

Network Working Group
Request for Comments: 4388
Category: Standards Track

R. Woundy
Comcast Cable
K. Kinnear
Cisco Systems
February 2006

Dynamic Host Configuration Protocol (DHCP) Leasequery

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

A Dynamic Host Configuration Protocol version 4 (DHCPv4) server is the authoritative source of IP addresses that it has provided to DHCPv4 clients. Other processes and devices that already make use of DHCPv4 may need to access this information. The leasequery protocol provides these processes and devices a lightweight way to access IP address information.

Table of Contents

1. Introduction	2
2. Terminology	5
3. Background	7
4. Design Goals	7
4.1. Broadcast ARP Is Undesirable	7
4.2. SNMP and LDAP Are Not Appropriate	8
4.3. DHCP Relay Agent Functionality Is Common	8
4.4. DHCP Servers Are a Reliable Source of Location Information	9
4.5. Minimal Additional Configuration Is Required	9
5. Protocol Overview	9
6. Protocol Details	12
6.1. Definitions Required for DHCPLEASEQUERY Processing	12
6.2. Sending the DHCPLEASEQUERY Message	14
6.3. Receiving the DHCPLEASEQUERY Message	15
6.4. Responding to the DHCPLEASEQUERY Message	16
6.5. Receiving a DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN Message	20
6.6. Receiving No Response to the DHCPLEASEQUERY Message	21
6.7. Lease Binding Data Storage Requirements	22
6.8. Using the DHCPLEASEQUERY Message with Multiple DHCP Servers	23
7. Security Considerations	23
8. IANA Considerations	24
9. Acknowledgements	24
10. References	25
10.1. Normative References	25
10.2. Informative References	25

1. Introduction

A DHCPv4 server contains considerable authoritative information concerning the IP addresses it has leased to DHCP clients. Sometimes devices or other processes may need access to this information. In some cases, these devices or processes already have the capability to send and receive DHCP packets, and so the leasequery protocol is designed to give these processes and devices a low-overhead way to access such information.

For example, access concentrators that act as DHCP relay agents sometimes derive information important to their operation by extracting data out of the DHCP packets they forward, a process known as "gleaning". Unfortunately, the typical access concentrator loses its gleaned information when the access concentrator is rebooted or is replaced. This memo proposes that when gleaned DHCP information is not available, the access concentrator/relay agent can obtain the

location information directly from the DHCP server(s) using the DHCPLEASEQUERY message.

To continue this example in more depth, in many broadband access networks, the access concentrator needs to associate an IP address lease to the correct endpoint location, which includes knowledge of the host hardware address, the port or virtual circuit that leads to the host, and/or the hardware address of the intervening subscriber modem. This is particularly important when one or more IP subnets are shared among many ports, circuits, and modems. Representative cable and DSL environments are depicted in Figures 1 and 2 below.

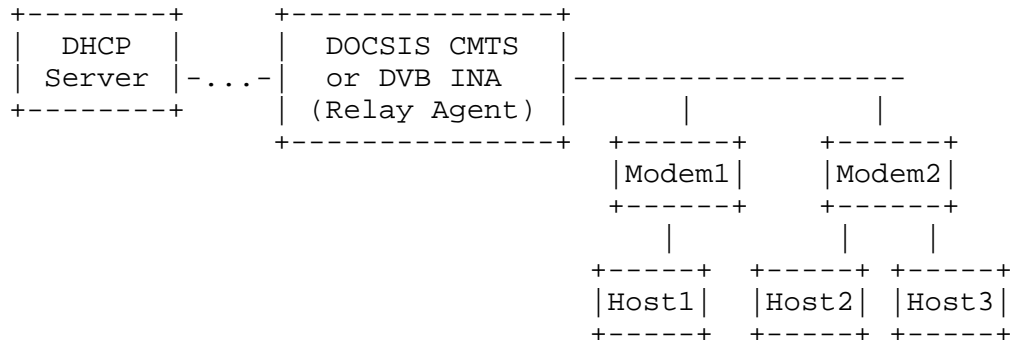


Figure 1: Cable Environment for DHCPLEASEQUERY

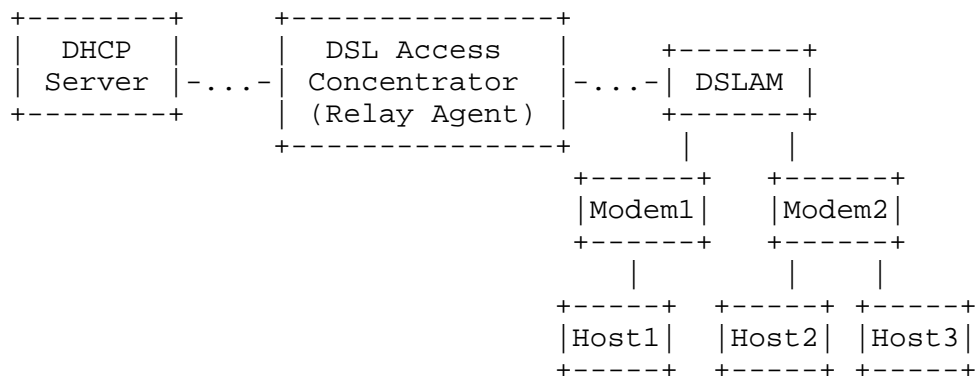


Figure 2: DSL Environment for DHCPLEASEQUERY

Knowledge of this location information can benefit the access concentrator in several ways:

1. The access concentrator can forward traffic to the access network using the correct access network port, down the correct virtual circuit, through the correct modem, to the correct hardware address.

2. The access concentrator can perform IP source address verification of datagrams received from the access network. The verification may be based on the datagram source hardware address, the incoming access network port, the incoming virtual circuit, and/or the transmitting modem.
3. The access concentrator can encrypt datagrams that can only be decrypted by the correct modem, using mechanisms such as [BPI] or [BPI+].

The access concentrator in this example obtains the location information primarily from "gleaning" information from DHCP server responses sent through the relay agent. When location information is not available from "gleaning", e.g., because the access concentrator has rebooted, the access concentrator can query the DHCP server(s) for location information using the DHCPLEASEQUERY message defined in this document.

The DHCPLEASEQUERY message is a new DHCP message type transmitted from a DHCP relay agent to a DHCP server. A DHCPLEASEQUERY-aware relay agent sends the DHCPLEASEQUERY message when it needs to know the location of an IP endpoint. The DHCPLEASEQUERY-aware DHCP server replies with a DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN message. The DHCPLEASEACTIVE response to a DHCPLEASEQUERY message allows the relay agent to determine the IP endpoint location and the remaining duration of the IP address lease. The DHCPLEASEUNASSIGNED is similar to a DHCPLEASEACTIVE message, but indicates that there is no currently active lease on the resultant IP address but that this DHCP server is authoritative for this IP address. The DHCPLEASEUNKNOWN message indicates that the DHCP server has no knowledge of the information specified in the query (e.g., IP address, MAC address, or Client-identifier option).

The DHCPLEASEQUERY message does not presuppose a particular use for the information it returns -- it is simply designed to return information for which the DHCP server is an authoritative source to a client that requests that information. It is designed to make it straightforward for processes and devices that already interpret DHCP packets to access information from the DHCP server.

This document specifies an extension specifically to the DHCPv4 protocol [RFC2131]. Given the nature of the DHCPv6 protocol [RFC3315], there is no effective way to make the DHCPLEASEQUERY message interaction common between DHCPv4 and DHCPv6 even should the desire to do so exist.

The DHCPLEASEQUERY message was the result of a set of specific real-world implementation needs that appeared many years after the DHCPv4 protocol was in wide use. Furthermore, at the time of this writing, the DHCPv6 protocol has yet to be widely deployed. The needs of access concentrators in yet to be determined DHCPv6 deployment scenarios are difficult to estimate. If a DHCPLEASEQUERY-like function is necessary in DHCPv6, many of the ideas of this document will probably be applicable, while others may not. We have been cautioned against designing protocol capabilities for which there is only an imagined consumer, and that is all that exists today in the realm of DHCPLEASEQUERY for DHCPv6.

Thus, this document applies only to DHCPv4, and for clarity we have not appended DHCPv4 to every appearance of several common terms. In this document, all references to IP addresses should be taken to mean IPv4 addresses, and all references to DHCP servers and DHCP clients should be taken to mean DHCPv4 servers and DHCPv4 clients.

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].

This document uses the following terms:

- o "access concentrator"

An access concentrator is a router or switch at the broadband access provider's edge of a public broadband access network. This document assumes that the access concentrator includes the DHCP relay agent functionality.

- o "DHCP client"

A DHCP client is an Internet host using DHCP to obtain configuration parameters such as a network address.

- o "DHCP relay agent"

A DHCP relay agent is a third-party agent that transfers Bootstrap Protocol (BOOTP) and DHCP messages between clients and servers residing on different subnets, per [RFC951] and [RFC1542].

- o "DHCP server"

A DHCP server is an Internet host that returns configuration parameters to DHCP clients.

- o "downstream"

Downstream is the direction from the access concentrator towards the broadband subscriber.

- o "gleaning"

Gleaning is the extraction of location information from DHCP messages, as the messages are forwarded by the DHCP relay agent function.

- o "location information"

Location information is information needed by the access concentrator to forward traffic to a broadband-accessible host. This information includes knowledge of the host hardware address, the port or virtual circuit that leads to the host, and/or the hardware address of the intervening subscriber modem.

- o "MAC address"

In the context of a DHCP packet, a MAC address consists of the following fields: hardware type "htype", hardware length "hlen", and client hardware address "chaddr".

- o "stable storage"

Every DHCP server is assumed to have some form of what is called "stable storage". Stable storage is used to hold information concerning IP address bindings (among other things) so that this information is not lost in the event of a server failure that requires restart of the server.

- o "upstream"

Upstream is the direction from the broadband subscriber towards the access concentrator.

3. Background

The focus of this document is to enable processes and devices that wish to access information from the DHCP server in a lightweight and convenient manner. It is especially appropriate for processes and devices that already interpret DHCP packets.

One important motivating example is that the DHCPLEASEQUERY message allows access concentrators to send DHCPLEASEQUERY messages to DHCP servers to obtain location information of broadband access network devices.

This document assumes that many access concentrators have an embedded DHCP relay agent functionality. Typical access concentrators include DOCSIS Cable Modem Termination Systems (CMTSS) [DOCSIS], DVB Interactive Network Adapters (INAs) [EUROMODEM], and DSL Access Concentrators.

The DHCPLEASEQUERY message is an extension to the DHCP protocol [RFC2131].

The DHCPLEASEQUERY message is a query message only and does not affect the state of the IP address or the binding information associated with it.

4. Design Goals

The goal of this document is to provide a lightweight mechanism for processes or devices to access information contained in the DHCP server. It is designed to allow processes and devices that already process and interpret DHCP messages to access this information in a rapid and lightweight manner.

Some of this information might be acquired in a different way, and the following sections discuss some of these alternative approaches.

4.1. Broadcast ARP Is Undesirable

The access concentrator can transmit a broadcast Address Resolution Protocol (ARP) Request [RFC826], and observe the origin and contents of the ARP Reply, to reconstruct the location information.

The ARP mechanism is undesirable for three reasons:

1. the burden on the access concentrator to transmit over multiple access ports and virtual circuits (assuming that IP subnets span multiple ports or virtual circuits),

2. the burden on the numerous subscriber hosts to receive and process the broadcast, and
3. the ease by which a malicious host can misrepresent itself as the IP endpoint.

4.2. SNMP and LDAP Are Not Appropriate

Access concentrator implementations typically do not have Simple Network Management Protocol (SNMP) management client interfaces nor Lightweight Directory Access Protocol (LDAP) client interfaces (although they typically do include SNMP management agents). This is one reason why this document does not leverage the proposed DHCP Server MIB [DHCPMIB].

The DHCP Server MIB effort [DHCPMIB] grew out of traffic engineering and troubleshooting activities at large DHCP installations, and is primarily intended as a method of gathering performance statistics about servers the load presented to them.

Despite the presence in the proposed DHCPv4 server MIB of objects that report configuration and status information, the MIB is intended to provide more generic, server-wide aggregated or summarized data. DHCPLEASEQUERY is intended to provide detailed, specific information about individual leases at a level that would be difficult or impossible to shoehorn into a MIB.

From an implementation standpoint, the DHCPLEASEQUERY message is not required to be supported by all DHCPv4 servers. Since it appears that defining optional MIB objects and objects for optional features in a MIB is discouraged, trying to support DHCPLEASEQUERY functionality optionally through a MIB would be similarly discouraged from an SNMP MIB standpoint.

4.3. DHCP Relay Agent Functionality Is Common

Access concentrators commonly act as DHCP relay agents. Furthermore, many access concentrators already glean location information from DHCP server responses, as part of the relay agent function.

The gleaning mechanism as a technique to determine the IP addresses valid for a particular downstream link is preferred over other mechanisms (ARP, SNMP, LDAP) because of the lack of additional network traffic, but sometimes gleaning information can be incomplete. The access concentrator usually cannot glean information from any DHCP unicast (i.e., non-relayed) messages due to performance reasons. Furthermore, the DHCP-gleaned location information often

does not persist across access concentrator reboots (due to lack of stable storage), and almost never persists across concentrator replacements.

4.4. DHCP Servers Are a Reliable Source of Location Information

DHCP servers are the most reliable source of location information for access concentrators, particularly when the location information is dynamic and not reproducible by algorithmic means (e.g., when a single IP subnet extends behind many broadband modems). DHCP servers participate in all IP lease transactions (and therefore in all location information updates) with DHCP clients, whereas access concentrators sometimes miss some important lease transactions.

An access concentrator can be configured with the IP addresses of multiple different DHCP servers, so that no one DHCP server is a single point of failure.

4.5. Minimal Additional Configuration Is Required

Access concentrators can usually query the same set of DHCP servers used for forwarding by the relay agent, thus minimizing configuration requirements.

5. Protocol Overview

In the following discussion of the DHCPLEASEQUERY message, the client of the message is assumed to be an access concentrator. Note that access concentrators are not the only allowed (or required) consumers of the information provided by the DHCPLEASEQUERY message, but they do give readers a concrete feel for how the message might be used.

The access concentrator initiates all DHCPLEASEQUERY message conversations. This document assumes that the access concentrator gleans location information in its DHCP relay agent function. However, the location information is usually unavailable after the reboot or replacement of the access concentrator.

Suppose the access concentrator is a router, and further suppose that the router receives an IP datagram to forward downstream to the public broadband access network. If the location information for the downstream next hop is missing, the access concentrator sends one or more DHCPLEASEQUERY message(s), each containing the IP address of the downstream next hop in the "ciaddr" field.

This query will then be answered by returning the information current when this client's lease was last granted or renewed, allowing the access concentrator to forward the IP datagram.

An alternative approach is to send in a DHCPLEASEQUERY message with the "ciaddr" field empty and the MAC address (i.e., "htype", "hlen", and "chaddr" fields) with a valid MAC address or a Client-identifier option (option 61) appearing in the options area. In this case, the DHCP server must return an IP address in the ciaddr if it has any record of the client described by the Client-identifier or MAC address. In the absence of specific configuration information to the contrary (see Section 6.4), it SHOULD be the IP address with the latest client-last-transaction-time associated with the client described by the MAC address or Client-identifier option.

The DHCP servers that implement this protocol always send a response to the DHCPLEASEQUERY message: either a DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN. The reasons why a DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN message might be generated are explained in the specific query regimes, below.

Servers that do not implement the DHCPLEASEQUERY message SHOULD simply not respond.

The DHCPLEASEQUERY message can support three query regimes: A server that implements the DHCPLEASEQUERY message must implement all three query regimes.

- o Query by IP address:

For this query, the requester supplies only an IP address in the DHCPLEASEQUERY message. The DHCP server will return any information that it has on the most recent client to have been assigned that IP address.

The DHCP server replies with a DHCPLEASEUNASSIGNED or DHCPLEASEACTIVE message if the IP address in the DHCPLEASEQUERY message corresponds to an IP address about which the server has definitive information (i.e., it is authorized to lease this IP address). The server replies with a DHCPLEASEUNKNOWN message if the server does not have definitive information concerning the address in the DHCPLEASEQUERY message.

- o Query by MAC address:

For this query, the requester supplies only a MAC address in the DHCPLEASEQUERY message. The DHCP server will return any information that it has on the IP address most recently accessed by a client with that MAC address. In addition, it may supply additional IP addresses that have been associated with that MAC address in different subnets. Information about these bindings

can then be found using the Query by IP Address, described above.

The DHCP server replies with a DHCPLEASEACTIVE message if the MAC address in the DHCPLEASEQUERY message corresponds to a MAC address with an active lease on an IP address in this server. The server replies with a DHCPLEASEUNKNOWN message if the server does not presently have an active lease by a client with this MAC address in this DHCP server.

- o Query by Client-identifier option:

For this query, the requester supplies only a Client-identifier option in the DHCPLEASEQUERY message. The DHCP server will return any information that it has on the IP address most recently accessed by a client with that Client-identifier. In addition, it may supply additional IP addresses that have been associated with Client-identifier in different subnets. Information about these bindings can then be found using the Query by IP Address, described above.

The DHCP server replies with a DHCPLEASEACTIVE message if the Client-identifier in the DHCPLEASEQUERY message currently has an active lease on an IP address in this DHCP server. The server replies with a DHCPLEASEUNKNOWN message if the server does not have an active lease by a client with this Client-identifier.

For many DHCP servers, the query by IP address is likely to be the most efficient form of leasequery. This is the form of DHCPLEASEQUERY that SHOULD be used if possible.

The DHCPLEASEUNASSIGNED or DHCPLEASEACTIVE message reply must always contain the IP address in the "ciaddr" field. The DHCPLEASEACTIVE message SHOULD contain the physical address of the IP address lease owner in the "htype", "hlen", and "chaddr" fields. The Parameter Request List (option 55) can be used to request specific options to be returned about the IP address in the ciaddr. The reply often contains the time until expiration of the lease, and the original contents of the Relay Agent Information option [RFC3046]. The access concentrator uses the "chaddr" field and Relay Agent Information option to construct location information, which can be cached on the access concentrator until lease expiration.

Any DHCP server that supports the DHCPLEASEQUERY message SHOULD save the information from the most recent Relay Agent Information option (option 82) [RFC3046] associated with every IP address that it serves. It is assumed that most clients that generate the DHCPLEASEQUERY message will ask for the Relay Agent Information

option (option 82) in the Parameter Request List (option 55), and so supporting the DHCPLEASEQUERY message without having the Relay Agent Information option around to return to the client is likely to be less than helpful.

A server that implements DHCPLEASEQUERY SHOULD also save the information on the most recent Vendor class identifier, option 60, associated with each IP address, since this option is also likely to be requested by clients sending the DHCPLEASEQUERY message.

6. Protocol Details

6.1. Definitions Required for DHCPLEASEQUERY Processing

The operation of the DHCPLEASEQUERY message requires the definition of the following new and extended values for the DHCP packet beyond those defined by [RFC2131] and [RFC2132]. See also Section 8, IANA Considerations.

1. The message type option (option 53) from [RFC2132] requires four new values: one for the DHCPLEASEQUERY message itself and one for each of its three possible responses DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, DHCPLEASEUNKNOWN. The values of these message types are shown below in an extension of the table from section 9.6 of [RFC2132]:

Value	Message Type
-----	-----
10	DHCPLEASEQUERY
11	DHCPLEASEUNASSIGNED
12	DHCPLEASEUNKNOWN
13	DHCPLEASEACTIVE

2. There is a new option, the client-last-transaction-time:

client-last-transaction-time

This option allows the receiver to determine the time of the most recent access of the client. It is particularly useful when DHCPLEASEACTIVE messages from two different DHCP servers need to be compared, although it can be useful in other situations. The value is a duration in seconds from the current time into the past when this IP address was most recently the subject of communication between the client and the DHCP server.

This MUST NOT be an absolute time. This MUST NOT be an absolute number of seconds since Jan. 1, 1970. Instead, this MUST be an integer number of seconds in the past from the time the DHCPLEASEACTIVE message is sent that the client last dealt with this server about this IP address. In the same way that the IP Address Lease Time option (option 51) encodes a lease time that is a number of seconds into the future from the time the message was sent, this option encodes a value that is a number of seconds into the past from when the message was sent.

The code for the this option is 91. The length of the this option is 4 octets.

Code	Len	Seconds in the past			
91	4	t1	t2	t3	t4

3. There is a second new option, the associated-ip option:

associated-ip

This option is used to return all of the IP addresses associated with the DHCP client specified in a particular DHCPLEASEQUERY message.

The code for this option is 92. The minimum length for this option is 4 octets, and the length MUST always be a multiple of 4.

Code	Len	Address 1				Address 2		
92	n	a1	a2	a3	a4	a1	a2	...

6.2. Sending the DHCPLEASEQUERY Message

The DHCPLEASEQUERY message is typically sent by an access concentrator. The DHCPLEASEQUERY message uses the DHCP message format as described in [RFC2131], and uses message number 10 in the DHCP Message Type option (option 53). The DHCPLEASEQUERY message has the following pertinent message contents:

- o The giaddr MUST be set to the IP address of the requester (i.e., the access concentrator). The giaddr is independent of the "ciaddr" field to be searched -- it is simply the return address of the DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN message from the DHCP server.

Note that this use of the giaddr is consistent with the definition of giaddr in [RFC2131], where the giaddr is always used as the return address of the DHCP response message. In some (but not all) contexts in RFC 2131, the giaddr is used as the "key" to access the appropriate address pool. The DHCPLEASEQUERY message is one of those cases where the giaddr MUST NOT be used as such a "key".

- o The Parameter Request List option (option 55) SHOULD be set to the options of interest to the requester. The interesting options are likely to include the IP Address Lease Time option (option 51), the Relay Agent Information option (option 82), and possibly the Vendor class identifier option (option 60). In the absence of a Parameter Request List option, the server SHOULD return the same options it would return for a DHCPREQUEST message that didn't contain a DHCPLEASEQUERY message, which includes those mandated by Section 4.3.1 of [RFC2131] as well as any options that the server was configured to always return to a client.

Additional details concerning different query types are:

- o Query by IP address:

The values of htype, hlen, and chaddr MUST be set to zero.

The "ciaddr" field MUST be set to the IP address of the lease to be queried.

The Client-identifier option (option 61) MUST NOT appear in the packet.

- o Query by MAC address:

The values of htype, hlen, and chaddr MUST be set to the value of the MAC address to search for.

The "ciaddr" field MUST be set to zero.

The Client-identifier option (option 61) MUST NOT appear in the packet.

- o Query by Client-identifier option:

There MUST be a Client-identifier option (option 61) in the DHCPLEASEQUERY message.

The "ciaddr" field MUST be set to zero.

The values of htype, hlen, and chaddr MUST be set to zero.

The DHCPLEASEQUERY message SHOULD be sent to a DHCP server which is known to possess authoritative information concerning the IP address. The DHCPLEASEQUERY message MAY be sent to more than one DHCP server, and in the absence of information concerning which DHCP server might possess authoritative information concerning the IP address, it SHOULD be sent to all DHCP servers configured for the associated relay agent (if any are known).

Any device expecting to use a DHCPLEASEQUERY message SHOULD ensure that the Relay Agent Info option that it uses contains information that unambiguously identifies the device.

6.3. Receiving the DHCPLEASEQUERY Message

A server that implements the DHCPLEASEQUERY message MUST implement all three query regimes: query by IP address, query by MAC address, and query by Client-identifier.

A DHCPLEASEQUERY message MUST have a non-zero giaddr. The DHCPLEASEQUERY message MUST have exactly one of the following: a non-zero ciaddr, a non-zero htype/hlen/chaddr, or a Client-identifier option.

The DHCP server that receives a DHCPLEASEQUERY message MUST base its response on the particular data item used in the query.

The `giaddr` is used only for the destination address of any generated response and, while required, is not otherwise used in generating the response to the `DHCPLEASEQUERY` message. It **MUST NOT** be used to restrict the processing of the query in any way, and **MUST NOT** be used to locate a subnet to which the `ciaddr` (if any) must belong.

Note that this use of the `giaddr` is consistent with the definition of `giaddr` in [RFC2131], where the `giaddr` is always used as the return address of the DHCP response message. In some (but not all) contexts in RFC 2131, the `giaddr` is used as the "key" to access the appropriate address pool. The `DHCPLEASEQUERY` message is one of those cases where the `giaddr` **MUST NOT** be used as such a "key".

6.4. Responding to the `DHCPLEASEQUERY` Message

There are three possible responses to a `DHCPLEASEQUERY` message:

- o `DHCPLEASEUNASSIGNED`

The server **MUST** respond with a `DHCPLEASEUNASSIGNED` message if this server has information about the IP address, but there is no active lease for the IP address. The `DHCPLEASEUNASSIGNED` message is only returned for a query by IP address, and indicates that the server manages this IP address, but there is no currently active lease on this IP address.

- o `DHCPLEASEUNKNOWN`

The `DHCPLEASEUNKNOWN` message indicates that the server does not manage the IP address or the client specified in the `DHCPLEASEQUERY` message does not currently have a lease on an IP address.

When responding with a `DHCPLEASEUNKNOWN`, the DHCP server **MUST NOT** include other DHCP options in the response.

- o `DHCPLEASEACTIVE`

The `DHCPLEASEACTIVE` message indicates that the server not only knows about the IP address and client specified in the `DHCPLEASEACTIVE` message, but also knows that there is an active lease by that client for that IP address.

The server **MUST** respond with a `DHCPLEASEACTIVE` message when the IP address returned in the "`ciaddr`" field is currently leased.

6.4.1. Determining the IP address about Which to Respond

Since the response to a DHCPLEASEQUERY request can only contain full information about one IP address -- the one that appears in the "ciaddr" field -- determination of which IP address about which to respond is a key issue. Of course, the values of additional IP addresses for which a client has a lease must also be returned in the associated-ip option (Section 6.1, #3). This is the only information returned not directly associated with the IP address in the "ciaddr" field.

In the event that an IP address appears in the "ciaddr" field of a DHCPLEASEQUERY message, if that IP address is one managed by the DHCP server, then that IP address MUST be set in the "ciaddr" field of a DHCPLEASEUNASSIGNED message.

If the IP address is not managed by the DHCP server, then a DHCPLEASEUNKNOWN message must be returned.

If the "ciaddr" field of the DHCPLEASEQUERY is zero, then the DHCPLEASEQUERY message is a query by Client-identifier or MAC address. In this case, the client's identity is any client that has proffered an identical Client-identifier option (if the Client-identifier option appears in the DHCPLEASEQUERY message), or an identical MAC address (if the MAC address fields in the DHCPLEASEQUERY message are non-zero). This client matching approach will, for the purposes of this section, be described as "Client-identifier or MAC address".

If the "ciaddr" field is zero in a DHCPLEASEQUERY message, then the IP address placed in the "ciaddr" field of a DHCPLEASEACTIVE message MUST be that of an IP address for which the client that most recently used the IP address matches the Client-identifier or MAC address specified in the DHCPLEASEQUERY message.

If there is only a single IP address that fulfills this criteria, then it MUST be placed in the "ciaddr" field of the DHCPLEASEACTIVE message.

In the case where more than one IP address has been accessed by the client specified by the MAC address or Client-identifier option, then the DHCP server MUST return the IP address returned to the client in the most recent transaction with the client unless the DHCP server has been configured by the server administrator to use some other preference mechanism.

If, after all of the above processing, no value is set in the "ciaddr" field of the DHCPLEASEUNASSIGNED or DHCPLEASEACTIVE message, then a DHCPLEASEUNKNOWN message MUST be returned instead.

6.4.2. Building a DHCPLEASEUNASSIGNED or DHCPLEASEACTIVE Message Once the "ciaddr" Field Is Set

Once the "ciaddr" field of the DHCPLEASEUNASSIGNED is set, the processing for a DHCPLEASEUNASSIGNED message is complete. No other options are returned for the DHCPLEASEUNASSIGNED message.

For the DHCPLEASEACTIVE message, the rest of the processing largely involves returning information about the IP address specified in the "ciaddr" field.

The IP address in the "ciaddr" field of the DHCPLEASEUNASSIGNED or DHCPLEASEACTIVE message MUST be one for which this server is responsible (or a DHCPLEASEUNKNOWN message would have already been returned early in the processing described in the previous section).

The MAC address of the DHCPLEASEACTIVE message MUST be set to the values that identify the client associated with the IP address in the "ciaddr" field of the DHCPLEASEUNASSIGNED message.

If the Client-identifier option (option 61) is specified in the Parameter Request List option (option 55), then the Client-identifier (if any) of the client associated with the IP address in the "ciaddr" field SHOULD be returned in the DHCPLEASEACTIVE message.

In the case where more than one IP address has been involved in a DHCP message exchange with the client specified by the MAC address and/or Client-identifier option, then the list of all of the IP addresses MUST be returned in the associated-ip option, whether or not that option was requested as part of the Parameter Request List option.

If the IP Address Lease Time option (option 51) is specified in the Parameter Request List and if there is a currently valid lease for the IP address specified in the ciaddr, then the DHCP server MUST return this option in the DHCPLEASEACTIVE message with its value equal to the time remaining until lease expiration. If there is no valid lease for the IP address, then the server MUST NOT return the IP Address Lease Time option (option 51).

A request for the Renewal (T1) Time Value option or the Rebinding (T2) Time Value option in the Parameter Request List of the DHCPLEASEQUERY message MUST be handled like the IP Address Lease Time option is handled. If there is a valid lease and these times are not

yet in the past, then the DHCP server SHOULD return these options (when requested) with the remaining time until renewal or rebinding, respectively. If these times are already in the past, or if there is not currently a valid lease for this IP address, the DHCP server MUST NOT return these options.

If the Relay Agent Information (option 82) is specified in the Parameter Request List, then the information contained in the most recent Relay Agent Information option received from the relay agent associated with this IP address MUST be included in the DHCPLEASEACTIVE message.

The DHCPLEASEACTIVE message SHOULD include the values of all other options not specifically discussed above that were requested in the Parameter Request List of the DHCPLEASEQUERY message and that are acceptable to return based on the list of "non-sensitive options", discussed below.

DHCP servers SHOULD be configurable with a list of "non-sensitive options" that can be returned to the client when specified in the Parameter Request List of the DHCPLEASEQUERY message. Any option not on this list SHOULD NOT be returned to a client, even if requested by that client.

The DHCP server uses information from its lease binding database to supply the DHCPLEASEACTIVE option values. The values of the options that were returned to the DHCP client would generally be preferred, but in the absence of those, options that were sent in DHCP client requests would be acceptable.

In some cases, the Relay Agent Information option in an incoming DHCPREQUEST packet is used to help determine the options returned to the DHCP client that sent the DHCPREQUEST. When responding to a DHCPLEASEQUERY message, the DHCP server MUST use the saved Relay Agent Information option just like it did when responding to the DHCP client in order to determine the values of any options requested by the DHCPLEASEQUERY message. The goal is to return the same option values to the DHCPLEASEQUERY as those that were returned to the DHCPDISCOVER or DHCPREQUEST from the DHCP client (unless otherwise specified, above).

In the event that two servers are cooperating to provide a high-availability DHCP server, as supported by [RFC2131], they would have to communicate some information about IP address bindings to each other. In order to properly support the DHCPLEASEQUERY message, these servers MUST ensure that they communicate the Relay Agent Information option information to each other in addition to any other IP address binding information.

6.4.3. Sending a DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN Message

The server expects a giaddr in the DHCPLEASEQUERY message, and unicasts the DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN message to the giaddr. If the "giaddr" field is zero, then the DHCP server MUST NOT reply to the DHCPLEASEQUERY message.

6.5. Receiving a DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN Message

When a DHCPLEASEACTIVE message is received in response to the DHCPLEASEQUERY message, it means that there is a currently active lease for this IP address in this DHCP server. The access concentrator SHOULD use the information in the "htype", "hlen", and "chaddr" fields of the DHCPLEASEACTIVE as well as any Relay Agent Information option information included in the packet to refresh its location information for this IP address.

When a DHCPLEASEUNASSIGNED message is received in response to the DHCPLEASEQUERY message, that means that there is no currently active lease for the IP address present in the DHCP server, but that this server does in fact manage that IP address. In this case, the access concentrator SHOULD cache this information in order to prevent unacceptable loads on the access concentrator and the DHCP server in the face of a malicious or seriously compromised device downstream of the access concentrator. This caching could be as simple as simply setting a bit saying that a response was received from a server that knew about this IP address but that there was no current lease. This would, of course, need to be cleared when the access concentrator next "gleaned" that a lease for this IP address came into existence.

In either case, when a DHCPLEASEUNASSIGNED or DHCPLEASEACTIVE message is received in response to a DHCPLEASEQUERY message, it means that the DHCP server that responded is a DHCP server that manages the IP address present in the ciaddr, and the Relay Agent SHOULD cache this information for later use.

When a DHCPLEASEUNKNOWN message is received by an access concentrator that has sent out a DHCPLEASEQUERY message, it means that the DHCP server contacted supports the DHCPLEASEQUERY message but that the DHCP server does not have definitive information concerning the IP address contained in the "ciaddr" field of the DHCPLEASEQUERY message. If there is no IP address in the "ciaddr" field of the DHCPLEASEQUERY message, then a DHCPLEASEUNKNOWN message means that

the DHCP server does not have definitive information concerning the DHCP client specified in the "hlen", "htype", and "chaddr" fields or the Client-identifier option of the DHCPLEASEQUERY message.

The access concentrator SHOULD cache this information, but only for a relatively short lifetime, approximately 5 minutes.

Having cached this information, the access concentrator SHOULD only infrequently direct a DHCPLEASEQUERY message to a DHCP server that responded to a DHCPLEASEQUERY message for a particular "ciaddr" field with a DHCPLEASEUNKNOWN.

6.6. Receiving No Response to the DHCPLEASEQUERY Message

When an access concentrator receives no response to a DHCPLEASEQUERY message, there are several possible reasons:

- o The DHCPLEASEQUERY or a corresponding DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN was lost during transmission or the DHCPLEASEQUERY arrived at the DHCP server but it was dropped because the server was too busy.
- o The DHCP server doesn't support DHCPLEASEQUERY.

In the first of the cases above, a retransmission of the DHCPLEASEQUERY would be appropriate, but in the second of the two cases, a retransmission would not be appropriate. There is no way to tell these two cases apart (other than, perhaps, because of a DHCP server's response to other DHCPLEASEQUERY messages indicating that it does or does not support the DHCPLEASEQUERY message).

An access concentrator that utilizes the DHCPLEASEQUERY message SHOULD attempt to resend DHCPLEASEQUERY messages to servers that do not respond to them using a backoff algorithm for the retry time that approximates an exponential backoff. The access concentrator SHOULD adjust the backoff approach such that DHCPLEASEQUERY messages do not arrive at a server that is not otherwise known to support the DHCPLEASEQUERY message at a rate of more than approximately one packet every 10 seconds, and yet (if the access concentrator needs to send DHCPLEASEQUERY messages) not less than one DHCPLEASEQUERY per 70 seconds.

In practice, this approach would probably best be handled by a per-server timer that is restarted whenever a response to a DHCPLEASEQUERY message is received, and expires after one minute. The per-server timer would start off expired, and in the expired state only one DHCPLEASEQUERY message would be queued for the associated server.

All DHCPLEASEQUERY messages SHOULD use the exponential backoff algorithm specified in Section 4.1 of [RFC2131].

Thus, in the initial state, the per-server timer is expired, and a single DHCPLEASEQUERY message is queued for each server. After the first response to a DHCPLEASEQUERY message, the per-server timer is started. At that time, multiple DHCPLEASEQUERY messages can be sent in parallel to the DHCP server, though the total number SHOULD be limited to 100 or 200, to avoid swamping the DHCP server. Each of these messages uses the [RFC2131] exponential backoff algorithm. Every time a response to any of these messages is received, the per-server timer is reset and starts counting again up to one minute. In the event the per-server timer goes off, then all outstanding messages SHOULD be dropped except for a single DHCPLEASEQUERY message that is used to poll the server at approximately 64-second intervals until such time as another (or the first) response to the DHCPLEASEQUERY is received.

In the event that there is no DHCPLEASEQUERY traffic for one minute, then the per-server timer will expire. After that time, there will only be one DHCPLEASEQUERY message allowed to be outstanding to that server until a response to that message is received.

6.7. Lease Binding Data Storage Requirements

DHCP server implementations that implement the DHCPLEASEQUERY capability MUST save the most recent Relay Agent Information option from the most recent DHCPREQUEST packet for two reasons. First, it is almost certain to be requested by in the dhcp-parameter-request-list option in any DHCPLEASEQUERY request. Second, the saved Relay Agent Information option may be necessary to determine the value of other options given to the DHCP client, if these are requested by the dhcp-parameter-request list in the DHCPLEASEQUERY request.

This is a list of the information that is required to successfully implement

- o relay-agent-info option from client packet: MUST store with binding.
- o client-last-transaction-time of last client interaction: MUST store with binding.
- o vendor-class-id: SHOULD store with binding.

These data storage requirements are minimally larger than those required for normal operation of the DHCP protocol, as required to properly implement [RFC2131].

6.8. Using the DHCPLEASEQUERY Message with Multiple DHCP Servers

When using the DHCPLEASEQUERY message in an environment where multiple DHCP servers may contain authoritative information about the same IP address (such as when two DHCP servers are cooperating to provide a high-availability DHCP service), multiple, possibly conflicting, responses might be received.

In this case, some information in the response packet SHOULD be used to decide among the various responses. The client-last-transaction-time (if it is available) can be used to decide which server has more recent information concerning the IP address returned in the "ciaddr" field.

7. Security Considerations

Access concentrators that use DHCP gleanings, refreshed with DHCPLEASEQUERY messages, will maintain accurate location information. Location information accuracy ensures that the access concentrator can forward data traffic to the intended location in the broadband access network, can perform IP source address verification of datagrams from the access network, and can encrypt traffic that can only be decrypted by the intended access modem (e.g., [BPI] and [BPI+]). As a result, the access concentrator does not need to depend on ARP broadcasts across the access network, which is susceptible to malicious hosts that masquerade as the intended IP endpoints. Thus, the DHCPLEASEQUERY message allows an access concentrator to provide considerably enhanced security.

DHCP servers SHOULD prevent exposure of location information (particularly the mapping of hardware address to IP address lease, which can be an invasion of broadband subscriber privacy) by employing the techniques detailed in [RFC3118], "Authentication for DHCP Messages".

This RFC describes how a DHCP client interacts with a DHCP server. Access concentrators that send the DHCPLEASEQUERY message are essentially DHCP clients for the purposes of the DHCPLEASEQUERY message, even though they perform the functions of a DHCP relay agent as well. Thus, [RFC3118] is an appropriate mechanism for DHCPLEASEQUERY messages.

Since [RFC3118] discusses the normal DHCP client interaction, consisting of a DHCPDISCOVER, DHCPOFFER, DHCPREQUEST, and DHCPACK, it is necessary to transpose the operations described in [RFC3118] to the DHCPLEASEQUERY domain. The operations described in [RFC3118] for

DHCPDISCOVER are performed for DHCPLEASEQUERY, and the operations described for DHCPPOFFER are performed for DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, and DHCPLEASEUNKNOWN messages.

Access concentrators SHOULD minimize potential denial of service attacks on the DHCP servers by minimizing the generation of DHCPLEASEQUERY messages. In particular, the access concentrator SHOULD employ negative caching (i.e., cache DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE, and DHCPLEASEUNKNOWN responses to DHCPLEASEQUERY messages) and ciaddr restriction (i.e., don't send a DHCPLEASEQUERY message with a ciaddr outside of the range of the attached broadband access networks). Together, these mechanisms limit the access concentrator to transmitting one DHCPLEASEQUERY message (excluding message retries) per legitimate broadband access network IP address after a reboot event.

DHCP servers supporting the DHCPLEASEQUERY message SHOULD ensure that they cannot be successfully attacked by being flooded with large quantities of DHCPLEASEQUERY messages in a short time.

In some environments, it may be appropriate to configure a DHCP server with the IP addresses of the relay agents for which it may respond to DHCPLEASEQUERY messages, thereby allowing it to respond only to requests from only a handful of relay agents. This does not provide any true security, but may be useful to thwart unsophisticated attacks of various sorts.

8. IANA Considerations

IANA has assigned six values for this document. See Section 6.1 for details. There are four new messages types, which are the value of the message type option (option 53) from [RFC2132]. The value for DHCPLEASEQUERY is 10, the value for DHCPLEASEUNASSIGNED is 11, the value for DHCPLEASEUNKNOWN is 12, and the value for DHCPLEASEACTIVE is 13. Finally, there are two new DHCP option defined; the client-last-transaction-time option -- option code 91, and the associated-ip option -- option code 92.

9. Acknowledgements

Jim Forster, Joe Ng, Guenter Roeck, and Mark Stapp contributed greatly to the initial creation of the DHCPLEASEQUERY message.

Patrick Guelat suggested several improvements to support static IP addressing. Thomas Narten made many suggestions for improvements. Russ Housley pressed effectively for increased security capabilities, and Ted Hardie suggested ways to minimize undesired information leakage. Bert Wijnen suggested we clarify our focus to DHCPv4 and

distinguish our approach from that of the DHCP MIB. R. Barr Hibbs, one of the authors of the DHCP MIB, supplied information to effectively distinguish that effort from DHCPLEASEQUERY.

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.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.
- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

10.2. Informative References

- [RFC826] 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.
- [RFC951] Croft, W. and J. Gilmore, "Bootstrap Protocol", RFC 951, September 1985.
- [RFC1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, October 1993.
- [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.
- [BPI] SCTE Data Standards Subcommittee, "Data-Over-Cable Service Interface Specifications: DOCSIS 1.0 Baseline Privacy Interface Specification SCTE 22-2 2002", 2002, available at <http://www.scte.org/standards/>.

- [BPI+] CableLabs, "Data-Over-Cable Service Interface Specifications: Baseline Privacy Plus Interface Specification CM-SP-BPI+_I12-050812", August 2005, available at <http://www.cablemodem.com/>.
- [DHCPMIB] Hibbs, R., Waters, G., "Dynamic Host Configuration Protocol (DHCP) Server MIB", Work in Progress, February 2004.
- [DOCSIS] SCTE Data Standards Subcommittee, "Data-Over-Cable Service Interface Specifications: DOCSIS 1.0 Radio Frequency Interface Specification SCTE 22-1 2002", 2002, available at <http://www.scte.org/standards/>.
- [EUROMODEM] ECCA, "Technical Specification of a European Cable Modem for digital bi-directional communications via cable networks", Version 1.0, May 1999.

Authors' Addresses

Rich Woundy
Comcast Cable
27 Industrial Ave.
Chelmsford, MA 01824

Phone: (978) 244-4010
EMail: richard_woundy@cable.comcast.com

Kim Kinnear
Cisco Systems
1414 Massachusetts Ave
Boxborough, MA 01719

Phone: (978) 936-0000
EMail: kinnear@cisco.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).

