

March 1997

Multicast address allocation extensions to the
Dynamic Host Configuration Protocol
<[draft-ietf-dhc-mdhcp-00.txt](#)>

Status of this memo

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``[1id-abstracts.txt](#)'' listing contained in the Internet-Drafts Shadow Directories on [ftp.is.co.za](#) (Africa), [nic.nordu.net](#) (Europe), [munnari.oz.au](#) (Pacific Rim), [ds.internic.net](#) (US East Coast), or [ftp.isi.edu](#) (US West Coast).

Abstract

The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts on a TCP/IP network. The multicast extensions to DHCP add additional capability of dynamic allocation of the multicast addresses and additional configuration options.

1. Introduction

The multicast extensions to DHCP (MDHCP) provide configuration parameters to the multicast applications. MDHCP is built on a client-server model, where designated DHCP server allocate multicast addresses and deliver parameters associated with the address to dynamically configured hosts. Throughout the remainder of this document, the term "server" refers to a host providing multicast address(es) and parameters through DHCP, and the term "client" refers to a host requesting multicast address(es) and parameters from a DHCP server. MDHCP server is used at times, to

indicate a DHCP server capable of handling MDHCP extensions to the

DHCP protocol and the MDHCP client is used to indicate the MDHCP capable DHCP client. MDHCP is not a separate protocol, but is simply extensions to the DHCP protocol.

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

Like DHCP, MDHCP should be a mechanism rather than a policy. MDHCP 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.

The MDHCP client is not required to obtain IP address from a DHCP server in order to use MDHCP protocol.

The design goals specified in the DHCP RFC also apply to MDHCP.

1.1 Requirements

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

- o "MUST"

This word or the adjective "REQUIRED" means that the item is an absolute requirement of this specification.

- o "MUST NOT"

This phrase means that the item is an absolute prohibition of this specification.

- o "SHOULD"

This word or the adjective "RECOMMENDED" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.

- o "SHOULD NOT"

This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable

or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

- o "MAY"

This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

1.2 Terminology

This document uses the following terms:

- o "DHCP client"

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

- o "DHCP server"

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

- o "MDHCP client"

A MDHCP client is a DHCP client that supports MDHCP extensions.

- o "MDHCP server"

A MDHCP server is a DHCP server that supports MDHCP extensions.

- o "BOOTP relay agent"

A BOOTP relay agent or relay agent is an Internet host or router

that passes DHCP messages between DHCP clients and DHCP servers. DHCP is designed to use the same relay agent behavior as specified in the BOOTP protocol specification.

- o "binding"

A binding is a collection of configuration parameters, including at least an IP address, associated with or "bound to" a DHCP client. Bindings are managed by DHCP servers.

1.3 Motivation

The current mechanisms used for multicast address allocation are fairly adhoc and are specific to the applications. For example, the mbone tools listen to the SAP messages to determine the multicast addresses that are currently in use. Since there are no unified mechanisms for allocating multicast address, each class of

applications or tools need to implement their own solution. This not only creates many different groups of address but also leads to potential conflict in address usage. The router that filters multicast packets based on the administrative scopes must be aware of different pools of addresses used by different applications. If an organization intends to implement a multicast allocation and use policy, each of every application must be aware of it and know how the addresses are to be used within this policy. The protocols that attempt to guess multicast addresses by listening to SAP messages may at times lead to conflicting use of multicast addresses as well. The conflicts may arise due to use of different TTLs, Scoping, packet losses, and temporary shutdown of the originating system. Consider an example where a site does not allow any multicast to enter but allows all the out going multicasts. In that case, an application internal to the site has no way of determining which multicast address are in use outside the site. Similarly, if some announcement packets are lost, the application may incorrectly conclude that a multicast address is not in use. Finally, since the multicast address allocation scheme is adhoc, the organization cannot clearly define the policy for multicasts.

Therefore, there is a need to provide a mechanism for requesting and assigning a multicast address. If this mechanism is based on client/server paradigm, the client is not responsible for ensuring uniqueness of multicast address. Moreover, the organization may be able to enforce a multicast policy through which all the multicast addresses are assigned. An example of a policy would be to assign multicast addresses from range X to be used within an organization,

range Y to used for an entire corporation etc. The applications should not have to be manually configured to determine these policies.

The proposed protocol does not address the specific means that are needed at a DHCP server to determine the address to be allocated. For the administratively scoped addresses, the DHCP server may have a block of address that it can assign to the application. Currently, it is not clear how will it determine the addresses to be allocated for the Internet (mbone). If and when these mechanisms become available, the DHCP server could act as a proxy in obtaining those addresses as well.

1.4 Protocol Summary

From the client's point of view, MDHCP is an extension of the DHCP mechanisms. The MDHCP servers assigns multicast addresses to the hosts to be used within a specific scope, and valid for a specific period. A client may request multiple multicast addresses.

The client requests a multicast address(es) to be used for a specific multicast scope available to it, and for a specific lease period. The MDHCP server would ideally assign the address from the requested scope or may allocate it for a different scope. However, if it allocates the address from a different scope, it will provide

this information as an option. The DHCP server MUST provide a TTL value. The multicast packets using the assigned address MUST NOT use a TTL value larger then the one provided. The lease period is defined by the duration of the lease and the time at which the lease becomes effective. Since the client may want to extend lease at a later time, the DHCP server SHOULD make every attempt at allocating an address which is not currently allocated to any other client. The DHCP server MUST NOT allocate the same addresses to different clients with overlapping lease period. The multicast scope list is one of the DHCP configuration parameters.

The scope list may be obtained through the DHCP option described in [3], or may be obtained with some other means. Similarly, the MDHCP server address (unicast or multicast) may also be obtained by the option described in [3] or by some other means.

The MDHCP protocol uses M flag and a set of options defined below.

2 MDHCP messages and options.

The following options and flags are used by MDHCP extensions.

2.1 M flag

A new flag (M) is defined to differentiate the MDHCP messages from DHCP messages. All the messages (DHCPDISCOVER, DHCP OFFER etc.) use M flag (this is a new flag) defined below to indicate multicast address negotiations. The second bit of the flag field (bit 1) defines M (multicast) flag. The M bit must be set for all the message exchanges pertinent to the multicast address assignment. The client MUST obtain an IP address prior to requesting a multicast address. Therefore, B flag MUST not be set when M flag is set.

```

                                1 1 1 1 1 1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+
|B|M|           MBZ           |
+---+---+---+---+---+---+---+---+---+---+
```

B: BROADCAST flag
M: Multicast address request flag.

MBZ: MUST BE ZERO (reserved for future use)

2.2 Multicast Scope Option

This 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 DHCP server. If this option is not specified, the DHCP server MAY allocate an address from a DEFAULT scope or reject the request.

```


                                1 1 1 1 1 1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+
| code           | length=4     |
+---+---+---+---+---+---+---+---+---+---+
| Scope id      |                 |
+---+---+---+---+---+---+---+---+---+---+
| Scope id      |                 |
+---+---+---+---+---+---+---+---+---+---+
```

The client may obtain the scope list through the option described in [3] or using some other means. The scope id is the numeric representation of the scope as described in [3].

The 'code' for this option is TBD.

2.3 Start time Option

The start time is used in a client request (DHCPDISCOVER or DHCPREQUEST) to allow the client to request the starting time for the use of the assigned address. This option allows client to request a multicast address for use at a future time.




The time is the Coordinated Universal Time(UTC) in unit of seconds and is specified as a 32-bit integer and is specified in the network time format.

The 'code' for this option is TBD.

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

2.4 Multicast TTL Option

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.



2.5 Multicast Block Size option

In some cases, an application may require a group of consecutive addresses to be assigned. This option is used by a client to request n consecutive addresses. It is also used by the DHCP server to indicate number of consecutive addresses assigned starting at the

address specified in ``yiaddr'' field.

```

                                1 1 1 1 1 1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+
| code           | length=1         |
+---+---+---+---+---+---+---+---+---+---+
| n              |
+---+---+---+---+---+---+---+---+---+---+
```

2.6 Client Port Option

In order to facilitate implementations outside the operating system kernel, and to allow two separate client implementations: one for DHCP and one for MDHCP, if this option is specified, the MDHCP server MUST use the source port number used in the DHCPDISCOVER, DHCPREQUEST, DHCPINFORM, and DHCPRELEASE as the destination port number in the response messages.

```

      0 1 2 3 4 5 6 7
+---+---+---+---+---+---+
| code           |
+---+---+---+---+---+---+
```

2.7 Cookie Option

The MDHCP server may issue a cookie along with the multicast address(es) so that a different user may use the cookie to renew lease on address(es).

```

                                1 1 1 1 1 1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+
| code           | length=2         |
+---+---+---+---+---+---+---+---+---+---+
| cookie         |
+---+---+---+---+---+---+---+---+---+---+
```

This option is useful when the "owner" of the address leaves a multicast group and some other member decides to either renew or terminate the lease. If a different member of the group from the one who was assigned the multicast address wants to modify the terms of multicast address, it must use this cookie as its client identifier option. For example, if host X was issued a multicast address, who decides to leave the multicast group that is using the assigned the address. Then, another participant in the group

determines that the work group must continue beyond the lease time for the multicast address, it may renew the lease by specifying the cookie to uniquely identify the group of multicast addresses. Note that cookie is not used as a security capability but is used to simplify the client and server implementations.

3. MDHCP protocol

The client needs to obtain the IP address of the MDHCP server (this may be a unicast or a multicast address for MDHCP group), and the multicast scope list. This list may be obtained as part of the normal DHCP protocol using the options specified in [3] or by some other means.

The client selects one of multicast scopes and requests multicast address(es) from the MDHCP servers. The fields and options that are different from the normal DHCP message exchange are summarized in Table 1 to 3. details on rest of the parameters, please consult DHCP RFC[1]. The mutlicast addresses are renewed or released using the DHCP exchanges for network addresses as defined in the DHCP RFC[1].

Note that all the messages in this exchange have their M flag set and B flag not set.

The MDHCP Client MUST provide client identifier option when sending messages for multicast address assignment. The client generates a unique key and uses that as a client identifier in the DHCPDISCOVER message. When the server responds to this with DHCPOFFER, it also provides a cookie along with it. This cookie is generated on the server and it uniquely idenfies the transaction associated with the multicast addresss(es) being offered to this client. For all the subsequent messages, client uses this cookie as a client identifier. Each client may be running several different multicast enabled applications, and each application may require separate multicast address(es). Client MUST use separate unique client identifier when requesting separate multicast address(es) for each application.

A client implementation may choose to use hardware address, hardware type and application instance number to generate unique client identifier

Field	DHCPOFFER	DHCPACK	DHCPNAK
'ciaddr'	'ciaddr' from DHCPDISCOVER or 0	'ciaddr' from DHCPREQUEST or 0	0
'yiaddr'	Starting address of the multicast block assigned to client	Starting address of the multicast block assigned to client	0

'siaddr'	Server's IP address reachable from the client.	Server's IP address 0 reachable from the client.
----------	--	--

'chaddr'	'chaddr' from client DHCPDISCOVER message	'chaddr' from client DHCPREQUEST message	'chaddr' from client DHCPREQUEST message (unused)
'file'	may contain options	may contain options	
'options'	options	options	

Option	DHCPOFFER	DHCPACK	DHCPNAK
-----	-----	-----	-----
IP address lease time	MUST	MUST (DHCPREQUEST)	MUST NOT
Server identifier	MUST	MUST	MUST
Multicast TTL	MUST	MUST	MUST NOT
Multicast Block size	MAY	MAY	MUST NOT
Cookie	MUST	MAY	MUST NOT

Table 1: Fields and options that are different in multicast DHCP server messages.

Field	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE, DHCPRELEASE
-----	-----	-----	-----
'flags'	Set 'M' Bit. BROADCAST bit 0	set 'M' Bit BROADCAST bit 0	set 'M' bit BROADCAST bit 0
'ciaddr'	client's network addr reachable from the server	client's network addr reachable from the server	0
'chaddr'	may contain hardware address	may contain hardware address	may contain hardware address
'options'	options	options	(unused)
Option	DHCPDISCOVER	DHCPREQUEST	DHCPDECLINE, DHCPRELEASE
-----	-----	-----	-----
Requested IP address	MAY	MUST (in SELECTING or INIT-REBOOT) MUST NOT (in BOUND or RENEWING)	MUST (DHCPDECLINE), MUST NOT (DHCPRELEASE)
Start time	MAY	MAY	MUST NOT
Client identifier	MUST	MUST	MAY

Table 2: Fields and options that are different in multicast DHCP client messages

	INIT-REBOOT	SELECTING	RENEWING	REBINDING
multi/unicast	multicast	multicast if multicast DISCOVER	unicast	multicast
		unicast otherwise		
server-ip	MUST NOT	MUST	MUST NOT	MUST NOT
requested-ip	MUST	MUST	MUST NOT	MUST NOT
ciaddr	IP addr	IP addr	IP addr	IP address

Table 3: Client messages from different states

3.1 DHCPDISCOVER Message.

If the unicast address of a MDHCP server is known and it supports the desired multicast scope, the MDHCP client SHOULD send a DHCPDISCOVER address to the MDHCP server. If the MDHCP server fails to allocate address(es) or fails to respond, the DHCP client SHOULD send a multicast DHCPDISCOVER message to the group address (multicast) of the MDHCP server. In both cases, if the client uses non-standard DHCP port number, it MUST specify the client port option. The client MUST also specify its IP address in the ciaddr field so that the MDHCP server and respond to the client request with a unicast message. The B flag must not be set and M flag MUST be set.

The client MUST include client identifier option.

In addition, the DHCPDISCOVER option SHOULD include the following options:

- o DHCP Scope,
- o Start time,

- o Lease time (duration)

If any of these options are not specified, the DHCP server may assume default values.

[3.2](#) DHCPOFFER Message.

The DHCP server may respond to a DHCPDISCOVER message with a unicast DHCPOFFER to the client. This message MUST include an available multicast address using the 'yiaddr' field. The MDHCP server SHOULD reserve the offered address. When allocating the address, the server MUST make every effort to ensure that the address is not in use for the lease period.

The server MUST include configuration parameters such as DHCP scope, start and lease time, in the DHCPOFFER message, if different from the ones requested. The MDHCP server must specify a cookie value in this message and this cookie MUST be specified in all the subsequent messages exchanged between the MDHCP clients and server pertaining to associated address(es). The MDHCP server MUST use the

cookie to identify the addresses instead of the client IP address.

[3.3](#) DHCPREQUEST

The client will select a multicast address(es) from a DHCPOFFER response. The client SHOULD send a unicast DHCPREQUEST message indicating the selected multicast address(es) to the MDHCP server, when the DHCPOFFER was in response to a unicast DHCPDISCOVER message, and using a multicast message, when the DHCPOFFER was in response to a multicast address. It MUST include multicast address option field in the response. If the number of address selected are different from the number of offered address, the client MUST also include the multicast block size option.

The M flag MUST be set and B flag MUST NOT be set.

[3.4](#) DHCPACK.

If the multicast address(es) are still available, the MDHCP server MUST reserve the address and send a DHCPACK message. Any configuration parameters in the DHCPACK message SHOULD NOT conflict with the ones in earlier DHCPOFFER message. The M flag MUST be set and B flag MUST NOT be set.

[3.5](#) DHCPNACK

The server MAY choose to mark the multicast address in DHCP OFFER unavailable to the client. In that case it will send DHCP NACK message. The M flag MUST be set and B flag MUST NOT be set.

[3.6](#) Renewing and termination of lease

The client may choose to release address(es) before the lease time has expired. The usual DHCP messages are used for this purpose.

The M flag MUST be set and B flag MUST not be set. Moreover, the client port option SHOULD be specified, if the client is using a port different from the standard DHCP port. The cookie MUST be specified with RENEW and RELEASE messages.

[4.](#) Examples of usage

The MDHCP server is not required to be co-located with a DHCP server. Therefore, in a typical deployment, there may be fewer MDHCP servers than the DHCP servers. We consider specific examples of DHCP configurations and the use of MDHCP protocol extensions.


[4.1](#) One MDHCP server

There is one MDHCP server which is configured to allocate multicast

addresses to a client and there may be many DHCP servers. The DHCP servers should be configured to provide the address of the MDHCP server capable of allocating multicast address to the MDHCP client, and should include a multicast scope list supported by the MDHCP server. The client may obtain the DHCP server address and scope list through DHCP client configuration procedure (and may use DHCP INFORM message). The client then selects a multicast scope from which the multicast address is to be requested and sends out a unicast DHCP DISCOVER address and includes multicast scope, start time, and lease time information using DHCP options. It may also specify multicast block size. The MDHCP server responds with an DHCP OFFER for multicast address and includes a TTL value to be used with this address. The client sends out a DHCP REQUEST message and includes the selected. If the address is still available, the server responds with an DHCP ACK message, else responds with a NACK message.

Since the DHCP messages are directly send to the MDHCP server, the

server is capable of interpreting M flag and therefore, there will be no conflict between the interpretation of DHCP and MDHCP messages.






Figure 1: Timeline diagram of messages exchanged between MDHCP client and servers when allocating multicast address(es) using unicast messages to a MDHCP capable server.

4.2 One or more MDHCP servers

If one or more MDHCP servers are available to a MDHCP client for

the purpose of assigning multicast addresses, the DHCP scope list option SHOULD specify an administratively scoped group address used by the MDHCP servers to receive DHCPDISCOVER messages. Each scope in the scope list MUST be supported by at least one server listening to the group multicast address used by MDHCP servers.

The client SHOULD select a scope and send out a DHCPDISCOVER, DHCPREQUEST messages to the group multicast address. The multicast DHCPREQUEST message is only received by the MDHCP capable DHCP servers, and therefore, there is no conflict between the MDHCP and DHCP messages. Further, the messages for renewing and releasing lease are sent directly to the MDHCP servers only, and therefore, there is conflict between DHCP and MDHCP message interpretation by a non-MDHCP capable server.

A summary of fields of MDHCP in messages that are different from the corresponding DHCP [1] messages are specified in Tables 1 to 3.

In some cases, the client may be aware of the unicast address of an MDHCP capable server, and may also be aware of the group multicast address of the MDHCP capable servers. In that case, the client SHOULD first try to use the unicast address, and if unsuccessful, SHOULD try the group multicast address for MDHCP servers.




Figure 2: Timeline diagram of messages exchanged between MDHCP client and servers when allocating multicast address(es) using group multicast address for MDHCP capable servers.

5. Acknowledgements

The authors would like to thank Thomas Pfenning, Rajeev Byrisetty, and Ramesh Vyaghrapuri for their assistance in creating this document.

6. References

- [1] Droms, R., "Dynamic Host Configuration Protocol", [RFC1541](#), October 1993

- [2] Alexander, S., and R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 1533](#), Lachman Technology, Inc., Bucknell University, October 1993.

- [3] Patel, B., and Shah, M., ``Multicast address allocation extensions options''
<[draft-ietf-dhc-multopt-00.txt](#)>

- [4] Meyer, D., ``Administratively scoped IP Multicast'', Internet draft,
<[draft-ietf-mboned-admin-ip-space-01.txt](#)>

7. Author's Address

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

Phone: 503 264 2422
EMail: baiju@ibeam.intel.com

Munil Shah
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Phone:206 703 3924
Email:munils@microsoft.com

This document will expire on Sept, 1997