

DHC Working Group
INTERNET DRAFT
Expires: May 2004
Internet Draft

Ted Lemon
Nominum
Bill Sommerfeld

Document: <[draft-ietf-dhc-3315id-for-v4-00.txt](#)>

Category: Standards Track

October, 2003

Node-Specific Client Identifiers for DHCPv4

Status of this Memo

This document is a submission by the Dynamic Host Configuration Working Group of the Internet Engineering Task Force (IETF). Comments should be submitted to the dhcwg@ietf.org mailing list.

Distribution of this memo is unlimited.

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at:

<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at:

<http://www.ietf.org/shadow.html>.

Abstract

This document specifies the format that is to be used for encoding DHCPv4 ([RFC2131](#)/[RFC2132](#)) client identifiers, so that those identifiers will be interchangeable with identifiers used in the DHCPv6 protocol ([RFC3315](#)).

Introduction

This document specifies the way in which DHCPv4 clients should identify themselves. DHCPv4 client implementations that conform to this specification use a DHCPv6-style DHCP Unique Identifier (DUID) encapsulated in a DHCPv4 client identifier option. This supersedes the behaviour specified in [RFC2131](#) and [RFC2132](#).

The reason for making this change is that as we make the transition from IPv4 to IPv6, there will be network devices that must use both DHCPv4 and DHCPv6. For reasons we will explain later, users of

these devices will have a smoother network experience if the devices identify themselves consistently, regardless of the version of DHCP they are using at any given moment. This change also addresses certain limitations in the functioning of [RFC2131](#)/2132-style DHCP client identifiers.

This document first describes the problem to be solved. It then states the new technique that is to be used to solve the problem.

Finally, it describes the specific changes that one would have to make to [RFC2131](#) and [RFC2132](#) in order for those documents not to contradict what is described in this document.

[1.0](#) Applicability

This document updates [RFC2131](#) and [RFC2132](#). DHCPv4 servers implementations SHOULD conform to this document. DHCPv4 clients on network devices that are expected to support DHCPv6 in the future SHOULD conform to this document. This document makes no changes to the behavior of DHCPv6 clients or servers.

DHCPv4 clients and servers that are implemented according to this document should be implemented as if the changes specified in section ??? have been made to [RFC2131](#) and [RFC2132](#).

[2.0](#) Problem Statement

[2.1](#). Client identities are ephemeral

[RFC2132](#) recommends that client identifiers be generated by using the permanent link-layer address of the network interface that the client is trying to configure. In cases where a network interface is removed from the client computer and replaced with a different network interface with a different permanent link-layer address, the identity of the client changes. The client loses its IP address and any other resources associated with its old identifier - for example, its domain name as published through the DHCP server.

[2.2](#). Clients can accidentally present multiple identities

Consider a DHCP client that has two network interfaces connected to the same link, one of which is wired and one of which is wireless. There are three interesting cases here:

- (a) Each network interface is attached to a different link.
- (b) Both network interface are connected to the same link.
- (c) Only one network interface is attached to a link.

Case (a) is problematic, and is beyond the scope of this document. Even the full implications of cases (b) and (c) are beyond the scope of this document. However, it seems safe to point out that cases (b) and (c) are very common in practice, because many devices such as laptop computers that are popular at the time of this writing have both wireless and wired network interfaces that are installed simultaneously. Both wired and wireless have advantages - wired has the advantage of speed, and wireless the advantage of mobility.

So it seems likely that there will be devices that are in states (b) and (c) frequently, and indeed frequently make transitions between these states. If the DHCP client that configures these devices uses the link-layer address of each device as an identifier, the two devices will appear to the DHCP server to be different nodes, and thus will be assigned different IP addresses, and, in state (b), only one of the two devices will be reachable

through the domain name registered by the DHCP server. Furthermore, if a device in state (b) makes the transition to state (c), it is quite possible that the lease for the device that has lost connectivity will remain valid for some time, and will be the one that got the registered domain name. In this case, the client will not be reachable through its registered domain name.

2.3. [RFC2131](#)/2132 and [RFC3315](#) identifiers are incompatible

The 'client identifier' option is used by DHCP clients and servers to identify clients. In some cases, the value of the 'client identifier' option is used to mediate access to resources (for example, the client's domain name, as published through the DHCP server). [RFC2132](#) and [RFC3315](#) specify different methods for deriving client identifiers. These methods guarantee that the DHCPv4 and DHCPv6 identifier will be different. This means that mediation of access to resources using these identifiers will not work correctly in cases where a node may be configured using DHCPv4 in some cases and DHCPv6 in other cases.

2.4. [RFC2131](#) does not require the use of a client identifier

[RFC2131](#) allows the DHCP server to identify clients either by using the client identifier option sent by the client, or, if the client did not send one, the client's link-layer address. Like the client identifier format recommended by [RFC2131](#), this suffers from the problems previously described in (2) and (3).

3. Solutions

The solution to problem (1) is to use a DHCP client identifier that is persistent - not tied to a particular piece of removable network hardware. Then, when network hardware is swapped in and out, the client identifier does not change, and thus the client has a consistent IP address and consistent use of whatever resources have been associated with its identifier.

The solution to problem (2) is the same. Although it is beyond the scope of this document to say how a DHCP client supporting two network interfaces in this way would provide a smooth user experience, it does seem safe to say that it will need to use a persistent DHCP client identifier that is the same across the interfaces that it configures.

It is also worth noting that in case (a), it is harmless for the device to use the same client identifier on both interfaces - in this case, the DHCP server or servers serving these two links will see the two network interfaces as distinct because they are connected to different links, even though they use the same

identifier.

The solution to problem (3) is to use the DHCP Unique Identifier as defined in [RFC3315](#) as a client identifier. The DUID provides several different ways of producing persistent DHCP client identifiers, at least one of which is likely to be appropriate for any particular sort of network device. So it turns out that this also solves problems (1) and (2).

The solution to problem (4) is to deprecate the use of the contents of the chaddr field in the DHCP packet as a means of identifying the client.

4. Recommendations

4.1. DHCP Client behavior

DHCP clients conforming to this specification MUST use stable DHCP node identifiers in the dhcp-client-identifier option. DHCP clients MUST NOT use client identifiers based solely on layer two addresses that are hard-wired to the layer two device (e.g., the ethernet MAC address) as suggested in [RFC2131](#). DHCP clients MUST send a 'client identifier' option containing a DHCP Unique Identifier, as defined in [RFC3315](#).

The general format of the DHCPv4 'client identifier' option is defined in [section 9.14 of RFC2132](#). To send a

To send a DUID in a DHCPv4 'client identifier' option, the type of the 'client identifier' option should be 255. The type field is immediately followed by the DUID. The format of the 'client identifier' option is as follows:

Code	Len	Type	DHCP Unique Identifier			
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
61	n	255	d1	d2	d3	d4 ...
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

Any DHCPv4 or DHCPv6 client that conforms to this specification SHOULD provide a means by which an operator can learn what DUID the client has chosen. Such clients SHOULD also provide a means by which the operator can configure the DUID. A device that is normally configured with both a DHCPv4 and DHCPv6 client SHOULD automatically use the same DUID for DHCPv4 and DHCPv6 without any operator intervention.

4.2. DHCPv4 Server behavior

This document does not require any change to DHCPv4 or DHCPv6 servers that follow [RFC2131](#) and [RFC2132](#). However, some DHCPv4 servers can be configured not to conform to [RFC2131](#) and [RFC2132](#), in the sense that they ignore the 'client identifier' option and use the client's hardware address instead. Some DHCP servers do not take into account the possibility that the same client identifier may be used on separate links, and thus will behave incorrectly when a DHCP client acquires leases on two different IP addresses on two different links at the same time.

DHCP servers that conform to this specification MUST use the 'client identifier' option to identify the client if the client sends it. DHCP servers MUST assume that when a lease on an IP address is bound to a particular DHCP client identifier, all other still-valid leases on IP addresses bound to that client identifier are still in use.

4.3. Changes to [RFC2131](#)

In [section 2 of RFC2131](#), on page 9, the text that reads "; for example, the 'client identifier' may contain a hardware address, identical to the contents of the 'chaddr' field, or it may contain another type of identifier, such as a DNS name" is deleted.

In [section 4.2 of RFC2131](#), the text "The client MAY choose to explicitly provide the identifier through the 'client identifier' option. If the client supplies a 'client identifier', the client MUST use the same 'client identifier' in all subsequent messages, and the server MUST use that identifier to identify the client. If the client does not provide a 'client identifier' option, the server MUST use the contents of the 'chaddr' field to identify the client." is replaced by the text "The client MUST explicitly provide a client identifier through the 'client identifier' option."

In the same section, the text "Use of 'chaddr' as the client's unique identifier may cause unexpected results, as that identifier may be associated with a hardware interface that could be moved to a new client. Some sites may choose to use a manufacturer's serial number as the 'client identifier', to avoid unexpected changes in a clients network address due to transfer of hardware interfaces among computers. Sites may also choose to use a DNS name as the 'client identifier', causing address leases to be associated with the DNS name rather than a specific hardware box." is replaced by the text "The DHCP client MUST NOT rely on the 'chaddr' field to identify it."

In [section 4.4.1 of RFC2131](#), the text "The client MAY include a different unique identifier" is replaced with "The client MUST include a unique identifier".

In sections [3.1](#), item 4 and 6, 3.2 item 3 and 4, and 4.3.1, where [RFC2131](#) says that 'chaddr' may be used instead of the 'client identifier' option, the text that says "or 'chaddr'" and "'chaddr' or" is deleted.

Note that these changes do not relieve the DHCP server of the obligation to use 'chaddr' as an identifier if the client does not send a 'client identifier' option. Rather, they oblige clients that conform with this document to send a 'client identifier' option, and not rely on 'chaddr' for identification. DHCP servers MUST use 'chaddr' as an identifier in cases where 'client identifier' is not sent, in order to support old clients that do not conform with this document.

4.4. Changes to [RFC2132](#)

The text in [section 9.14](#), beginning with "The client identifier MAY consist of" through "that meet this requirement for uniqueness." is replaced with "the client identifier consists of a type field whose value is normally 255, followed by a two-byte type field, followed by the contents of the identifier. The two-byte type field and the format of the contents of the identifier are defined in RFC3315, [section 9](#)." The text "its minimum length is 2" in the following paragraph is deleted.

5. Security Considerations

This document raises no new security issues. Potential exposure to attack in the DHCPv4 protocol are discussed in [section 7](#) of the DHCP protocol specification [[RFC2131](#)] and in Authentication for DHCP messages [[RFC3118](#)]. Potential exposure to attack in the DHCPv6 protocol is discussed in [section 23 of RFC3315](#).

6. IANA Considerations

This document defines no new name spaces that need to be administered by the IANA. This document deprecates all 'client identifier' type codes other than 255, and thus there is no need for the IANA to track possible values for the type field of the 'client identifier' option.

7. Author's Addresses

Ted Lemon
Nominum
2385 Bay Road
Redwood City, CA 94063 USA
+1 650 381 6000
mellon@nominum.com

Bill Sommerfeld

8. Full Copyright Statement

"Copyright (C) The Internet Society July 2003. All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any

way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

9. Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

