Dynamic Host Configuration Working Group
INTERNET DRAFT

Rich Woundy
Comcast Cable

Kim Kinnear
Cisco Systems

October 2003
Expires April 2004

# DHCP Lease Query
<draft-ietf-dhc-leasequery-06.txt>

## Status of this Memo



## Copyright Notice



## Abstract

A DHCP server contains considerable authoritative information concerning the IP addresses it has leased to DHCP clients. Other processes and devices, many that already send and receive DHCP format packets, sometimes need to access this information. The leasequery protocol is designed to give these processes and devices a lightweight way to access information that may be critical to their operation.

## 1. Introduction

A DHCP server contains considerable authoritative information concerning the IP addresses it has leased to DHCP clients. Other processes and devices, many that already send and receive DHCP format packets, sometimes need to access this information. The leasequery protocol is designed to give these processes and devices a lightweight way to access information that may be critical to their operation.

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.

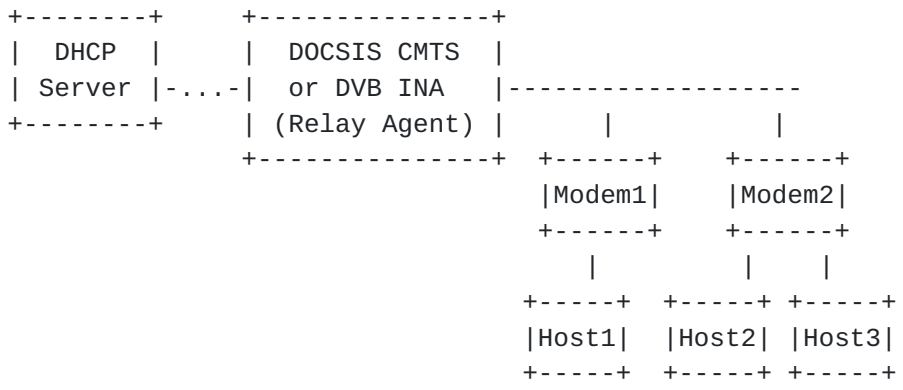


Figure 1: Cable Environment for DHCPLEASEQUERY

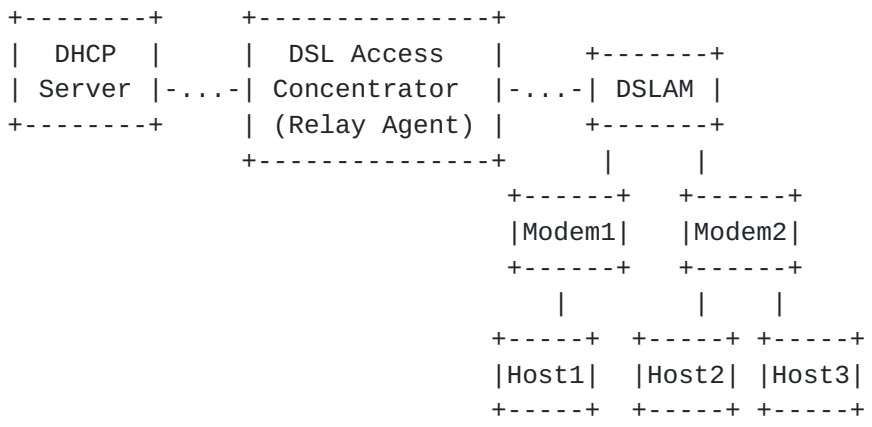


Figure 2: DSL Environment for DHCPLEASEQUERY

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 which 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 DHCPLEASEKNOWN, 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 DHCPLEASEKNOWN 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-id 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 which requests that information. It is designed to make it straightforward for processes and devices which already interpret DHCP packets to access information from the DHCP server.

## 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 [RFC 2119].

This document uses the following terms:

- "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.

- "DHCP client"

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

- "DHCP relay agent"

  A DHCP relay agent is a third-party agent that transfers BOOTP and DHCP messages between clients and servers residing on different subnets, per [RFC 951] and [RFC 1542].

- "DHCP server"

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

- "downstream"

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

- "gleaning"

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

- "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.

- "MAC address"

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

- "primary DHCP server"

  The primary DHCP server in a DHCP Failover environment is configured to provide primary service to a set of DHCP clients for a particular set of subnet address pools.

- "secondary DHCP server"

  The secondary DHCP server in a DHCP Failover environment is configured to act as backup to a primary server for a particular set of subnet address pools.

- "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 which requires restart of the server.

- "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 which wish to access information from the DHCP server in a lightweight and convenient manner. It is especially appropriate for processes and devices which 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 optional extension to the DHCP protocol [RFC 2131].

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 which 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 ARP Request [RFC 826], 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 Client Functionality is Lacking

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

### 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 as 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 reader 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.

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 SHOULD 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 MUST be the IP address most recently used by the client described by the MAC address or Client-identifier option (or the client described by both, if both appear).

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

Servers which do not implement the DHCPLEASEQUERY message fall into two classes. Those that simply do not know about the DHCPLEASEQUERY message will simply not respond to it, so clients which send the DHCPLEASEQUERY message MUST be prepared to deal with this behavior. Servers which are aware of the DHCPLEASEQUERY message but do not implement it SHOULD respond with a DHCPUNIMPLEMENTED message but MAY simply not respond.

The DHCPLEASEQUERY message can support three query regimes:

- 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 DHCPLEASEKNOWN or DHCPLEASEACTIVE message if the IP address in the DHCPLEASEQUERY message corresponds to an IP address about which the server has definitive information (ie., 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.

  A server which implements the DHCPLEASEQUERY message MUST implement this capability.

- 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 addition IP addresses which 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 an 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.

A server which implements the DHCPLEASEQUERY message SHOULD implement this capability. If it does not, it SHOULD respond with a DHCPUNIMPLEMENTED message when it receives a query by MAC address.

- Query by Client-identifier option:

  For this query, the requester supplies only a client-id 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-id. In addition, it may supply addition IP addresses which have been associated with client-id 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-id 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-id.

  A server which implements the DHCPLEASEQUERY message SHOULD implement this capability. If it does not, it SHOULD respond with a DHCPUNIMPLEMENTED message when it receives a query by Client-identifier option address.

Generally, the query by IP address is likely to be the most efficient and widely implemented form of leasequery, and it SHOULD be used if at all possible. Use of the other two query formats SHOULD be minimized, as they can potentially place a large load on some servers.

The DHCPLEASEKNOWN or DHCPLEASEACTIVE message reply MUST always contain the IP address in the ciaddr field. The DHCPLEASEACTIVE message SHOULD contains 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 [RFC 3046]. The access concentrator uses the "chaddr" and Relay Agent Information option to construct location information, which can be cached on the access concentrator until lease expiration.

Any DHCP server which supports the DHCPLEASEQUERY message SHOULD save the information from the most recent Relay Agent Information option (option 82) [RFC 3046] associated with every IP address which it serves. It is assumed that most clients which 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 which 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 a likely candidate 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 [RFC 2131] and [RFC 2132]. See also Section 8, IANA considerations.

1. The message type option (option 53) from [RFC 2132] requires five new values: one for the DHCPLEASEQUERY message itself and and one for each of its four possible responses DHCPLEASEKNOWN, DHCPLEASEACTIVE, DHCPLEASEUNKNOWN, and DHCPUNIMPLEMENTED. The values of these message types are shown below in a reproduction of the table from [RFC 2132]:

| Value | Message Type |
|---|---|
| 1 | DHCPDISCOVER |
| 2 | DHCPOFFER |
| 3 | DHCPREQUEST |
| 4 | DHCPDECLINE |
| 5 | DHCPACK |
| 6 | DHCPNAK |
| 7 | DHCPRELEASE |
| 8 | DHCPINFORM |
| TBD | DHCPLEASEQUERY |
| TBD | DHCPLEASEKNOWN |
| TBD | DHCPLEASEUNKNOWN |
| TBD | DHCPLEASEACTIVE |
| TBD | DHCPUNIMPLEMENTED |

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 which is a number of seconds into the future from the time the message was sent, this option encodes a value which is a number of seconds into the past from when the message was sent.

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

```
 Code   Len      Seconds in the past
+-----+-----+-----+-----+-----+-----+
| TBD |  4  |  t1 |  t2 |  t3 |  t4 |
+-----+-----+-----+-----+-----+-----+
```

3. There in 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 TBD. 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
+-----+-----+-----+-----+-----+-----+-----+-----+--
| TBD |  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 [RFC 2131], and uses message number TBD in the DHCP Message Type option (option 53). The DHCPLEASEQUERY message has the following pertinent message contents:

- 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 for the DHCPLEASEKNOWN, DHCPLEASEACTIVE or DHCPLEASEUNKNOWN message from the DHCP server.

- 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 will

return the same options it would return for a DHCPREQUEST message which didn't contain a DHCPLEASEQUERY message, which includes those mandated by [RFC 2131, Section 4.3.1] as well as any options which 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 0.

  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 0.

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).

## 6.3. Receiving the DHCPLEASEQUERY Message

A DHCPLEASEQUERY message MUST have a non-zero giaddr. The

DHCPLEASEQUERY message MUST have exactly one of: a non-zero ciaddr, a non-zero "htype"/"hlen"/"chaddr", or a Client-identifier.

The DHCP server which 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 locate a subnet to which the ciaddr (if any) must belong.

### 6.4. Responding to the DHCPLEASEQUERY Message

There are four possible responses to a DHCPLEASEQUERY message:

- DHCPLEASEKNOWN

  The server MUST respond with a DHCPLEASEKNOWN message if this server has information about the IP address, but there is no active lease for the IP address. The DHCPLEASEKNOWN 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.

- 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 SHOULD NOT include other DHCP options in the response.

- DHCPLEASEACTIVE

  The DHCPLEASEACTIVE message indicates that the server not only knows about the IP address and client specified in the DHCPLEASEACTIVE message but also 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.

- DHCPUNIMPLEMENTED

  The DHCPUNIMPLEMENTED response to the DHCPLEASEQUERY message

indicates that the particular form of DHCPLEASEQUERY used is not implemented in this DHCP server. It may mean that the DHCPLEASEQUERY message as a whole is not implemented by this DHCP server although it is usually used to indicate that a query by Client-identifier or MAC address is not implemented by a DHCP server that otherwise supports a DHCPLEASEQUERY by IP address.

The DHCPUNIMPLEMENTED message can apply to any unimplemented messages, and MAY be used to respond to messages other than DHCPLEASEQUERY.

### 6.4.1. Determining the IP address to 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 to 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, #4). 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 DHCPLEASEKNOWN 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 which 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 which 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 DHCPLEASEKNOWN or DHCPLEASEACTIVE message, then a DHCPLEASEUNKNOWN message MUST be returned instead.

### 6.4.2. Building a DHCPLEASEKNOWN or DHCPLEASEACTIVE message once the "ciaddr" field is set

Once the "ciaddr" field of the DHCPLEASEKNOWN or DHCPLEASEACTIVE message is set, the processing for a DHCPLEASEKNOWN message is complete.

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 DHCPLEASEKNOWN or DHCPLEASEACTIVE message MUST be one for which this server is responsible (or a DHCPLEASEUNKNOWN message would be 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 which identify the client associated with the IP address in the "ciaddr" field of the DHCPLEASEKNOWN 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 SHOULD be returned in the associated-ip option (option TBD), if 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 and if the DHCP server has saved the information contained in the most recent Relay Agent Information option, the DHCP server MUST include that information in a Relay Agent Information option 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. 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 order to accommodate DHCPLEASEQUERY messages sent to a DHCP Failover secondary server [FAILOVER] when the primary server is down, the primary server MUST communicate the Relay Agent Information option (option 82) values to the secondary server via the DHCP Failover BNDUPD messages.

### 6.4.3. **Sending a DHCPLEASEKNOWN, DHCPLEASEACTIVE, or DHCPLEASEUNKNOWN** message

The server expects a giaddr in the DHCPLEASEQUERY message, and unicasts the DHCPLEASEKNOWN, 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 DHCPLEASEKNOWN, 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 DHCPLEASEKNOWN 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 cacheing could be as simple as simply setting a bit saying that a response was received from a server which 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 existance.

In either case, when a DHCPLEASEKNOWN or DHCPLEASEACTIVE message is received in response to a DHCPLEASEQUERY message, it means that the DHCP server which responded is a DHCP server which 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 which 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 any 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, and only infrequently direct a DHCPLEASEQUERY message to a DHCP server that responded to a DHCPLEASEQUERY message for a particular "ciaddr" field with a DHCPLEASEUNKNOWN.

When a DHCPUNIMPLEMENTED message is received by an access concentrator, it means that the particular aspect of DHCPLEASEQUERY processing requested is not implemented in the responding server. It may or may not be the case that other aspects of DHCPLEASEQUERY processing are not implemented in that server.

### 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:

- The DHCPLEASEQUERY or a corresponding DHCPLEASEKNOWN, DHCPLEASEACTIVE or DHCPLEASEUNKNOWN were lost during transmission or the DHCPLEASEQUERY arrived at the DHCP server but it was dropped because the server was too busy.

- 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 which utilizes the DHCPLEASEQUERY message SHOULD attempt to resend DHCPLEASEQUERY messages to servers which 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 which 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 exponetial backoff algorithm specified in RFC 2131, section 4.1 [RFC 2131].

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 message 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 RFC 2131 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 which 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 recieved.

### 6.7. 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 failover [FAILOVER] is operating), 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 gleaning, 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 which 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 which 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 leveraging DHCP authentication [RFC 3118]. With respect to authentication, the access concentrator acts as the "client". The use of "Authentication Protocol 0" (using simple unencoded authentication token(s) between the access concentrator and the DHCP server) is straightforward. Alternatively, use of IPsec would also be

a way to ensure security between the relay agent and the DHCP server.

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 cacheing (i.e. cache DHCPLEASEKNOWN, 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.

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 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 seven values for this document. See Section 6.1 for details. There are five new messages types, which are the value of the message type option (option 53) from [RFC 2132]. The value for DHCPLEASEQUERY is TBD, the value for DHCPLEASEKNOWN is TBD, the value for DHCPLEASEACTIVE is TBD, the value for DHCPLEASEUNKNOWN is TBD and the value for DHCPUNIMPLEMENTED is TBD. Finally, there are two new DHCP option defined; the client-last-transaction-time option -- option code TBD, and the associated-ip option -- option code TBD.

## 9. Acknowledgments

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.

## 10. References

### 10.1. Normative References

[RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", RFC 2119, March 1997.

[RFC 2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC 3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

## 10.2. Informative References

[RFC 826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", RFC 826, November 1982.

[RFC 951] Croft, B., Gilmore, J., "Bootstrap Protocol (BOOTP)", RFC 951, September 1985.

[RFC 1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, October 1993.

[RFC 2132] Alexander, S., Droms, R., "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

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

[BPI] CableLabs, "Baseline Privacy Interface Specification", SP-BPI-I02-990319, March 1999, available at http://www.cablemodem.com/.

[BPI+] CableLabs, "Baseline Privacy Plus Interface Specification", SP-BPI+-I04-000407, April 2000, available at http://www.cablemodem.com/.

[DHCPMIB] Hibbs, R., Waters, G., "Dynamic Host Configuration Protocol (DHCP) Server MIB", draft-ietf-dhc-server-mib-06.txt, February 2002.

[DOCSIS] CableLabs, "Data-Over-Cable Service Interface Specifications: Cable Modem Radio Frequency Interface Specification SP-RFI-I05-991105", November 1999.

[EUROMODEM] ECCA, "Technical Specification of a European Cable Modem for digital bi-directional communications via cable networks", Version 1.0, May 1999.

[FAILOVER] Droms, R., Kinnear, K., Stapp, M., Volz, B., Gonczi, S., Rabil, G., Dooley, M., Kapur, A., "DHCP Failover Protocol",

draft-ietf-dhc-failover-12.txt, March 2003.

## 11. Author's information

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: kkinnear@cisco.com

## 12. Intellectual Property Statement

Director.

## 13. Full Copyright Statement