

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 10, 2009

R. Johnson
J. Jumarasamy
K. Kinnear
M. Stapp
Cisco Systems, Inc.
March 9, 2009

Subnet Allocation Option
draft-ietf-dhc-subnet-alloc-09.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 10, 2009.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

Subnet Allocation Option

March 2009

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Internet-Draft

Subnet Allocation Option

March 2009

Abstract

This document defines a new DHCP option which is passed between the DHCP Client and the DHCP Server to request dynamic allocation of a subnet, give specifications of subnet(s) allocated, and report usage statistics. This memo documents the current usage of the option in agreement with [\[RFC3942\]](#), which declares that any pre-existing usages of option numbers in the range 128 - 223 should be documented and the working group will try to officially assign those numbers to those options.

Table of Contents

1.	Introduction	5
2.	Conventions	6
3.	Subnet Allocation Option format	7
3.1.	Subnet-Request Suboption format	7
3.2.	Subnet-Information Suboption format	9
3.2.1.	Subnet Prefix Information block format	10
3.3.	Subnet-Name Suboption format	12
3.4.	Suggested-Lease-Time Suboption format	12
4.	Requesting allocation of a subnet	14
4.1.	Client DHCPDISCOVER message	14
4.2.	Server DHCPOFFER message	14
4.3.	Client DHCPREQUEST message	15
4.4.	Server DHCPACK message	15
5.	Client renewal of subnet lease	17
5.1.	Client RENEW DHCPREQUEST message	17
5.2.	Server RENEW DHCPREQUEST response	17
5.3.	Client DHCPRELEASE message	18
5.4.	Server DHCPFORCERENEW message	18
6.	Client requesting previously allocated subnet information . .	19
6.1.	Client Initial DHCPDISCOVER Message	19
6.2.	Server Initial DHCPOFFER Response	19
6.3.	Client Additional DHCPDISCOVER Messages	19

6.4.	Server Additional DHCP OFFER Responses	20
7.	DHCP Server Subnet Allocation method	21
8.	Examples	22
8.1.	Example 1:	22
8.2.	Example 2:	23
9.	Differences from DHCPv6 Prefix Delegation	27
10.	Security Considerations	28
11.	IANA Considerations	29
12.	Intellectual Property Rights and Copyright	30
13.	References	31
13.1.	Normative References	31
13.2.	Informative References	31

Authors' Addresses	32
------------------------------	--------------------

1. Introduction

There is a need for a DHCP Client to be able to allocate a subnet from a DHCP Server. Alternate methods of allocation, such as AAA are not appropriate for various reasons and the most straight-forward way to accomplish this allocation is via DHCP. A DHCP option to support this may be utilized by a simple Dialin Aggregation box, or even to implement a Hierarchical chain of DHCP Servers, each one in turn leasing one or more subnets to the next DHCP Server down the chain.

This new DHCP option [[RFC2132](#)], the Subnet Allocation option is specified in such a way as to use one DHCP Option number, while using suboption numbers for each function. The Subnet-Request suboption tells what types of subnets are needed and how many. The Subnet-Information suboption gives the actual subnet number(s) and allows for extra flags to convey additional information about each subnet. The "Subnet-Name" suboption allows a method of passing additional information about the requested subnet(s), such as department name, user name, customer number, etc. The "Suggested-Lease-Time" suboption provides a method for the DHCP Server to suggest a lease time for addresses allocated from the offered subnet. The DHCP Server has the option of not supplying all subnets requested or even

returning different sized subnets than were requested. Additionally, usage statistics may be included in RENEW messages from the DHCP Client back to the DHCP Server.

[2.](#) Conventions

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 [[RFC2119](#)].

This document also uses the following terms:

DHCP Client: DHCP Client or "Client" is an Internet host using DHCP to obtain configuration parameters such as a network address.


DHCP Server: A DHCP Server or "Server" is an Internet host that returns configuration parameters to DHCP Clients.

[illegible]

Code = Subnet Allocation option code (1 octet): 220
 Len = Length of the entire option including all sub-options
 (1 octet)
 Flags = Various flags: (None currently defined)

One or more sub-options may be specified as described below.

3.1. Subnet-Request Suboption format



Len = Length of the suboption (always 2 for this suboption)
 (1 octet)
 Flags = Flags field. (all unused bits must be zero)

- 'h' = Hierarchical flag
 - 1 : Client will be allocating addresses from this subnet.
 - 0 : Client will be relaying DHCP requests to the Server from this subnet.
- 'i' = Information flag
 - 1 : Client is seeking information about previously allocated subnets.
 - 0 : Client is seeking a new subnet allocation.

Prefix = Network prefix length requested
 (0 means no suggested length is given) (1 octet)

The DHCP Server SHOULD NOT include the Subnet Request suboption in any replies to the DHCP Client. Enough information will be present in the Subnet-Information suboption, such that the Subnet Request suboption is not needed in replies to the Client.

The DHCP Server SHOULD allocate a subnet with prefix length [\[RFC4632\]](#) less than or equal to the "Prefix" field length specified in the request. In other words, a subnet containing a number of addresses at least the size indicated by the prefix length requested and

possibly more. The DHCP Server MAY allocate a subnet with a prefix

length greater than that specified in the request (subnet with a number of addresses less than requested), but this is not encouraged.

A "Prefix" field size of 0 MAY be specified by the DHCP Client. In this case, no suggested prefix length is given.

Multiple Subnet-Request suboptions in a DHCP packet indicate that multiple subnets are being requested. Note that multiple occurrences of this suboption MUST NOT be concatenated in the sense of [[RFC3396](#)].

Each Subnet-Request suboption MUST result in no more than one Subnet-Information suboption in the DHCP OFFER message from the Server, and may result in no Subnet-Information suboptions.

The Client MAY also include the Subnet-Information suboption in order to request a particular subnet. In this case, the Client SHOULD only include one Subnet-Request suboption, since it would otherwise be unclear which Subnet-Information suboption referred to which Subnet-Request suboption. Multiple subnet requests can be made by sending multiple DHCPDISCOVER packets.

Setting the 'h' flag to 1 indicates the Client will be allocating addresses from the allocated subnet(s) itself. This can be thought of as a "Hierarchical DHCP" design in that control of allocation for the subnet(s) will be passed to the Client.

Setting the 'h' flag to 0 indicates the Client wants the DHCP Server to retain control over allocation of addresses from the subnet(s). Any address allocation requests on the subnet will be relayed back to the DHCP Server.

Setting the 'i' flag to 1 indicates the Client is NOT seeking allocation of any subnets, but is simply seeking information from the Server as to what subnet(s) have been allocated (or reserved) for this Client. If the 'i' flag is set to 1, then the "Prefix" field SHOULD be set to 0 and MUST be ignored by the Server. In this case, the conversation between the Client and the Server will only progress to the DHCP OFFER packet from the Server, giving the information, as described in [section 6](#) below.

3.2. Subnet-Information Suboption format

```

      0               1               2
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           2           |   Len   |  Flags   :c:s| SP1, SP2, ...
+---+---+---+---+---+---+---+---+---+---+---+---+

```

Len = Length of the sub-option (min. length of 8) (1 octet)

Flags = Various flags which apply to ALL Subnet Prefix
Information fields specified in this suboption.
Unused flags must be zero.

'c' : Client flag (explained below)

1 : This information is in response to a client request
(or has been echoed back by the client, when asking
for the next previously allocated subnet from the
server)

0 : Otherwise

's' : Server flag (explained below)

1 : Server has additional subnet information for this
client

0 : Otherwise

SP1, SP2, ... Subnet Prefix information blocks as specified below
(variable sized)

The "Client flag" ('c') is set to 1 if this Subnet-Information suboption is in response to a Client request for information from the Server as to what subnet(s) have been allocated. This flag is used in response to a Subnet-Request suboption with the 'i' flag set and should be 0 in other Server responses. Note, the flag is echoed back from the Client to the Server when requesting the "next previously allocated subnet". Another way to think of the 'c' bit would be that it indicates that the subnet information contained in this suboption does not represent a new allocation by the Server or a new request for allocation by the Client, but instead represents previously allocated subnet information.

The "Server flag" ('s') is set to 1 if the Server has additional subnet information for the Client.

The Subnet-Information suboption is used by the DHCP Server in order to return information as to what subnets are offered (in the case of a DHCPOFFER packet) or allocated (in the case of a DHCPACK packet).

As is indicated above, multiple subnets may be returned in one Subnet-Information suboption.

The Subnet-Information suboption is also used by the DHCP Client in order to request allocation of specific subnets in a DHCPREQUEST packet. In this case, the "Network", "Prefix", and "Flags" fields contained in the associated Subnet Prefix Information blocks MUST NOT be changed from the information which was received in the DHCPOFFER packet from the server. The Client MAY, however, use multiple Subnet-Information suboptions in order to request subnets which were originally specified by the Server inside one Subnet-Information suboption.

[3.2.1.](#) Subnet Prefix Information block format

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Network                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Prefix   |   Flags |h|d|   Stat-len   |   Optional stats...   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Addr = IPv4 network number (4 octets)

Prefix = Prefix length (1 octets)

Flags = Flags field (Undefined bits must be zero) (1 octet)

'd' = Deprecate flag (explained below)

1 : Deprecation of this subnet is requested

0 : No deprecation is needed

'h' = Hierarchical flag (explained below)

1 : Client will be allocating addresses from this subnet

0 : Client will be relaying DHCP requests to the server
from this subnet.

Stat-len = Length of the optional statistics information field
(zero if no statistics are included)

The 'd' flag may only be returned by the Server to the Client (inside a DHCPACK packet, in response to a DHCP RENEW). It's presence means

that the Client should prepare to give up the subnet. For example, if the Client is assigning addresses from this subnet to other clients, it should cease doing so immediately and should not renew any leases when clients ask for renewal. As soon as all addresses in the subnet are unallocated, the Client should send a DHCPRELEASE message to the Server, including a Subnet Prefix Information suboption for the subnet in order to release the subnet. The format of this message is described farther below.

The 'h' flag tells the Client how the Server intends for the Client to use the allocated subnet. It is interpreted in the same manner as that in the Subnet-Request suboption. In response to a Subnet-Request, the Server should normally specify the 'h' flag in the same manner as it was in the Subnet-Request suboption from the Client. The Server MAY, however, change the 'h' flag from that specified in the Subnet-Request suboption if it has been configured to override the Client's request.

If any usage statistics information is to be included, then the "Stat-len" field specifies the number of bytes of statistics information which is included. See below for more information. If no statistics information is included, then this byte MUST be zero. The "Stat-len" field SHOULD always be zero when this suboption is sent by the DHCP Server. The usage statistics information is intended for use only to report usage statistics from the Client to the Server.

[3.2.1.1](#). Subnet Usage Statistics

The Subnet-Information suboption may also include usage statistics information. If this information is included, then the "Stat-len" (Statistics length) field MUST be set to the number of bytes of statistics information which is being included. The statistics information MUST be in the following form and order:

```

      0                               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|                               |
|           High water         |
+---+---+---+---+---+---+---+---+
|                               |
|           Currently in use    |
|                               |

```

```

+---+---+---+---+---+---+---+---+---+---+
|           Unusable           |
+---+---+---+---+---+---+---+---+---+---+

```

"High water" refers the to "high water mark" of allocated addresses within the subnet. I.e., the largest number of addresses which were ever allocated at one time from the subnet.

"Currently in use" refers to the number of addresses currently allocated in the subnet.

"Unusable" refers to the number of addresses which are currently unusable for any reason (such as a client returning a DHCPDECLINE, or finding the address already in use).

Additional statistics may be added to this option in the future. If

so, they MUST be appended to the end of the option. All statistics fields MUST remain in the same order. Use the all ones value (0xFFFF) in order to skip reporting a number for a particular field. Fewer fields may be included than what is specified in any current RFC, but all fields which are included MUST remain in order specified here.

[3.3.](#) Subnet-Name Suboption format

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+---+---+---+---+---+---+---+---+---+---+
|           3           |   Len   | Name ...
+---+---+---+---+---+---+---+---+---+---+

```

Len = length of the sub-option (min. length of 1) (1 octet)

The Subnet-Name suboption may be used in order to pass a subnet name to the server for use during allocation. This name may be used for any purpose but is intended to tell the server something extra about the needed subnet; for example, "sales department", "customer 1002", "address pool F00", or some such. The "name" should NOT be NULL terminated since the "len" field already specifies the length of the name. The "Name" in this suboption is simply a number of octets and need not be ASCII text.

3.4. Suggested-Lease-Time Suboption format

[illegible]

Len = length of the sub-option (always 4 for this suboption) (1
octet)

The Suggested-Lease-Time suboption MAY be included by the Server in order to suggest the lease time to be used by the Client when allocating addresses from the subnet allocated. The four (4) octet value of the lease time is in the same format as that of the "IP Address Lease Time" option (option 51), as described in [\[RFC2132\]](#).

If included, this suboption **MUST** appear only once. (Inclusion of multiple such suboptions would result in ambiguity as to which applied to which subnet.) If different suggested lease times are

needed, the Server SHOULD, instead, reply with only one offered subnet and should set the "Server flag" in the Subnet-Information suboption to indicate to the Client that it should send another subnet request to gather the others.

The Client SHOULD use a lease time, when allocating addresses from the subnet, which is the lesser of the remaining lease time of the subnet itself and the Suggested-Lease-Time suboption.

[4.](#) Requesting allocation of a subnet

[4.1.](#) Client DHCPDISCOVER message

The DHCP Client creates a DHCPDISCOVER message including the Subnet Allocation option, and its set of suboptions, to request allocation of a subnet. The DHCP Client should include the Subnet-Request suboption, specifying the prefix length of the subnet requested. The 'h' bit should be set to 1 if the Client intends to control allocation of addresses within the subnet itself, or 0 if the Server should retain control of addresses within the subnet. More than one

Subnet Allocation option may appear in a DHCPDISCOVER message, however the client SHOULD limit the number of requests, noting that the DHCP replies will need to include the Subnet-Information suboption, which takes up more space than the Subnet-Request suboption.

If more than one subnet is being requested, multiple Subnet-Request suboptions MAY be included or multiple DHCPDISCOVER messages MAY be sent instead. The prefix length field of each Subnet-Request suboption MUST be either 0, or in the range of 1 to 30 inclusive.

The DHCP "IP address lease time" option (code 51) MAY be included in the DHCPDISCOVER message to specify the lease time the Client is requesting for the subnet. If not present, no suggested lease time is given.

The DHCP "Client ID" option (code 61) MAY be included in the DHCPDISCOVER message as it may be used by the Server in performing the subnet allocation.

[4.2.](#) Server DHCPOFFER message

Upon receiving a DHCPDISCOVER containing the Subnet Allocation option, the DHCP Server should respond with a DHCPOFFER message including the Subnet-Information suboption in order to specify the subnet(s) which it is willing to allocate to the Client in order to fulfill the request(s).

The Server need not reserve the subnets which are being offered, but SHOULD strive to not offer the same subnets to another DHCP Client until a reasonable time period (implementation dependent) has passed. (This is similar to normal DHCP address allocation.)

The Server MUST NOT include the Subnet-Request suboption in the DHCPOFFER. The same information is already present in the Subnet Information suboption(s) which SHOULD be included in the DHCPOFFER.

The Server SHOULD also include the IP address lease time option (option 51) in the DHCPOFFER message. This gives the lease time for all subnets given in all Subnet-Request suboptions contained in the DHCPOFFER message. The Server MAY also include the Renewal and/or

Rebinding options in order to further control the "T1" and "T2" lease timers of the client. There MUST be only one IP address lease time, rebind, and/or renew option each in the DHCP OFFER message.

The Server MAY set the "Server flag" ('s') to 1 to indicate that it would like to allocate one or more additional subnet(s) to the Client. This indicates that the Client should send another DHCPDISCOVER message specifying a zero prefix length field, P, in order to request the additional subnet allocation(s) information. This may be necessary if the subnets are to be allocated with different lease times, for example.

The "Client flag" ('c') MUST be set to 0 to indicate this is a Server response to a client request for a new subnet allocation and not a response to a request for information about already allocated subnets.

The Server SHOULD set the DHCP yiaddr value to all zeros (0.0.0.0) and the Client MUST ignore fields having to do with address assignment if the packet contains a Subnet Allocation option. In other words, a DHCP packet exchange can not provide subnet allocation and address assignment simultaneously.

[4.3.](#) Client DHCPREQUEST message

When sending a DHCPREQUEST, the Client MUST NOT modify any fields of all Subnet-Information suboptions received from the Server. However, the Client MAY choose not to include some Subnet-Information suboptions when issuing the DHCPREQUEST. Subnet-Request suboptions MUST NOT be included in the DHCPREQUEST message, only the Subnet-Information suboption(s) should be included.

[4.4.](#) Server DHCPACK message

The DHCP Server, upon receipt of the Client's DHCPREQUEST message, MAY refuse allocation of any subnets (for example, if they have been allocated elsewhere in the meantime), however since the Server should have set aside the subnets offered for a short period of time, and since the Client should have requested the subnets within a short period of time after receiving the offer(s) from the server(s), this last minute rejection should be rare. The DHCP Server MUST NOT change the network number(s) or prefix length(s), however it MAY remove some Subnet-Information suboptions from the list.

The Server SHOULD include the IP address lease time option specifying the lease period for all subnet(s) in the DHCPACK. All subnets allocated in one DHCP message will have the same lease time and only one IP address lease time option must appear in the DHCP message.

If the Server has internal information which states that the Client should be allocated more subnets than were requested, the Server MAY set the 's' bit in the last Subnet-Information suboption to indicate that the Client needs to request more subnets (as described above).

The allocable unit is the tuple (network number, prefix length). Multiple subnets may be allocated in one DHCPACK, however since there can be only one Lease-time option, each subnet allocated is assigned the same lease time. Each subnet lease tuple (network number, prefix length) MAY be renewed or released individually.

[5.](#) Client renewal of subnet lease

[5.1.](#) Client RENEW DHCPREQUEST message

The Client MUST renew all subnets allocated with a lease time in much the same manner as renewing an allocated IP address. Renewal timers need not be set in exactly the same manner, however. If Renewal and/or Rebinding options were included in the DHCPACK of the subnet allocation, then these "T1" and "T2" timers should be used just as they would be in the case of address allocation timers.

The DHCPREQUEST message MUST include a Subnet-Information suboption for which the Client is seeking renewal of the lease. This Subnet-Information suboption may optionally include subnet usage statistics, as described above in the Subnet-Information suboption format section (3.2).

The subnet network number field ("Network") and subnet prefix length field ("Prefix") MUST agree with the values as they were originally allocated to the Client by the Server. In any of the statistics fields (High, Current, Unusable), a value of all ones (0xffff) SHOULD be used if the Client has no information to report for a statistic.

[5.2.](#) Server RENEW DHCPREQUEST response

The Server MAY respond to a subnet RENEW with either a DHCPACK or DHCPNAK response. If a DHCPNAK response is given the Client MUST immediately stop using the subnet(s) specified and, if possible, notify all Clients with addresses allocated from this subnet that their addresses are no longer valid. The Client MAY, of course, send a DHCPDISCOVER message containing the Subnet Allocation option and the Subnet-Request suboption in order to acquire another subnet for use. In general, the Server should ask the Client to deprecate subnets by using the 'd' bit of the Subnet-Information suboption in a DHCPACK message (see below).

If a DHCPACK response is given, the "Deprecate" ('d') bit of the flags field in the Subnet-Information suboption may also be set. This indicates the DHCP Client should "prepare to stop using this subnet". If the Client is allocating IP addresses for other clients

from this subnet (e.g. via DHCP), the Client SHOULD immediately stop allocating such addresses. Once all allocated addresses in the subnet have been released, the Client SHOULD send a DHCPRELEASE message, including the Subnet-Information suboption (with optional usage statistics) in order to release the subnet(s) back to the Server.

[5.3.](#) Client DHCPRELEASE message

The DHCP Client should send a DHCPRELEASE message in order to release allocated subnet(s) back to the Server when it is finished using them. This message MUST NOT include the Subnet-Request suboption, but MUST include one or more Subnet-Information suboptions, and optionally including usage statistics.

The Client MUST release the same subnet(s) of the same prefix length ("Prefix"), as was originally allocated. The Client MAY release a subset of the subnets which were allocated originally. In other words, the allocable unit is the tuple (network number, prefix length). Multiple subnets may be allocated in one DHCPACK, however each subnet MAY be released individually.

[5.4.](#) Server DHCPFORCERENEW message

The DHCP Server may issue a DHCPFORCERENEW [[RFC3203](#)] message containing the Subnet Allocation option and the Subnet-Information suboption. Upon receiving this message, the DHCP Client MUST issue a DHCPREQUEST message to the DHCP Server in order to renew the lease on the subnet mentioned. No other subnets allocated to the Client are effected. As is the case with all DHCP RENEW messages, the Client may include subnet usage information in the Subnet-Information suboption in order to report subnet usage statistics, or set the "Stat-len" field to 0 if no statistics are to be reported.

If the Server responds to this DHCPREQUEST with a DHCPNAK message, then the Client MUST immediately stop using the subnet(s) and, if possible, notify all Clients with addresses allocated from this/these subnet(s) that their addresses are no longer valid. The Client MAY, of course, send a DHCPDISCOVER message containing the Subnet Allocation option and the Subnet-Request suboption in order to

acquire another subnet for use.

[6.](#) Client requesting previously allocated subnet information

A DHCP Client may request from the DHCP Server a list of what subnets are currently allocated to the Client. This may be used to recover from a restart if the Client does not have local storage in order to retain the information itself. (For example of this, see the [section 8.2](#) below.)

[6.1.](#) Client Initial DHCPDISCOVER Message

The DHCP Client DHCPDISCOVER message, when used to discover already allocated subnet information, should contain a Subnet-Request suboption with the "Prefix" field set to 0 and with the 'i' flag set to 1 to indicate that the Client is seeking already allocated subnet information from the Server. No Subnet-Information suboptions should be included in this message. Note, no Subnet-Information suboption is included in this message, since the client would not know of any subnet to request at this point.

This DHCPDISCOVER message MAY be unicast to a particular DHCP Server, or broadcast in the normal fashion.

[6.2.](#) Server Initial DHCPOFFER Response

Any DHCP Server which has allocated subnets to the Client should respond to the DHCPDISCOVER message with a DHCPOFFER message. The

DHCPOFFER message should contain one or more Subnet-Information suboption(s) telling the prefix of the subnet(s) allocated to the Client.

The Server SHOULD, internally, retain an ordered list of subnets which are allocated to each Client. In response to an initial DHCPDISCOVER message requesting allocated subnet information (i.e., one with the 'i' flag set to 1, but not carrying a Subnet-Information suboption), the server returns in the DHCPOFFER message the subnet(s) information for the first subnet(s) from this list. If the end of the list has been reached, then the 's' bit of the last Subnet-Information suboption included in the message MUST be set to 0. If there are more subnets in the list, the 's' bit MUST be set to 1 to indicate to the Client that more information is available. Since this information is in response to a client request for previously allocated subnet information, the 'c' bit MUST be set to 1.

[6.3.](#) Client Additional DHCPDISCOVER Messages

The Client, upon receiving any Server DHCPOFFER messages containing Subnet Information suboption information with the 'c' ("Client") bit set, should gather the network number ("Network") and prefix length

("Prefix") information from the message.

If the 's' bit is set in the last of the Subnet-Information suboptions included in the message, then the client SHOULD construct a new DHCPDISCOVER message containing the Subnet Allocation option and the last Subnet-Information suboption from the Server's message. This message SHOULD then be sent back to the same DHCP Server originating the DHCPOFFER message. The 'c' and 's' bits MUST retain the same settings they had from the Server's DHCPOFFER message and the network number ("Network") and prefix length ("Prefix") fields MUST be unaltered as well.

If the 's' bit in all of the Subnet-Information suboptions from the Server was 0, then it indicates the Server has no more information about subnets allocated to the Client.

[6.4.](#) Server Additional DHCPOFFER Responses

The Server, upon receiving from a Client an additional DHCPDISCOVER

message for allocated subnet information retrieval, with the 'i' flag set to 1 and containing one or more Subnet Information suboptions with the 'c' and the 's' bits set, MUST use the network number ("Network") and prefix length ("Prefix") fields contained in the last such Subnet Information suboption in order to locate the position in the internal table of allocated subnets for this Client, and then return an DHCPOFFER message containing a Subnet-Information suboption giving information about the next set of subnets allocated to this Client. If this finishes the list in the table for this Client, then the 's' bit MUST be set to 0 to indicate there is no more information. Any Subnet Information suboptions encountered without both the 'c' and 's' bits set should be ignored by the Server.

[7.](#) DHCP Server Subnet Allocation method

The actual method of allocating subnets on the DHCP Server, as well as the configuration of what networks may be subnetted and how, is left up to the implementation.

8. Examples

Only the Subnet Allocation option and accompanying suboptions are displayed in these examples. All other fields in the DHCP messages are described in [[RFC2131](#)].

8.1. Example 1:

A DHCP Client requesting a subnet with prefix length 24 from which the Client will allocate addresses to other clients. The Server responds with an allocation of exactly the size requested:

The Client sends a DHCPDISCOVER message including a Subnet Allocation option with the Subnet-Request suboption:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      5      |      0      |      1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      | 0 | 0 |      24      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The Server responds with a DHCPOFFER message including a Subnet Allocation option with a Subnet-Information suboption, offering the subnet 10.0.1.0/24.

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11     |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      0      | 0 | 0 |      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      1      |      0      |      24      |      | 0 | 0 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The Client sends a DHCPREQUEST including a Subnet Allocation option with a Subnet-Information suboption:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11     |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      0      | 0 | 0 |      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      1      |      0      |      24      |      | 0 | 0 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The Server responds with a DHCPACK message including a Subnet Allocation option with a Subnet-Information suboption:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      0      | 0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      1      |      0      |      24      |      | 0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+

```

Later the Client sends a DHCPRELEASE message including a Subnet Allocation option with a Subnet-Information suboption:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      0      | 0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      1      |      0      |      24      |      | 0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+

```

8.2. Example 2:

A DHCP Client requesting two subnets, each with prefix length 24:

The Client sends a DHCPDISCOVER message including a Subnet Allocation option with a Subnet-Request suboption:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      9      |      0      |      1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      | 0|0|      24      |      1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      | 0|0|      24      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Server responds with a DHCPOFFER message including a Subnet Allocation option with a Subnet-Information suboption:

The DHCPOFFER specifies 1 subnet of size 24 and 1 subnet of size 28.

Internet-Draft

Subnet Allocation Option

March 2009

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      18      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      15      |      |0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |      10      |      0      |      3      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |      28      |      |0|0|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Client sends a DHCPREQUEST message including a Subnet Allocation option with a Subnet-Information suboption:

The Client decides that the subnet of size 28 is not sufficient so it doesn't include that subnet into the DHCPREQUEST message.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      |0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Server responds with a DHCPACK message including a Subnet Allocation option with a Subnet-Information suboption:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      |0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Later the Client sends a DHCPREQUEST message in order to renew the

lease on the one subnet and includes subnet usage information. It reports that a maximum of 10 addresses were allocated from the subnet since the last report, 7 addresses are currently allocated, and 2 addresses were found to be unusable.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      17      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      14      |      |0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      6      |      0      |      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      7      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Server responds with a DHCPACK message, however it signals to the Client that the subnet should be deprecated.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      |0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|1|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Client reloads at this point and upon completion of the reload sends a DHCPDISCOVER asking for information about all subnets which were allocated to it.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      5      |      0      |      1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      |1|0|      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Server responds with a DHCP OFFER, giving the subnet information of the one subnet which is allocated to the Client. Also the Server specifies that the one allocated subnet should be immediately deprecated. Note that the 's' ("Server") bit is 0 thus indicating that there is no more information available for this Client.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      |1|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|1|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+

```

The Client responds with a DHCP RELEASE message after having deprecated the subnet:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      220      |      11      |      0      |      SIS      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      8      |      |0|0|      10      |      0      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      2      |      0      |      24      |      |0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      0      |
+---+---+---+---+---+

```

9. Differences from DHCPv6 Prefix Delegation

The following differences may be noticed between Subnet Allocation as described in this document and DHCPv6 Prefix Delegation as described in [[RFC3633](#)]:

- o This option does not use anything like an "IA_PD" as is used in DHCPv6.
- o If the Server can not allocate a subnet, it remains silent, instead of returning a special response saying nothing is available.
- o DHCPv6 Prefix Delegation has no mechanism for returning subnet/prefix usage statistics.
- o DHCPv6 has no equivalent to the "subnet deprecation" flag as described here.
- o DHCPv6 Prefix Delegation makes no mention of what Client actions should result from receiving a DHCPNAK during a RENEW of a

delegation.

- o DHCPv6 has no equivalent of the subnet allocation "Network name" suboption, which may be used by the Server for various purposes, such as to specify a pool to use when allocating a subnet.
- o DHCPv6 Prefix Delegation corresponds to "Hierarchical Subnet Allocation" (setting the 'h' flag in the Prefix Information suboption). There is no V6 equivalent of clearing the 'h' flag, in which the Server retains authority over allocation of addresses from the subnet.
- o DHCPv6 Prefix Delegation has nothing to correspond to the Suggested-Lease-Time suboption.

10. Security Considerations

Potential exposures to attack are discussed in [section 7](#) of the DHCP protocol specification [[RFC2131](#)]. The Subnet Allocation option can be used to hoard all allocable subnets on a network.

Implementations should consider using the DHCP Authentication option [[RFC3118](#)] in order to provide a higher level of security if it is deemed necessary in their environment. Potential exposures to attack are discussed in [section 7](#) of the DHCP protocol specification in [[RFC2131](#)].

11. IANA Considerations

IANA is requested to assign DHCP option number 220 for this option, in accordance with [[RFC3942](#)].

No assignment of values for the suboption codes need be made at this time. New values may only be defined by IETF Consensus, as described

in [[RFC5226](#)]. Basically, this means that they are defined by RFCs approved by the IESG.

12. Intellectual Property Rights and Copyright

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

Internet-Draft

Subnet Allocation Option

March 2009

[13.](#) References

[13.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), March 1997.
- [RFC3942] Volz, B., "Reclassifying Dynamic Host Configuration Protocol version 4 (DHCPv4) Options", [RFC 3942](#), November 2004.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

[13.2.](#) Informative References

- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", [RFC 3118](#), June 2001.
- [RFC3203] T'Joens, Y., Hublet, C., and P. De Schrijver, "DHCP reconfigure extension", [RFC 3203](#), December 2001.
- [RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", [RFC 3396](#), November 2002.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", [BCP 122](#), [RFC 4632](#), August 2006.

Internet-Draft

Subnet Allocation Option

March 2009

Authors' Addresses

Richard A. Johnson
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
US

Phone: +1 408 526 4000
Email: raj@cisco.com

Jay Kumarasamy
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
US

Phone: +1 408 526 4000
Email: jayk@cisco.com

Kim Kinnear
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
US

Phone: +1 408 526 4000
Email: kkinnear@cisco.com

Mark Stapp
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134

US

Phone: +1 408 526 4000

Email: mjs@cisco.com

Johnson, et al.

Expires September 10, 2009

[Page 32]