MALLOC working group
Internet Draft
November 1998
Expires: May 1999

Baiju V. Patel, Intel Corp.
Munil Shah, Microsoft Corp.
Stephen R. Hanna, Sun Microsystems, Inc.

draft-ietf-malloc-mdhcp-01.txt

Multicast address allocation based on the Dynamic Host Configuration Protocol

Status of this Memo

Abstract

This document defines a protocol, MDHCP, that allows hosts to request multicast addresses from multicast address allocation servers. MDHCP is similar to DHCP, but not dependent on it.

1. Introduction

Multicast address allocation based on the Dynamic Host Configuration Protocol (MDHCP) is a protocol similar to DHCP ([1], [2]) that allows hosts to request multicast address allocation services from multicast address allocation servers. This protocol is part of the Multicast Address Allocation Architecture defined in [5]. However, it may be used separately from the rest of that architecture as appropriate.

1.1. Protocol Overview

MDHCP is built on a client-server model, where hosts request address allocation services from address allocation servers. See Appendix A for examples of typical protocol exchanges.

1.2. 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.

Throughout the rest of this document, the words "server" or "MDHCP server" refer to a host providing multicast address allocation services via MDHCP. The words "client" or "MDHCP client" refer to a host requesting multicast address allocation services via MDHCP.

1.3. Motivation and Protocol Requirements

For multicast applications to be deployed everywhere, there is a need to define a protocol that any host may use to allocate multicast addresses. Here are the requirements for such a protocol.

Quick response: The host should be able to allocate a multicast address and begin to use it promptly.

Low network load: Hosts that are not allocating or deallocating multicast addresses at the present time should not need to send or receive any network traffic.

Support for intermittently connected or power managed systems: Hosts should be able to be disconnected from the network, powered off, or otherwise inaccessible except during the brief period during which they are allocating a multicast address.

Multicast address scopes: The protocol must be able to allocate both the administratively scoped and globally scoped multicast addresses.

Efficient use of address space: The multicast address space is fairly small. The protocol should make efficient use of this scarce resource.

Authentication: Because multicast addresses are scarce, it is important to protect against hoarding of these addresses. One way to do this is by authenticating clients.

Policy neutral: Allocation policies (such as who can allocate addresses) should not be dictated by the protocol.

1.4. Relationship with DHCP

In order to allow code reuse, MDHCP is based on a subset of DHCP. However, it has been carefully designed to ensure that there are no dependencies or interactions between the two protocols. MDHCP may be deployed without concern for impacts on existing DHCP servers or clients.

As stated above, MDHCP is based on a subset of DHCP. The message format and behavior of the protocols are similar, but there are differences. This specification has been designed to stand on its own, independent of the DHCP specifications. Implementers of MDHCP do not need to read the DHCP specifications, although they may find them useful.

Where there are conflicts between the MDHCP and DHCP specifications (and there are several), the MDHCP specifications apply to MDHCP and the DHCP specifications apply to DHCP. Remember, the protocols are similar but independent.

2. Protocol Description

The MDHCP protocol is a client-server protocol. In general, the client unicasts or multicasts a message to one or more servers, which optionally respond with messages unicast to the client.

Messages are always sent via UDP. A reserved port number dedicated for MDHCP is used on the server (port number 2535, as assigned by IANA). Any port number may be used on client machines. When an MDHCP server sends a message to an MDHCP client, it MUST use a destination port number that matches the source port number provided by the client in the message that caused the server to send its message.

Like DHCP, MDHCP is a mechanism rather than a policy. MDHCP allows local system administrators to exercise control over configuration parameters where desired. For example, MDHCP servers may be configured to limit the number of multicast addresses allocated to a single client. Properly enforcing such a limit requires cryptographic security, which will be addressed in a supplementary document.

All MDHCP messages have the same format. This format is similar to that of a DHCP message. However, many of the fields are unused. Each message includes a message type and a type-length-value encoded options field.

The next few sections describe the MDHCP message format and message types. A full list of MDHCP options is provided in section 3.

2.1. Message Format

The format of an MDHCP message is similar to that of a DHCP message. However, many of the fields are unused.

Figure 1 gives the format of an MDHCP message and Table 1 describes each of the fields in the MDHCP 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 MDHCP messages.

Any message whose UDP data is too short to hold this format (at least 32 bytes) MUST be ignored.

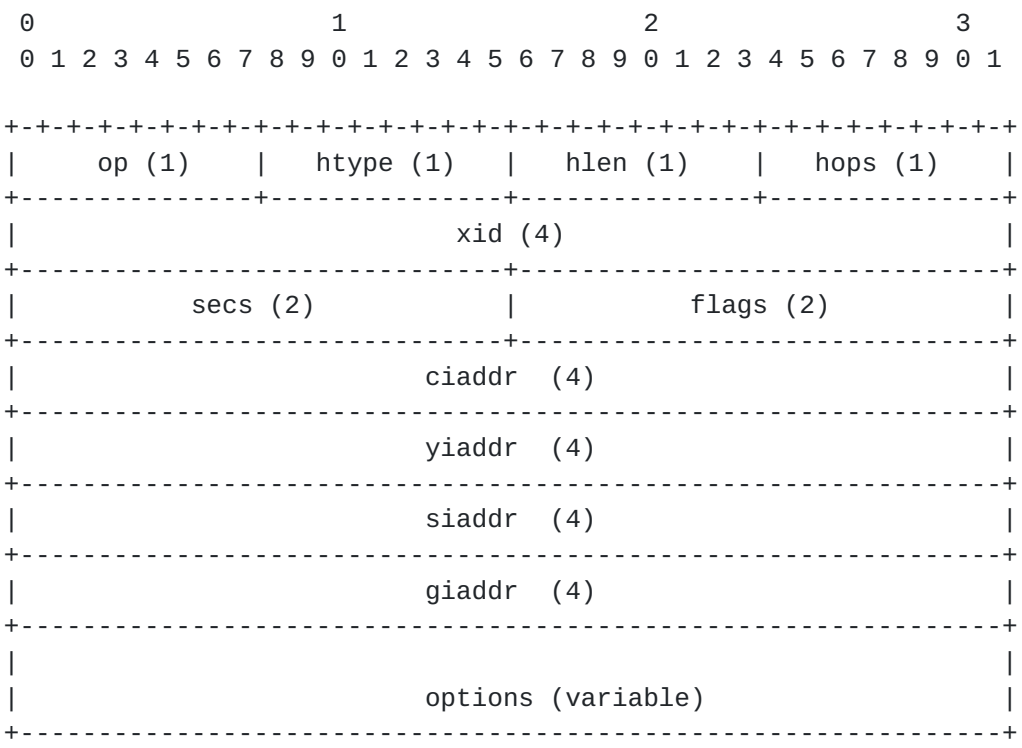

Figure 1: Format of an MDHCP message

FIELD	OCTETS	DESCRIPTION
op	1	Message op code / message type. 1 = BOOTREQUEST, 2 = BOOTREPLY
htype	1	Should be zero. Ignored.
hlen	1	Should be zero. Ignored.
hops	1	Must be zero.
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	Must be zero.
flags	2	Must be 64 (decimal).
ciaddr	4	Must be zero.
yiaddr	4	First allocated multicast address (must be zero for Messages from clients).
siaddr	4	Must be zero.
giaddr	4	Must be zero.
options	var	Optional parameters field. See section 3 for a list of defined options.

Table 1: Description of fields in an MDHCP message

2.1.1. MDHCP Message Fields

All multi-octet quantities are in network byte-order.

The op field of each MDHCP message sent from a client to a server MUST contain BOOTREQUEST (1). BOOTREPLY MUST be used in the op field of each MDHCP message sent from a server to a client.

The htype and hlen fields SHOULD be zero. They SHOULD be ignored by MDHCP clients and servers.

The hops, secs, ciaddr, siaddr, and giaddr fields MUST have the value zero (0). Messages containing any other value MUST be ignored.

The flags field MUST have the value sixty-four decimal (64). Messages containing any other value MUST be ignored.

The yiaddr field MUST be set to zero (0) by MDHCP clients. Servers use this field for returning the first allocated Multicast address, as described in section 2.5.

2.1.2. The options field

The first four octets of the 'options' field of an MDHCP message MUST have the (decimal) values 99, 130, 83 and 99, respectively. This sequence is known as the "magic cookie".

The remainder of the 'options' field consists of a list of tagged parameters that are called "options". Options may be fixed length or variable length. All options begin with a tag (or 'code') octet, which uniquely identifies the option. Fixed-length options without data consist of only a tag octet. Only options 0 and 255 are fixed length. All other options are variable-length with a length octet

following the tag octet. The value of the length octet does not include the two octets specifying the tag and length. The length octet is followed by "length" octets of data. In the case of some variable-length options, the length field is a constant but must still be specified. Any options defined subsequent to this document MUST contain a length octet even if the length is fixed or zero.

The option field MUST contain the magic cookie defined above, followed by any number of options that are not an end option, and ending with an end option (code 255) followed by any number of pad options (code 0). Any message whose options field does not conform to this syntax MUST be ignored.

Anyone sending an MDHCP message SHOULD include only options listed in section 3, but MAY include other MDHCP options that are defined in the future. Anyone receiving an MDHCP message MUST ignore unrecognized options. New MDHCP options may only be defined by submitting a standards track RFC.

2.2. Message Types

The "MDHCP message type" option MUST be included in every MDHCP message. This option defines the "type" of the MDHCP message.

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

Here are descriptions of the message types a client may send and the way a server should respond. Table 2, which appears at the end of this section, summarizes which options are allowed with each message type.

2.2.1. MDHCPDISCOVER

The MDHCPDISCOVER message is used by a MDHCP client that wants to discover MDHCP servers that can probably satisfy a request. MDHCP clients MAY employ other methods to find MDHCP servers, such as caching an IP address that worked in the past or obtaining a DNS name or IP address from DHCP or prior configuration. Using the MDHCPDISCOVER message has the particular advantage that it allows clients to automatically move to another server if one fails.

The MDHCP client begins by sending a multicast MDHCPDISCOVER message to an MDHCP server multicast address. Any servers that wish to assist the client respond by sending a unicast MDHCPOFFER message to the client. If a server can process the request with a shorter lease time

or later start time than the client requested, it MAY send an MDHCPOFFER message with the lease time that it can offer.

After a suitable delay, the client selects the server it wants to use and moves into the request phase. The time over which the client collects messages and the mechanism used to select one MDHCPOFFER are implementation dependent. If no MDHCPOFFER messages are received after an appropriate delay, the client SHOULD resend its MDHCPDISCOVER message.

For more details about the MDHCP Server Multicast Address, see section 2.9.

2.2.2. MDHCPREQUEST

The MDHCPREQUEST message is used by an MDHCP client that wants to allocate or extend the lease of a multicast address.

The MDHCP client sends out an MDHCPREQUEST message. If this request was previously sent as an MDHCPDISCOVER message, the MDHCPREQUEST message SHOULD be multicast to the MDHCP server multicast address so that all MDHCP servers know which server was selected. Otherwise, the MDHCPREQUEST message SHOULD be unicast to the MDHCP server that the client wants to use.

If the selected server can process the request successfully, it SHOULD unicast an MDHCPACK message to the client. Otherwise, it SHOULD unicast an MDHCPNAK to the client. If a server can process the request with a shorter lease time or later start time than the client requested, it MAY send an MDHCPACK message with the lease time that it can offer.

If the server responds with an MDHCPNAK or fails to respond within a reasonable (implementation-dependent) delay, the client MAY try to find another server by sending an MDHCPDISCOVER request with another xid.

2.2.3 MDHCPRELEASE

If a client wants to deallocate a multicast address before its lease expires, the client unicasts an MDHCPRELEASE message to the server from which it allocated the address. The server does not respond to this message.

2.2.2. MDHCPINFORM

The MDHCPINFORM message is used by an MDHCP client that wants to acquire configuration parameters.

The MDHCP client sends out an MDHCPINFORM message. The message may be unicast to a particular MDHCP server or multicast to the MDHCP server multicast address.

If a server receives an MDHCPINFORM message and it can process the request successfully, it SHOULD unicast an MDHCPACK message to the client. The MDHCPACK message SHOULD include the Multicast Scope List option and MAY include the Current Time option. Otherwise, it SHOULD ignore the MDHCPINFORM message.

If no MDHCPACK messages are received after an appropriate delay, the client may resend its MDHCPINFORM message to the MDHCP server multicast address.

2.2.4. Options Allowed

Table 2 summarizes which options are allowed with each message type.

Option	MDHCPOFFER	MDHCPACK	MDHCPNAK
Requested IP Address	MUST NOT	MUST NOT	MUST NOT
IP address lease time	MUST	MUST	MUST NOT
MDHCP Message Type	MUST	MUST	MUST
Server Identifier	MUST	MAY	MAY
Client identifier	MUST	MUST	MUST
Multicast Scope	MUST	MUST	MUST NOT
Start Time	MAY	MAY	MUST NOT
Multicast TTL	MAY	MAY	MUST NOT
Number of Addresses Requested	MAY	MAY	MUST NOT
Requested Language	MUST NOT	MUST NOT	MUST NOT
Multicast Scope List	MAY	MAY	MUST NOT
List of Address Ranges Allocated	MAY	MAY	MUST NOT
Current Time	MAY	MAY	MUST NOT

Option	MDHCPDISCOVER	MDHCPREQUEST	MDHCPRELEASE
Requested IP Address	MAY	MAY	MUST
IP address lease time	MAY	MAY	MUST NOT
MDHCP Message Type	MUST	MUST	MUST
Server Identifier	MUST NOT	MUST (if multicast)	MUST NOT
Client identifier	MUST	MUST	MUST
Multicast Scope	SHOULD	SHOULD	MUST NOT
Start Time	MAY	MAY	MUST NOT
Multicast TTL	MUST NOT	MUST NOT	MUST NOT
Number of Addresses			

Requested	MAY	MAY	MAY
Requested Language	MAY	MAY	MUST NOT
Multicast Scope List	MAY	MAY	MUST NOT
List of Address Ranges			
Allocated	MAY	MAY	MAY
Current Time	MAY	MAY	MUST NOT

Option	MDHCPINFORM
Requested IP Address	MUST NOT
IP address lease time	MUST NOT
MDHCP Message Type	MUST
Server Identifier	MUST NOT
Client identifier	MUST
Multicast Scope	MUST NOT
Start time	MUST NOT
Multicast TTL	MUST NOT
Number of Addresses	
Requested	MUST NOT
Requested Language	MAY
Multicast Scope List	MUST NOT
List of Address Ranges	
Allocated	MUST NOT
Current Time	MUST NOT

Table 2: Options allowed in MDHCP messages

2.3. Retransmission

MDHCP clients are responsible for all message retransmission. The client MUST adopt a retransmission strategy that incorporates a randomized exponential backoff algorithm to determine the delay between retransmissions. The delay between retransmissions SHOULD be chosen to allow sufficient time for replies from the server to be delivered based on the characteristics of the internetwork between the client and the server. For example, in a 10Mb/sec Ethernet internetwork, the delay before the first retransmission SHOULD be 4 seconds randomized by the value of a uniform random number chosen from the range -1 to +1. Clients with clocks that provide resolution granularity of less than one second may choose a non-integer randomization value. The delay before the next retransmission SHOULD be 8 seconds randomized by the value of a uniform number chosen from the range -1 to +1. The retransmission delay SHOULD be doubled with subsequent retransmissions up to a maximum of 64 seconds. The client MAY provide an indication of retransmission attempts to the user as an indication of the progress of the process. The client MAY halt retransmission at any point.

2.4. Associating Client and Server Messages

Messages between clients and servers are associated with one another using the client identifier option and xid field. Each client MUST choose a client identifier that is unique within a multicast address allocation domain. For each transaction initiated by a client, the client MUST generate an xid value that is unique for that client identifier and likely to be unique across all client identifiers. For instance, a client might start with a random xid and increment from there. The client identifier option and xid field MUST be included in each message sent by the client or the server.

The client MUST check the client identifier option and xid field in each incoming message to ensure that they match its client identifier and an outstanding transaction. If not, the message MUST be discarded. The server MUST check the client identifier option and xid field in each incoming message to establish the proper context for the message. If the message is inappropriate for its context, it SHOULD be discarded.

A transaction can be an attempt to allocate a multicast address (consisting of MDHCPDISCOVER, MDHCPOFFER, MDHCPREQUEST, MDHCPACK, and MDHCPNAK messages), an attempt to extend a lease (consisting of MDHCPREQUEST, MDHCPACK, and MDHCPNAK messages), an attempt to release a previously allocated multicast address (consisting of a single MDHCPRELEASE message), or an attempt to acquire configuration parameters (consisting of MDHCPINFORM and MDHCPACK messages).

2.5. Allocating Multiple Addresses

An MDHCP client may request the allocation of more than one multicast address in a single request by including the Number of Addresses Requested option in the MDHCPDISCOVER and MDHCPREQUEST messages. An MDHCP server may include this option in an MDHCPOFFER or MDHCPACK message to indicate its willingness to supply more than one address. Finally, an MDHCP client may include this option in an MDHCPRELEASE message to release a set of addresses or in an MDHCPREQUEST message to renewing a lease for a set of addresses.

When the Number of Addresses Requested option is included in an MDHCPOFFER or MDHCPACK message, it MUST be accompanied by a List of Address Ranges Allocated option listing the address ranges offered or allocated. This is in addition to the normal requirement that the yiaddr field be set to the first multicast address allocated.

When the Number of Addresses Requested option is included in an MDHCPREQUEST message for the purposes of renewing a lease for a set of addresses or in an MDHCPRELEASE message, it MUST be accompanied by

a List of Address Ranges Allocated option listing the address ranges affected. This is in addition to the normal requirement that the Requested IP Address option be used to specify the first multicast address allocated.

2.6. Multicast Scopes

RFC 2365 [3] provides for dividing the multicast address space into a number of administratively scopes. Routers should be configured so that each scope corresponds to a particular partition of the network into disjoint regions. Messages sent to a multicast address that falls within a certain administrative scope should only be delivered to hosts that have joined that multicast group *and* fall within the same region as the sender. For instance, packets sent to an address in the organization-local scope should only be delivered to hosts that have joined that group and fall within the same organization as the sender.

Different sets of scopes may be in effect at different places in the network and at different times. Before attempting to allocate an address from an administrative scope (other than global or link-level scope, which are always in effect), an MDHCP client SHOULD determine that the scope is in effect at its location at this time. Several techniques that an MDHCP client may use to determine the set of administrative scopes in effect (the scope list) are: manual configuration, configuration via MDHCP (using the Multicast Scope List option), or listening to MZAP messages [6].

If an MDHCP client is unable to determine its scope list using one of these techniques, it MAY temporarily assume a scope list consisting of IPv4 Local Scope and IPv4 Allocation Scope, both with a maximum TTL of 16. Using this temporary scope list, it MAY attempt to contact an MDHCP server that can provide a scope list for it.

When an MDHCP client requests an address, it SHOULD specify the administrative scope from which the address should be allocated. This scope is indicated with the Multicast Scope option. If no scope is specified, the server MAY allocate an address from some default scope or refuse to allocate an address. In any case, the server MUST include the Multicast Scope option in all MDHCPOFFER and MDHCPACK messages.

2.7. Multicast TTL

Another way to limit propagation of multicast messages is by setting the TTL field before sending them. This technique has several disadvantages in comparison to administratively scoped multicast addresses, but it is currently in widespread usage.

An MDHCP client MAY request a multicast address for use with a specific TTL by including a Multicast TTL option in an MDHCPDISCOVER or MDHCPREQUEST message. The server SHOULD respond to this option by returning only addresses which support the specified TTL (or greater). If the client does not include this option, the server MUST assume that 255 is indicated.

When sending an MDHCPOFFER or MDHCPACK message with a multicast address in it, an MDHCP server MUST indicate the maximum TTL that may be used with the address by including a Multicast TTL option. This is known as the "maximum TTL" associated with that address.

When using a multicast address allocated with MDHCP, all hosts SHOULD NOT set the TTL field to a number larger than the maximum TTL associated with that address.

2.8. Locating MDHCP Servers

There are several ways for an MDHCP client to locate an MDHCP server. For instance, the client may obtain a DNS name or IP address from DHCP or manual configuration.

One particularly convenient technique is for the client to send an MDHCPDISCOVER message to an MDHCP Server Multicast Address and wait for MDHCPOFFER responses. This technique is described in more detail in the next section.

2.9. MDHCP Server Multicast Address

Each multicast scope has an associated MDHCP Server Multicast Address. This address has been reserved by the IANA as the address with a relative offset of -1 from the last address of a multicast scope.

An MDHCP client looking for servers that can provide multicast allocation services MAY send an MDHCPDISCOVER message to an MDHCP Server Multicast Address. Any MDHCP servers listening to this address SHOULD respond with a unicast MDHCPOFFER message to the client if they wish to offer a response. Clients may also send MDHCPINFORM messages in the same manner, with servers responding with MDHCPACK messages if they can supply the requested information.

If a client receives no response to a message sent to an MDHCP Server Multicast Address (after retransmission), it MAY send the message to a larger scope and repeat this process as necessary. However, the client MUST NOT send an MDHCP message to the MDHCP Server Multicast Address associated with the global scope.

The MDHCP Server Multicast Address used by a client MAY be established by configuration (manually or via DHCP). If a client has no such configuration, it SHOULD start with the MDHCP Server Multicast Address associated with IPv4 Local Scope. If no response is received on this address, it SHOULD try the MDHCP Server Multicast Address associated with the scope from which it is trying to allocate an address, if the scope is a "small" scope. If this does not succeed, it SHOULD try the MDHCP Server Multicast Address associated with Allocation Scope.

This technique allows MDHCP servers to provide services for scopes in which they do not reside. However, such servers MUST make special efforts to ensure that their services meet clients' needs for largely conflict-free allocation and accurate scope list information. In particular, AAP traffic does not extend outside of the scope that is being managed so coordinating with other servers may be difficult. Also, establishing which scope the client is in may be difficult. If an MDHCP server is not prepared to provide services for scopes in which it does not reside, it SHOULD ignore requests whose scope does not match (or is enclosed by) the scope of the MDHCP Server Multicast Address on which the request was received.

2.10. Clock Skew

The Current Time option is used to detect and handle clock skew between MDHCP clients and servers. This option SHOULD be included in any MDHCP message that includes an absolute time (such as the Start Time option). It MAY be included in any MDHCPDISCOVER, MDHCPOFFER, MDHCPREQUEST, or MDHCPACK message.

Clock skew is a situation where two systems have clocks that are not synchronized. MDHCP servers SHOULD expect and tolerate a small amount of clock skew with their clients (such as an hour) by ensuring that multicast addresses are allocated for an extra hour or so on either side of the lease given to the client. However, large amounts of clock skew require special handling.

The Current Time option contains the sender's opinion of the current time in UTC at or about the time the message was assembled. Because of delays in transmission and processing, this value will rarely match the receiver's opinion of the current time at the time the option is processed by the receiver.

If an MDHCP server detects substantial clock skew, it SHOULD ignore the client. The server MAY log a message.

2.11. Using MDHCP Without Administrative Scoping

MDHCP can be used in an environment that does not have administrative scoping enabled. In such an environment, TTL scoping SHOULD be used. The Multicast Scope List option MAY be used to distribute information about TTL scopes that are enforced with TTL thresholds. MDHCP clients MAY include the Multicast TTL option in MDHCPDISCOVER or MDHCPREQUEST packets. MDHCP servers MUST include the Multicast TTL option in MDHCPOFFER and MDHCPACK packets when the TTL to be used is less than 255.

Multicast MDHCPDISCOVER, MDHCPREQUEST, and MDHCPINFORM packets must be carefully managed in a TTL scoped environment to avoid flooding them around the world. MDHCP clients MAY be configured (manually or via DHCP) with an MDHCP Server Multicast Address and appropriate TTL. Alternatively, the MDHCP Server Multicast Addresses associated with the IPv4 Local Scope or IPv4 Allocation Scope may be blocked so that packets sent to them with TTL of 16 are not sent outside a certain boundary. In this case, MDHCP clients need not be configured manually or via DHCP, since their default behavior in choosing an MDHCP Server Multicast Address without configuration is to choose the MDHCP Server Multicast Address associated with the IPv4 Local Scope and, if this fails, IPv4 Allocation Scope, using TTL no more than 16.

3. MDHCP Options

The following options are defined for use in MDHCP messages. The options are listed in numerical order.

3.1. Pad Option

The pad option can be used to cause subsequent fields to align on word boundaries.

The code for the pad option is 0, and its length is 1 octet.

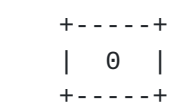

3.2. Requested IP Address

This option is used in a client request (MDHCPDISCOVER or MDHCPREQUEST) to allow the client to request that a particular multicast address be assigned.

The code for this option is 50, and its length is 4.

```
 Code   Len          Address
+-----+-----+-----+-----+-----+-----+
|  50 |  4  |  a1 |  a2 |  a3 |  a4 |
+-----+-----+-----+-----+-----+-----+
```

3.3. IP Address Lease Time

This option is used in a client request (MDHCPDISCOVER or MDHCPREQUEST) to allow the client to request a lease time for the multicast address. In a server reply (MDHCPOFFER), an MDHCP server uses this option to specify the lease time it is willing to offer.

The time is in units of seconds, and is specified as a 32-bit unsigned integer.

The code for this option is 51, and its length is 4.

```
 Code   Len         Lease Time
+-----+-----+-----+-----+-----+-----+
|  51 |  4  |  t1 |  t2 |  t3 |  t4 |
+-----+-----+-----+-----+-----+-----+
```

3.4. MDHCP Message Type

This option is used to convey the type of the MDHCP message. The code for this option is 53, and its length is 1. Legal values for this option are:

Value	Message Type
1	MDHCPDISCOVER
2	MDHCPOFFER
3	MDHCPREQUEST
4	reserved
5	MDHCPACK
6	MDHCPNAK
7	MDHCPRELEASE
8	MDHCPINFORM

```
 Code   Len  Type
+-----+-----+-----+
|  53 |  1  | 1-8 |
+-----+-----+-----+
```

3.5. Server Identifier

This option is used in MDHCPOFFER and MDHCPREQUEST messages, and may optionally be included in the MDHCPACK and MDHCPNAK messages. MDHCP

servers include this option in the MDHCPOFFER in order to allow the client to distinguish between lease offers. MDHCP clients use the contents of the 'server identifier' field as the destination address for any MDHCP messages unicast to the MDHCP server. MDHCP clients also indicate which of several lease offers is being accepted by including this option in an MDHCPREQUEST message.

The identifier is the IP address of the selected server.

The code for this option is 54, and its length is 4.

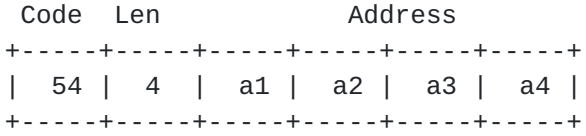

3.6. Client Identifier

This option is used by MDHCP clients to specify their unique identifier. MDHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain.

Identifiers SHOULD be treated as opaque objects by MDHCP servers.

The client identifier MAY consist of type-value pairs similar to the of a hardware type and hardware address. In this case the type field SHOULD be one of the ARP hardware types defined in STD2 [8]. A hardware type of 0 (zero) should be used when the value field contains an identifier other than a hardware address (e.g. a fully qualified domain name).

For correct identification of clients, each client's client-identifier MUST be unique among the client-identifiers used on the subnet to which the client is attached. Vendors and system administrators are responsible for choosing client-identifiers that meet this requirement for uniqueness.

The code for this option is 61, and its minimum length is 2.

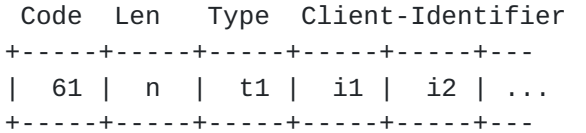

3.7. Multicast Scope

The multicast scope option is used by the client to indicate the

multicast scope for the requested multicast address. It is also used to indicate the scope of the assigned address by the MDHCP server. If this option is not specified, the MDHCP server MAY allocate an address from a default scope or reject the request.

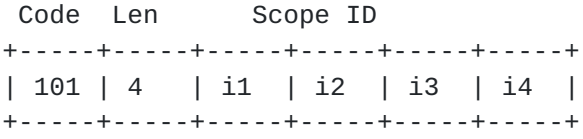

The client may obtain the scope list through the Multicast Scope List option or using some other means. The scope id is the first multicast address in the scope.

The code for this option is 101 and the length is 4.

3.8. Start Time

The Start Time option is used in a client request (MDHCPDISCOVER or MDHCPREQUEST) to allow the client to request the starting time for the use of the assigned address. This option allows a client to request a multicast address for use at a future time.

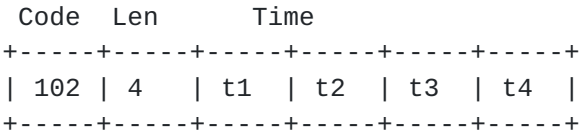

The time value is an unsigned 32 bit integer in network byte order giving the number of seconds since 00:00 UTC, 1st January 1970. This is consistent with the time format used in AAP [9] and can be converted to an NTP timestamp by adding decimal 2208988800. This time format will not wrap until the year 2106.

If the Start Time option is present, the IP Address Lease Time option specifies the duration of the lease beginning at the Start Time option value.

If the Start Time option is present, the Current Time option MUST also be present, as described in section 2.10.

The code for this option is 102 and the length is 4.

3.9. Multicast TTL

This option specifies the TTL value to be used with the multicast address. The TTL is specified as an octet with a value between 1 and 255. The implied value of this option is 255 when not included.

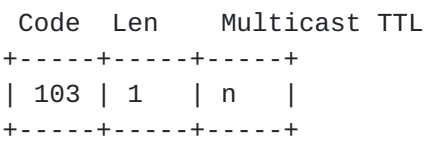

The code for this option is 103 and the length is 1.

3.10. Number of Addresses Requested

This option specifies the minimum and desired number of addresses requested by the client. These values are unsigned 16 bit integers stored in network byte order. The minimum MUST be less than or equal to the desired number. If a packet is received where this is not the case, the desired value MUST be used for both.

The client MAY obtain more than one address either by repeating the protocol for every address or by requesting several addresses at the same time via this option. When the client is requesting only one address, this option SHOULD not be included. An MDHCP server receiving an MDHCPDISCOVER or MDHCPREQUEST packet including this option MUST include between minimum and desired number of addresses in any MDHCPOFFER or MDHCPACK response.

The server MAY use this option to indicate to the client the number of addresses it has allocated to the client. In this case, the minimum and desired values MUST be equal.

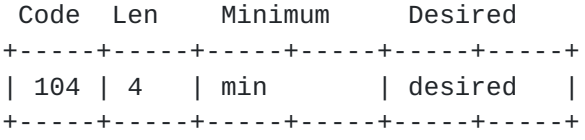

The code for this option is 104 and the length is 4.

3.11. Requested Language

This option specifies the language in which the MDHCP client would like strings such as zone names to be returned. It is an RFC 1766 [11] language tag. The proper way to handle this tag with respect to zone names is discussed further in the definition of the Multicast Scope List option.

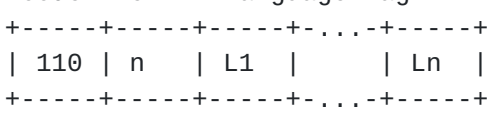

The code for this option is 110.

3.11. Multicast Scope List

The format of the multicast scope list option is:

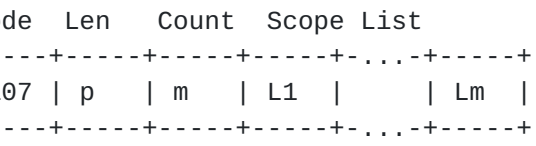

The scope list is a list of m tuples, where each tuple is of the form,

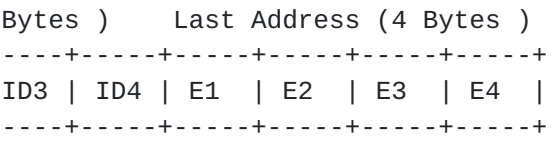

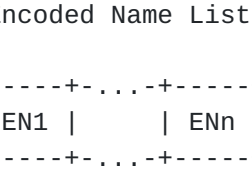

where Scope ID is the first multicast address in the scope, Last Address is the last multicast address in the scope, TTL is the multicast TTL value for the multicast addresses of the scope, and Name Count is the number of encoded names for this zone (which may be zero). Each encoded name is of the form

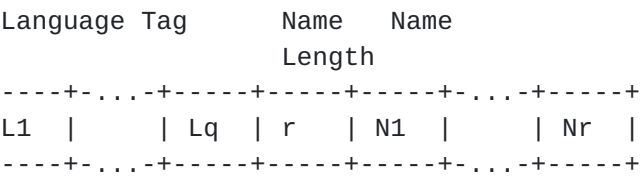

where Name Flags is a flags field with flags defined below, Lang Length is the length of the Language Tag in octets (which may be zero), Language Tag is a language tag indicating the language of the zone name (as described in [11]), Name Length is the length of the Name in octets, and Name is a UTF-8 [10] string indicating the name given to the scope zone.

The high bit of the Name Flags field is set if the following name should be used if no name is available in a desired language. Otherwise, this bit is cleared. All remaining bits in the octet SHOULD be set to zero and MUST be ignored.

The scope IDs of entries in the list SHOULD be unique and the scopes SHOULD be listed from smallest to largest.

If the client has not specified a Requested Language option in its request, the MDHCP server SHOULD return all zone names for each zone. If the client has specified a Requested Language option in its request, the MDHCP server MUST return no more than one zone name for each zone. For each zone, the MDHCP server SHOULD first look for a zone name that matches the requested language tag (using a case-insensitive ASCII comparison). If any names match, one of them should be returned. Otherwise, the MDHCP server SHOULD choose another zone name to return (if any are defined). It SHOULD give preference to zone names that are marked to be used if no name is available in a desired language.

If the scope list is too long to fit into a single Multicast Scope List option (255 bytes), the server SHOULD split it at 255 bytes and place the remaining bytes in one or more subsequent Multicast Scope List options. The client SHOULD carefully scan the entire set of options and reassemble the scope list by concatenating all Multicast Scope List options before attempting to parse them.

The code for this option is 107.

Example:

There are two scopes supported by the multicast address allocation server: Inside abcd.com with addresses 239.192.0.0-239.195.255.255, and world with addresses 224.0.1.0-238.255.255.255. Then this option will be given as:

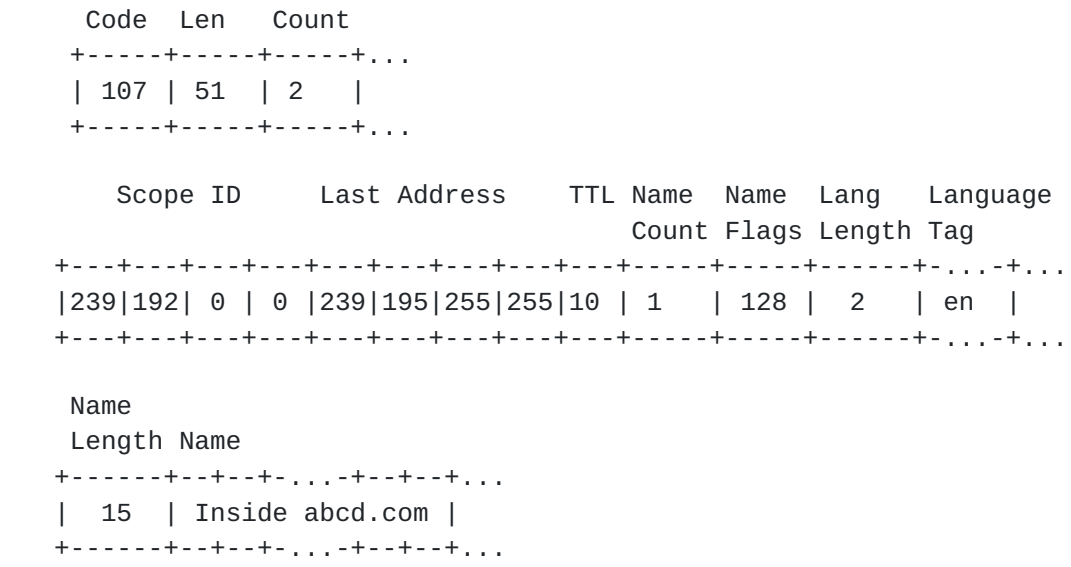

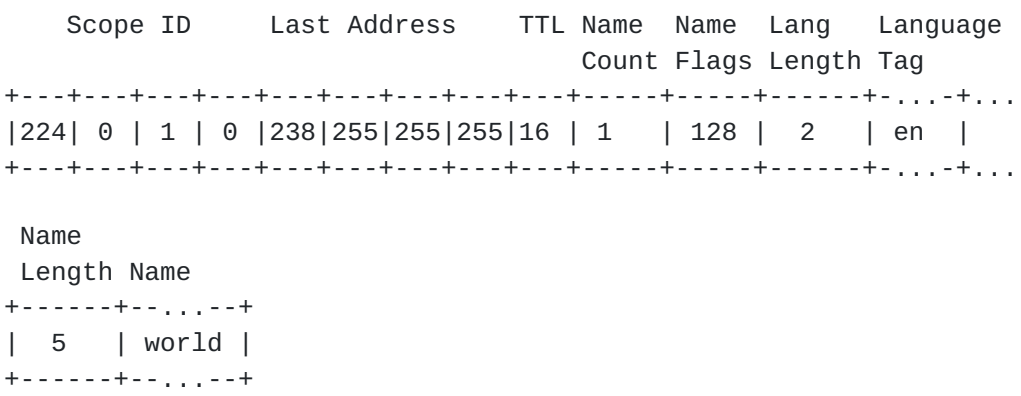

3.12. List of Address Ranges Allocated

This option is used by the server to provide the list of all the address ranges allocated to the client when the client requests more than one address. When a client requests only one address, the server uses the 'yiaddr' field to specify the allocated address. When a client requests more than one address, additional address ranges are listed via this option.

This option is also used by the client when requesting a lease extension for more than one address or releasing more than one address.

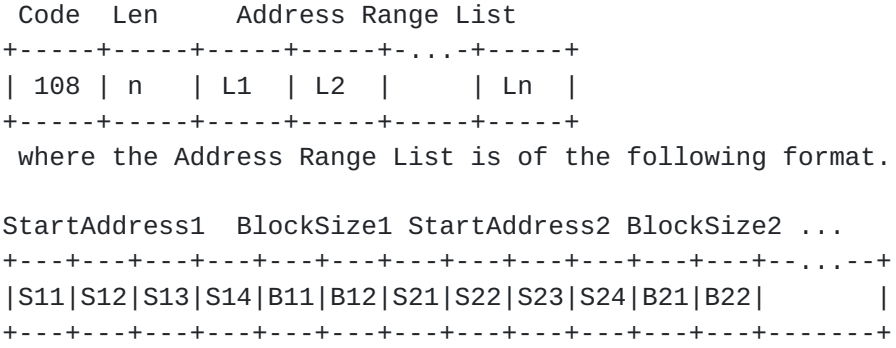

The code for this option is 108 and the minimum length is 6.

3.13. Current Time

The current time may be used to express what the sender thinks the current time is. This option may be used to detect clock skew and MUST be used if the Start Time option is used, as described in section 2.10.

```
 Code  Len      Time
+-----+-----+-----+-----+-----+-----+
| 109 | 4   | t1  | t2  | t3  | t4  |
+-----+-----+-----+-----+-----+-----+
```

The time value is an unsigned 32 bit integer in network byte order giving the number of seconds since 00:00 UTC, 1st January 1970. This is consistent with the time format used in AAP [9] and can be converted to an NTP [4] timestamp by adding decimal 2208988800. This time format will not wrap until the year 2106.

The code for this option is 109 and the length is 4.

3.14. End Option

The end option marks the end of valid information in the vendor field. Subsequent octets should be filled with pad options.

The code for this option is 255, and the length is 1.

```
 Code
+-----+
| 255 |
+-----+
```

4. Open Issues and Action Items

We refer to Allocation Scope, but this has not been assigned by IANA yet. If this is still the case when we want to go to Proposed Standard, we will have to remove these references.

5. Security Considerations

There are several security risks associated with multicast address allocation.

First, multicast addresses are a scarce commodity. Therefore, it may be desirable or necessary to provide access controls on allocation services. Enforcing such controls requires client authentication with replay prevention.

Second, there are many possibilities for denial of service. Allocating excessive numbers of addresses may be addressed by providing access controls as described above. Deallocating or modifying other clients' addresses may be handled in a similar way. Installing a false server can be addressed by requiring clients to authenticate servers.

Third, information about who has allocated what may disclose confidential information and may be useful in other attacks such as sending extra data to your enemies' multicast groups. Providing confidentiality via encryption addresses those problems somewhat, although traffic analysis is still possible.

There are already efforts underway in the DHC Working Group that should be helpful with these problems. The authors are monitoring this work and hope to apply the solutions developed there in this context.

6. Acknowledgments

The authors would like to thank Rajeev Byrisetty, Steve Deering, Peter Ford, Mark Handley, Van Jacobson, David Oran, Thomas Pfenning, Dave Thaler, Ramesh Vyaghrapuri and the participants of the IETF for their assistance with this protocol.

Much of the text in this document is based on or copied from [1] and [2]. The authors of this document would like to express their gratitude to the authors of these previous works. Any errors in this document are solely the fault of the authors of this document.

7. References

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

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

[3] Meyer, D., "Administratively Scoped IP Multicast", RFC 2365, July 1998.

[4] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation and Analysis", RFC 1305, March 1992.

[5] Handley, M., D. Thaler, and D. Estrin, "The Internet Multicast Address Allocation Architecture", Internet Draft, draft-handley-malloc-arch-00.txt, December 1997.

[6] Handley, M., "Multicast-Scope Zone Announcement Protocol (MZAP)", Internet Draft, draft-ietf-mboned-mzap-00.txt, December 1997.

[7] Croft, W., and J. Gilmore, "Bootstrap Protocol", RFC 951, Stanford University and Sun Microsystems, September 1985.

[8] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1700, USC/Information Sciences Institute, July 1992.

[9] Handley, M., "Multicast Address Allocation Protocol (AAP)", Internet Draft, draft-handley-aap-00.txt, December 1997.

[10] Yergeau, F., "UTF-8, a transformation format of ISO 10646", RFC 2279, January 1998.

[11] Alvestrand, H., "Tags for the Identification of Languages", RFC 1766, March 1995.

8. Authors' Addresses

Baiju V. Patel
Intel Corp.
2111 NE 25th Ave.
Hillsboro, OR 97124

Phone: 503 264 2422
EMail: baiju.v.patel@intel.com

Munil Shah
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Phone: 425 703 3924
Email: munils@microsoft.com

Stephen R. Hanna
Sun Microsystems, Inc.
2 Elizabeth Drive, M/S UCHL03-205
Chelmsford, MA 01824

Phone: +1.978.442.0166
Email: steve.hanna@sun.com

APPENDIX A: Examples

This appendix includes several examples of typical MDHCP protocol exchanges.

1. Unicast Address Allocation

In this example, an MDHCP client wants to allocate a multicast address from the global scope for use during the next two hours. It knows (through prior configuration or communication) the scopes that

are in effect at its location and the unicast address of an MDHCP server that provides allocation services for the global scope.

The client begins by unicasting an MDHCPREQUEST packet to the server. This packet includes the IP Address Lease Time, MDHCP Message Type, Client Identifier, and Multicast Scope options.

The server responds with an MDHCPACK packet containing the requested address. The address is stored in the yiaddr field. This packet includes the IP Address Lease Time, MDHCP Message Type, Client Identifier, and Multicast Scope options.

At this time, the client is said to have a two hour "lease" on the multicast address, indicating that (with high likelihood) this address will not be allocated to any other clients in the same scope for use during this period. The client may then proceed to distribute the address to others and conduct a multicast session on the address.

If the client had not received a response from the server, it would probably have retransmitted its MDHCPREQUEST packet for a while. If it still received no response, it could try another server or move to multicast discovery with the MDHCPDISCOVER message.

The following figure illustrates this exchange.

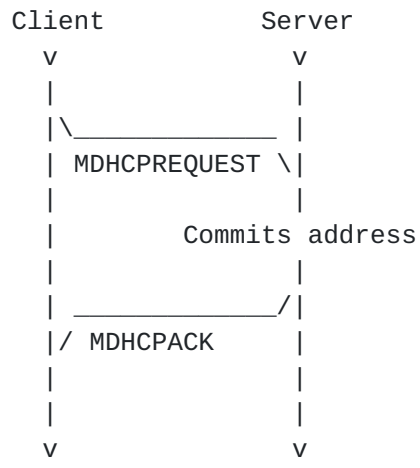

Figure 2: Timeline diagram of messages exchanged in Unicast Address Allocation example

2. Lease Extension

This is a continuation of the previous example. The client has already allocated a multicast address from the global scope for use during the next two hours. Half way through this two hour period, it decides that it wants to extend its lease for another hour.

The client unicasts an MDHCPREQUEST packet to the server from which it allocated the address. This packet includes the Requested IP Address, IP Address Lease Time, MDHCP Message Type, and Client Identifier options. The time included in the IP Address Lease Time is two hours, since the client wants the lease to expire two hours from the current time.

The server responds with an MDHCPACK packet indicating that the lease extension has been granted. The address is stored in the yiaddr field. This packet includes the IP Address Lease Time, MDHCP Message Type, Client Identifier, and Multicast Scope options.

If the server did not want to grant the requested lease extension, it would have responded with an MDHCPACK packet with the remaining lease time indicated in the IP Address Lease Time option. It would *not* send an MDHCPNAK packet, since that would cancel the lease immediately.

The following figure illustrates this exchange.

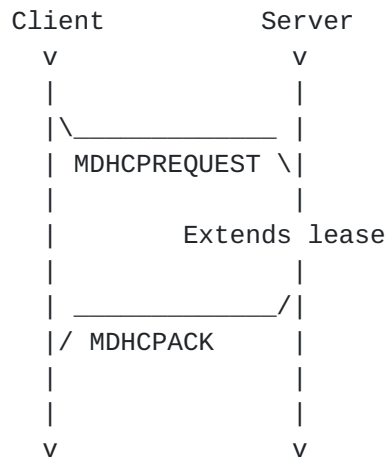

Figure 3: Timeline diagram of messages exchanged in Lease Extension example

3. Address Release

This is a continuation of the previous example. The client has already allocated a multicast address and extended its lease for another two hours. Half an hour later, the client finishes its use of the multicast address and wants to release it so it can be reused.

The client unicasts an MDHCPRELEASE packet to the server from which it allocated the address. This packet includes the Requested IP Address, MDHCP Message Type, and Client Identifier options. When the server receives this packet, it cancels the client's lease on the address.

Since the server does not acknowledge the MDHCPRELEASE packet, there is no provision for the client retransmitting it. However, this is not very harmful. If an MDHCPRELEASE packet is lost, the server will keep the multicast address reserved for the client's use until the end of its lease. At that point, the address will once again be available for use by others.

The following figure illustrates this exchange.

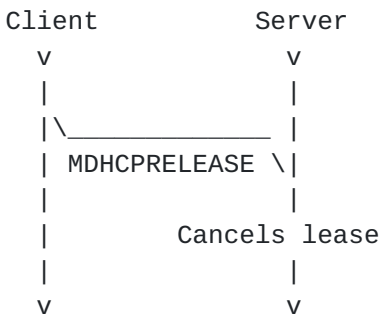

Figure 4: Timeline diagram of messages exchanged in Address Release example

4. Multicast Discovery and Address Allocation

Let's revisit the first example, but say that the MDHCP client did not know the unicast address of the MDHCP server that it wanted to use.

The client begins by multicasting an MDHCPDISCOVER packet to an MDHCP Server Multicast Address. This packet includes the IP Address Lease Time, MDHCP Message Type, Client Identifier, and Multicast Scope options.

Any servers that receive the MDHCPDISCOVER packet and can satisfy this request temporarily reserve an address for the client and unicast an MDHCPOFFER packet to the client. These packets contain the IP Address Lease Time, MDHCP Message Type, Server Identifier, Client Identifier, and Multicast Scope options. The reserved address is also stored in the yiaddr field.

After a suitable delay, the client multicasts an MDHCPREQUEST packet to the MDHCP Server Multicast Address. This packet contains all of the options included in the MDHCPDISCOVER packet, but also includes the Server Identifier option, indicating which server it has selected for the request.

The server whose Server Identifier matches the one specified by the client responds with an MDHCPACK packet containing the same information as the MDHCPOFFER packet. All the other servers that had sent MDHCPOFFER packets stop reserving an address for the client and forget about the whole exchange.

The client now has a two hour "lease" on the multicast address, just like it did in the first example.

If the client had not received a response from the server, it would have retransmitted its MDHCPREQUEST packet for a while. If it still

received no response, it could try another server or return to multicast discovery with a new MDHCPDISCOVER message.

The following figure illustrates this exchange.

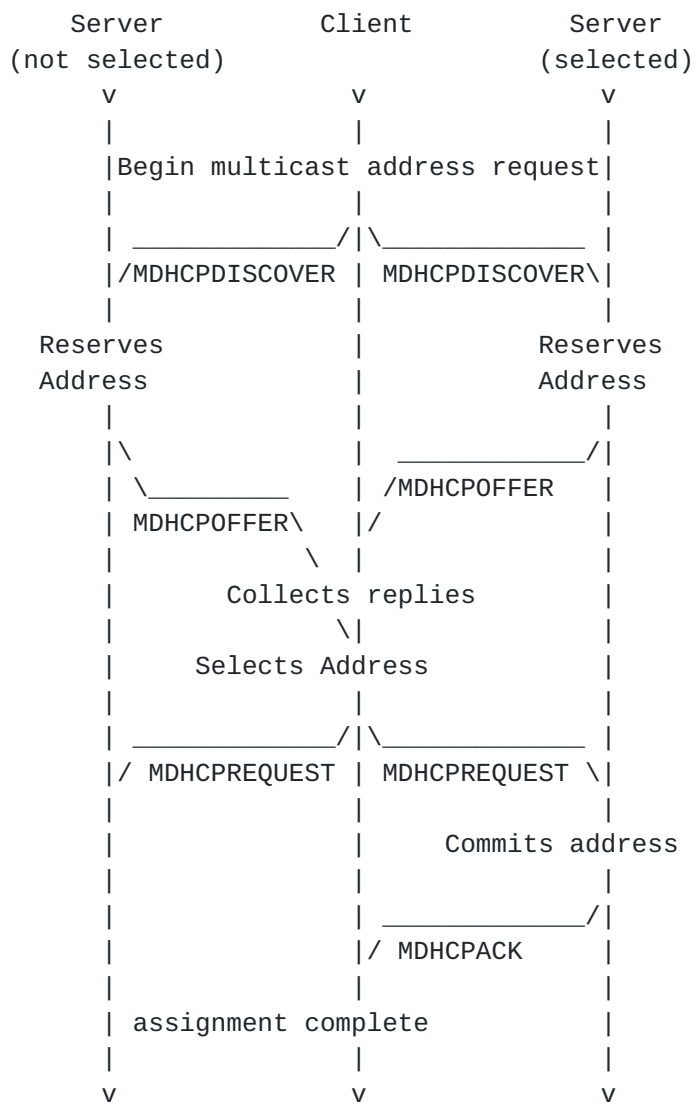

Figure 5: Timeline diagram of messages exchanged in Multicast Discovery and Address Allocation example

APPENDIX B: Change Log

CHANGES FROM draft-ietf-malloc-mdhcp-00.txt

Port number assigned.

Removed unused chaddr, sname, and file fields.

Split MDHCP option space from DHCP.

Changed technique for choosing MDHCP Server Multicast Addresses when initial requests fail.

Added Requested Language option.

Added language tags and multiple names per zone to the Multicast Scope List option.

CHANGES FROM draft-ietf-dhc-mdhcp-03.txt

Many changes to make this document no longer dependent on the DHCP spec. This should make the document easier to read and understand.

Removed MDHCPDECLINE.

Added Current Time option to deal with clock skew.

Scopes are now identified by the first multicast address in the scope instead of using a scope ID.

Changed Total Addresses Requested option to Number of Addresses Requested. Changed this option to have minimum and desired fields.

Clarified that servers MAY send MDHCPOFFER or MDHCPACK messages with shorter lifetimes or later start times than those requested by the client. This is consistent with DHCP and provides a simple way to achieve the minimum/maximum lifetime functionality described in the malloc abstract API.