Internet Engineering Task Force
INTERNET-DRAFT

R. Droms
Bucknell University
December, 1992

Dynamic Host Configuration Protocol

1 Abstract

The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts on a TCP/IP network.

2 Status of this memo

This draft document is a product of the IETF Dynamic Host Configuration Working Group; it will be submitted to the RFC editor as a standards document. Distribution of this memo is unlimited. It is available in both ASCII and PostScript formats. The figures are described in PostScript and are not included with the ASCII version of the document. Copies of the figures are available from the author. Please respond with comments to the `host-conf@sol.cs.bucknell.edu` mailing list. This document will expire on June 1, 1993.

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress." Please check the `1id-abstracts.txt` listing contained in the internet-drafts Shadow Directories on `nic.ddn.mil`, `nnsc.nsf.net`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `munnari.oz.au` to learn the current status of any Internet Draft.

Contents

List of Figures

List of Tables

3 Introduction

The Dynamic Host Configuration Protocol (DHCP) provides configuration parameters to Internet hosts. DHCP consists of three components: a protocol for delivering host–specific configuration parameters from a DHCP server to a host; a mechanism for allocation of network addresses to hosts; and a protocol through which a collection of DHCP servers can cooperatively allocate network addresses from a shared pool of network addresses.

DHCP is built on a client–server model, where designated DHCP server hosts allocate network addresses and deliver configuration parameters to dynamically configured hosts. Throughout the remainder of this document, the term "server" refers to a host providing initialization parameters through DHCP, and the term "client" refers to a host requesting initialization parameters from a DHCP server.

A host should not act as a DHCP server unless explicitly configured to do so by a system administrator. The diversity of hardware and protocol implementations in the Internet would preclude reliable operation if random hosts were allowed to respond to DHCP requests. For example, IP requires the setting of many parameters within the protocol implementation software. Because IP can be used on many dissimilar kinds of network hardware, values for those parameters cannot be guessed or assumed to have correct defaults. Also, distributed address allocation schemes depend on a polling/defense mechanism for discovery of addresses that are already in use. IP hosts may not always be able to defend their network addresses, so that such a distributed address allocation scheme cannot be guaranteed to avoid allocation of duplicate network addresses.

DHCP supports three mechanisms for IP address allocation. In "automatic allocation", DHCP assigns a permanent IP address to a host. In "dynamic allocation", DHCP assigns an IP address to a host for a limited period of time (or until the host explicitly relinquishes the address). In "manual allocation", a host's IP address is assigned by the network administrator, and DHCP is used simply to convey the assigned address to the host. A particular network will use one or more of these mechanisms, depending on the policies of the network administrator.

Dynamic allocation is the only one of the three mechanisms that allows automatic reuse of an address that is no longer needed by the host to which it was assigned. Thus, dynamic allocation is particularly useful for assigning an address to a host that will be connected to the network only temporarily or for sharing a limited pool of IP addresses among a group of hosts that do not need permanent IP addresses. Dynamic allocation may also be a good choice for assigning an IP address to a new host being permanently connected to a network where IP addresses are sufficiently scarce that it is important to retire them when old hosts are retired. Manual allocation allows DHCP to be used to eliminate the error–prone process of manually configuring hosts with IP addresses in environments where (for whatever reasons) it is desirable to manage IP address assignment outside of the DHCP mechanisms.

The format of DHCP messages is based on the format of BOOTP [7] messages, to capture the BOOTP relay agent behavior described as part of the BOOTP specification [7, 21] and to allow interoperability of existing BOOTP clients with DHCP servers. Using BOOTP relaying agents eliminates the necessity of having a DHCP server on each physical network segment.

3.1 Related Work

There are several Internet protocols and related mechanisms that address some parts of the dynamic host configuration problem. RARP [9] (through the extensions defined in the DRARP draft RFC [5]) explicitly addresses the problem of network address discovery, and includes an automatic IP address assignment mechanism. TFTP [20] provides for transport of a boot image from a boot server. ICMP [15] provides for informing hosts of additional routers via "ICMP redirect" messages. ICMP also can provide subnet mask information through the "ICMP mask request" message and other information through the (obsolete) "ICMP information request" message. Hosts can locate routers through the ICMP router discovery mechanism [8].

BOOTP is a transport mechanism for a collection of configuration information. BOOTP is also extensible, and official extensions [16, 17] have been defined for several configuration parameters. Morgan has proposed extensions to BOOTP for dynamic IP address assignment [14]. NIP, used by the Athena project at MIT, is a distributed mechanism for dynamic IP address assignment [19]. RLP [1] provides for location of higher level services. Sun Microsystems diskless workstations use a boot procedure that employs RARP, TFTP and an RPC mechanism called "bootparams" to deliver configuration information and operating system code to diskless hosts. (Sun Microsystems, Sun Workstation and SunOS are trademarks of Sun Microsystems, Inc.) Some Sun networks also use DRARP and an auto–installation mechanism to automate the configuration of new hosts in an existing network.

In other related work, the path MTU discovery algorithm can determine the MTU of an arbitrary internet path [13]. Comer and Droms have proposed the use of ARP as a transport protocol for resource location and selection [6]. Finally, the Host Requirements RFCs [3, 4] mention specific requirements for host reconfiguration and suggest a scenario for initial configuration of diskless hosts.

3.2 Problem definition and issues

DHCP is designed to supply hosts with the configuration parameters defined in the Host Requirements RFCs. After obtaining parameters via DHCP, a host should be able to exchange packets with any other host in the Internet. The parameters supplied by DHCP are listed in Appendix A.

Not all of these parameters are required for a newly initialized host. A client and server may

negotiate for the transmission of only those parameters required by the client or specific to a particular subnet.

DHCP allows but does not require the configuration of host parameters not directly related to the IP protocol. DHCP also does not address registration of newly configured hosts with DNS[11, 12].

DHCP is not intended for use in configuring routers.

3.3 Requirements

The following list gives general requirements for DHCP.

- DHCP should be a mechanism rather than a policy. DHCP must allow local system administrators control over configuration parameters where desired; e.g., local system administrators should be able to enforce local policies concerning allocation and access to local resources where desired.
- Hosts should require no manual configuration. Each host should be able to discover appropriate local configuration parameters without user intervention and incorporate those parameters into its own configuration.
- Networks should require no hand configuration for individual hosts. Under normal circumstances, the network manager should not have to enter any per–host configuration parameters.
- DHCP should not require a server on each subnet. To allow for scale and economy, DHCP must work across routers or through the intervention of BOOTP/DHCP relay agents.
- A DHCP host must be prepared to receive multiple responses to a request for configuration parameters. Some installations may include multiple, overlapping DHCP servers to enhance reliability and increase performance.
- DHCP must coexist with statically configured, non–participating hosts and with existing network protocol implementations.
- DHCP must interoperate with the BOOTP relay agent behavior as described by RFC 951 and by Wimer's Internet Draft.
- DHCP must interoperate with existing BOOTP clients.

The following list gives requirements specific to the transmission of the network layer parameters. DHCP must:

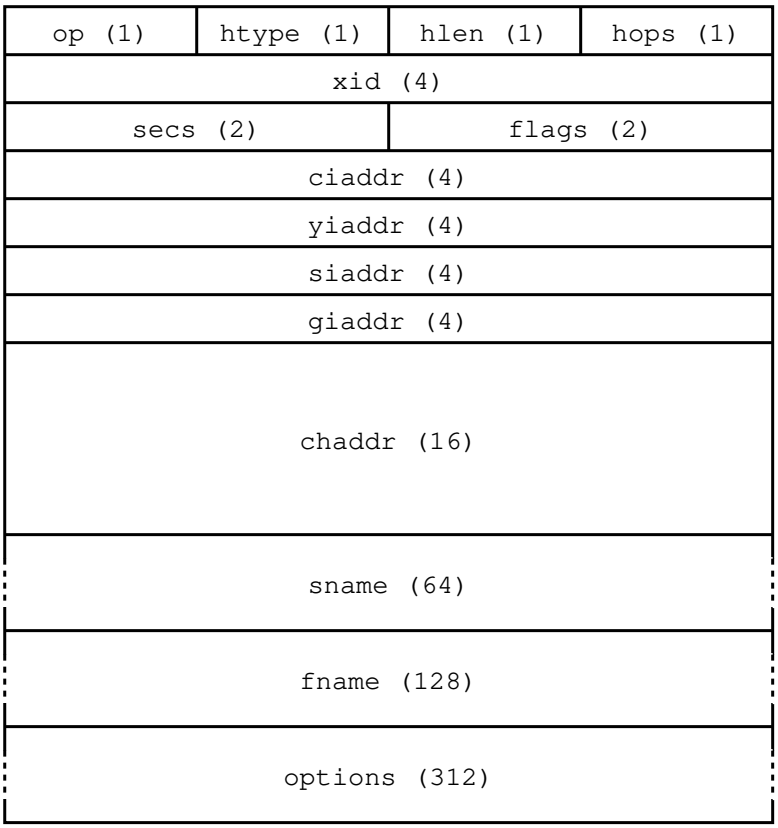

Figure 1: Format of a DHCP message

- Guarantee that any specific network address will not be in use by more than one host at a time,
- Retain host configuration across host reboot. A host should, whenever possible, be assigned the same configuration parameters (e.g., network address) in response to each request,
- Retain host configuration across server reboots, and, whenever possible, a host should be assigned the same configuration parameters despite restarts of the DHCP mechanism,
- Allow automatic assignment of configuration parameters to new hosts to avoid hand configuration for new hosts,
- Support fixed or permanent allocation of configuration parameters to specific hosts.

4 Protocol Summary

From the client's point of view, DHCP is an extension of the BOOTP mechanism. This behavior allows existing BOOTP clients to interoperate with DHCP servers without requiring any change

to the clients' initialization software. A separate document (currently an Internet Draft) details the interactions between BOOTP and DHCP clients and servers. There are some new, optional transactions that optimize the interaction between DHCP clients and servers that are described in sections 5 and 6.

Figure 1 gives the format of a DHCP message and table 1 describes each of the fields in the DHCP message. The numbers in parentheses indicate the size of each field in octets. The names for the fields given in the figure will be used throughout this document to refer to the fields in DHCP messages.

There are two primary differences between DHCP and BOOTP. First, DHCP provides the mechanism for a client to acquire all of the IP configuration parameters that it needs in order to operate. Second, DHCP defines mechanisms through which DHCP servers coordinate the dynamic allocation of network addresses to requesting clients.

DHCP extends the use of the 'htype', 'hlen' and 'chaddr' fields to allow the use of client identifiers that are not hardware addresses. Note that 'htype', 'hlen' and 'chaddr' are historic names and those fields may be used for any type of client identifier, other than just hardware addresses such as an Ethernet address. The hardware type values defined in the ARP section of the "Assigned Numbers" RFC are reserved for use when the client identifier is a hardware address. Two additional hardware type values, specifying a DNS name or a machine–specific serial number permanently stored with the client, are defined in table 1. Other client identifier types may be defined as needed for use with DHCP. New client identifier types should be registered with the IANA [18], and will be included in future revisions of the DHCP Options Internet Draft [2].

DHCP introduces a small change in terminology intended to clarify the meaning of one of the fields. What was the "vendor extensions" field in BOOTP has been re-named the "options" field in DHCP. Similarly, the tagged data items that were used inside the BOOTP "vendor extensions" field, which were formerly referred to as "vendor extensions," are now termed simply "options."

DHCP also carries forward the interpretation of the 'siaddr' field as the address of the server to use in the next step of the client's bootstrap process [21]. A DHCP server returns its own address in the 'server identifier' option.

In addition, the options field is now variable length, with the minimum extended to 312 octets. This brings the minimum size of a DHCP message up to 576 octets, the minimum IP datagram size a host must be prepared to accept [3, Sec. 3.2.2, p. 56]. DHCP clients may negotiate the use of larger DHCP messages through the 'Maximum DHCP message size' option. The options field may be further extended into the 'file' and 'sname' fields.

4.1 Components of the Protocol

DHCP provides three distinct services:

- A persistent, dynamic repository of configuration information for clients.
- Dynamic allocation of configuration resources such as network layer addresses.
- Distribution of configuration information among protocol servers.

This document will separately describe the mechanisms through which DHCP provides the first two of these services. A separate document will describe the operation of the DHCP distributed database.

4.2 Configuration parameters repository

The first service provided by DHCP is to provide persistent storage of network parameters for network clients. The model of DHCP persistent storage is that the DHCP service stores a key–value entry for each client, where the key is some unique identifier (an IP subnet number and a unique identifier within the subnet) and the value contains the configuration parameters for the client.

For example, the key might be the pair (IP–subnet–number, hardware–address), allowing for serial or concurrent reuse of a hardware address on different subnets, and for hardware addresses that may not be globally unique. Alternately, the key might be the pair (IP–subnet–number, hostname), allowing the server to assign parameters intelligently to a host that has been moved to a different subnet or has changed hardware addresses (perhaps because the network interface failed and was replaced).

A client can query the DHCP service to retrieve its configuration parameters. The client interface to the configuration parameters repository consists of protocol messages to request configuration parameters and responses from the server carrying the configuration parameters.

4.3 Dynamic allocation of network addresses

The second service provided by DHCP is the allocation of temporary or permanent network (IP) addresses to hosts. The basic mechanism for the dynamic allocation of network addresses is simple: a client requests the use of an address for some period of time. The allocation mechanism (the collection of DHCP servers) guarantees not to reallocate that address within the requested time and attempts to return the same network address each time the client requests an address. In

this document, the period over which a network address is allocated to a client is referred to as a "lease" [10]. The client may extend its lease with subsequent requests. The client may issue a message to release the address back to the server when the client no longer needs the address. The client may ask for a permanent assignment by asking for an infinite lease. Even when assigning "permanent" addresses, a server may choose to give out lengthy but non–infinite leases to allow detection of the fact that the host has been retired.

In some environments it will be necessary to reassign network addresses due to exhaustion of available addresses. In such environments, the allocation mechanism will reuse addresses whose lease has expired. The server should use whatever information is available in the configuration information repository to choose an address to reuse. For example, the server may choose the least recently assigned address. As a consistency check, the allocation mechanism may probe the reused address, e.g., with an ICMP echo request, before allocating the address, and the client will probe the newly received address, e.g., with ARP.

5 The Client–Server Protocol

DHCP uses the BOOTP message format defined in RFC 951 and given in table 1 and figure 1. The 'op' field of each DHCP message sent from a client to a server contains BOOTREQUEST. BOOTREPLY is used in the 'op' field of each DHCP message sent from a server to a client.

The first four octets of the 'options' field of the DHCP message contain the (decimal) values 99, 130, 83 and 99, respectively (this is the same magic cookie as is defined in RFC 1048). The remainder of the 'options' field consists a list of tagged parameters that are called "options". All of the "vendor extensions" listed in RFC 1048 are also DHCP options. A separate document, currently an Internet Draft, gives the complete set of options defined for use with DHCP.

Several options have been defined so far. One particular option – the "DHCP message type" option – must be included in every DHCP message. This option defines the "type" of the DHCP message. Additional options may be allowed, required, or not allowed, depending on the DHCP message type.

Throughout this document, DHCP messages that include a 'DHCP message type' option will be referred to by the type of the message; e.g., a DHCP message with 'DHCP message type' option type 1 will be referred to as a "DHCPDISCOVER" message.

5.1 Client–server interaction – allocating a network address

The following summary of the protocol exchanges between clients and servers refers to the DHCP messages described in table 2. The timeline diagram in figure 2 shows the timing relationships

FIELD	OCTETS	DESCRIPTION
op	1	Message op code / message type. 1 = BOOTREQUEST, 2 = BOOTREPLY
htype	1	Hardware address type, see ARP section in "Assigned Numbers" RFC; e.g., '1' = 10mb ethernet. In addition to ARP identifiers, '0' = "other type identifier" and '128' = DNS name.
hlen	1	Hardware address length (e.g. '6' for 10mb ethernet).
hops	1	Client sets to zero, optionally used by relay–agents when booting via a relay–agent.
xid	4	Transaction ID, a random number chosen by the client, used by the client and server to associate messages and responses between a client and a server.
secs	2	Filled in by client, seconds elapsed since client started trying to boot.
–	2	Unused.
ciaddr	4	Client IP address; filled in by client in DHCPREQUEST if unwilling to accept new IP address from DHCP server.
yiaddr	4	'your' (client) IP address; filled by server if client doesn't know its own address ('ciaddr' was 0).
siaddr	4	Server IP address; returned in DHCPOFFER, DHCPACK and DHCPNAK by server.
giaddr	4	Relay agent IP address, used in booting via a relay–agent.
chaddr	16	Client hardware address.
sname	64	Optional server host name, null terminated string.
file	128	Boot file name, null terminated string; "generic" name or null in DHCPDISCOVER, fully qualified directory–path name in DHCPOFFER.
options	312	Optional parameters field. See the options documents for a list of defined options.

Table 1: Description of fields in a DHCP message

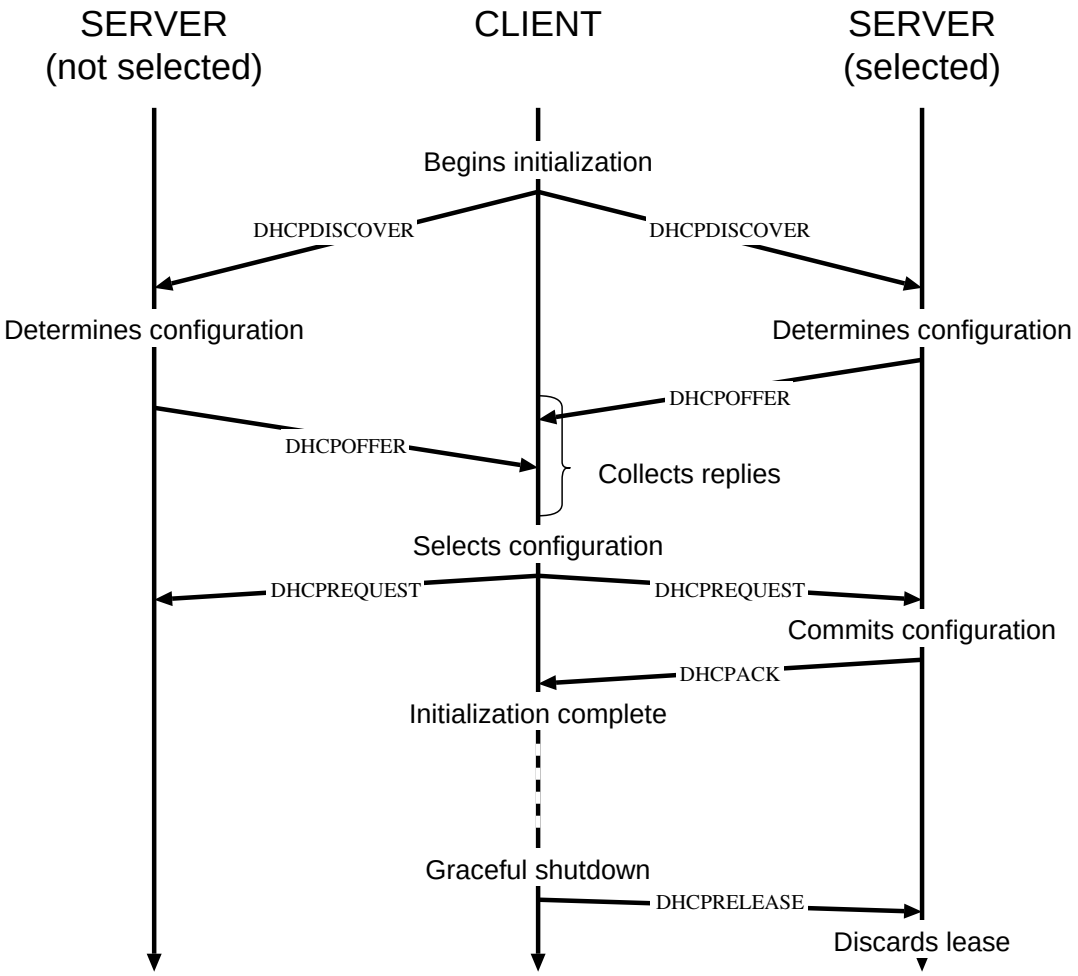

Figure 2: Timeline diagram of messages exchanged between DHCP client and servers when allocating a new network address

in a typical client–server interaction. If the client already knows its address, some steps may be omitted; this abbreviated interaction is described in section 5.2.

1. The client broadcasts a DHCPDISCOVER message on its local physical subnet. The DHCPDISCOVER message may include options that suggest values for the network address and lease duration. DHCP/BOOTP relay agents pass the message on to DHCP servers not on the same physical subnet.

2. Each server may respond with a DHCPOFFER message that includes an available network address in the 'yiaddr' field. Other configuration parameters may be returned as options. Servers need not reserve the offered network address, although the protocol will work more efficiently if the server avoids allocating the offered network address to another client. The server unicasts the DHCPOFFER message to the client (using the DHCP/BOOTP relay agent if necessary) if possible, or may broadcast the message to the all–one's broadcast address on the client's subnet.

3. The client receives one or more DHCPOFFER messages from one or more servers. The client may choose to wait for multiple responses. The client chooses one server from which to request configuration parameters, based on the configuration parameters offered in the DHCPOFFER messages. The client broadcasts a DHCPREQUEST message with the 'ciaddr' field filled in with the network address from the DHCPOFFER message sent by the selected server. The client MUST include the 'server identifier' option to indicate which server it has selected, and may include other options specifying desired configuration values. This DHCPREQUEST message is broadcast and relayed through DHCP/BOOTP relay agents. To help ensure that any DHCP/BOOTP relay agents forward the DHCPREQUEST message to the same set of DHCP servers that received the original DHCPDISCOVER message, the DHCPREQUEST message must use the same value in the DHCP message header's 'secs' field and be sent to the same IP broadcast address as the original DHCPDISCOVER message. The client times out and retransmits the DHCPDISCOVER message if the client receives no DHCPOFFER messages.

4. The servers receive the DHCPREQUEST broadcast from the client. Those servers not selected by the DHCPREQUEST message use the message as notification that the client has declined that server's offer. The server selected in the DHCPREQUEST message commits the binding for the client to persistent storage and responds with a DHCPACK message containing the configuration parameters for the requesting client. The server generates a unique value to identify the lease and places it in a 'lease identifier cookie' option included with the DHCPACK message. The 'yiaddr' field in the DHCPACK messages is filled in with the selected network address.

   If the selected server is unable to satisfy the DHCPREQUEST message (e.g., the requested network address has been allocated), the server responds with a DHCPNAK message.

Message		Use
DHCPDISCOVER	–	Client broadcast to locate available servers.
DHCPOFFER	–	Server to client in response to DHCPDISCOVER with offer of configuration parameters.
DHCPREQUEST	–	Client broadcast to servers requesting offered parameters from one server and implicitly declining offers from all others.
DHCPACK	–	Server to client with configuration parameters, including committed network address.
DHCPNAK	–	Server to client refusing request for configuration parameters (e.g., requested network address already allocated).
DHCPDECLINE	–	Client to server indicating configuration parameters (e.g., network address) invalid.
DHCPRELEASE	–	Client to server relinquishing network address and cancelling remaining lease.

Table 2: DHCP messages

A server may choose to mark addresses offered to clients in DHCPOFFER messages as unavailable. The server should mark an address offered to a client in a DHCPOFFER message as available if the server receives no DHCPREQUEST message from that client.

5. The client receives the DHCPACK message with configuration parameters. The client performs a final check on the parameters (e.g., ARP for allocated network address), and notes the duration of the lease and the lease identification cookie specified in the DHCPACK message. At this point, the client is configured. If the client detects a problem with the parameters in the DHCPACK message, the client sends a DHCPDECLINE message to the server and restarts the configuration process. The client should wait a minimum of ten seconds before restarting the configuration process to avoid excessive network traffic in case of looping.

   If the client receives a DHCPNAK message, the client restarts the configuration process.

   The client times out and retransmits the DHCPREQUEST message if the client receives neither a DHCPACK or a DHCPNAK message.

6. The client may choose to relinquish its lease on a network address by sending a DHCPRELEASE message to the server. The client identifies the lease to be released by including the 'lease identifier' option in the DHCPRELEASE message.

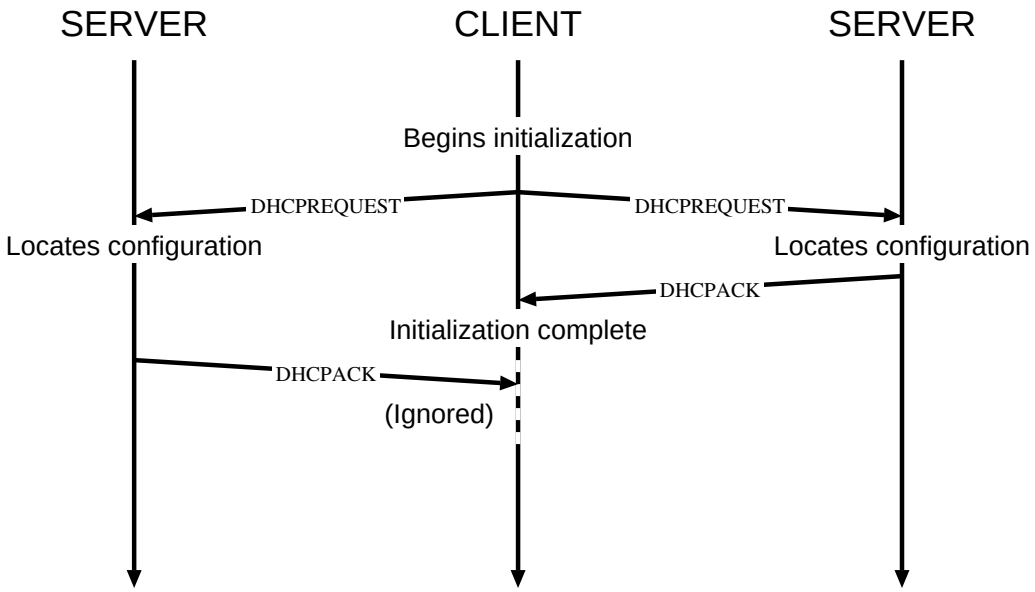

Figure 3: Timeline diagram of messages exchanged between DHCP client and servers when reusing a previously allocated network address

5.2 Client–server interaction – reusing a previously allocated network address

If a client remembers and wishes to reuse a previously allocated network address (allocated either by DHCP or some means outside the protocol), a client may choose to omit some of the steps described in the previous section. The timeline diagram in figure 3 shows the timing relationships in a typical client–server interaction for a client reusing a previously allocated network address.

1. The client broadcasts a DHCPREQUEST message on its local subnet. The DHCPREQUEST message includes the client's network address in the 'ciaddr' field and any other suggested configuration values as options. DHCP/BOOTP relay agents pass the message on to DHCP servers not on the same subnet.

2. Servers with knowledge of the client's configuration parameters respond with a DHCPACK message to the client.

   If the client's request is invalid (e.g., the client has moved to a new subnet), servers may respond with a DHCPNAK message to the client.

3. Client receives the DHCPACK message with configuration parameters. The client performs a final check on the parameters, and notes the duration of the lease and the lease identification cookie specified in the DHCPACK message. At this point, the client is configured.

   If the client detects a problem with the parameters in the DHCPACK message, the client sends a DHCPDECLINE message to the server and restarts the configuration process in INIT state,

requesting a new network address.

If the client receives a DHCPNAK message, it cannot reuse its remembered network address. It must instead request a new address by restarting the configuration process, this time using the (non–abbreviated) procedure described in section 5.1. This action corresponds to the client moving to the INIT state in the DHCP state diagram, which will be described in section 6.4.

The client times out and retransmits the DHCPREQUEST message if the client receives neither a DHCPACK or a DHCPNAK message. If the client receives no response to repeated DHCPREQUEST messages (how many?), the client restarts the configuration process in INIT state.

4. The client may choose to relinquish its lease on a network address by sending a DHCPRELEASE message to the server. The client identifies the lease to be released with the lease identification cookie.

   Note that in this case, where the client retains its network address locally, the client will not normally relinquish its lease during a graceful shutdown. Only in the case where the client explicitly needs to relinquish its lease, e.g., the client is about to be moved to a different subnet, will the client send a DHCPRELEASE message.

5.3 Interpretation and representation of time values

A client acquires a lease for a network address for a fixed period of time (which may be infinite). Throughout the protocol, times are to be represented in units of seconds. The time value of all 1s is reserved to represent "infinity". The minimum lease duration is one hour.

As clients and servers may not have synchronized clocks, times are represented in DHCP messages as relative times, to be interpreted with respect to the client's local clock. Representing relative times in units of seconds in an unsigned 32 bit word gives a range of relative times from 0 to approximately 100 years, which is sufficient for the relative times to be measured using DHCP.

The algorithm for lease duration interpretation given in the previous paragraph assumes that client and server clocks are stable relative to each other. If there is drift between the two clocks, the server may consider the lease expired before the client does. To compensate, the server may return a shorter lease duration to the client than the server commits to its local database of client information.

5.4 Host parameters in DHCP

Not all clients require initialization of all parameters listed in Appendix A. Two techniques are used to reduce the number of parameters transmitted from the server to the client. First, most

of the parameters have defaults defined in the Host Requirements RFCs; if the client receives no parameters from the server that override the defaults, a client uses those default values. Second, in its initial DHCPDISCOVER or DHCPREQUEST message, a client may provide the server with a list of specific parameters the client is interested in.

The parameters returned to a client may still exceed the space allocated to options in a DHCP message. In this case, two additional options flags (which must appear in the 'options' field of the message) indicate that the 'file' and 'sname' fields are to be used for options.

There are two ways that the client can inform the server which configuration parameters the client is interested in. First, it can include the 'parameter request vector' option in the DHCPDISCOVER or DHCPREQUEST message. In the 'parameter request vector' data, a one bit in position `n` in the vector represents an explicit request for the option parameter with tag `n`. Second, the client can include the 'parameter request list' option. The data portion of this option explicitly lists options by tag number.

In addition, the client may suggest values for the network address and lease time in the DHCPDISCOVER message. The client may include the 'requested IP address' option to suggest that a particular IP address be assigned, and may include the 'IP address lease time' option to suggest the lease time it would like. The client may include the 'maximum DHCP message size' option to let the server know how large the server may make its DHCP messages. No other options representing "hints" at configuration parameters are allowed in a DHCPDISCOVER message. The 'ciaddr' field is to be filled in

If a server receives a DHCPREQUEST message with an invalid 'ciaddr', the server responds to the client with a DHCPNAK message and may choose to report the problem to the system administrator. The server may include an error message in the 'message' option.

The client should specify the largest acceptable DHCP message with the 'DHCP message size' option to ensure that the server can transmit all the appropriate parameters in a single DHCP message.

5.5 Use of DHCP in clients with multiple interfaces

A host with multiple network interfaces must use DHCP through each interface independently to obtain configuration information parameters for those separate interfaces.

5.6 When clients should use DHCP

A host should use DHCP to reacquire or verify its IP address and network parameters whenever the local network parameters may have changed; e.g., at system boot time or after a disconnection from the local network, as the local network configuration may change without the host's or user's knowledge.

If a host has knowledge of a previous network address and is unable to contact a local DHCP server, the host may continue to use the previous network address until the lease for that address expires. If the lease expires before the host can contact a DHCP server, the host must immediately discontinue use of the previous network address and may inform local users of the problem.

6 Specification of the DHCP client–server protocol

In this section, we assume that a DHCP server has a block of network addresses from which it can satisfy requests for new addresses. Each server also maintains a database of allocated addresses and leases in local permanent storage.

6.1 Constructing and sending DHCP messages

DHCP clients and servers both construct DHCP messages by filling in fields in the fixed format section of the message and appending tagged data items in the variable length option area. The options area includes first a four–octet 'magic cookie' (which was described in section 5), followed by the options. The last option must always be the 'end' option.

DHCP uses UDP as its transport protocol. DHCP messages from a client to a server are sent to the 'DHCP server' port (67), and DHCP messages from a server to a client are sent to the 'DHCP client' port (68).

DHCP messages broadcast by a client prior to that client obtaining its IP address must have the source address field in the IP header set to 0.

If the 'giaddr' field in a DHCP message from a client is non–zero, the server sends any return messages to the 'DHCP client' port on the DHCP relaying agent whose address appears in 'giaddr'. If the 'giaddr' field is zero, the client is on the same subnet, and the server sends any return messages to either the client's network address, if that address was supplied in the 'ciaddr' field, or to the client's hardware address or to the local subnet broadcast address.

6.2 DHCP server administrative controls

DHCP servers are not required to respond to every DHCPDISCOVER and DHCPREQUEST message they receive. For example, a network administrator, to retain more precise control over the hosts attached to the network, may choose to configure DHCP servers to respond only to hosts that have been previously registered through some external mechanism. The DHCP specification describes only the interactions between clients and servers when the clients and servers choose to interact; it is beyond the scope of the DHCP specification to describe all of the administrative controls that system administrators might want to use. Specific DHCP server implementations may incorporate any controls or policies desired by a network administrator.

DHCP servers are also not required to furnish the same configuration parameters to every client on a particular network or subnet. For example, a DHCP server may return the IP address of different bootstrap servers in the 'siaddr' field depending on the type of DHCP client.

6.3 DHCP server behavior

A DHCP server processes incoming DHCP messages from a client based on the current state of the binding for that client. A DHCP server can receive the following messages from a client:

- DHCPDISCOVER
- DHCPREQUEST
- DHCPDECLINE
- DHCPRELEASE

Table 3 gives the use of the fields and options in a DHCP message by a server. The remainder of this section describes the action of the DHCP server for each possible incoming message.

6.3.1 DHCPDISCOVER message

When a server receives a DHCPDISCOVER message from a client, the server chooses a network address for the requesting client. If no address is available, the server may choose to report the problem to the system administrator and may choose to reply to the client with a DHCPNAK message. If the server chooses to respond to the client, it may include an error message in the 'message' option. If an address is available, the new address should be chosen as follows:

Field	DHCPOFFER	DHCPACK	DHCPNAK
'op'	BOOTREPLY	BOOTREPLY	BOOTREPLY
'htype'	(From "Assigned Numbers" RFC; 0 implies "other type identifier")		
'hlen'	(Hardware adress length in octets)		
'hops'	0	0	0
'xid'	'xid' from client DHCP-DISCOVER message	'xid' from client DHCP-REQUEST message	'xid' from client DHCP-REQUEST message
secs	0	0	0
'ciaddr'	0	0	0
'yiaddr'	IP address offered to client	IP address assigned to client	0
'siaddr'	IP address of server	IP address of server	IP address of server
'giaddr'	0	0	0
'chaddr'	'chaddr' from client DHCPDISCOVER message	'chaddr' from client DHCPREQUEST message	'chaddr' from client DHCPREQUEST message
'sname'	Server host name or options	Server host name or options	(unused)
'file'	Client boot file name or options	Client boot file name or options	(unused)
'options'	options	options	

Option	DHCPOFFER	DHCPACK	DHCPNAK
Requested IP address	MUST NOT	MUST NOT	MUST NOT
IP address lease time	MUST	MUST	MUST NOT
Use 'file'/'sname' fields	MAY	MAY	MUST NOT
DHCP message type	DHCPOFFER	DHCPACK	DHCPNAK
Lease identifier cookie	MUST	MUST	MUST NOT
Parameter request vector	MUST NOT	MUST NOT	MUST NOT
Parameter request list	MUST NOT	MUST NOT	MUST NOT
Message	MAY	MAY	MAY
All others	MAY	MAY	MUST NOT

Table 3: Fields and options used by DHCP servers

- The client's previous address as recorded in the client's binding, if that address is in the server's pool of available addresses and not already allocated, else
- The address requested in the 'Requested IP Address' option, if that address is valid and not already allocated, else
- A new address allocated from the server's pool of available addresses.

While not required for correct operation of DHCP, the server should arrange to avoid reusing the selected network address soon after offering the address to the client. The server may choose to record the address as offered to the client.

The server must also choose an expiration time for the lease, as follows:

- If the client has not requested a specific lease in the DHCPDISCOVER message and the client already has an assigned network address, the server returns the existing lease expiration time.
- If the client has not requested a specific lease in the DHCPDISCOVER message and the client does not have an assigned network address, the server assigns a locally configured default lease time.
- If the client has requested a specific lease in the DHCPDISCOVER message (regardless of whether the client has an assigned network address), the server may choose either to return the requested lease (if the lease is acceptable to local policy) or select another lease.

Once the network address and lease have been determined, the server constructs a DHCPOFFER message with the offered initialization parameters:

- The client's network address and subnet mask.
- The expiration time for the client's lease.
- Parameters requested by the client.
- Parameters with non–default values on the client's subnet.

The server inserts the 'xid' field from the DHCPDISCOVER message into the 'xid' field of the DHCPOFFER message and sends the DHCPOFFER message to the requesting client.

6.3.2 DHCPREQUEST message

A DHCPREQUEST message may come from a client responding to a DHCPOFFER message from a server, or from a client verifying a previously allocated IP address. If the DHCPREQUEST message contains a 'server identifier' option, the message is in response to a DHCPREQUEST message.

Consider first the case of a DHCPREQUEST message in response to a DHCPOFFER message. If the server is identified in the 'server identifier' option in the DHCPREQUEST message, the server checks to confirm that the requested parameters are acceptable. Usually, the requested parameters will match those returned to the client in the DHCPOFFER message; however, the client may choose to request a different lease duration. Also, there is no requirement that the server cache the parameters from the DHCPOFFER message. The server must simply check that the parameters requested in the DHCPREQUEST are acceptable. If the parameters are acceptable, the server records the new client binding and returns a DHCPACK message to the client.

If the requested parameters are unacceptable, e.g., the requested lease time is unacceptable to local policy, the server sends a DHCPNAK message to the client. The server may choose to return an error message in the 'message' option.

If a different server is identified in the 'server identifier' field, the client has selected a different server from which to obtain configuration parameters. The server may discard any information it may have cached about the client's request, and may free the network address that it had offered to the client.

Note that the client may choose to collect several DHCPOFFER messages and select the "best" offer. The client indicates its selection by identifying the offering server in the DHCPREQUEST message. If the client receives no acceptable offers, the client may choose to try another DHCPDISCOVER message. Therefore, the servers may not receive a specific DHCPREQUEST from which they can decide whether or not the client has accepted the offer. Because the servers have not committed any network address assignments on the basis of a DHCPOFFER, servers are free to reuse offered network addresses in response to subsequent requests. As an implementation detail, servers should try to arrange to avoid reusing offered addresses and may use an implementation–specific timeout mechanism to decide when to reuse an offered address.

In the second case, when there is no 'server identifier' option, the client is verifying a previously allocated IP address. The server checks to confirm that the requested parameters are acceptable. If the parameters specified in the DHCPREQUEST message match the previous parameters, the server returns a DHCPACK message to the requesting client. Otherwise, the server returns a DHCPNAK message to the client.

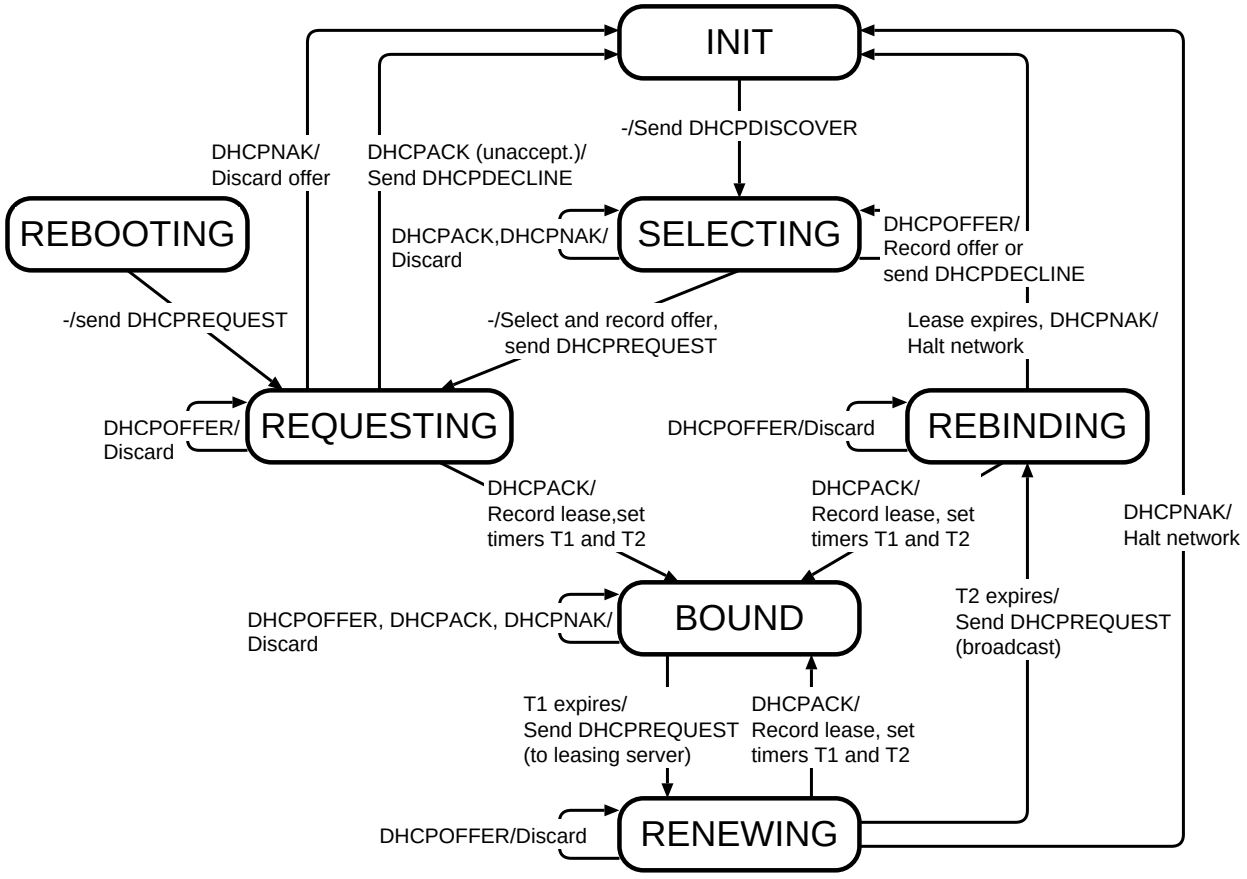

Figure 4: State–transition diagram for DHCP clients

6.3.3 DHCPDECLINE message

If the server receives a DHCPDECLINE message, the client has discovered through some other means that the suggested network address is already in use. The server marks the network address as not allocated and may notify the local system administrator of a possible configuration problem.

6.3.4 DHCPRELEASE message

Upon receipt of a DHCPRELEASE message, the server marks the network address as not allocated. The server should retain a record of the client's initialization parameters for possible reuse in response to subsequent requests from the client.

6.4 DHCP client behavior

Figure 4 gives a state–transition diagram for a DHCP client. A client can receive the following messages from a server:

- DHCPOFFER
- DHCPACK
- DHCPNAK

Table 4 gives the use of the fields and options in a DHCP message by a client. The remainder of this section describes the action of the DHCP client for each possible incoming message.

6.4.1 Initialization and allocation of network address

The client begins in INIT state and forms a DHCPDISCOVER message. The client should wait a random time between one and ten seconds to desynchronize the use of DHCP at startup. The client sets 'ciaddr' to all 0s. The client may request specific parameters by including the 'parameter request vector' or 'parameter request list' option. The client may suggest a network address and/or lease time by including the 'requested IP address' and 'IP address lease time' options. The client generates and records a random transaction identifier and inserts that identifier into the 'xid' field. The client records its own local time for later use in computing the lease expiration. The client then broadcasts the DHCPDISCOVER on the local hardware broadcast address to the all-ones IP broadcast address and 'DHCP server' UDP port.

> AUTHOR'S NOTE: The client must implement a timeout and retransmission with exponential backoff algorithm for the receipt of the DHCPOFFER message.

If the 'xid' of an arriving DHCPOFFER message does not match the 'xid' of the most recent DHCPDISCOVER message, the DHCPOFFER message is silently discarded. Any arriving DHCPACK messages are silently discarded.

The client collects DHCPOFFER messages over a period of time, selects one DHCPOFFER message from the (possibly many) incoming DHCPOFFER messages (e.g., the first DHCPOFFER message or the DHCPOFFER message from the previously used server) and extracts the server address from the DHCPOFFER message. The time over which the client collects messages and the mechanism used to select one DHCPOFFER are implementation dependent. The client may perform a check on the suggested address to ensure that the address is not already in use. For example, if the client is on a network that supports ARP, the client may issue an ARP request for the suggested request. When broadcasting an ARP request for the suggested address, the client

Field	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE, DHCPRELEASE
'op'	BOOTREQUEST	BOOTREQUEST	BOOTREQUEST
'htype'	(From "Assigned Numbers" RFC; 0 implies "other type identifier")		
'hlen'	(Hardware address length in octets)		
'hops'	0	0	0
'xid'	selected by client	selected by client	selected by client
'secs'	(opt.)	(opt.)	0
'ciaddr'	0	requested address	0
'yiaddr'	0	0	0
'siaddr'	0	0	0
'giaddr'	0	0	0
'chaddr'	client's hardware address or identifier	client's hardware address or identifier	client's hardware address or identifier
'sname'	options (opt.)	options (opt.)	(unused)
'file'	options (opt.)	options (opt.)	(unused)
'options'	options	options	(unused)

Option	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE, DHCPRELEASE
Requested IP address	MAY	MUST NOT	MUST NOT
IP address lease time	MAY	MAY	MUST NOT
Use 'file'/'sname' fields	MAY	MAY	MAY
DHCP message type	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE/ DHCPRELEASE
Lease identifier cookie	MUST NOT	MUST NOT	MUST NOT
Parameter request vector	MAY	MAY	MUST NOT
Parameter request list	MAY	MAY	MUST NOT
Message	MUST NOT	MUST NOT	MUST NOT
All others	MUST NOT	MUST NOT	MUST NOT

Table 4: Fields and options used by DHCP clients

must fill in its own hardware address as the sender's hardware address, and 0 as the sender's IP address, to avoid confusing ARP caches in other hosts on the same subnet. If the network address appears to be in use, the client sends a DHCPDECLINE message to the server and waits for another DHCPOFFER. As the client does not have a valid network address, the client must broadcast the DHCPDECLINE message.

If the parameters are acceptable, the client records the address of the server that supplied the parameters from the 'server identifier' field and sends that address in the 'server identifier' field of a DHCPREQUEST broadcast message. Once the DHCPACK message from the server arrives, the client is initialized and moves to BOUND state. The DHCPREQUEST message contains the same 'xid' as the DHCPOFFER message. The client records the lease expiration time as the sum of the time at which the original request was sent and the duration of the lease from the DHCPOFFER message.

6.4.2 Initialization with known network address

The client begins in REBOOTING state and sends a DHCPREQUEST message with the 'ciaddr' field set to the client's network address. The client may request specific configuration parameters by including the 'parameter request vector' or 'parameter request list' options. The client generates and records a random transaction identifier and inserts that identifier into the 'xid' field. The client records its own local time for later use in computing the lease expiration. The client then broadcasts the DHCPREQUEST on the local hardware broadcast address to the 'DHCP server' UDP port.

> AUTHOR'S NOTE: The client must implement a timeout and retransmission with exponential backoff algorithm for the receipt of the DHCPOFFER message.

Once a DHCPACK message with an 'xid' field matching that in the client's DHCPREQUEST message arrives from any server, the client is initialized and moves to BOUND state. The client records the lease expiration time as the sum of the time at which the DHCPREQUEST message was sent and the duration of the lease from the DHCPACK message.

6.4.3 Initialization with a known DHCP server address

When the DHCP client knows the address of a DHCP server, in either INIT or REBOOTING state, the client may use that address in the DHCPDISCOVER or DHCPREQUEST rather than the IP broadcast address. If the client receives no response to DHCP messages sent to the IP address of a known DHCP server, the DHCP client reverts to using the IP broadcast address.

6.4.4 Reacquisition and expiration

The client maintains two times, T_1 and T_2, that specify the times at which the client tries to extend its lease on its network address. T_1 is the time at which the client enters the RENEWING state and attempts to contact the server that originally issued the client's network address. T_2 is the time at which the client enters the REBINDING state and attempts to contact any server.

At time T_1 after the client accepts the lease on its network address, the client moves to RENEWING state and sends (via unicast) a DHCPREQUEST message to the server to extend its lease. The client generates a random transaction identifier and inserts that identifier into the 'xid' field in the DHCPREQUEST. The client records the local time at which the DHCPREQUEST message is sent for computation of the lease expiration time.

Any DHCPACK messages that arrive with an 'xid' that does not match the 'xid' of the client's DHCPREQUEST message are silently discarded. When the client receives a DHCPACK from the server, the client computes the lease expiration time as the sum of the time at which the client sent the DHCPREQUEST message and the duration of the lease in the DHCPACK message. The client has successfully reacquired its network address, returns to BOUND state and may continue network processing.

If no DHCPACK arrives before time T_2 ($T_2 > T_1$) before the expiration of the client's lease on its network address, the client moves to REBINDING state and sends (via broadcast) a DHCPREQUEST message to extend its lease. The client sets the 'ciaddr' field in the DHCPREQUEST to its current network address.

> AUTHOR'S NOTE: The client must implement a timeout and retransmission with exponential backoff algorithm to retry the unicast and broadcast DHCPDISCOVER messages.

Times T_1 and T_2 are configurable by the server through options. T_1 defaults to (0.5 * duration_of_lease). T_2 defaults to (0.875 * duration_of_lease). Times T_1 and T_2 should be chosen with some random "fuzz" around a fixed value, to avoid synchronization of client reacquisition.

If the lease expires before the client receives a DHCPACK, the client moves to INIT state, must immediately stop any other network processing and request network initialization parameters as if the client were uninitialized. If the client then receives a DHCPACK allocating that client its previous network address, the client may continue network processing. If the client is given a new network address, it may not continue using the previous network address and must notify the local users of the problem.

6.4.5 DHCPRELEASE

If the client no longer requires use of its assigned network address (e.g., the client is gracefully shut down), the client sends a DHCPRELEASE message to the server. Note that the correct operation of DHCP does not depend on the transmission of DHCPRELEASE messages.

7 Security Considerations

This document does not address security issues.

8 Acknowledgments

Greg Minshall, Leo McLaughlin and John Veizades have patiently contributed to the the design of DHCP through innumerable discussions, meetings and mail conversations. Jeff Mogul first proposed the client–server based model for DHCP. Steve Deering searched the various IP RFCs to put together the list of network parameters supplied by DHCP. Walt Wimer contributed a wealth of practical experience with BOOTP and wrote a document clarifying the behavior of BOOTP/DHCP relay agents. Jesse Walker analyzed DHCP in detail, pointing out several inconsistencies in earlier specifications of the protocol. Steve Alexander reviewed Walker's analysis and the fixes to the protocol based on Walker's work. And, of course, all the members of the Dynamic Host Configuration Working Group of the IETF have contributed to the design of the protocol through discussion and review of the protocol design.

9 Author's Address

Ralph Droms
Computer Science Department
323 Dana Engineering
Bucknell University
Lewisburg, PA 17837

(717) 524–1145

droms@bucknell.edu

10 References

This document references several other Internet Drafts. According to the IETF policy stated in section 2, Internet Drafts should not be used as reference material. Any references to Internet Drafts will be deleted or changed to reference the appropriate RFCs before this document is published as an RFC.

11 Expiration date

This document will expire on June 1, 1993.

References

[1] M. Acetta. Resource Location Protocol. RFC 887, NIC, December 1983.

[2] Steve Alexander and Ralph Droms. DHCP Options and BOOTP Vendor Extensions. Internet Draft, NIC, November 1992.

[3] R. Braden (Ed.). Requirements for Internet Hosts – Communication Layers. RFC 1122, NIC, October 1989.

[4] R. Braden (Ed.). Requirements for Internet Hosts – Application and Support. RFC 1123, NIC, October 1989.

[5] David Brownell. Dynamic Reverse Address Resolution Protocol (DRARP). RFC DRAFT, NIC, 1989.

[6] D. Comer and R. Droms. Uniform Access to Internet Directory Services. *Proc. of ACM SIGCOMM '90* (Special issue of Computer Communications Review), 20(4):50–59, 1990.

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

[8] S. Deering. ICMP Router Discovery Messages. RFC 1256, NIC, September 1991.

[9] R. Finlayson, T. Mann, J. Mogul, and M. Theimer. A Reverse Address Resolution Protocol. RFC 903, NIC, June 1984.

[10] C. G. Gray and D. R. Cheriton. Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency. In *Proc. of the Twelfth ACM Symposium on Operating Systems Design*, 1989.

[11] P. Mockapetris. Domain Names – Concepts and Facilities. RFC 1034, NIC, November 1987.

[12] P. Mockapetris. Domain Names – Implementation and Specification. RFC 1035, NIC, November 1987.

[13] J. Mogul and S. Deering. Path MTU Discovery. RFC 1191, NIC, November 1990.

[14] R. L. Morgan. Dynamic IP Address Assignment for Ethernet Attached Hosts. RFC DRAFT, NIC, 1989.

[15] J. Postel. Internet Control Message Protocol. RFC 792, NIC, September 1981.

[16] P. Prindeville. BOOTP Vendor Information Extensions. RFC 1048, NIC, February 1988.

[17] J. Reynolds. BOOTP Vendor Information Extensions. RFC 1084, NIC, December 1988.

[18] J. Reynolds and J. Postel. Assigned Numbers. RFC 1340, NIC, July 1992.

[19] Jeffrey Schiller and Mark Rosenstein. A Protocol for the Dynamic Assignment of IP Addresses for use on an Ethernet. (Available from the Athena Project, MIT), 1989.

[20] K. Sollins. The TFTP Protocol (Revision 2). RFC 783, NIC, June 1981.

[21] W. Wimer. Clarifications and Extensions for the Bootstrap Protocol. Internet Draft, NIC, 1991.

A Host Configuration Parameters

IP–layer parameters, per host:		
Be a router	on/off	HRC 3.1
Non–local source routing	on/off	HRC 3.3.5
Policy filters for non–local source routing	(list)	HRC 3.3.5
Maximum reassembly size	integer	HRC 3.3.2
Default TTL	integer	HRC 3.2.1.7
PMTU aging timeout	integer	MTU 6.6
MTU plateau table	(list)	MTU 7
IP–layer parameters, per interface:		
IP address	(address)	HRC 3.3.1.6
Subnet mask	(address mask)	HRC 3.3.1.6
MTU	integer	HRC 3.3.3
All–subnets–MTU	on/off	HRC 3.3.3
Broadcast address flavor	all 0s/all 1s	HRC 3.3.6
Perform mask discovery	on/off	HRC 3.2.2.9
Be a mask supplier	on/off	HRC 3.2.2.9
Perform router discovery	on/off	RD 5.1
Router solicitation address	(address)	RD 5.1
Default routers, list of:		
router address	(address)	HRC 3.3.1.6
preference level	integer	HRC 3.3.1.6
Static routes, list of:		
destination	(host/subnet/net)	HRC 3.3.1.2
destination mask	(address mask)	HRC 3.3.1.2
type–of–service	integer	HRC 3.3.1.2
first–hop router	(address)	HRC 3.3.1.2
ignore redirects	on/off	HRC 3.3.1.2
PMTU	integer	MTU 6.6
perform PMTU discovery	on/off	MTU 6.6
Link–layer parameters, per interface:		
Trailers	on/off	HRC 2.3.1
ARP cache timeout	integer	HRC 2.3.2.1
Ethernet encapsulation	(RFC 894/RFC 1042)	HRC 2.3.3
TCP parameters, per host:		
TTL	integer	HRC 4.2.2.19
Keep–alive interval	integer	HRC 4.2.3.6
Keep–alive data size	0/1	HRC 4.2.3.6

Key:

HRC = Requirements for Internet Hosts – Communication Layers (RFC 1122)

MTU = Path MTU Discovery (RFC 1191, Proposed Standard)

RD = Router Discovery (RFC 1256, Proposed Standard)