

Dynamic Host Configuration Working
Group
Internet-Draft
Updates: [2131](#) (if approved)
Intended status: Standards Track
Expires: August 30, 2009

D. Hankins
ISC
February 26, 2009

Dynamic Host Configuration Protocol DHCPINFORM Message Clarifications
draft-ietf-dhc-dhcpinform-clarify-02

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 30, 2009.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

The DHCPINFORM message within the DHCPv4 protocol has in operation

Internet-Draft

DHCPINFORM Clarify

February 2009

diverged incompatibly from the current defined standard. This document seeks to provide clarification of actual behaviour and guidance for some situations that were previously omitted.

Table of Contents

1.	Introduction	3
2.	Requirements Language	4
3.	Client Behaviour	4
4.	Server Behaviour	4
5.	Security Considerations	7
6.	IANA Considerations	8
7.	Acknowledgements	8
8.	References	8
8.1.	Normative References	8
8.2.	Informative References	8
Appendix A.	RFC Interpretation Notes	9
	Author's Address	11

Internet-Draft

DHCPINFORM Clarify

February 2009

1. Introduction

The most recent DHCPv4 Standard [[RFC2131](#)] added a new DHCPv4 message: DHCPINFORM. The intent of the DHCPINFORM message was for clients that used manually entered fixed IPv4 addresses to still be able to get some configuration state dynamically. Since that time, however, we have seen this message used by normal DHCPv4 dynamically addressed clients; clients that have previously succeeded in receiving configuration through DHCPDISCOVER, DHCPOFFER, DHCPREQUEST, and finally DHCPACK messages.

These clients are attempting DHCPINFORM messages in order to obtain additional configuration state that was not present in their lease binding. The discovery is that DHCPINFORM can be used to reach extra DHCP servers, other than the one that gave an address, which may have more configuration options available but aren't in a position to give addresses. This extra configuration state is often required by applications that were not running at system startup, when the DHCP client was initialized.

Some of these DHCPINFORM clients have surfaced which run with stripped down user privileges, but still perform some network related functions. This software does not have the capacity to determine its IPv4 address(es), nor does it know what interface(s) are present on the system, or their hardware addresses. But it can send and receive DHCP packets. Consequently, the 'ciaddr' and 'chaddr' fields have been witnessed to be empty, even though they appear to be required to be filled by [RFC 2131](#). Clarification is sought for server behaviour when ciaddr is zero.

Another set of DHCP clients set the 'chaddr' field to a fixed magic value, rather than the client's hardware address, identifying them as part of a vendor's product. Although the 'chaddr' contents were never defined by any IETF RFC to be a valid place to store 'Vendor Identifying Information', their implementors believed this field was unused by the DHCP protocol in specific regards to DHCPINFORM because

a server would determine the client's hardware address through normal UDP unicast methods; IP forwarding leading to ARP [[RFC0826](#)] processing or similar.

We also wish to clarify a DHCPv4 server's behaviour when it receives a DHCPINFORM via a relay (when 'giaddr' is non-zero). [Section 4.1](#) of the DHCPv4 specification [[RFC2131](#)] seems to include DHCPINFORM->DHCPACK exchanges by describing generic behaviour for all DHCP OFFER and DHCPACK replies, and it requires that if giaddr is non-zero that it "MUST" be used. But this advice does not work in practice (due to BOOTP Relay Agent [[RFC1542](#)] requirements to use 'yiaddr' field contents). As a result, it also does not describe

current operational deployments of the DHCPINFORM message exchange.

[2.](#) Requirements Language

In this document, the key words "MAY", "MUST", "SHALL", "MUST NOT", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [BCP 14](#), [RFC 2119](#) [[RFC2119](#)].

[3.](#) Client Behaviour

Clients are still required to fulfill the DHCPv4 requirements for DHCPINFORM messages ([[RFC2131](#)], Sections [4.4.1](#) and [4.4.3](#)). But the following are clarified as in addition, or to overlay those requirements:

- o Clients MUST set 'ciaddr' to a working IPv4 address which they can use to receive replies. This address SHOULD be an address that is currently assigned to the interface upon which the client is transmitting its DHCPINFORM, except in the condition where the DHCP client is unable to determine a valid IP address for its host, in which case the client MUST set 'ciaddr' to all-zero.
- o Clients MUST set 'chaddr', 'htype', and 'hlen' to the hardware address of the interface upon which the DHCPINFORM message is being transmitted, except in the condition where the DHCP client is unable to determine this address, in which case all three fields MUST be set all-zero.

- o Clients MUST set the 'flags' field to zero. This means that the client MUST NOT set the 'BROADCAST' flag, and MUST be capable of receiving IP unicasts.
- o Clients SHOULD direct their DHCPINFORM via unicast UDP to the IPv4 address contained in the Server Identifier [[RFC2132](#)] option, if they have a currently active binding from previous DHCPREQUEST message exchanges.

4. Server Behaviour

DHCPv4 server behaviour in processing DHCPINFORM messages is a more difficult question to answer, due to inconsistent client behaviour and conflicting directions in [RFC 2131](#). The following is intended to be a more complete reference.

First, upon receiving a DHCPINFORM, a DHCPv4 Server MUST determine

the client's "relevant IPv4 address" according to the following in order of priority:

1. The Subnet Selection Option [[RFC3011](#)], if it is present and supported.
2. The 'ciaddr' field, if it is non-zero.
3. The Relay Agent Link Selection Sub-Option [[RFC3046](#)], if it is present in a Relay Agent Information Option [[RFC3046](#)] in the DHCPINFORM packet (never cached from a previous exchange).
4. The 'giaddr' field, if it is non-zero.
5. The IPv4 source address field, if it is non-zero.
6. The DHCPv4 Server's address on the interface on which the DHCPINFORM was received.

The DHCPv4 server checks to see if the "relevant IPv4 address" is within a range or subnet over which it holds authority, or if it is configured to respond. It will manufacture a DHCPACK response with

configuration values appropriate for the "relevant IPv4 address", possibly in addition to configuration values appropriate for the 'ciaddr' field contents if they are non-zero (configuration values granted to a specific address, or range of addresses of which the non-zero 'ciaddr' is a member). This MAY involve inspecting that address's current lease, but MUST NOT modify it in any way (such as by extending the lease time or granting the address).

In the manufactured response:

- o The 'htype', 'hlen', 'chaddr', 'ciaddr', 'xid', 'flags' (with the exception noted below), and 'giaddr' fields MUST be copied from the client's DHCPINFORM.
- o The 'hops' field MUST be zero.
- o The 'secs' field MUST be zero.
- o The 'yiaddr' field MUST be zero.
- o The 'siaddr' field MUST be zero.
- o The 'sname' and 'file' fields MAY be used exclusively for 'option overloading', but MUST be zeroed otherwise.

- o The 'options' field MUST be filled as described in [RFC 2131 Section 4.3.1](#).

Next, the DHCPv4 server MUST determine the "reply address and port" according to the following conditions in order:

1. If the 'ciaddr' field is non-zero, the server selects its contents as an IPv4 address and port 68 ('DHCP client').
2. If the 'giaddr' field is non-zero, the server selects its contents as an IPv4 address and port 67 ('DHCP server').
3. If the IPv4 source address field is non-zero, the server selects its contents as an IPv4 address and port 68 ('DHCP client')

4. The server selects the limited broadcast address (all-ones) and port 68 ('DHCP client').

At this point, the DHCPv4 server verifies that it holds configuration authority over the reply address (or link in case of limited broadcast address) it has selected to transmit the reply to. If the server has not been configured to hold authority over this address, it MUST NOT reply. It SHOULD increment a counter visible to the operator but SHOULD NOT log an error (unless a mechanism is used to suppress repeated log messages). See the Security section ([Section 5](#)) for the rationale behind this direction.

Note very carefully that a DHCPv4 server will send replies directly to a DHCPv4 client by way of 'ciaddr' even if the DHCPINFORM message was relayed. Note that this means DHCPINFORM processing is intentionally broken in deployments where the client's address space is unreachable by the DHCPv4 server. In such cases, the server should probably be configured not to reply to DHCPINFORMs.

Now, the server performs an exception to assist relay agents. If it selected the 'giaddr' as the destination address and port, then it MUST set the 'BROADCAST' bit in the flags field true, no matter what its value was in the client's DHCPINFORM message. Otherwise, the response could not be delivered; a BOOTP Relay Agent [[RFC1542](#)] is required to direct unicast server replies to the 'chaddr' and 'yiaddr' field contents, but 'chaddr' is not reliably filled, and 'yiaddr' is required to be all-zero. Setting the broadcast flag assists the relay agent in locating the client.

Having selected a destination IPv4 address and port number, the last step is to select a destination link layer address.

For the all-ones limited broadcast address, the DHCPv4 server MUST

use the all-ones broadcast hardware address.

For all other (unicast) destination selections, the DHCPv4 server MUST use its host operating system's usual methods to determine hardware addressing, as by IP forwarding and subsequent address resolution (such as through ARP [[RFC0826](#)]). Note that the DHCPv4 server MAY have seeded its ARP cache from a previous stateful exchange with the client (from 'chaddr' contents while processing a

DHCPREQUEST message, due to the requirement of DHCPv4 servers to unicast some replies before clients will process ARP), and some DHCPv4 software MAY still use 'chaddr' contents to direct replies to directly connected clients. Consequently, DHCPINFORM can not be reasonably expected to instigate an immediate ARP broadcast, nor can 'chaddr' contents be used for any purpose other than to carry the unicast hardware address with which a client might reasonably be reached.

5. Security Considerations

As with all DHCP messages, DHCPINFORM and DHCPACK replies contain no capacity for encryption, and all packet contents must be presumed readable in the clear. In particular, and as outlined above, in some circumstances the packets may be broadcast and so more easily intercepted than most other messages.

Authentication for DHCPv4 Messages [[RFC3118](#)] does exist, but is not well deployed. Care should be taken in the degree to which configuration parameters provided by DHCPv4 are trusted, as the replies can be easily spoofed by any eavesdropper. Again noting that packets may be broadcast under some circumstances, the BOOTP header Transaction Id field ("XID") is insufficient protection from man-in-the-middle attacks.

A relay agent receives replies via unicast UDP messages from a DHCP server, and may broadcast these packets on the inside-facing network. If an outside attacker was aware of this relay agent and its unicast address, this facility could be used to produce broadcast storms on the network. Care should be taken to ensure that the relay agent is not open to this kind of attack, possibly making use of Relay Agent Authentication [[RFC4030](#)] to ensure that a DHCPv4 server can not be induced to sending bogus replies to the relay.

This protocol uses the 'ciaddr' field contents to direct replies, which may be set blindly by the client to any value, regardless of IP source address validation or related filter restrictions. If an attacker were to identify a number of DHCPv4 servers which reply to addresses not under their authority to configure, and those servers

had enough large DHCPv4 options in configuration to request, it could

represent a significant amplification vector in straight packet-load Denial-of-Service attacks. For this reason, servers MUST NOT make replies to addresses not explicitly configured under their authority to configure.

6. IANA Considerations

This document has no action for IANA.

7. Acknowledgements

This document has been reviewed and improved by the comments of several people, but the author would like to take a moment to thank Alfred Hoenes, who has submitted revised text for this document.

8. References

8.1. 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.
- [RFC3011] Waters, G., "The IPv4 Subnet Selection Option for DHCP", [RFC 3011](#), November 2000.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", [RFC 3046](#), January 2001.

8.2. Informative References

- [RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, [RFC 826](#), November 1982.
- [RFC1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", [RFC 1542](#), October 1993.

- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", [RFC 3118](#), June 2001.
- [RFC4030] Stapp, M. and T. Lemon, "The Authentication Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option", [RFC 4030](#), March 2005.

[Appendix A](#). RFC Interpretation Notes

This section will self-destruct as (if) we near last-call. It is a list of [RFC 2131](#) notations I've used as a guide to navigate this maze.

- o [Section 4.1](#): "If the BROADCAST bit is cleared to 0, the message SHOULD be sent as an IP unicast to the IP address specified in the 'yiaddr' field and the link-layer address specified in the 'chaddr' field." But in other sections we say that 'yiaddr' is set zero. So any message via a relay has to be broadcast in response. I don't know if relays check for 'yiaddr' equal zero and downgrade to broadcast, so I think it's best to set this bit just to help them out (this is like DHCPNAK replies, where 'yiaddr' is also zero).
- o [Section 4.1](#) also has a lengthy paragraph that's been brought up on the DHCWG mailing list, which seems to indicate that for all "DHCP OFFER and DHCPACK" messages, giaddr is always first, followed by ciaddr, (then depending on broadcast bit) followed by yiaddr, followed by broadcast. DHCPACK is certainly the message that is used to reply to DHCPINFORMs! But it really isn't clear in my mind that this section was updated in step with the addition of the DHCPINFORM message; parts of it seem very clearly to be presuming client's broadcasts, so parts can only be interpreted for DHCPREQUEST when not RENEWING. However! It could be true that we should send to giaddr first, in which case we will always be broadcasting replies to DHCPINFORM via relays? Setting yiaddr? None of this is specified, and it doesn't work without those clues.
- o [Section 4.4.1](#), Table 5, uses the same column for both DHCPINFORM and DHCPDISCOVER. It is clear that both DHCPINFORM and DHCPDISCOVER make the same use of chaddr/htype/hlen. It is also clear that 'ciaddr' is zero on DHCPDISCOVER - but very clearly non-zero ("the client's network address") on DHCPINFORM. Since this is in the same column, and uses a wording that is similar to other columns (which have an "or" in them), it may be overlooked

if you weren't looking closely enough. There is no normative language that reinforces this, but it seems like a non-zero

'ciaddr' was not one of [RFC 2131](#)'s intentions.

- o [Section 3.4](#): "Servers receiving a DHCPINFORM message construct a DHCPACK message with any local configuration parameters appropriate for the client without: allocating a new address, checking for an existing binding, filling in 'yiaddr' or including lease time parameters." ...snip... "The server SHOULD check the network address in a DHCPINFORM message for consistency, but MUST NOT check for an existing lease. The server forms a DHCPACK message containing the configuration parameters for the requesting client and sends the DHCPACK message directly to the client." This is kind of problematic. Our DHCP software lets you scope configuration parameters in a tree hierarchy, and this includes right on the lease itself. So the MUST NOT (and the non-normative language before) that keeps us from checking for an existing lease (very vague language) also means we may give different answers to the same client at DORA time versus DHCPINFORM time. The client actually over-writes its config with the DHCPINFORM values and becomes broken. I think these two validations in [RFC 2131](#) can be simplified to one validation, which is even simpler: The server validates that the client's address is one which it is responsible for configuring.
- o Again [Section 3.4](#): "The servers SHOULD unicast the DHCPACK reply to the address given in the 'ciaddr' field of the DHCPINFORM message." That 'SHOULD' kind of makes you wonder what /else/ you would do.
- o [Section 4.3.5](#): "The server responds to a DHCPINFORM message by sending a DHCPACK message directly to the address given in the 'ciaddr' field of the DHCPINFORM message. The server MUST NOT send a lease expiration time to the client and SHOULD NOT fill in 'yiaddr'." So, a non-normative indication for 'ciaddr', followed by a normative SHOULD NOT for 'yiaddr' (conflicts with non-normative language in 3.4 above which makes it sound like yiaddr is zeroed). Curious. [Section 4.3.1](#) Table 3 seems to indicate 'yiaddr' is always set on DHCPACK to the "IP address assigned to client", with no reservation for message type (other fields in this table make distinctions for DHCPINFORM).

- o [Section 4.3.1](#) Table 3 lists in the DHCPACK column a lot of strange values when processing a DHCPINFORM packet. Namely, "'xid' from client DHCPREQUEST message". That is a strange thing to hang on to from the previous DORA exchange, and it's supposed that a client might not even do the DORA exchange (or might do it with a different server). Obviously this was just overlooked, but it brings everything in this table into question. We should remove those questions.

Hankins

Expires August 30, 2009

[Page 10]

Internet-Draft

DHCPINFORM Clarify

February 2009

Author's Address

David W. Hankins
Internet Systems Consortium, Inc.
950 Charter Street
Redwood City, CA 94063
US

Phone: +1 650 423 1307
Email: David_Hankins@isc.org

Hankins

Expires August 30, 2009

[Page 11]