

Network Working Group
Request for Comments: 4361
Updates: 2131, 2132, 3315
Category: Standards Track

T. Lemon
Nominum
B. Sommerfield
Sun Microsystems
February 2006

Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document specifies the format that is to be used for encoding Dynamic Host Configuration Protocol Version Four (DHCPv4) client identifiers, so that those identifiers will be interchangeable with identifiers used in the DHCPv6 protocol. This document also addresses and corrects some problems in RFC 2131 and RFC 2132 with respect to the handling of DHCP client identifiers.

Table of Contents

1. Introduction	2
2. Terminology	2
3. Applicability	2
4. Problem Statement	3
4.1. Client identities are ephemeral.	3
4.2. Clients can accidentally present multiple identities.	3
4.3. RFC 2131/2132 and RFC 3315 identifiers are incompatible. ...	4
4.4. RFC 2131 does not require the use of a client identifier. ..	4
5. Requirements	4
6. Implementation	6
6.1. DHCPv4 Client Behavior	6
6.2. DHCPv6 Client Behavior	7
6.3. DHCPv4 Server Behavior	7
6.4. Changes from RFC 2131	8
6.5. Changes from RFC 2132	9

7. Notes on DHCP Clients in Multi-stage Network Booting	9
8. Security Considerations	10
9. References	10
9.1. Normative References	10
9.2. Informative References	10

1. Introduction

This document specifies the way in which Dynamic Host Configuration Protocol Version 4 [RFC2131] clients should identify themselves. DHCPv4 client implementations that conform to this specification use a DHCP Unique Identifier (DUID) as specified in Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [RFC3315]. The DUID is encapsulated in a DHCPv4 client identifier option, as described in "DHCP Options and BOOTP Vendor Extensions" [RFC2132]. The behaviour described here supersedes the behavior 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. 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. Most obviously, DNS updates made by the DHCP server on behalf of the client will be handled more correctly. This change also addresses certain limitations in the functioning of RFC 2131/2132-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 RFC 2131 and RFC 2132 in order for those documents not to contradict what is described in this document.

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 RFC 2119 [RFC2119].

3. Applicability

This document updates RFC 2131 and RFC 2132. This document also specifies behavior that is required of DHCPv4 and DHCPv6 clients that are intended to operate in a dual-stack configuration. Finally, this document recommends behavior for host configurations where more than one DHCP client must operate in sequence in order to fully configure

the client (e.g., a network boot loader and the operating system it loads).

DHCPv4 clients and servers that are implemented according to this document should be implemented as if the changes specified in sections 6.3 and 6.4 have been made to RFC 2131 and RFC 2132. DHCPv4 clients should, in addition, follow the behavior specified in section 6.1. DHCPv6 clients should follow the behavior specified in section 6.2. DHCPv4 servers should additionally follow the behavior specified in section 6.3.

4. Problem Statement

4.1. Client identities are ephemeral.

RFC 2132 recommends that client identifiers be generated by using the permanent link-layer address of the network interface that the client is trying to configure. One result of this recommendation is that when the network interface hardware on a client computer is replaced, 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 DHCPv4 server).

4.2. Clients can accidentally present multiple identities.

Consider a DHCPv4 client that has two network interfaces, one of which is wired and one of which is wireless. The DHCPv4 client will succeed in configuring either zero, one, or two network interfaces. Under the current specification, each network interface will receive a different IP address. The DHCPv4 server will treat each network interface as a completely independent DHCPv4 client, on a completely independent host.

Thus, when the client presents some information to be updated in a network directory service, such as the DNS, the name that is presented will be the same on both interfaces, but the identity that is presented will be different. What will happen is that one of the two interfaces will get the name, and will retain that name as long as it has a valid lease, even if it loses its connection to the network, while the other network interface will never get the name. In some cases, this will achieve the desired result; when only one network interface is connected, sometimes its IP address will be published. In some cases, the one connected interface's IP address will not be the one that is published. When there are two interfaces, sometimes the correct one will be published, and sometimes not.

This is likely to be a particular problem with modern laptops, which usually have built-in wireless ethernet and wired ethernet. When the user is near a wired outlet, he or she may want the additional speed and privacy provided by a wired connection, but that same user may unplug from the wired network and wander around, still connected to the wireless network. When a transition like this happens, under the current scheme, if the address of the wired interface is the one that gets published, this client will be seen by hosts attempting to connect to it as if it has intermittent connectivity, even though it actually has continuous network connectivity through the wireless port.

Another common case of a duplicate identity being presented occurs when a boot monitor such as a Pre-Boot Execution Environment (PXE) loader specifies one DHCP client identifier, and then the operating system loaded by the boot loader specifies a different identity.

4.3. RFC 2131/2132 and RFC 3315 identifiers are incompatible.

The 'client identifier' option is used by DHCPv4 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 DHCPv4 server). RFC 2132 and RFC 3315 specify different methods for deriving client identifiers. These methods guarantee that the DHCPv4 and DHCPv6 identifiers 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.

4.4. RFC 2131 does not require the use of a client identifier.

RFC 2131 allows the DHCPv4 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 RFC 2131, this suffers from the problems previously described in sections 4.2 and 4.3.

5. Requirements

In order to address the problems stated in section 4, DHCPv4 client identifiers must have the following characteristics:

- They must be persistent, in the sense that a particular host's client identifier must not change simply because a piece of network hardware is added or removed.

- It must be possible for the client to represent itself as having more than one network identity, for example, so that a client with two network interfaces can express to the DHCPv4 server that these two network interfaces are to receive different IP addresses, even if they happen to be connected to the same link.
- In cases where the DHCPv4 client is expressing more than one network identity at the same time, it must nevertheless be possible for the DHCPv4 server to determine that the two network identities belong to the same host.
- In some cases it may be desirable for a DHCP client to present the same identity on two interfaces, so that if they both happen to be connected to the same network, they will both receive the same IP address. In such cases, it must be possible for the client to use exactly the same identifier for each interface.
- DHCPv4 servers that do not conform to this specification, but that are compliant with the older client identifier specification, must correctly handle client identifiers sent by clients that conform to this specification.
- DHCPv4 servers that do conform to this specification must interoperate correctly with DHCPv4 clients that do not conform to this specification, except that when configuring such clients, behaviors such as those described in section 2 may occur.
- The use by DHCPv4 clients of the chaddr field of the DHCPv4 packet as an identifier must be deprecated.
- DHCPv4 client identifiers used by dual-stack hosts that also use DHCPv6 must use the same host identification string for both DHCPv4 and DHCPv6. For example, a DHCPv4 server that uses the client's identity to update the DNS on behalf of a DHCPv4 client must register the same client identity in the DNS that would be registered by the DHCPv6 server on behalf of the DHCPv6 client running on that host, and vice versa.

In order to satisfy all but the last of these requirements, we need to construct a DHCPv4 client identifier out of two parts. One part must be unique to the host on which the client is running. The other must be unique to the network identity being presented. The DHCP Unique Identifier (DUID) and Identity Association Identifier (IAID) specified in RFC 3315 satisfy these requirements.

In order to satisfy the last requirement, we must use the DUID to identify the DHCPv4 client. So, taking all the requirements together, the DUID and IAID described in RFC 3315 are the only possible solution.

By following these rules, a compliant DHCPv4 client will interoperate correctly with both compliant and non-compliant DHCPv4 servers. A non-compliant DHCPv4 client will also interoperate correctly with a compliant DHCPv4 server. If either server or client is not compliant, the goals stated in the document are not met, but there is no loss of functionality.

6. Implementation

Here we specify changes to the behavior of DHCPv4 clients and servers. We also specify changes to the wording in RFC 2131 and RFC 2132. DHCPv4 clients, servers, and relay agents that conform to this specification must implement RFC 2131 and RFC 2132 with the wording changes specified in sections 6.3 and 6.4.

6.1. DHCPv4 Client Behavior

DHCPv4 clients conforming to this specification **MUST** use stable DHCPv4 node identifiers in the dhcp-client-identifier option. DHCPv4 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 RFC 2131, except as allowed in section 9.2 of RFC 3315. DHCPv4 clients **MUST** send a 'client identifier' option containing an Identity Association Unique Identifier, as defined in section 10 of RFC 3315, and a DHCP Unique Identifier, as defined in section 9 of RFC 3315. These together constitute an RFC 3315-style binding identifier.

The general format of the DHCPv4 'client identifier' option is defined in section 9.14 of RFC 2132.

To send an RFC 3315-style binding identifier in a DHCPv4 'client identifier' option, the type of the 'client identifier' option is set to 255. The type field is immediately followed by the IAID, which is an opaque 32-bit quantity. The IAID is immediately followed by the DUID, which consumes the remaining contents of the 'client identifier' option. The format of the 'client identifier' option is as follows:

Code	Len	Type	IAID				DUID			
+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+
61	n	255	i1	i2	i3	i4	d1	d2	...	
+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+	+---+

Any DHCPv4 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 by both a DHCPv4 and DHCPv6 client SHOULD automatically use the same DUID for DHCPv4 and DHCPv6 without any operator intervention.

DHCPv4 clients that support more than one network interface SHOULD use the same DUID on every interface. DHCPv4 clients that support more than one network interface SHOULD use a different IAID on each interface.

A DHCPv4 client that generates a DUID and that has stable storage MUST retain this DUID for use in subsequent DHCPv4 messages, even after an operating system reboot.

6.2. DHCPv6 Client Behavior

Any 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 by both a DHCPv4 and DHCPv6 client SHOULD automatically use the same DUID for DHCPv4 and DHCPv6 without any operator intervention.

6.3. DHCPv4 Server Behavior

This document does not require any change to DHCPv4 or DHCPv6 servers that follow RFC 2131 and RFC 2132. However, some DHCPv4 servers can be configured not to conform to RFC 2131 and RFC 2132, in the sense that they ignore the 'client identifier' option and use the client's hardware address instead.

DHCPv4 servers that conform to this specification MUST use the 'client identifier' option to identify the client if the client sends it.

DHCPv4 servers MAY use administrator-supplied values for chaddr and htype to identify the client in the case where the administrator is assigning a fixed IP address to the client, even if the client sends a client identifier option. This is ONLY permitted in the case where the DHCPv4 server administrator has provided the values for chaddr and htype, because in this case if it causes a problem, the administrator can correct the problem by removing the offending configuration information.

6.4. Changes from RFC 2131

In section 2 of RFC 2131, 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 RFC 2131, 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. The client MUST use the same 'client identifier' option for all messages."

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 client's 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 RFC 2131, the text "The client MAY include a different unique identifier" is replaced with "The client MUST include a unique identifier".

In section 3.1, items 4 and 6; section 3.2 item 3 and 4; and section 4.3.1, where RFC 2131 says that 'chaddr' may be used instead of the 'client identifier' option, the text "or 'chaddr'" and "'chaddr' or" is deleted.

Note that these changes do not relieve the DHCPv4 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. DHCPv4 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.

6.5. Changes from RFC 2132

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 four-byte IA_ID field, followed by the DUID for the client as defined in RFC 3315, section 9." The text "its minimum length is 2" in the following paragraph is deleted.

7. Notes on DHCP Clients in Multi-stage Network Booting

In some cases a single device may actually run more than one DHCP client in sequence, in the process of loading an operating system over the network. In such cases, it may be that the first-stage boot uses a different client identifier, or no client identifier, than the subsequent stage or stages.

The effect of this, under the DHCPv4 protocol, is that the two (in some cases more than two!) boot stages will present different identities. A DHCPv4 server will therefore allocate two different IP addresses to the two different boot stages.

Some DHCP servers work around this problem for the common case where the boot Programmable Read Only Memory (PROM) presents no client identifier, and the operating system DHCP client presents a client identifier constructed from the Message Authentication Code (MAC) address of the network interface -- both are treated as the same identifier. This prevents the consumption of an extra IP address.

A compliant DHCPv4 client does not use a client identifier constructed from the MAC address of the network interface, because network interfaces are not stable. So a compliant DHCPv4 client cannot be supported by a simple hack like the one described previously; this may have some significant impact at some sites.

We cannot state the solution to this problem as a set of requirements, because the circumstances in which this occurs vary too widely. However, we can make some suggestions.

First, we suggest that DHCP clients in network boot loaders request short lease times, so that their IP addresses are not retained. Such clients should send a DHCPRELEASE message to the DHCP server before moving on to the next stage of the boot process. Such clients should provide a way for the operating system DHCP client to configure a DUID to use in subsequent boots. DHCP clients in the final stage should, where possible, configure the DUID used by the boot PROM to be the same as the DUID used by the operating system.

Second, implementors of DHCPv4 clients that are expected to only be used in a multi-stage network boot configuration, that are not expected ever to network boot using DHCPv6, and that have a MAC address that cannot be easily changed may not need to implement the changes described in this specification. There is some danger in making this assumption--the first solution suggested is definitely better. A compromise might be to have the final-stage DHCP client detect whether it is running on legacy hardware; if it is, it uses the old identifier; if it is not, it follows the scheme described in the previous paragraph.

8. Security Considerations

This document raises no new security issues. Potential exposure to attack in the DHCPv4 protocol is 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 RFC 3315.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, 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.

9.2. Informative References

- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

Authors' Addresses

Ted Lemon
Nominum
2385 Bay Road
Redwood City, CA 94063 USA

Phone: +1 650 381 6000
EMail: mellon@nominum.com

Bill Sommerfeld
Sun Microsystems
1 Network Drive
Burlington, MA 01824

Phone: +1 781 442 3458
EMail: sommerfeld@sun.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

