IETF 125 – 6lo

Generic Address Assignment Option for 6LoWPAN Neighbor Discovery

~~draft-ietf-6lo-nd-gaao-08~~

draft-ietf-6lo-nd-gaao-09

L. Iannone, D. Lou, A. Rashid,

IETF 125 – Shenzhen

Since IETF 124

Generic Address Assignment Option for 6LoWPAN Neighbor Discovery
draft-ietf-6lo-nd-gaao-09

Status | IESG evaluation record | IESG writeups | Email expansions | History

Versions:

00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09

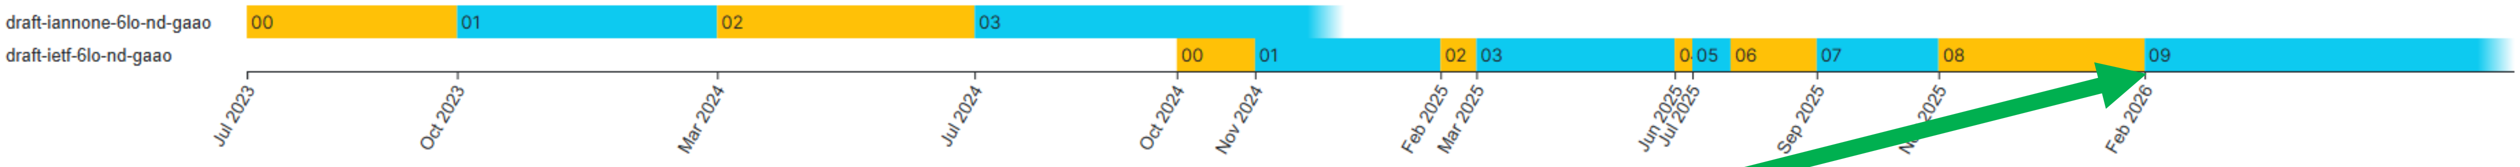

draft-ietf-6lo-nd-gaao-09.txt
February 2026
Main changes: Revised Introduction to address Lorenzo’s comments (Kudos to Adnan)

Main Changes -08: Structure

Old	New
Table of Contents	Table of Contents
1. Introduction 2	1. Introduction 2
	1.1. Limitations of DHCPv6 in Constrained LLNs 3
	1.2. Algorithmic and Distributed Address Assignment 4
2. Terminology 4	2. Terminology 4

Old	New
1. Introduction	1. Introduction
Low Power and Lossy Networks (LLNs) have required adapting the design of Internet protocols to more constrained environments, by taking into consideration energy saving, limited memory capacity, and duty cycling of the LLN devices, as well as low-power and lossy transmissions. Since the wireless interface is a major energy drain, protocols aiming at being deployed over LLNs must be designed in such a way to reduce as much as possible transmissions and idle listening, allowing to turn off the radio interface or put the interface or the whole node in the sleeping mode.	Low Power and Lossy Networks (LLNs) require adaptations of Internet protocols to operate efficiently under constraints such as limited energy, low data rates, constrained memory, and duty-cycled radio operation. In many LLN deployments, the wireless interface is the dominant source of energy consumption. As a result, protocol design must minimize transmissions, idle listening, and the number of nodes involved in control-plane operations.
IPv6 Neighbor Discovery has been also adapted to the LLN environment in [RFC6775], later updated by [RFC8505], [RFC8929], and [RFC9010]. The target has been to design protocols that reduce energy consumption, especially in LLNs, though in general their design could be applied in any context targeting lowering carbon emissions. In particular, interface address assignment relies on address auto-configuration (such as Stateless Address Autoconfiguration (SLAAC)[RFC4862]), since the use of Dynamic Host Configuration Protocol (DHCP [RFC8415]) is not adapted, from an energy and bandwidth perspective, to LLN deployments. Indeed, LLN environments aim at avoiding as much as possible asynchronous multicast operations, because that would keep nodes awake and listening. Furthermore, it is also preferable to reduce as much as possible the number of nodes involved in control plane operations, because of energy and bandwidth constraints typical of LLN. DHCP can still be used in Internet-of-Things (IoT) deployments where energy and bandwidth are not an issue.	IPv6 Neighbor Discovery (ND) was optimized for LLNs in [RFC6775] and later extended by [RFC8505], [RFC8929], [RFC9010], and [RFC9685]. These specifications reduce multicast usage, limit control-plane participation, and introduce explicit address registration mechanisms to better support energy-constrained and duty-cycled devices.
To avoid multicast operations and to limit the number of nodes involved in address assignment in LLNs, mechanisms to register self-generated addresses have been designed ([RFC6775], [I-D.ietf-6lo-prefix-registration], [RFC8505], [RFC9685]).	In general IPv6 networks, address and prefix assignment are well supported by Stateless Address Auto-Configuration (SLAAC) [RFC4862] and DHCPv6 [RFC9915]. However, these mechanisms do not fully align with the architectural and operational goals of RFC8505-based 6LoWPAN deployments, particularly in scenarios requiring:
Recent use cases show, however, that there are some advantages in assigning addresses in an algorithmically managed way. In particular, in some scenarios, routing and forwarding can be simplified ([RFC9453], [I-D.ietf-6lo-path-aware-semantic-addressing], [SHENOY21], [BLESS22], [RIDOUX05]), hence reducing the power consumption and memory footprint. Algorithmic address assignment has its own pros and cons, as well as deployment requirements. However, they have the common benefit of being easily distributed. In other words, it is not necessary to have a centralized approach, like DHCP, rather address assignment is distributed by construction and a node can obtain an address from one of its neighbors who simply runs a distributed algorithm.	* Strict minimization of multicast traffic,
This situation highlights an existing gap that this document tries to fill: 6LoWPAN nodes have no means to directly request an address (or address prefix) from routers that are their direct neighbors. Currently, either auto-configuration is used, or DHCP has to be deployed. The former is energy efficient, but makes it hard to implement solutions like	* Avoidance of centralized infrastructure, * Localized control-plane interactions, * Algorithmically structured address assignment to support routing optimizations. This document does not attempt to replace SLAAC or DHCPv6 in general IPv6 networks. Instead, it addresses a specific gap in 6LoWPAN LLNs operating under RFC8505-based Neighbor Discovery optimizations, where nodes may need to request addresses or prefixes directly from neighboring routers without introducing DHCPv6 infrastructure.

Minor Changes:

- Re-organized Introduction

Main Changes -08: Explicit sub-section on DHCPv6

```
1.1.  Limitations of DHCPv6 in Constrained LLNs

   DHCPv6 relies on a client-server model and typically uses multicast
   (e.g., Solicit messages sent to FF02::1:2).  In IEEE 802.15.4
   [IEEE802154] and similar LLNs, IPv6 multicast is commonly mapped to
   link-layer broadcast.  Such broadcasts cause all nodes on the channel
   to wake up and process the frame.  In duty-cycled networks, this
   increases:

   *  Radio wake-ups,

   *  Idle listening time,

   *  Channel contention,

   *  Overall energy consumption.

   Furthermore, DHCPv6 requires a reachable server, often via relay
   agents, which may introduce multi-hop control paths and centralized
   state management.  In lossy multi-hop LLNs, longer control paths
   increase failure probability and recovery cost.

   While DHCPv6 can be efficient in traditional IPv6 networks, including
   support for long lifetimes and reduced renewal frequency, its
   architectural model does not align with the distributed, strictly
   localized control-plane design promoted by RFC8505-based 6LoWPAN
   Neighbor Discovery.
```

Discussed the inefficiencies of DHCPv6 in the context of constrained LLNs:

1. Drawbacks caused by the use of broadcasts.
2. Centralized architecture
3. Requirement of using relays leading to multi-hop control paths.

Main Changes -08: Algorithmic Address Assignment

```
[I-D.ietf-6lo-path-aware-semantic-addressing], [SHENOY21], [BLESS22],
and [RIDOUX05].  The latter, on the opposite, allows the use of
sophisticated assignment algorithms, but remains inefficient from an
energy and bandwidth consumption viewpoint.

This document proposes a new Neighbor Discovery Option, namely the
Generic Address Assignment Option (GAAO), in order for a node to
issue an address/prefix request to neighboring routers.  GAAO
complements the Extended Address Registration Option (EARO), defined
in [RFC8505], further extended in [I-D.ietf-6lo-prefix-registration]
and [RFC9685].

2.  Terminology
```

```
1.2.  Algorithmic and Distributed Address Assignment

   Recent work has demonstrated the benefits of algorithmically
   structured addressing in constrained networks (e.g., [RFC9453],
   [I-D.ietf-6lo-path-aware-semantic-addressing], [SHENOY21], [BLESS22],
   [RIDOUX05]).  Such approaches can simplify routing, reduce forwarding
   state, and improve scalability.  These schemes often require routers
   to assign addresses or prefixes according to a distributed Address
   Assignment Function (AAF).

   Existing mechanisms do not provide a standardized way for a 6LoWPAN
   Node (6LN) to explicitly request an address or prefix from a
   neighboring 6LoWPAN Router (6LR) within the optimized ND framework
   defined by [RFC8505].

   This document specifies a new Neighbor Discovery option, the Generic
   Address Assignment Option (GAAO), that enables a node to request an
   address or prefix directly from a neighboring router using ND
   messages.  The mechanism:

   *  Operates strictly at 1-hop,

   *  Avoids introducing DHCPv6 infrastructure,

   *  Aligns with RFC8505 registration procedures,

   *  Supports distributed algorithmic address assignment.

   GAAO complements the Extended Address Registration Option (EARO)
   defined in [RFC8505] and its extensions, integrating address/prefix
   assignment into the existing optimized ND framework for 6LoWPAN.
```

- Sub-section focusing on the motivation, requirements, and main benefits of GAAO:
- Generic support for algorithmic address assignment solutions
- 1-hop operation
- No multi-hop infrastructure
- Alignment with RFC 8505

Provided explanations to Lorenzo on the 6man mailing list

[IPv6]Re: [6lo] Re: Request for review of draft-ietf-6lo-nd-gaao prior to IETF Last Call

Adnan Rashid <adnanrashidpk@gmail.com> | Fri, 13 February 2026 12:37 UTC | Show header

Dear Lorenzo,

Thank you for your thoughtful and detailed review. Your question, whether a new address assignment mechanism is truly needed given the existence of SLAAC and DHCPv6, is both valid and important.

After considering your comments, we agree that the current draft did not sufficiently articulate the scope and motivation of the work. In particular, it did not clearly distinguish this mechanism from general IPv6 address assignment approaches.

To clarify, the intent of GAAO is *NOT *to replace SLAAC or DHCPv6 in general IPv6 deployments. Rather, it is scoped specifically to 6LoWPAN LLNs operating under RFC6775/RFC8505 Neighbor Discovery optimizations, where the architectural goals include:

- Strict minimization of multicast traffic
- Avoidance of centralized infrastructure
- Localized (1-hop) control-plane interactions
- Support for distributed, algorithmic address assignment

Regarding DHCPv6 efficiency, while DHCPv6 can indeed be efficient in traditional IPv6 networks (e.g., single-RTT exchanges with long lifetimes), its model relies on multicast discovery (FF02::1:2) and a client-server architecture. In IEEE 802.15.4 and similar LLNs, IPv6 multicast is typically mapped to link-layer broadcast, which causes all nodes on the channel to wake and process the frame. In duty-cycled networks, this increases radio wake-ups, idle listening time, and energy consumption.

Additionally, DHCPv6 typically requires a reachable server (possibly via relays), introducing multi-hop control paths and centralized state. In contrast, GAAO operates strictly at 1-hop within the existing ND framework

https://mailarchive.ietf.org/arch/msg/ipv6/9AYDF_XigB4sjqJu1RWYGNShINI/

Next Steps

- **Publication requested March 2026**
  - **Now waiting for Éric ;-)**

THANKS!