

Network Working Group
Internet-Draft
Updates: [3315](#) (if approved)
Intended status: Standards Track
Expires: April 25, 2014

T. Lemon
Nominum, Inc.
October 22, 2013

Techniques for Tracking Inventory Using DHCPv6 DUID
draft-lemon-dhc-inventory-02

Abstract

In the years since DHCPv4 gained widespread popularity, one of the uses to which organizations have put it is inventory tracking: associating identifiers scanned from packaging with records in an inventory database. This document describes various means for accomplishing the same purpose using DHCPv6. This document also updates [RFC3315](#) by clarifying the meaning of some normative language regarding the DUID-LL and DUID-LLT DUID types.

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 25, 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

Internet-Draft

Tracking Inventory with DHCPv6 DUID

October 2013

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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	2
1.1.	Terminology	3
1.2.	Scope of applicability	3
2.	General Mechanism	4
2.1.	Entering a new device into inventory	4
2.2.	Distributing the new device	5
2.3.	First appearance on the network	5
3.	Using DUID-LL or DUID-LLT	7
4.	Using the DHCPv6 Client Link-layer Address Option	8
5.	Which algorithm to use	8
6.	New labeling requirements	8
7.	Acknowledgments	9
8.	Security Considerations	9
9.	IANA Considerations	10
10.	References	10
10.1.	Normative References	10
10.2.	Informative References	10
	Author's Address	10

[1.](#) Introduction

In DHCPv4 [RFC2131], the link-layer address of a DHCP client is commonly used to identify the client. The link-layer address appears in the chaddr field of the DHCP packet, and is also frequently included in the Client Identifier option.

Not coincidentally, the link-layer addresses of network devices are almost always present as bar codes and machine-readable text on the outside of the boxes in which these devices are delivered. This is true of most mobile phones, laptop computers, desktop computers, network routers and switches, and so on.

Services providers and enterprises have taken advantage of these two facts in their inventory tracking systems: when a new device arrives, the bar codes are scanned into a database, and an inventory tracking number is assigned to the device. When the device is assigned to a user or to a use, that information can be added to the database.

This means that, for example, when a network router is installed, the inventory tracking system can be updated both with the physical location of the router and with its intended purpose: for example, "router between backbone and first floor network." This information can in turn be used to provision the router: to send it a configuration. When a router is replaced, the provisioning system can then automatically configure the new router simply by knowing that it is the "router between backbone and first floor network"; DHCPv4 takes care of noticing that the device is new on the network, and that can trigger the provisioning of the device.

Unlike DHCPv4, DHCPv6 [[RFC3315](#)] does not directly make use of a device's link-layer address as an identifier. This is because the link-layer address is specific to an interface, and it was considered useful to be able to notice that requests being issued on multiple interfaces related to the same device. It was also considered useful that the device's identifier remain stable when network hardware was added or removed.

Consequently, the inventory management solution in DHCPv6 is somewhat more complicated than that in DHCPv4. This document describes several mechanisms that are available to administrators to address this concern.

[1.1.](#) 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](#)].

[1.2.](#) Scope of applicability

This document is intended to provide guidance to implementors of DHCP servers, network device inventory management systems and provisioning systems as to how to connect information sent over the network by devices with the physical devices that sent the information.

Lemon

Expires April 25, 2014

[Page 3]

Internet-Draft

Tracking Inventory with DHCPv6 DUID

October 2013

Nothing in this document should be construed as a requirement for such systems; rather, it is intended as helpful advice. The normative language in the document applies to implementations that attempt to follow the advice given in this document, and is not intended to apply to systems that solve the same problem in different ways.

[2.](#) General Mechanism

This mechanism takes as its input two pieces of information: one of the link-layer addresses of a device, and the DUID of the device. If the DUID is known, the link-layer address MUST be ignored. If the DUID is unknown, the link-layer address is used to find the the inventory record for the device, and then the the DUID is added to the inventory record.

[2.1.](#) Entering a new device into inventory

We assume that when a new device arrives, the box has one bar code on the side for the link layer address of each network interface on the device. The person responsible for receiving the device scans each bar code off of the box. This person then generates an inventory control tag for the device, and scans that into the system as well. The inventory control tag is affixed to the device in a location where it can be easily scanned or read in the future.

Suppose a new router arrives. It has two network interfaces: one with a link-layer address of 00:53:01:1f:24:32 and one with a link-layer address of 00:53:02:05:49:ad. The device is assigned an

inventory control tag number of 11029938. This will produce several rows in a database table listing link-layer addresses:

link-layer address	inventory number
00:53:01:1f:24:32	11029938
00:53:02:05:49:ad	11029938

Table 1: Link Layer address to Inventory Table

It will also produce one additional row in a database table listing inventory items:

inventory number	description	DUID	user?	use
11029938	router	NULL	false	NULL

inventory number	description	DUID	user?	use
------------------	-------------	------	-------	-----

Table 2: Inventory Items Table

Note that the DUID and use fields are NULL at this point, because the device hasn't yet been assigned a user, and has never been connected to the network. The user field in this example is a flag indicating whether the device will be assigned to the user (true), or is infrastructure equipment (false).

2.2. Distributing the new device

Eventually the new device is moved from inventory to its intended use: either on a machine room rack somewhere, or to a user's desk, for example. When this happens, the inventory record is updated; the link layer address records are not:

inventory number	description	DUID	user?	use
11029938	router	NULL	false	bb::first floor

Table 3: Inventory Items Table (distributed)

In this example the router is marked with a token that will be meaningful to the provisioning system: "backbone::first floor".

[2.3.](#) First appearance on the network

Now the device is plugged into a rack in a distribution closet; one network interface is plugged into the backbone network; the other is plugged into the cascade of switches that support the first floor network. The device is powered on.

When the device is powered on, it first does router solicitation and gets a prefix on the backbone network (we assume that there is no router on the first floor network, so it doesn't get a prefix there). The prefix is marked with the 'M' bit, indicating that this is a managed network, so the router issues a DHCP Solicit message.

The solicit messages is received by the DHCP server. The DHCP server does not have a record of the DUID presented by the router, so it logs it as unknown in the provisioning system and does nothing further. The DHCP server provides both the link-layer address of the router's interface on the backbone network, and the DUID that the router presented.

If the DHCP server and router are connected to the same physical link, the DHCP server can acquire the router's link-layer address from the link-layer framing of the DHCP Solicit message. Otherwise, the server must obtain the link-layer address using the techniques described in [Sections 3](#) or, if possible, 4.

In this example we'll assume that the link-layer address the router sent is 00:53:01:1f:24:32, and that the DUID is 00:03:00:01:00:53:02:05:49:ad.

The provisioning system takes the log entry for the unknown device and does a lookup in the Inventory Items table for the DUID that the router presented. The DUID is not in the table, so the provisioning system gets an empty result table, indicating that the DUID is currently unknown to the provisioning system.

The provisioning system then looks up the link-layer address in the Link Layer Address to Inventory table. This produces a result table with a single row:

link-layer address	inventory number
00:53:01:1f:24:32	11029938

Table 4: Link Layer address to Inventory Result Table

The provisioning system now uses the inventory number to find the inventory table entry and update it with the DUID; after this is done, the record looks like this:

inventory number	description	DUID	user?	use
11029938	router	00:03:00:01 00:53:02:05:49:ad	false	bb::first floor

Table 5: Inventory Items Table (finished)

The provisioning system now has enough information to configure the DHCP server with an IP address specific to the router, and to configure the router itself with information about prefixes on the first floor network. How this is done is beyond the scope of this document.

3. Using DUID-LL or DUID-LLT

[RFC3315](#) defines the DHCP Unique Identifier (DUID) and describes several different formats suited to various uses. Two of those formats, DUID-LL and DUID-LLT, include the link-layer address of the client. [RFC3315](#) states:

Clients and servers MUST treat DUIDs as opaque values and MUST

only compare DUIDs for equality. Clients and servers MUST NOT in any other way interpret DUIDs. Clients and servers MUST NOT restrict DUIDs to the types defined in this document, as additional DUID types may be defined in the future.

This text is specifically intended to exclude the possibility that the DHCP server might treat some portion of the DUID, rather than the entire DUID, as a unique identifier for the client. However, the text is stated so unequivocally that it is often interpreted to mean that it's not permissible to look at the contents of the option for any other reason; this was not the original intent of the requirement.

We therefore update the above paragraph from [RFC3315](#) as follows:

Clients and servers MUST NOT use any part of a DUID as a unique identifier. Clients and servers MUST use the entire contents of the DUID as an opaque token for the purpose of uniquely identifying the client. Clients and servers MUST NOT restrict DUIDs to the types defined in this document, as additional DUID types may be defined in the future. Clients and servers MAY use the semantic contents of the DUID to generate a one-time mapping between a link-layer address known to be configured in a specific device, and that device's DUID.

This change to [RFC3315](#) allows DHCP servers or provisioning systems to use the link-layer address from a DUID-LL or DUID-LLT as input to the process described in [Section 2](#) for mapping the DUID to a specific device in an inventory database.

It is important to note that the usual reason for using a DUID-LLT, as opposed to a DUID, is that the network interface used to generate the DUID-LLT is not permanently installed in the device. This means that there is no assurance that a device that came with a removable network interface will not have a new interface installed when it generates its DUID. In that case, the device will present an unknown link-layer address to the DHCP server in the DUID-LLT.

For this reason, nodes that contain both removable and fixed interfaces MUST use the link-layer address of a fixed interface when

address of a fixed interface to generate the DUID SHOULD use DUID-LL, not DUID-LLT, since there is no benefit to the additional timestamp in DUID-LLT.

4. Using the DHCPv6 Client Link-layer Address Option

The DHCPv6 Client Link-layer Address option [[RFC6939](#)] is a new DHCPv6 extension which allows the DHCPv6 relay agent to include the client's link-layer address as an option in the RELAY-FORW message. The DHCP server can use the provided link-layer address as a key in the lookup described in [Section 2](#). Note that the link-layer address will come from the RELAY-FORW message, but the DUID to be mapped will come from the inner encapsulated packet—for example, a DHCP Solicit or other client-sourced packet.

5. Which algorithm to use

Since the use of DUID-LL and DUID-LLT is not required, it is best not to rely on these DUIDs as a source for the client's link-layer address. If the client is connected to the same link as the server, the server SHOULD use the link-layer address presented by the client for the inventory table lookup. If the client is configured through a relay, and the relay provides the Client Link-layer Address option, the server SHOULD use the contents of that option to identify the client.

6. New labeling requirements

The extra work involved in matching link-layer addresses to DUIDs is only necessary because network equipment boxes typically only have bar code labels for link-layer addresses and not for DUIDs. It would greatly simplify inventory management for dual-stack and IPv6-only sites if these boxes were additionally labeled with DUIDs.

However, this is not as simple as it sounds. The problem is the DUID, like a link-layer address, is simply a sequence of octets. A bar code containing a DUID would therefore be difficult to reliably distinguish from a link-layer address. It might be possible for the person receiving the box to scan only the DUID. However, this is an extra bit of training that would be required, and of course it is error-prone.

For this reason, manufacturers of DHCPv6-capable network hardware with predetermined DUIDs are strongly encouraged to add a bar code label to the box containing the equipment with a DUID that matches the following ABNF [[RFC5234](#)]:

```
label = "DUID:" hex-sequence
hex-sequence = hex-octet * ( ":" hex-sequence )
hex-octet = hex-digit hex-digit
hex-digit = %x30 - %x39 / %x41 - %x46
```

For example, a device with a DUID in the DUID-LL format might have a bar code that reads as follows:

```
DUID:00:03:00:01:08:00:2b:4c:3d:9f
```

[7.](#) Acknowledgments

This document was motivated by my realization during a private conversation with Leaf Yeh that although this technique for mapping client link-layer addresses to inventory tracking systems is well-known to some experts in the DHCPv4 and DHCPv6 user community, it has not been documented by the IETF, and that readers of [RFC2131](#) and [RFC3315](#) might therefore be unaware that this usage pattern exists.

Thanks to Sten Carlsen, Niall O'Reilly and Simon Hobson for their careful review and discussion of this work.

[8.](#) Security Considerations

DHCP servers rely on information provided by the DHCP client to identify the client. In DHCPv6, the server typically relies on the DUID to uniquely identify the client; unless the DHCP packet is authenticated in some way, one clients can masquerade as another by presenting that client's DUID instead of its own.

This document proposes using one of the client's link-layer addresses as a means of identifying the client when it first connects to the network. This mechanism presents the same sorts of risks as does using the DUID to identify the client. Existing mitigation strategies in [RFC3315](#) will work equally well to prevent clients from presenting fraudulent link-layer identifiers.

In addition, there are several factors that make this less of an issue with the mechanisms described in this document. First, the link-layer address is only used as an identifier the first time the client connects to the network; like leap-of-faith authentication, this presents a very brief window of opportunity for an attack using the link-layer address.

Secondly, because the attack has to occur the first time the client connects to the network, it's more complicated to effect the attack:

the attacker has to snoop the link-layer address from the wire, and then attempt a DHCP transaction before the legitimate client starts its DHCP transaction. Provisioning software can detect such an attack, because two conflicting records for the same client will appear in the log in quick succession. If the attacker happens to guess the same DUID that the client chooses, this attack has no effect at all, since the correct information will be entered into the database.

[9.](#) IANA Considerations

The IANA is hereby absolved of any requirement to take any action in relation to this document.

[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.
- [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.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [RFC6939] Halwasia, G., Bhandari, S., and W. Dec, "Client Link-Layer Address Option in DHCPv6", [RFC 6939](#), May 2013.

[10.2.](#) Informative References

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.

Author's Address

Ted Lemon

Nominum, Inc.
2000 Seaport Blvd
Redwood City, CA 94063
USA

Phone: +1-650-381-6000
Email: Ted.Lemon@nominum.com

Lemon

Expires April 25, 2014

[Page 10]