

Dynamic Host Configuration (DHC)  
Internet-Draft  
Intended status: Standards Track  
Expires: September 6, 2020

B. Volz  
Cisco  
T. Mrugalski  
ISC  
CJ. Bernardos  
UC3M  
March 5, 2020

**Link-Layer Addresses Assignment Mechanism for DHCPv6**  
**draft-ietf-dhc-mac-assign-04**

**Abstract**

In certain environments, e.g. large scale virtualization deployments, new devices are created in an automated manner. Such devices typically have their link-layer (MAC) addresses randomized. With sufficient scale, the likelihood of collision is not acceptable. Therefore an allocation mechanism is required. This draft proposes an extension to DHCPv6 that allows a scalable approach to link-layer address assignments.

**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 <https://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 September 6, 2020.

**Copyright Notice**

Copyright (c) 2020 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 (<https://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>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Requirements</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Terminology</a>	<a href="#">3</a>
<a href="#">4.</a>	<a href="#">Deployment scenarios and mechanism overview</a>	<a href="#">4</a>
<a href="#">4.1.</a>	<a href="#">Proxy client mode scenario</a>	<a href="#">4</a>
<a href="#">4.2.</a>	<a href="#">Direct client mode scenario</a>	<a href="#">5</a>
<a href="#">4.3.</a>	<a href="#">Mechanism Overview</a>	<a href="#">5</a>
<a href="#">5.</a>	<a href="#">Design Assumptions</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">Information Encoding</a>	<a href="#">8</a>
<a href="#">7.</a>	<a href="#">Requesting Addresses</a>	<a href="#">8</a>
<a href="#">8.</a>	<a href="#">Renewing Addresses</a>	<a href="#">9</a>
<a href="#">9.</a>	<a href="#">Releasing Addresses</a>	<a href="#">10</a>
<a href="#">10.</a>	<a href="#">Option Definitions</a>	<a href="#">10</a>
<a href="#">10.1.</a>	<a href="#">Identity Association for Link-Layer Addresses Option</a>	<a href="#">10</a>
<a href="#">10.2.</a>	<a href="#">Link-Layer Addresses Option</a>	<a href="#">12</a>
<a href="#">11.</a>	<a href="#">Selecting Link-Layer Addresses for Assignment to an IA_LL</a>	<a href="#">14</a>
<a href="#">12.</a>	<a href="#">IANA Considerations</a>	<a href="#">15</a>
<a href="#">13.</a>	<a href="#">Security Considerations</a>	<a href="#">15</a>
<a href="#">14.</a>	<a href="#">Privacy Considerations</a>	<a href="#">15</a>
<a href="#">15.</a>	<a href="#">Acknowledgements</a>	<a href="#">16</a>
<a href="#">16.</a>	<a href="#">References</a>	<a href="#">16</a>
<a href="#">16.1.</a>	<a href="#">Normative References</a>	<a href="#">16</a>
<a href="#">16.2.</a>	<a href="#">Informative References</a>	<a href="#">16</a>
<a href="#">Appendix A.</a>	<a href="#">IEEE 802c Summary</a>	<a href="#">17</a>
	<a href="#">Authors' Addresses</a>	<a href="#">19</a>

## [1. Introduction](#)

There are several new deployment types that deal with a large number of devices that need to be initialized. One of them is a scenario where virtual machines (VMs) are created on a massive scale.

Typically the new VM instances are assigned a random link-layer (MAC) address, but that does not scale well due to the birthday paradox.

Another use case is IoT devices. The huge number of such devices would likely exhaust a vendor's OUI (Organizationally Unique Identifier) global address space, and while there is typically no need to provide global uniqueness for such devices, a link-layer assignment mechanisms allows for conflicts to be avoided inside an administrative domain. For those reasons, it is desired to have some



form of local authority that would be able to assign locally unique MAC addresses.

This document proposes a new mechanism that extends DHCPv6 operation to handle link-layer address assignments.

Since DHCPv6 ([\[RFC8415\]](#)) is a protocol that can allocate various types of resources (non-temporary addresses, temporary addresses, prefixes, but also many options) and has necessary infrastructure (numerous server and client implementations, large deployed relay infrastructure, supportive solutions, like leasequery and failover) to maintain such assignment, it is a good candidate to address the desired functionality.

While this document presents a design that should be usable for any link-layer address type, some of the details are specific to Ethernet / IEEE 802 48-bit MAC addresses. Future documents may provide specifics for other link-layer address types.

The IEEE originally set aside half of the 48-bit MAC Address space for local use (where the U/L bit is set to 1). In 2017, the IEEE specified an optional specification (IEEE 802c) that divides this space into quadrants (Standards Assigned Identifier, Extended Local Identifier, Administratively Assigned Identifier, and a Reserved quadrant) - more details are in [Appendix A](#).

The IEEE is also working to specify protocols and procedures for assignment of locally unique addresses (IEEE 802.1cq). This work may serve as one such protocol for assignment. For additional background, see [\[IEEE-802-Tutorial\]](#).

## 2. Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [\[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown here.

## 3. Terminology

The DHCPv6 terminology relevant to this specification from the DHCPv6 Protocol [\[RFC8415\]](#) applies here.

client	A device that is interested in obtaining link-layer addresses. It implements the basic DHCPv6 mechanisms needed by a DHCPv6 client as described in <a href="#">[RFC8415]</a> and supports the new options (IA_LL and LLADDR) specified
--------	--



in this document. The client may or may not support address assignment and prefix delegation as specified in [\[RFC8415\]](#).

**server** Software that manages link-layer address allocation and is able to respond to client queries. It implements basic DHCPv6 server functionality as described in [\[RFC8415\]](#) and supports the new options (IA\_LL and LLADDR) specified in this document. The server may or may not support address assignment and prefix delegation as specified in [\[RFC8415\]](#).

**address** Unless specified otherwise, an address means a link-layer (or MAC) address, as defined in IEEE802. The address is typically 6 octets long, but some network architectures may use different lengths.

**address block** A number of consecutive link-layer addresses. An address block is expressed as a first address plus a number that designates the number of additional (extra) addresses. A single address can be represented by the address itself and zero extra addresses.

#### **[4.](#) Deployment scenarios and mechanism overview**

This mechanism is designed to be generic and usable in many deployments, but there are two scenarios it attempts to address in particular: (i) proxy client mode, and (ii) direct client mode.

##### **[4.1.](#) Proxy client mode scenario**

This mode is used when an entity acts as a DHCP client and requests available DHCP servers to assign one or more MAC addresses (an address block), to be then assigned for use to the final end-devices. One relevant example of scenario of application of this mode is large scale virtualization. In such environments the governing entity is often called a hypervisor and is frequently required to spawn new VMs. The hypervisor needs to assign new MAC addresses to those machines. The hypervisor does not use those addresses for itself, but rather uses them to create new VMs with appropriate MAC addresses. It is worth pointing out the cumulative nature of this scenario. Over time, the hypervisor is likely to increase its MAC addresses usage. While some obsolete VMs will be deleted and their MAC addresses will become eligible for release or reuse, it is unexpected for all MAC addresses to be released.



#### **4.2. Direct client mode scenario**

This mode is used when an entity acts as a DHCP client and requests available DHCP servers to assign one or more MAC addresses (an address block) for its own use. This usage scenario is related to IoT (Internet of Things). With the emergence of IoT, a new class of cheap, sometimes short lived and disposable devices, has emerged. Examples may include various sensors (e.g. medical) and actuators or controllable LED lights. Upon first boot, the device uses a temporary MAC address, as described in [IEEE-802.11-02-109r0], to send initial DHCP packets to available DHCP servers. Then, such devices would typically request a single MAC address for each available network interface, which typically means one MAC address per device. Once the server assigns a MAC address, the device abandons its temporary MAC address and uses the assigned (leased) MAC address.

Note that a client that operates as above that does not have a globally unique link-layer address on any of its interfaces MUST NOT use a DUID-LLT or DUID-LL (link-layer based DUID). For more details, refer to [Section 11 of \[RFC8415\]](#).

#### **4.3. Mechanism Overview**

In all scenarios the protocol operates in fundamentally the same way. The device requesting an address, acting as a DHCP client, will send a Solicit message with a IA\_LL option to all available DHCP servers. That IA\_LL option MUST include a LLADDR option specifying the link-layer-type and link-layer-len and may specify a specific address or address block as a hint for the server. Each available server responds with an Advertise message with offered link-layer address or addresses. The client then picks the best server, as governed by [RFC8415], and will send a Request message. The target server will then assign the link-layer addresses and send a Reply message. Upon reception, the client can start using those link-layer addresses.

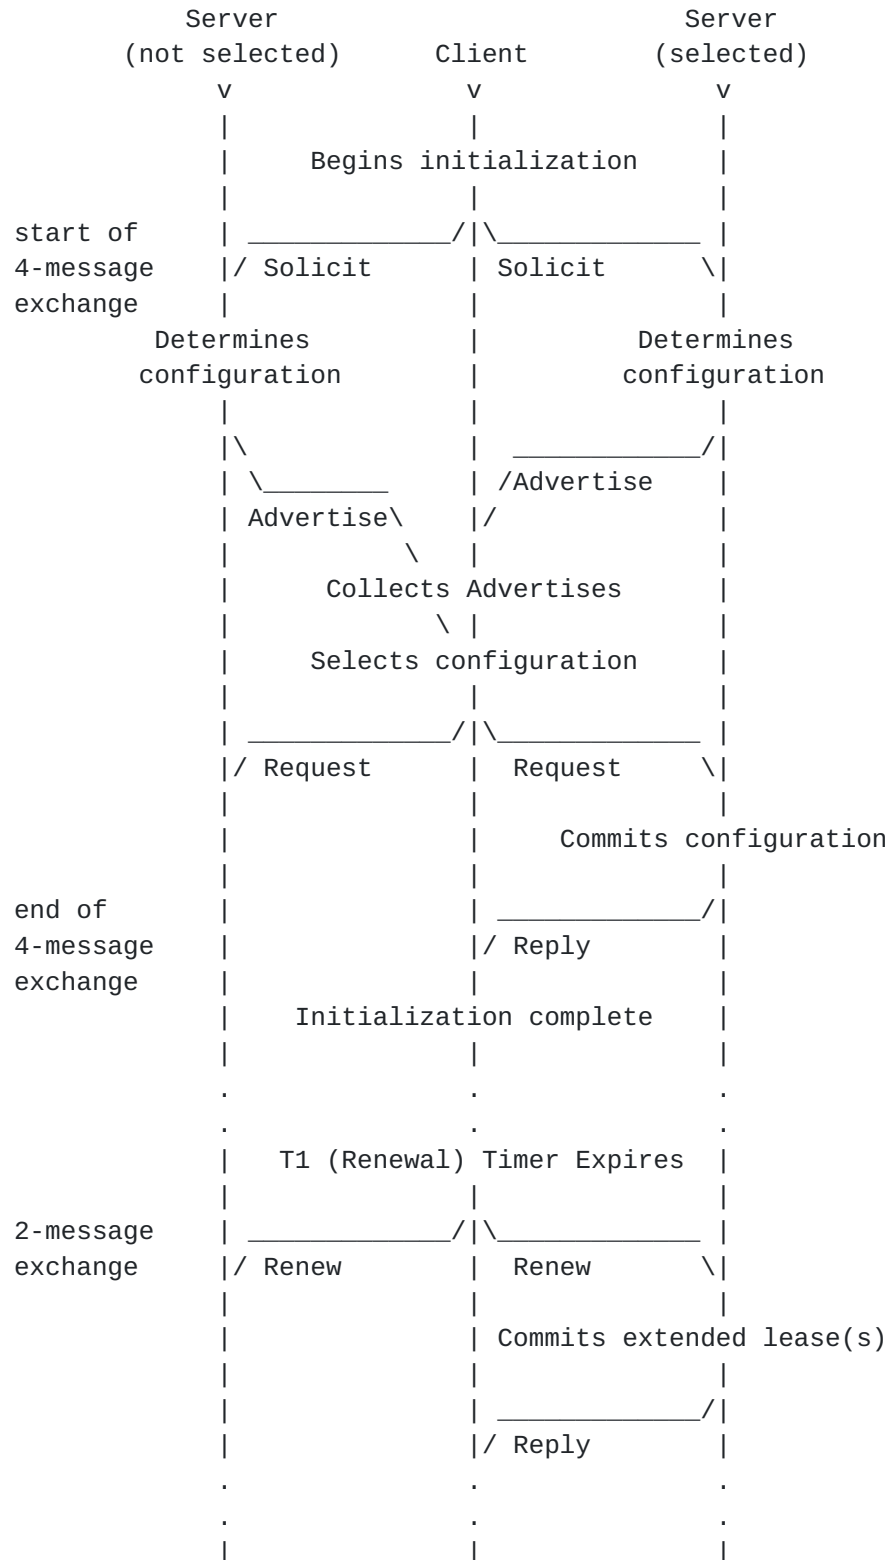
Normal DHCP mechanisms are in use. The client is expected to periodically renew the link-layer addresses as governed by T1 and T2 timers. This mechanism can be administratively disabled by the server sending "infinity" as the T1 and T2 values (see [Section 7.7 of \[RFC8415\]](#)).

The client can release link-layer addresses when they are no longer needed by sending a Release message (see [Section 18.2.7 of \[RFC8415\]](#)).





Figure 1, taken from [RFC8415], shows a timeline diagram of the messages exchanged between a client and two servers for the typical lifecycle of one or more leases








Figure 1: Timeline diagram of the messages exchanged between a client and two servers for the typical lifecycle of one or more leases

Confirm, Decline, and Information-Request messages are not used in link-layer address assignment.

Clients implementing this mechanism SHOULD use the Rapid Commit option as specified in [Section 5.1](#) and 18.2.1 of [\[RFC8415\]](#).

An administrator may make the address assignment permanent by specifying use of infinite lifetimes, as defined in [Section 7.7 of \[RFC8415\]](#). An administrator may also disable the need for the renewal mechanism by setting the T1 and T2 values to infinity.

Devices supporting this proposal MAY support the reconfigure mechanism, as defined in [Section 18.2.11 of \[RFC8415\]](#). If supported by both server and client, this mechanism allows the administrator to immediately notify clients that the configuration has changed and triggers retrieval of relevant changes immediately, rather than after the T1 timer elapses. Since this mechanism requires implementation of Reconfigure Key Authentication Protocol (See [Section 20.4 of \[RFC8415\]](#)), small footprint devices may choose to not support it.

DISCUSSION: A device may send its link-layer address in a LLADDR option to ask the server to register that address to the client (if available), making the assignment permanent for the lease duration. The client MUST be prepared to use a different address if the server chooses not to honor its hint.

## 5. Design Assumptions

One of the essential aspects of this mechanism is its cumulative nature, especially in the hypervisor scenario. The server-client relationship does not look like other DHCP transactions. This is especially true in the hypervisor scenario. In a typical environment, there would be one server and a rather small number of



hypervisors, possibly even only one. However, over time the number of MAC addresses requested by the hypervisor(s) will likely increase as new VMs are spawned.

Another aspect crucial for efficient design is the observation that a single client acting as hypervisor will likely use thousands of addresses. Therefore an approach similar to what is used for address or prefix assignment (IA container with all assigned addresses listed, one option for each address) would not work well. Therefore the mechanism should operate on address blocks, rather than single values. A single address can be treated as an address block with just one address.

The DHCPv6 mechanisms are reused to large degree, including message and option formats, transmission mechanisms, relay infrastructure and others. However, a device wishing to support only link-layer address assignment is not required to support full DHCPv6. In other words, the device may support only assignment of link-layer addresses, but not IPv6 addresses or prefixes.

## **6. Information Encoding**

A client MUST send a LLADDR option encapsulated in an IA\_LL option to specify the link-layer-type and link-layer-len values. For link-layer-type 1 (Ethernet / IEEE 802 48-bit MAC addresses), a client sets the link-layer-address field to:

1. 00:00:00:00:00:00 (all zeroes) if the client has no hint as to the starting address of the unicast address block. This address has the IEEE 802 individual/group bit set to 0 (individual).
2. Any other value to request a specific block of address starting with the specified address

A client sets the extra-addresses field to either 0 for a single address or to the size of the requested address block minus 1.

A client SHOULD set the valid-lifetime field to 0 (this field MUST be ignored by the server).

## **7. Requesting Addresses**

The link-layer addresses are assigned in blocks. The smallest block is a single address. To request an assignment, the client sends a Solicit message with an IA\_LL option in the message. The IA\_LL option MUST contain a LLADDR option as specified in [Section 6](#).



The server, upon receiving an IA\_LL option, inspects its content and may offer an address or addresses for each LLADDR option according to its policy. The server MAY take into consideration the address block requested by the client in the LLADDR option. However, the server MAY chose to ignore some or all parameters of the requested address block. In particular, the server may send a different starting address than requested, or grant a smaller number of addresses than requested. The server sends back an Advertise message an IA\_LL option containing an LLADDR option that specifies the addresses being offered. If the server is unable to provide any addresses it MUST return the IA\_LL option containing a Status Code option (see [Section 21.13 of \[RFC8415\]](#)) with status set to NoAddrsAvail.

The client MUST be able to handle an Advertise message containing a smaller number of addresses, or an address or addresses different than those requested.

The client waits for available servers to send Advertise responses and picks one server as defined in [Section 18.2.9 of \[RFC8415\]](#). The client then sends a Request message that includes the IA\_LL container option with the LLADDR option copied from the Advertise message sent by the chosen server.

Upon reception of a Request message with IA\_LL container option, the server assigns requested addresses. The server allocates block of addresses according to its configured policy. The server MAY assign a different block than requested in the Request message. It then generates and sends a Reply message back to the client.

Upon receiving a Reply message, the client parses the IA\_LL container option and may start using all provided addresses. It MUST restart its T1 and T2 timers using the values specified in the IA\_LL option.

The client MUST be able to handle a Reply message containing a smaller number of addresses, or an address or addresses different than those requested.

A client that has included a Rapid Commit option in the Solicit, may receive a Reply in response to the Solicit and skip the Advertise and Request steps above (see [Section 18.2.1 of \[RFC8415\]](#)).

## **8. Renewing Addresses**

Address renewals follow the normal DHCPv6 renewals processing described in [Section 18.2.4 of \[RFC8415\]](#). Once the T1 timer elapses, the client starts sending Renew messages with the IA\_LL option containing a LLADDR option for the address block being renewed. The server responds with a Reply message that contains the renewed





address block. The server SHOULD NOT alter the address block being renewed, unless its policy has changed. The server MUST NOT shrink or expand the address block - once a block is assigned and has a non-zero valid lifetime, its size, starting address, and ending address MUST NOT change.

If the requesting client needs additional MAC addresses -- e.g., in the hypervisor scenario because addresses need to be assigned to new VMs -- the simpler approach is for the requesting device to keep the address blocks as atomic once "leased". Therefore, if a client wants more addresses at a later stage, it SHOULD send an IA\_LL option with a different IAID to create another "container" for more addresses.

If the client is unable to Renew before the T2 timer elapses, it starts sending Rebind messages as described in 18.2.5 of [\[RFC8415\]](#).

## **9.    Releasing Addresses**

The client may decide to release a leased address block. A client MUST release the whole block in its entirety. A client releases an address block by sending a Release message that includes the IA\_LL option containing the LLADDR option for the address block to release. The Release transmission mechanism is described in [Section 18.2.7 of \[RFC8415\]](#).

## **10.   Option Definitions**

This mechanism uses an approach similar to the existing mechanisms in DHCP. There is one container option (the IA\_LL option) that contains the actual link-layer address or addresses, represented by an LLADDR option. Each such option represents an address block, which is expressed as a first address with a number that specifies how many additional addresses are included.

### **10.1.   Identity Association for Link-Layer Addresses Option**

The Identity Association for Link-Layer Addresses option (IA\_LL option) is used to carry one or more IA\_LL options, the parameters associated with the IA\_LL, and the address blocks associated with the IA\_LL.

The format of the IA\_LL option is:



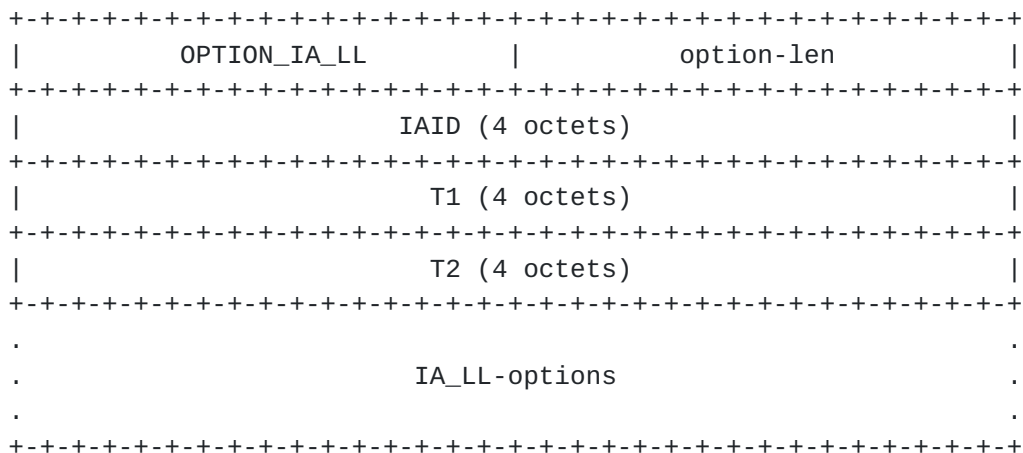


Figure 2: IA\_LL Option Format

option-code	OPTION_IA_LL (tbd1).
option-len	12 + length of IA_LL-options field.
IAID	The unique identifier for this IA_LL; the IAID must be unique among the identifiers for all of this client's IA_LLs. The number space for IA_LL IAIDs is separate from the number space for other IA option types (i.e., IA_NA, IA_TA, and IA_PD). A 4-octet field containing an unsigned integer.
T1	The time at which the client should contact the server from which the addresses in the IA_LL were obtained to extend the valid lifetime of the addresses assigned to the IA_LL; T1 is a time duration relative to the current time expressed in units of seconds. A 4-octet field containing an unsigned integer.
T2	The time at which the client should contact any available server to extend the valid lifetime of the addresses assigned to the IA_LL; T2 is a time duration relative to the current time expressed in units of seconds. A 4-octet field containing an unsigned integer.
IA_LL-options	Options associated with this IA_LL. A variable length field (12 octets less than the value in the option-len field).



An IA\_LL option may only appear in the options area of a DHCP message. A DHCP message may contain multiple IA\_LL options (though each must have a unique IAID).

The status of any operations involving this IA\_LL is indicated in a Status Code option (see [Section 21.13 of \[RFC8415\]](#)) in the IA\_LL-options field.

Note that an IA\_LL has no explicit "lifetime" or "lease length" of its own. When the valid lifetimes of all of the addresses in an IA\_LL have expired, the IA\_LL can be considered as having expired. T1 and T2 are included to give servers explicit control over when a client recontacts the server about a specific IA\_LL.

In a message sent by a client to a server, the T1 and T2 fields SHOULD be set to 0. The server MUST ignore any values in these fields in messages received from a client.

In a message sent by a server to a client, the client MUST use the values in the T1 and T2 fields for the T1 and T2 times, unless those values in those fields are 0. The values in the T1 and T2 fields are the number of seconds until T1 and T2.

As per [Section 7.7 of \[RFC8415\]](#)), the value 0xffffffff is taken to mean "infinity" and should be used carefully.

The server selects the T1 and T2 times to allow the client to extend the lifetimes of any address block in the IA\_LL before the lifetimes expire, even if the server is unavailable for some short period of time. Recommended values for T1 and T2 are .5 and .8 times the shortest valid lifetime of the address blocks in the IA that the server is willing to extend, respectively. If the "shortest" valid lifetime is 0xffffffff ("infinity"), the recommended T1 and T2 values are also 0xffffffff. If the time at which the addresses in an IA\_LL are to be renewed is to be left to the discretion of the client, the server sets T1 and T2 to 0. The client MUST follow the rules defined in [Section 14.2 in \[RFC8415\]](#).

If a client receives an IA\_LL with T1 greater than T2, and both T1 and T2 are greater than 0, the client discards the IA\_LL option and processes the remainder of the message as though the server had not included the invalid IA\_LL option.

## **[10.2. Link-Layer Addresses Option](#)**

The Link-Layer Addresses option is used to specify an address block associated with a IA\_LL. The option must be encapsulated in the



IA\_LL-options field of an IA\_LL option. The LLaddr-options fields encapsulates those options that are specific to this address block.

The format of the Link-Layer Addresses option is:

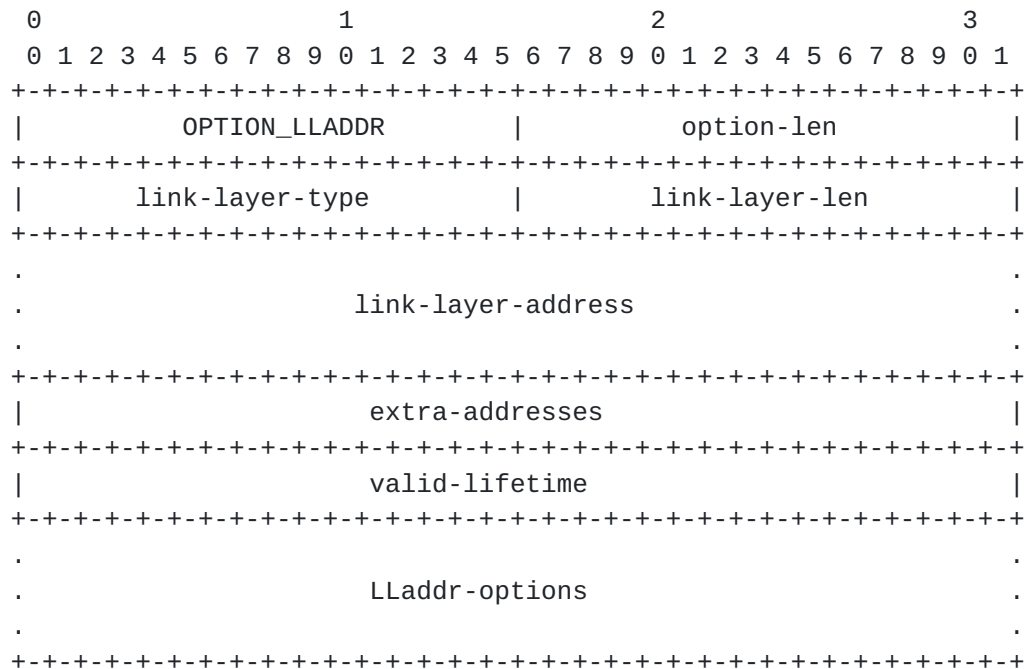


Figure 3: LLADDR Option Format

option-code      OPTION\_LLADDR (tbd2).

option-len      12 + link-layer-len field (typically 6) + length of LLaddr-options field. Assuming a typical link-layer address of 6 is used and there are no extra options, length should be equal to 18.

link-layer-type      The link-layer type MUST be a valid hardware type assigned by the IANA, as described in [RFC5494] and in the "Hardware Types" table at <https://www.iana.org/assignments/arp-parameters>. A 2-octet field containing an unsigned integer.

link-layer-len      Specifies the length of the link-layer-address field (typically 6, for a link-layer-type of 1 (Ethernet)). A 2-octet field containing an unsigned integer.

link-layer-address      Specifies the link-layer address that is being requested or renewed, or a special value to request any address. For a link-layer type of 1 (Ethernet / IEEE 802 48-bit MAC addresses), see [Section 6](#) for





details on these values. In responses from a server, this value specifies the first address allocated.

- extra-addresses Number of additional addresses that follow the address specified in link-layer-address. For a single address, 0 is used. For example: link-layer-address: 02:04:06:08:0a and extra-addresses 3 designates a block of 4 addresses, starting from 02:04:06:08:0a (inclusive) and ending with 02:04:06:08:0d (inclusive). In responses from a server, this value specifies the number of additional addresses allocated. A 4-octet field containing an unsigned integer.
- valid-lifetime The valid lifetime for the address(es) in the option, expressed in units of seconds. A 4-octet field containing an unsigned integer.
- LLaddr-options Any encapsulated options that are specific to this particular address block. Currently there are no such options defined, but there may be in the future.

In a message sent by a client to a server, the valid lifetime field SHOULD be set to 0. The server MUST ignore any received value.

In a message sent by a server to a client, the client MUST use the value in the valid lifetime field for the valid lifetime for the address block. The value in the valid lifetime field is the number of seconds remaining in the lifetime.

As per [Section 7.7 of \[RFC8415\]](#), the valid lifetime of 0xffffffff is taken to mean "infinity" and should be used carefully.

More than one LLADDR option can appear in an IA\_LL option.

## **11. Selecting Link-Layer Addresses for Assignment to an IA\_LL**

A server selects link-layer addresses to be assigned to an IA\_LL according to the assignment policies determined by the server administrator.

Link-layer addresses are typically specific to a link and the server SHOULD follow the steps in [Section 13.1 of \[RFC8415\]](#) to determine the client's link.

For Ethernet / IEEE 802 MAC addresses, a server MAY use additional options supplied by a relay agent or client to select the quadrant (see [Appendix A](#)) from which addresses are to be assigned. This MAY



include new options, such as those specified in [\[I-D.ietf-dhc-slap-quadrant\]](#).

## **12. IANA Considerations**

IANA is requested to assign the OPTION\_IA\_LL (tbd1) option code from the DHCPv6 "Option Codes" registry maintained at <http://www.iana.org/assignments/dhcpv6-parameters> and use the following data when adding the option to the registry:

Value:	tbd1
Description:	OPTION_IA_LL
Client ORO:	No
Singleton Option:	No
Reference:	this document

IANA is requested to assign the OPTION\_LLADDR (tbd2) option code from the DHCPv6 "Option Codes" registry maintained at <http://www.iana.org/assignments/dhcpv6-parameters> and use the following data when adding the option to the registry:

Value:	tbd2
Description:	OPTION_LLADDR
Client ORO:	No
Singleton Option:	No
Reference:	this document

## **13. Security Considerations**

See [\[RFC8415\]](#) for the DHCPv6 security considerations. See [\[RFC8200\]](#) for the IPv6 security considerations.

There is a possibility of the same link-layer address being used by more than one device if not all parties on a link use this mechanism to obtain a link-layer address from the space assigned to the DHCP server. It is also possible that a bad actor purposely uses a device's link-layer address.

## **14. Privacy Considerations**

See [\[RFC8415\]](#) for the DHCPv6 privacy considerations.

For a client requesting a link-layer address directly from a server, as the link-layer address assigned to a client will likely be used by



the client to communicate on the link, the address will be exposed to those able to listen in on this communication. For those peers on the link that are able to listen in on the DHCPv6 exchange, they would also be able to correlate the client's identity (based on the DUID used) with the assigned address. Additional mechanisms, such as the ones described in [[RFC7844](#)] can also be used.

## **15. Acknowledgements**

Thanks to the DHC Working Group participants that reviewed this document, provided comments, and support. With special thanks to Ian Farrer for his thorough reviews and shepherding of this document through the IETF process.

## **16. References**

### **16.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 8415](#), DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

### **16.2. Informative References**

- [I-D.ietf-dhc-slap-quadrant] Bernardos, C. and A. Mourad, "SLAP quadrant selection options for DHCPv6", [draft-ietf-dhc-slap-quadrant-04](#) (work in progress), February 2020.
- [IEEE-802-Tutorial] Thaler, P., "Emerging IEEE 802 Work on MAC Addressing", <<https://datatracker.ietf.org/meeting/96/materials/slides-96-edu-ieee802work-0/>>.



[IEEE-802.11-02-109r0]

Edney, J., Haverinen, H., Honkanen, J-P., and P. Orava,  
"Temporary MAC address for anonymity",  
<<https://mentor.ieee.org/802.11/dcn/02/11-02-0109-00-000i-temporary-mac-address-for-anonymity.ppt>>.

[IEEEStd802c-2017]

IEEE Computer Society, "IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture, Amendment 2: Local Medium Access Control (MAC) Address Usage, IEEE Std 802c-2017".

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", [RFC 2464](#), DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.

[RFC5494] Arkko, J. and C. Pignataro, "IANA Allocation Guidelines for the Address Resolution Protocol (ARP)", [RFC 5494](#), DOI 10.17487/RFC5494, April 2009, <<https://www.rfc-editor.org/info/rfc5494>>.

[RFC7844] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity Profiles for DHCP Clients", [RFC 7844](#), DOI 10.17487/RFC7844, May 2016, <<https://www.rfc-editor.org/info/rfc7844>>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

## **Appendix A. IEEE 802c Summary**

This appendix provides a brief summary of IEEE802c from [[IEEEStd802c-2017](#)].

The original IEEE 802 specifications assigned half of the 48-bit MAC address space to local use -- these addresses have the U/L bit set to 1 and are locally administered with no imposed structure.

In 2017, the IEEE issued the 802c specification which defines a new "optional Structured Local Address Plan (SLAP) that specifies different assignment approaches in four specified regions of the local MAC address space." Under this plan, there are 4 SLAP quadrants that use different assignment policies.





The first octet of the MAC address Z and Y bits define the quadrant for locally assigned addresses (X-bit is 1). In IEEE representation, these bits are as follows:




Figure 4: SLAP Bits

The SLAP quadrants are:

Quadrant	Y-bit	Z-bit	Local Identifier Type	Local Identifier
01	0	1	Extended Local	ELI
11	1	1	Standard Assigned	SAI
00	0	0	Administratively Assigned	AAI
10	1	0	Reserved	Reserved

SLAP Quadrants

Extended Local Identifier (ELI) derived MAC addresses are based on an assigned Company ID (CID), which is 24-bits (including the M, X, Y, and Z bits) for 48-bit MAC addresses. This leaves 24-bits for the locally assigned address for each CID for unicast (M-bit = 0) and also for multicast (M-bit = 1). The CID is assigned by the IEEE RA.

Standard Assigned Identifier (SAI) derived MAC addresses are assigned by a protocol specified in an IEEE 802 standard. For 48-bit MAC addresses, 44 bits are available. Multiple protocols for assigning SAIs may be specified in IEEE standards. Coexistence of multiple protocols may be supported by limiting the subspace available for assignment by each protocol.

Administratively Assigned Identifier (AAI) derived MAC addresses are assigned locally. Administrators manage the space as needed. Note that multicast IPv6 packets ([[RFC2464](#)]) use a destination address starting in 33-33 and this falls within this space and therefore should not be used to avoid conflict with IPv6 multicast addresses. For 48-bit MAC addresses, 44 bits are available.



The last quadrant is reserved for future use. While this quadrant may also be used for AAI space, administrators should be aware that future specifications may define alternate uses that could be incompatible.

#### Authors' Addresses

Bernie Volz  
Cisco Systems, Inc.  
1414 Massachusetts Ave  
Boxborough, MA 01719  
USA

Email: volz@cisco.com

Tomek Mrugalski  
Internet Systems Consortium, Inc.  
950 Charter Street  
Redwood City, CA 94063  
USA

Email: tomasz.mrugalski@gmail.com

Carlos J. Bernardos  
Universidad Carlos III de Madrid  
Av. Universidad, 30  
Leganes, Madrid 28911  
Spain

Phone: +34 91624 6236

Email: cjbc@it.uc3m.es

URI: <http://www.it.uc3m.es/cjbc/>

