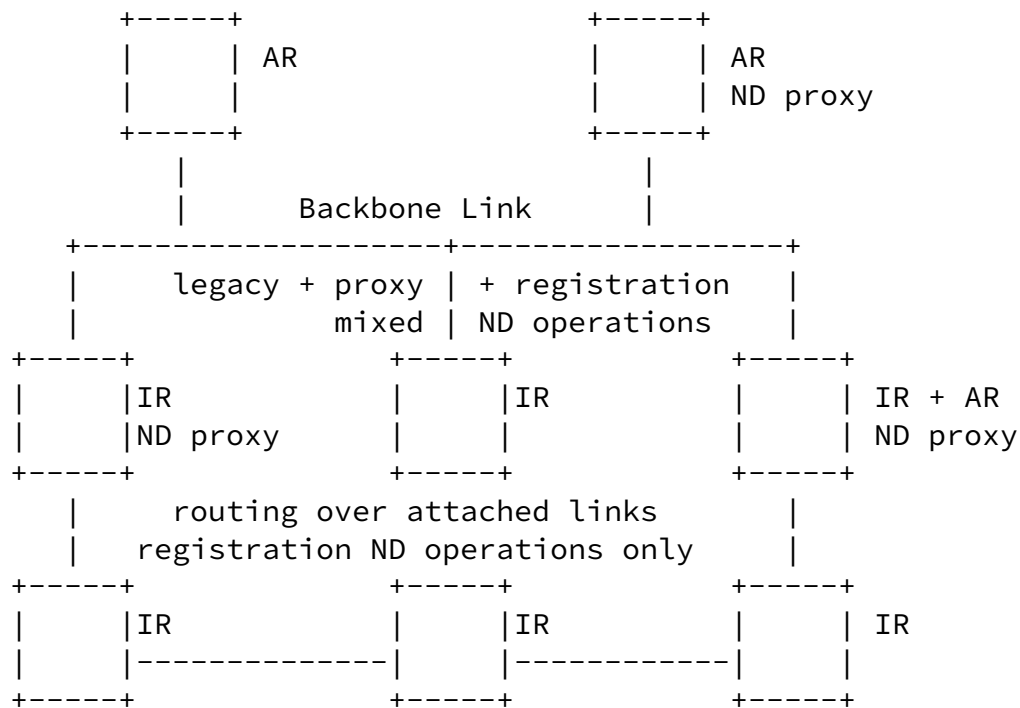
**Determination Interface (DIF):** The interface between ARs. It can be implemented using LISP, routing protocol extensions, or using IPv6 ND proxy extensions such as suggested by [I-D.thubert-6lowpan-backbone-router] .

### 3. Overview

The scope of this draft is a potentially large and potentially multi-link subnet [[I-D.ietf-ipv6-multilink-subnets](#)] formed by a high speed Backbone that federates additional links of heterogeneous MAC/PHY types, for instance an Ethernet switched fabric federating a Route-Over mesh that may interconnect thousands of LLN devices over multiple wireless hops.

In order to avoid floods of multicast packets inherent to a reactive discovery, a node - referred to as a Registering Node (RN) - needs to claim its addresses proactively, binding them with its location in the network and a Lower Layer Address (LLA), over confirmed exchanges with a neighborhood Intermediate Registrar (IR).

October 2013

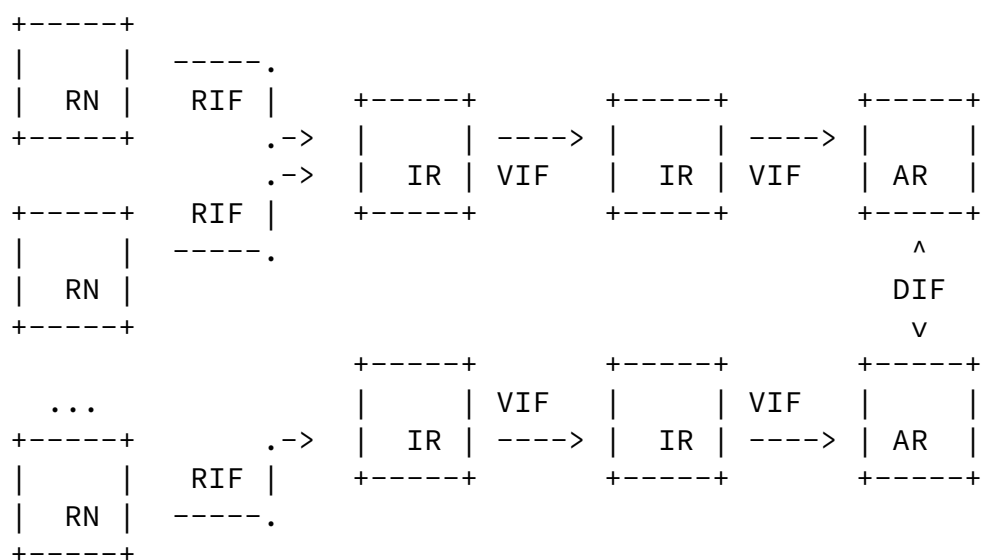


In the case of a meshed 6LoWPAN [[RFC6282](#)] [[RFC6775](#)] LLN topology , the neighborhood IR is a 6LoWPAN Router (6LR) and the AR is a 6LoWPAN Border Router (6LBR). When the topology grows, the IPv6 ND registration model as described in [[RFC6775](#)], with a single AR (the 6LBR), may not scale.

With this draft, all registrars maintain an abstract Binding Table of their registered addresses. The Binding Table operates as a distributed database of information related to addresses whether the address owner reside on the attached links or on the Backbone. ARs use extensions to the Neighbor Discovery Protocol to exchange that information across the Backbone either in the classical ND reactive fashion, or through a new pub/sub mechanism that is introduced by this specification.



With this specification, multiple IRs and one AR can be deployed so as to scale the IPv6 ND registration model yet avoiding any broadcast beyond one Layer-2 hop; IRs cover the whole multi-link subnet in a fashion that any node in the network has at least one neighborhood IR one Layer-2 hop away so it may perform an NDP registration with that IR using link local addresses regardless of the link type, wired or wireless.



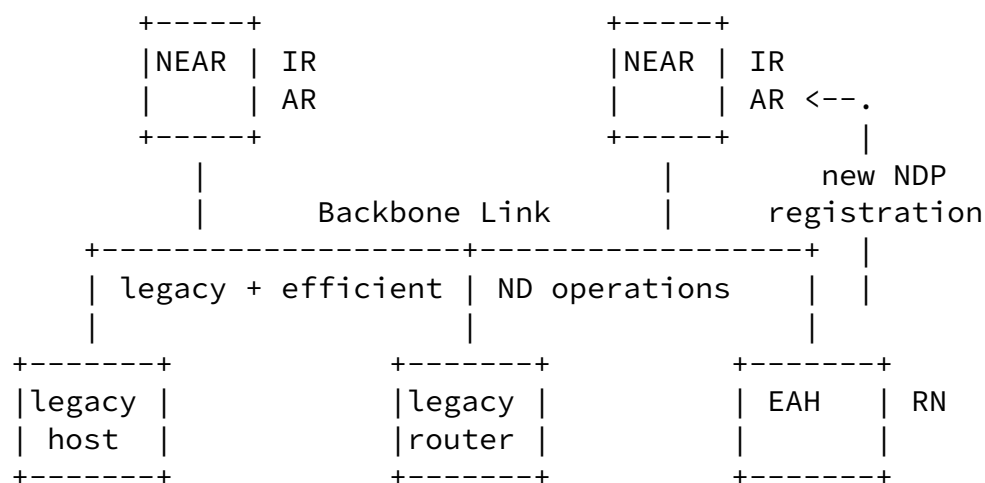
IRs and ARs form a DAG directed towards the AR(s); but how this DAG is set up is out of scope for this specification.

If more than one AR is deployed, a strategy (e.g. a distributed hash table (DHT) or a DNS-like hierarchy) and a method to distribute and synchronize the individual domains of authority between ARs, must be put in place. Such method is out of scope for this document. In the case where an overlap of domain is acceptable, a protocol must be put in place between ARs so as to resolve conflicts, and clean up stale states. Such a protocol is out of scope for this document.

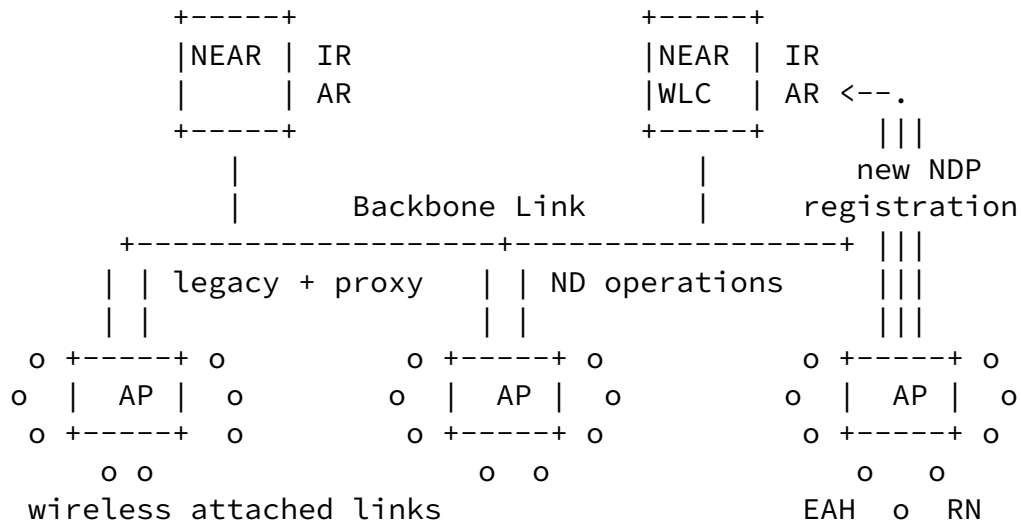
#### 4. General Context

#### 4.1. Efficient ND

[I-D.chakrabarti-nordmark-6man-efficient-nd] updates the specification of the RIF interface between the RN and the IR, that was initially defined in [RFC6775] for 6LoWPAN devices. The draft details the operation of a IPv6 ND-efficiency-aware Router (NEAR), that is the neighborhood IR to which an RN, which is called a Efficiency-Aware Host (EAH), registers. A NEAR is also an AR as it has the exclusive authority on the bindings for its registered EAHs.



In the case of a WIFI connection, the NEAR is a BBR for the wireless device, and may be collocated with a standalone AP or a Wireless LAN Controller.



This specification extends [I-D.chakrabarti-nordmark-6man-efficient-nd], by allowing the separation of IR and AR functions, which are collapsed inside the NEAR. This draft introduces the VIF interface between IRs, and between IRs and ARs, as well as the DIF interface between ARs with potentially overlapping domain, and other 6SMEs.

For the purpose of the new NDP registration, [I-D.chakrabarti-nordmark-6man-efficient-nd] defines an extended ARO option that is advertised by an EAH. The new ARO option includes a sequence counter called TID that enables a short term freshness assertion between rapid re-registrations of a mobile device, and a unique ID that is used for the Duplicate Address Detection (DAD).


A 6SME may maintain a state for a longer time than covered by the ARO TID, so a coarse age information is needed to compare old state information over the VIF and DIF interfaces. Additionally, the 6SME may qualify information with additional metadata to help resolve conflicts. For instance, in the case of a duplicated IPv6 address, additional meta information such as the protocol that was used to establish the state (SLAAC vs. DHCPv6), a device type (trusted server vs. unknown host), or a Secure ND cryptographic address ownership validation ([RFC3971], [RFC3972]) can help protect the address where it is assigned in a more trusted fashion even if a

rogue managed to grab the address while the more trusted owner was not able to defend it.

This specification proposes a new ND option that contains such information and complements the information in the ARO option for use on the VIF and DIF interfaces.

## 4.2. Proxying classical ND

A 6SME such as an IR, a AR or a BBR may proxy classical IPv6 NDP [RFC4861] on behalf of a virtual, a wireless, or a low power device so as to offload the device, to dampen the network load such as induced by the multicast operations of the proxied protocol, or simply to attract over the backbone and then relay its traffic to the mobile or sleeping device even if the device is not reachable at that particular time.



The 6SME may perform the proxy operation on behalf of an original device using the original device LLA, or may proxy the Layer-2 information with their own LLA and either rewrite it later in the packets, or route the packet again over an attached link, as exemplified by a MIPv6 [[RFC6275](#)] or a NEMO [[RFC3963](#)] Home Agent (HA).

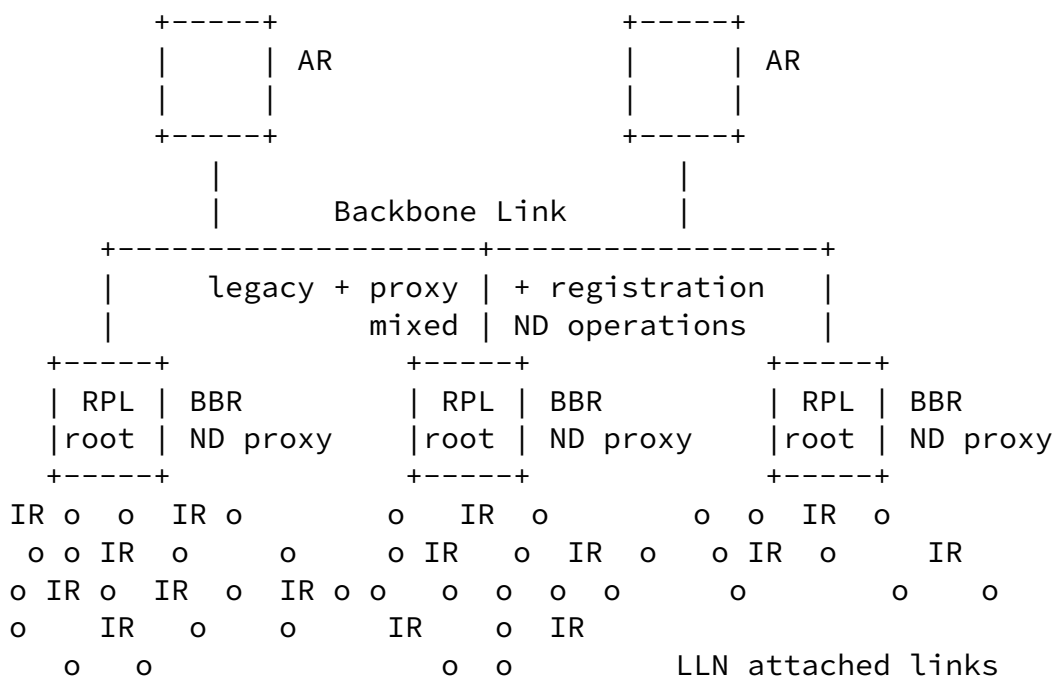
The HA is authoritative for any Mobile Node (MN) that successfully

registers to it through a Binding Update/Ack flow and its domain of authority is the subnet(s) on the Home Link, potentially overlapping with other HAs. ND proxy operations are used over the Home Link to resolve collisions.

The MN is thus an RN and the HA cumulates IR, AR and proxy functionalities. With NEMO [RFC3963], the model is conserved but the RN is now a Mobile Router that registers a prefix together with its own address, so the operation in the Backbone link is a mix of ND proxy and routing. Network Mobility Home Network Models [RFC3963] provides more information on that model.

#### 4.3. Federating Large LLNs

In the case of a large multi-link subnet, this specification expects that a Backbone link is deployed to interconnect all the ARs and legacy NDP devices. Each interconnected attached link, whether it is a WIFI access, a mesh network, a 6LoWPAN/RPL LLN or an overlaid tunnel, is anchored to the Backbone at a Backbone Router (BBR). The BBRs interconnect the multi-link subnet over the Backbone Link at layer 3, enabling connectivity within the subnet over IP.



If the LLN uses a Route-Over model based on RPL [RFC6550], the Backbone Router (BBR) that connects the LLN to the Backbone is the root for the RPL LLN. The BBR proxies the ND protocol over the backbone for the addresses that it has learnt through RPL as host routes, using its own LLA and location to attract traffic for the attached nodes and route it over the LLN.

Internet-Draft

WiND-SAIL

October 2013

Over the Backbone, this setup implements the "simple scenario" in [\[I-D.ietf-ipv6-multilink-subnets\]](#) whereby the router acts "as an asymmetric Neighbor Discovery proxy"; over the RPL-based LLN mesh, the setup implements the more "complex scenario" whereby "an arbitrary topology exists, and routers within the subnet communicate using some means of exchanging host routes".

[I-D.thubert-6lowpan-backbone-router] describes this mixed model, and how a Backbone Router perform ND proxy operation for their attached nodes over the The Backbone Link regardless of the mode of registration for the attached nodes. The operation described in the draft is compatible with that of a MIPv6 [\[RFC6275\]](#) Home Agent. This enables mobility support for wireless attached nodes that would move outside of the network delimited by the Backbone link and back. In any case, it is expected that the registration provides a sequence counter, a lifetime and a unique identifier of the attached nodes in such a fashion that they can be matched or compared across protocols.

This specifications indicates how the new ND option can be used in conjunction with ND proxy techniques over the Backbone to implement the DIF interface.

## [5.](#) New types and formats

This section introduces message formats for all messages used in this specification.

The specification expects that the protocol running on the LLN can provide a sequence number called Transaction ID (TID) that is associated to the registration. When a node registers to multiple registrars (IRs or ARs), it is expected that the same TID is used, to enable the registrar to correlate the registrations as being a single one, and differentiate that situation from a movement. Otherwise, the resolution makes it so that only the most recent registration was perceived from the highest TID is kept.

The specification expects that the protocol running on the LLN can provide a unique ID for the owner of the address that is being registered. The Owner Unique ID enables to differentiate a duplicate registration from a double registration. In case of a duplicate, the last registration loses. The Owner Unique ID can be as simple as a

EUI-64 burnin address, if the device manufacturer is convinced that there can not be a manuf error that would cause duplicate EUI64

Internet-Draft

WiND-SAIL

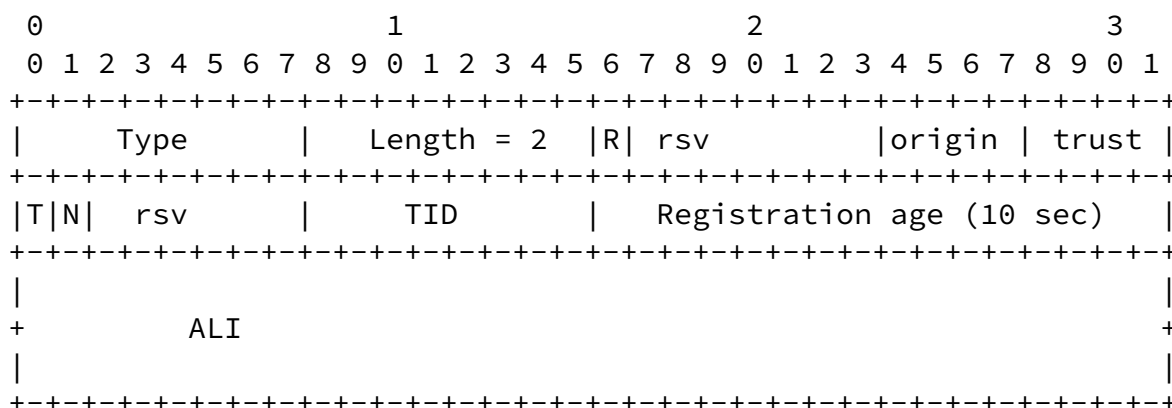
October 2013

addresses. Alternatively, the unique ID can be a hash of supposedly unique information from multiple orthogonal sources, for instance:

- o Burn in address.
- o configured address, id, security keys...
- o (pseudo) Random number, radio link metrics ...

In any fashion, it is recommended that the device stores the unique ID in persistent memory. Otherwise, it will be prevented to re-register after a reboot that would cause a loss of memory until the Backbone Router times out the registration.

The unique ID and the sequence number are placed in a new ND option that is used by the Backbone Routers over the Backbone link to detect duplicates and movements. The option format is as follows:



Option Fields

Type: 8-bit identifier of the type of option.

Length: 2

R: One bit flag. Set if the sender is relaying the option received from a downstream node, whether a RN or an IR.

origin : 4-bit unsigned integer. Indicates the origin of the entry.

0 - SLACC: Address was auto-configured on the RN [RF4862]

1 - DHCP: Address was assigned to the RN by DHCP

2 - LOCAL: Address was manually configured on the RN

3 - STATIC: Address was manually configured on the IR as a downstream address, i.e. an address assigned to a downstream node

4 - DATA: Address was gleaned on the first IR as the source of a data packet

trust : 4-bit unsigned integer. Indicates the level of trust the attaching node place in the entry

0 - NO\_TRUST: No particular trust associated with the entry

1 - L2L3\_MATCH: The layer-2 source MAC and Link-layer-Address claimed in the registration match

2 - TRUSTED\_BY\_POLICY: The address is trusted by policy on the attaching node

3 - AUTHENTICATED The address has been authenticated by a cryptographic protocol (CGA, etc.)>

T: One bit flag. Set if the next octet is a used as a TID following



follow [section 7](#) of RPL [[RFC6550](#)] for sequence counters. If the bit is not set, a unsigned char is expected.

N: One bit flag. Set if the device moved. If not set, the router will refrain from sending gratuitous NA(0) over the backbone, for instance after the DAD operation upon entry creation.

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

TID: 1-byte integer; a transaction id that is maintained by the device and incremented with each transaction. it is recommended that the device maintains the TID in a persistent storage. The TID is incremented at each registration.

Registration Age: 2-byte integer; the duration since the last update of TID in units of 10 seconds.

Attaching Location Identifier: A locally unique identifier for the IR interface attaching the registering host.

## [6.](#) Validation Interface Operations

The Validation Interface (VIF) is the interface between a child Intermediate Registrar (IR) and a parent registrar, whether an Intermediate Registrar or Authoritative Registrar (AR). An IR parent or upstream chain is defined as the set of IRs along the DAG starting from this IR, directed to (and including) the AR. An IR child or downstream chain is defined as the set of IRs from this IR an entry was learnt from including the IR attaching to the RN. The goal of VIF is to perform one or several of the followings:

- o Validate a new registration along the parent chain
- o Record a registration along the parent chain
- o Update a registration along the parent chain
- o Cancel a registration along the parent chain
- o Delete a record along the parent chain
- o Delete a record along the child chain

### 6.1. Child to Parent Operations

The first IR in the chain (attached to the RN) initiates validation and registration of an address registered by the RN or snooped on the interface attaching the RN to the node, by building a DAR message, including a SLLA and a SAIL option where:

- o Source of the DAR is an IP address of the IR interface to the next IR in the chain.
- o Destination of the DAR is an IP address of the next IR.
- o R bit set to zero. If the IR acts as an RN, the the bit is set to 1
- o Origin can take any of the values defined, based on how the address was assigned
- o If the registration was received from the RN (or gleaned) on an IR interface administratively trusted, the field "trust" is set to TRUSTED\_BY\_POLICY. Otherwise, if the registration carried CGA [[RFC3971](#)] credential that the IR successfully verified, the field "trust" is set to AUTHENTICATED. Otherwise, if the source mac of the registration message received by the IR is identical to the Link-Layer Address provided by the message in the SLLA option, the trust field is set to L2L3\_MATCH. Otherwise, it is set to NO\_TRUST.
- o An SLLA option MUST be included in the DAR message along with the SAIL option, that contains the Link-Layer address bound to the IP address bein registered.

While waiting for a response, a TENTATIVE entry is created on the IR.

Several attributes are stored next to the entry: the EUI64 provided by the RN, Link Layer Address provided in the SLLA option, the origin and trust values (computed by the IR, based on the registration, and local policies), the ALI the registration was received from, the lifetime. In the absence of a DAC response, DAR messages sent in the context of VIF from the IR are retransmitted after 250ms [DAR\_INTERVAL], up to 2 times [MAX\_DAR\_RETRANSMIT]. Upon receiving a negative response (duplicate address status) or when the maximum retransmit is exhausted, the entry is removed from the IR.

The IR can also initiate update and delete operations. An update is no different from an address validation: the DAR will carry the same address and EUI-64 as the one provided in a previous validation, while any other attribute such as SLLA option, Lifetime, Origin, Trust or ALI will eventually be different from the value previously provided

In order to cancel a registration, a DAR is sent, with a lifetime set to zero. It should carry a SAIL option to allow the receiving IR or AR to validate the delete.

## [6.2.](#) Parent to Child Operations

### [6.2.1.](#) Address validation and registration

Upon receiving a DAR message with a SAIL option, an IR will lookup in the local table to verify whether the address already exist.

If it does not exist, two cases arise.

1. The IR is not an AR. It creates the entry as "TENTATIVE", together with attributes such as EUI64, IR address it came from, origin and trust values. It then builds and sends a DAR message sourced with one address of its interface to the next IR, set destination address to the next IR address, sets the R bit to 1 and copies all other fields from the received DAR. It also starts a 250ms TENTATIVE\_TIMER timer of 250ms [DAR\_INTERVAL]. Should this timer expire, the DAR is re-transmitted up to 2 times, then the entry is deleted. If a negative response (DAC with status 1) is received, a DAC with status 1 is sent to the downstream IR, and the TENTATIVE entry is deleted. If a positive response (DAC with status 0) is received, the timer TENTATIVE\_TIMER is stopped, the entry state moved to REACHABLE and a DAC with status 0 is sent to downstream IR.

Internet-Draft

WiND-SAIL

October 2013

2. The IR is also the AR: it queries other ARs over the DIF interface. While waiting for a response, it may create the entry in "TENTATIVE" state. Upon confirmation from another AR that the entry exist elsewhere, the entry in TENTATIVE is deleted, and a DAC message is sent back to the source of the DAR message (previous IR), with a status set to 1 (Duplicate address). Upon receiving this DAC, each downstream IR deletes its own TENTATIVE entry, and sends a DAC, status 1, to the next child IR until it reaches the IR attaching the RN, which builds an NA with ARO option, and status set to duplicate address. If the DIF interface returns no conflict on the address, the entry state is moved to REACHABLE, and a DAC with status 0 is sent to the downstream IRs which move their TENTATIVE entry to REACHABLE. When the DAC reaches the attaching IR, it send an NA with ARO option, status 0 to the RN.

If the same address carried in the DAR exist on one of the IR or the AR, with a different EUI-64 interface identifier, the two entries attributes are compared. A trustlevel value is computed for each entry (as a function of the trust value, the origin and the R bit). The two trustlevel values are compared numerically as follows:

1. If the trustlevel of the existing entry is bigger or equal than the one carried by the DAR, the DAR is not propagated, and a DAC with status 1 (duplicate address) is sent back to downstream IR, up to the attaching IR which sends a NA with an ARO option, status 1 to the RN.
2. If the trustlevel of the existing entry is strictly smaller than the one carried by the DAR, it replaces it, and the DAR is propagated towards the upstream IR up to the AR. Again, DAR follow the rule of hop-by-hop retransmission and acknowledgment already described.
3. At the same time, if the entry being replaced was associated with a different IR than the one this DAR came from, another DAR, with the previous EUI64 value, and a lifetime set to zero is sent to downstream IR the previous registration came from. This message causes the downstream IR to remove the entry, provided that the EUI64 match, to build and send a DAR to the next IR, and to acknowledge the deletion with a DAC, status 0. If the EUI64 don't match, it means the entry has already been replaced, and the DAR

need not to be propagated from this IR. DAR retransmission follow the same pattern already described. DAC are not propagated. Upon receiving the DAR with lifetime set to zero, the attaching IR sends an unsolicited NA to the RN with an ARO option, status 1 (duplicate address).

### [6.2.2.](#) Registration update

An update message is a DAR that carries an address and EUI64 interface identifier matching an IR or AR table entry, and at least one of the following field different from the previously registered value: LifeTime, Origin, trust, ALI or IR child. Upon receiving an update, the IR or AR updates its entry and propagate to the upstream IR or AR. The AR will in return send a DAC message with a status of 0 to acknowledge the update.

If the attribute being updated is the IR address the DAR is coming from (child IR), the host has moved to a different downstream IR chain, and the entry along the previous chain must be cleaned up. A DAR message, with a lifetime set to zero is sent (and retried if not acknowledged) to the old downstream IR. This message causes the downstream IR to remove the entry. The downstream IR should propagate the DAR to the next IR in chain, and acknowledge it with a DAC.

### [6.3.](#) Registration deletion

A registration deletion can come from the IR attaching to the RN, because the RN left the link, from any IR as the result of an administrative action, or from the AR because the lifetime has expired or again following an administrative action. In all cases, a DAR message with a lifetime set to zero is sent either upstream or downstream, retried and acknowledged at each hop along the chain, if necessary. When the deletion is initiated on the IR attaching to the RN, a SAIL option MUST be provided to enable any upstream registrar to verify that the deletion is coming from the location the RN was attached to. For deletion following the parent chain, the ALI value carried in the SAIL option is compared with the ALI value registered for this address, and entry is deleted if the two match. For deletion following the child chain, this check is not required. Upon

deleting the entry, the IR builds and sends a DAC to acknowledge the deletion, then build and send a DAR to propagate the deletion, downstream or upstream.

## 7. Security Considerations

This specification expects that the link layer is sufficiently protected, either by means of physical or IP security for the Backbone Link or MAC sublayer cryptography. In particular, it is expected that the LLN MAC provides secure unicast to/from the Backbone Router and secure BBRoadcast from the Backbone Router in a way that prevents tempering with or replaying the RA messages.

The use of EUI-64 for forming the Interface ID in the link local address prevents the usage of Secure ND ([\[RFC3971\]](#) and [\[RFC3972\]](#)) and address privacy techniques. Considering the envisioned deployments and the MAC layer security applied, this is not considered an issue at this time.

## 8. IANA Considerations

Thubert & Levy-Abegnoli Expires April 18, 2014

[Page 18]

---

Internet-Draft

WiND-SAIL

October 2013

A new type is requested for an ND option.

## 9. Acknowledgments

TBD

## 10. References

### 10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S.E. and R.M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", [RFC 4429](#), April 2006.

- [RFC4443] Conta, A., Deering, S. and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W. and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC4862] Thomson, S., Narten, T. and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J. and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), September 2011.
- [RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP. and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", [RFC 6550](#), March 2012.
- [RFC6620] Nordmark, E., Bagnulo, M. and E. Levy-Abegnoli, "FCFS SAVI: First-Come, First-Served Source Address Validation Improvement for Locally Assigned IPv6 Addresses", [RFC 6620](#), May 2012.

- [RFC6775] Shelby, Z., Chakrabarti, S., Nordmark, E. and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", [RFC 6775](#), November 2012.

## [10.2.](#) Informative References

- [I-D.chakrabarti-nordmark-6man-efficient-nd]  
Chakrabarti, S., Nordmark, E. and M. Wasserman, "Efficiency aware IPv6 Neighbor Discovery Optimizations", Internet-Draft [draft-chakrabarti-nordmark-6man-efficient-nd-01](#), November 2012.

- [I-D.ietf-ipv6-multilink-subnets]  
Thaler, D. and C. Huitema, "Multi-link Subnet Support in IPv6", Internet-Draft [draft-ietf-ipv6-multilink-subnets-00](#), July 2002.
- [I-D.ietf-roll-terminology]  
Vasseur, J., "Terminology in Low power And Lossy Networks", Internet-Draft [draft-ietf-roll-terminology-12](#), March 2013.
- [I-D.thubert-6lowpan-backbone-router]  
Thubert, P., "6LoWPAN Backbone Router", Internet-Draft [draft-thubert-6lowpan-backbone-router-03](#), February 2013.
- [RFC3115] Dommety, G. and K. Leung, "Mobile IP Vendor/Organization-Specific Extensions", [RFC 3115](#), April 2001.
- [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A. and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", [RFC 3963](#), January 2005.
- [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.
- [RFC4389] Thaler, D., Talwar, M. and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", [RFC 4389](#), April 2006.
- [RFC4887] Thubert, P., Wakikawa, R. and V. Devarapalli, "Network Mobility Home Network Models", [RFC 4887](#), July 2007.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", [RFC 4903](#), June 2007.
- [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.

- [RFC6275] Perkins, C., Johnson, D. and J. Arkko, "Mobility Support in IPv6", [RFC 6275](#), July 2011.



## Authors' Addresses

Pascal Thubert  
Cisco Systems  
Village d'Entreprises Green Side  
400, Avenue de Roumanille  
Batiment T3  
Biot - Sophia Antipolis, 06410  
FRANCE

Phone: +33 4 97 23 26 34  
Email: pthubert@cisco.com

Eric Levy-Abegnoli  
Cisco Systems  
Village d'Entreprises Green Side  
400, Avenue de Roumanille  
Batiment T3  
Biot - Sophia Antipolis, 06410  
FRANCE

Phone: +33 4 97 23 26 34  
Email: elevyabe@cisco.com

