

Network Working Group

INTERNET DRAFT

Obsoletes: <[draft-ietf-dhc-sso-03.txt](#)> Ralph Droms, Bucknell University

Bill Sommerfeld, Integrated Systems

October 1999

Expires April 2000

Glenn Stump, IBM

Pratik Gupta, IBM

The Server Selection Option for DHCP **<[draft-ietf-dhc-sso-03.txt](#)>**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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/1id-abstracts.txt>

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

Abstract

This option is provided by DHCP servers to DHCP clients so that clients may optionally select from among multiple offers received from DHCP servers in the network offer from a DHCP. The information contained in this option is a 16-bit unsigned integer which represents a value indicating the priority of the offer in which this option is contained.

Several server profiles are presented in appendices showing different ways in which the option can be used by a set of servers. Regardless of the profile, the client behavior is the same.

1. Requirements 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](#) [3].

2. DHCP Terminology

- o "DHCP client"

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

- o "DHCP server"

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

3. The DHCP Server Selection Option

DHCP [1] provides a powerful mechanism for automating and centralizing the administration of IP host configuration and has become a critical service in many large IP networks. Because of its importance in networks and because it can create a single point of failure for network operations (from a DHCP client's perspective), many network administrators choose to deploy many DHCP in order to enhance availability and/or performance of DHCP services.

However, for networks with multiple DHCP servers, the DHCP protocol does not provide a means by which a DHCP client may select from among the DHCP server offers according to policy determined by the network administrator. The protocol specification [1] currently states that:

"DHCP clients are free to use any strategy in selecting a DHCP server among those from which the client receives a DHCP OFFER message."

Thus, currently, client "policy", of which there is no standardization, determines which offer is selected. In practice, most vendor's implementation of policy here is very basic (e.g., first offer received or first acceptable offer received) and is "hard-coded" (i.e., non-configurable).

A mechanism which enables a server-based policy for determining how clients select among DHCP offers would allow a network administrator better control of the manner in which addresses are allocated across the multiple DHCP servers. This type of control takes on increased importance considering that IP address spaces cannot generally be shared across DHCP servers, thus requiring network administrators to divide the available addresses among multiple DHCP servers.

This document specifies an option that can be configured in DHCP servers and which can influence how DHCP clients select an offer from among those servers. The option described in this document is a new DHCP option [2].

This document also presents several server profiles showing how this option can be used by administrators to effect various policies. Regardless of the profile in use, the operation of the client is the same -- the profiles listed are just different ways of assigning

priorities to dhcp servers.

3.1 Motivation

In general, the motivation for inventing the DHCP Server Selection option originates from shortcomings identified in networks where multiple DHCP servers are used enhance DHCP serving performance, availability, or both. Specifically, these problems are:

1. Server Segregation
2. Server Aggregation
3. Server Introduction and Deprecation
4. DNS IP Address Mapping Stability

There may well be others which are important - or will be important given new capabilities in the future - and not listed above. The DHCP Server Selection mechanism defined herein must be flexible enough to accommodate future requirements.

3.1.1 Server Segregation: Differentiating Servers Based on Varying Priorities

Multiple DHCP servers can also be segregated to differentiate the importance or roles of some DHCP server or servers vs. others. The simplest and most common example here is where one DHCP server is intended as a "primary" or "preferred" server for a subnet while some other server is intended as a "secondary" or "backup" server capacity in the case the primary fails. Typically, the preferred server(s) will have the most number of available addresses and the backups have some lesser number, which are only intended for use when the preferred server(s) fail or their addresses are depleted. Again, all servers are assumed to provide an equivalent set of options to a requesting client.

Segregated DHCP server deployments present two fundamental problems for today's DHCP implementations:

1. There is no standard means for a client to differentiate which of many equivalent offers to accept. Assuming that most client implementations will accept the first of many equivalent offers received, a "backup" server's address pool will be depleted each time its offers are served first to a client. This may happen frequently, for example, during times of peak loads on the preferred servers.
2. Once an address lease from an "alternate" DHCP server is selected (for any reason), that address will likely be re-selected on subsequent client reboot/init-reboots.

A mechanism which enables DHCP clients to select an offer based on an administrative preference by encouraging clients to always select a preferred DHCP server (or order of servers) over others who respond

in the network.

DISCUSSION:

Such a mechanism could also help in preventing DHCP service disruption due to the accidental introduction of rogue DHCP servers. That is, if all clients would be configured to choose offers with the DHCP Server Selection when present, and all DHCP Server Vendors would disable the option by default, then a DHCP offer from a rogue DHCP server accidentally started by a novice would essentially be ignored. Is this worth mentioning as a motivation?

3.1.2 Server Aggregation: Grouping Servers of Equal Priority

Multiple DHCP servers, typically two, which have essentially equivalent configuration characteristics - typically in terms of the options served and the number of addresses available -- be aggregated to enhance the responsiveness and availability of DHCP services to a particular subnet.

However, in the event that one of the DHCP servers responds more quickly than others the number of addresses available at one server (e.g., a "fast" server) may be disproportionately low. Thus, failure of another server in the aggregated set (e.g., the other, "slower" server) will cause a disproportionately high risk in availability of DHCP services (e.g., when a slow server, with more available addresses, fails or is removed from operation temporarily).

A mechanism which enables DHCP clients to select an offer based on server- administrated preference could allow clients to choose leases from among servers in a way which could maintain a prescribed balance of address availability across servers (by maintaining a balance in the relative address depletion rates).

3.1.3 DHCP Server Introduction or Deprecation

Currently, the process of changing the number of DHCP servers in a network cannot generally be done gradually or gracefully because, currently, there is no standard means of allowing servers to share IP address spaces. A mechanism which enables DHCP clients to select an offer based on an administrative preference could help in this process by identifying particular DHCP servers to (or from) which clients should "migrate" (rather than continuing to use an existing server).

3.1.4 DNS Address Mapping Stability

Regardless of whether multiple DHCP servers are aggregated, segregated, or a combination of both, a "mobile" DHCP client which...

1. moves frequently across many networks or subnetworks, and/or.
2. does not keep track of which leases are active beyond their

current lease.
single subnet over a series of connections and reconnections. This is due to the fact that, the client would not necessarily choose a lease based on an active or previous lease association since it is not instructed to do so and/or is not aware that one exists.

In network environments where either dynamic DNS updates are not employed or where there is a desire to minimize the number of updates to DNS mappings, a mechanism to identify a lease representing a previous address association can be used to, in effect, reduce the number of changes in a DNS host-name-to-address mapping (i.e., A-RR).

3.2 DHCP Server Selection Option Format

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

Code	Len	Priority
+-----+	+-----+	+-----+
TBD	2	p
+-----+	+-----+	+-----+

where: "p" is an unsigned 16 bit integer representing the priority value chosen for the server (x'0000'- x'FFFF', inclusive).

DISCUSSION:

Is 16 bits the right number of bits, as opposed to 8 or 32?

It is envisioned that this field can be used in many ways to accomplish various system behavior across servers. However, the values within this field must be administered consistently across all DHCP servers in "the system" (i.e., those serving a particular subnet) and across DHCP server vendors (so that the same basic capabilities are administered) in order to achieve the desired behaviors.

To that end, the policies for acceptable uses of the priority field will be directed through the specification of DHCP Server Selection option "profiles", which will be maintained - at least for now -- as appendices to this document. Each profile will list how the priority field value is derived, what DHCP serving behaviors are possible, and how these behaviors are achieved.

4. DHCP Client Behavior

A client supporting the DHCP Server Selection option MUST use the DHCP Server Selection option as the primary discriminator for selecting among multiple offers from servers. The highest priority value gets first preference. If two DHCP Server Selection priority values are equivalent, then the DHCP client can use other mechanisms as described in [RFC 2131](#) to choose among the offers.

DISCUSSION: Should the client be required to use the DHCP Server

Selection offer if present or only as a tie-breaker if the offers are equivalent (in terms of options offered)? In general, how should this option relate to other decision factors (implicit or explicit).

5. DHCP Server Behavior

When a DHCP server supports the DHCP Server Selection option, the server MUST select a value for the field according to the policy set forth by a selected DHCP Server Selection Option "profile". The set of standardized DHCP Server Selection profiles will be maintained by the IETF DHC Working Group. The current list of profiles under consideration is maintained in the appendices of this document.

It is intended that all of the DHCP servers serving a particular subnet or set of hosts MUST support the same profile in order to achieve the associated/desired DHCP service behavior for that subnet or set of hosts.

6. Security Considerations

DHCP currently provides no authentication or security mechanisms. Potential exposures to attack are discussed in [section 7](#) of the protocol specification [1].

The server selection option might be used to redirect DHCP clients to accept configuration parameters from a "bad guy" DHCP server.

7. References

- [1] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [2] S. Alexander, R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), March 1997.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," [RFC 2119](#), March 1997.

8. Author Information

Pratik Gupta
IBM, Inc.
4205 S.Miami Blvd
Research Triangle Park, NC 27709
Phone: (919)254-5616
email: pratik_gupta@vnet.ibm.com

Glenn Stump
IBM, Inc.
4205 S.Miami Blvd
Research Triangle Park, NC 27709
Phone: (919)254-5616

email: glennstump@vnet.ibm.com

Ralph Droms
323 Dana Engineering
Bucknell University
Lewisburg, PA 17837
Phone: (717) 524-1145
email: droms@bucknell.edu

Bill Sommerfeld
Integrated Systems, Inc.
1 Tara Boulevard suite 403
Nashua NH 03062
Phone: (603) 897 2060
email: sommerfeld@alum.mit.edu

Appendix A: Profile 0 [Rank (only)]

The DHCP Server Selection option specifies a mechanism for an administrator to determine the policy governing how a client chooses among multiple DHCP offers. In this profile, Profile 0, the selection criteria used to govern DHCP OFFER selection policy is a simple, relative "ranking" value. That is, a client will select the offer from the server with the highest designated priority.

[A.1](#) DHCP Server Selection Option Format

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

Code	Len	Priority
+-----+	+-----+	+-----+
TBD	2	r reserved
+-----+	+-----+	+-----+

where:

r = rank, or relative priority, of server (higher order 8 bits)

reserved = reserved for future use; must be x'00'

The priority field is actually a concatenation of the "rank", "address" availability, and "flags" field. The "flags" field, represented in the lower order byte, can be used to instruct the client to give preference to offers from servers which have a currently active or recently expired lease association. The priority field, represented in the high order byte, can be used to instruct the client to accept a lease from a server according to some other criteria (see "DHCP Server Behavior" below).

[A.2](#) DHCP Server Behavior

When a server supports the "rank" field, the value can be statically pre-configured or dynamically calculated and set to any integer between x'00' and x'FF', inclusive. The values can be coordinated

across DHCP servers to achieve the desired priority system behavior.

A.3 Application Notes

The DHCP Server Selection option allows the DHCP/network administrator to control DHCP services characteristics by influencing the way addresses are allocated across a set of DHCP servers. The administrator may wish to configure all of the servers to have equal serving priority or to configure some to have a greater priority than others. Further, not only can the network administrator prescribe an address allocation scheme across DHCP servers, but the option can also be used to enable the servers to return to the prescribed state in the event that a server failure resulted in an imbalance.

The following outlines how the DHCP Server Selection option can be used in each of these situations to achieve the desired service behavior. Note that these scenarios presume that there are no DHCP server-to-server coordination mechanisms which could be used to synchronize the server's databases and effectively share address spaces.

A.3.1 Server Segregation

DHCP servers which are assigned different rank values, "r", are viewed by DHCP clients as having varying priorities. That is, the DHCP servers are segregated according to a distinct priority system (set by the DHCP system administrator). Clients will choose an offer received from the server with the highest assigned rank value, "r".

Once a client can differentiate priorities among DHCP servers, the DHCP administrator can manipulate the priorities across DHCP servers to affect the DHCP serving system behavior, such as:

- primary-and-backup
- graceful addition or deletion of DHCP servers

A.3.2 Server Aggregation

DHCP servers which are assigned the same rank value, "r", are viewed by DHCP clients as equal. That is, the servers are an aggregated set of equivalent servers, and offers from equivalent servers can be selected using some other criteria (e.g., first offer received).

Appendix B: Profile 1 [Rank with Address Binding]

The DHCP Server Selection option specifies a mechanism for an administrator to determine the policy governing how a client chooses among multiple DHCP offers. In this profile, Profile 1, the selection criteria used to govern server selection policy is a simple, relative "ranking" system plus a mechanism which can be used to further differentiate equally-ranked servers based on a preference for offers representing a previous address binding. That is, a client will select an offer from a particular server because that

server has the highest designated priority, and, in the event that one or more servers have equal priority (i.e, are aggregated), the client will select the offer from the server having, in order of preference, an active address binding or a previous binding which is still available to the client.

B.1 DHCP Server Selection Option Format

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

Code	Len	Priority
TBD	2	r flags 0000

where:

r = rank, or relative priority, of server (higher order 8 bits)

flags = b `xAxP` where:

- x = reserved, must be x'0'
- A = offer contains currently active lease for client
- P = offer contain previously active lease available to client
- 0000 = bits reserved for future use; must be x'0'

Here, the priority field is actually a concatenation of the "rank", and the "address binding" fields. The use of the rank field is described in Profile 0. The address binding field indicates whether this DHCP OFFER represents a currently active (A-flag) or previous (P-flag) address lease binding.

B.2 DHCP Server Behavior

In addition to the setting the "rank" value as described in Profile 0, the address binding field MUST be set according to whether the OFFER contains an currently active address binding for the client (A-flag set, P-flag cleared), a previous binding (P-flag cleared, A-flag set), or no previous association (A- and P-flags cleared).

The "rank" and "flags" values used consistently (i.e. by all DHCP servers serving a subnet) enables a client to select a DHCP OFFER representing the highest rank and, when one or more offers are of equal rank, an active or (else) previous address binding.

B.3 Application Notes

The primary reason for using the "address binding" field is to minimize the number of DNS updates which are required due to changes in IP address.

Appendix C: Profile 2 [Rank with Address Balancing]

The DHCP Server Selection option specifies a mechanism for an

administrator to determine the policy governing how a client chooses among multiple DHCP offers. In this profile, Profile 2, the selection criteria used to govern server selection policy is a simple, relative "ranking" system plus a mechanism which can be used to further differentiate equally-ranked servers based on the availability of address leases at the servers. That is, a client will select an offer from a particular server because that server has the highest designated priority, and, in the event that one or more servers have equal priority (i.e., are aggregated), the client will select the offer from the server with the most available addresses (relative to the total number of addresses available for allocation by the server).

C.1 DHCP Server Selection Option Format

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

Code	Len	Priority
TBD	2	r v 0000

where:

r = rank, or relative priority, of server
(higher order 8 bits)

v = percentage of available leases remaining in pool,
calculated as:

$$v = (\text{\#remaining addresses} / \text{\#total addresses}) * 100 / 6$$

0000 = bits reserved for future use; must be x'0'

Here, the priority field is actually a concatenation of the "rank" field and the "address availability" field. The use of the rank field is described in Profile 0. The address availability field expresses the number of addresses currently available in a server relative to the number originally defined in the pool (expressed as a percentage and as calculated above).

C.2 DHCP Server Behavior

In addition to the setting the "rank" value as described in Profile 0, the address availability field MUST be set according to the relative number of addresses remaining "v", as expressed as an integer representing a percentage and calculated as follows:

$$v = \text{INT}[(\text{\#remaining addresses} / \text{\#total addresses}) * 100 / 6]$$

The "rank" and "v" values used consistently (i.e. by all DHCP servers serving a subnet) enables a client to select a lease from the highest ranked server with the best availability of addresses.

C.3 Application Notes

The important point to note here is that the client will select the highest rank server with "the best address availability", where "best

server" means the one with the largest percentage of addresses available remaining from its original pool. In other words, the "system" of DHCP servers using this Profile 2 will tend to maintain the original, prescribed "balance" of addresses allocated across the (equivalent) servers.

Maintaining this prescribed balance of address availability is important in cases one of the DHCP servers in an aggregated set responds more quickly than others and therefore depletes its addresses pools at a disproportionately higher rate than others in the set.. Thus, failure of one of the other servers in the set will cause a disproportionately high risk in availability of DHCP services.

To illustrate, suppose there is a set of two equivalent DHCP servers, A and B, which service a particular subnet. Each server is assigned half of the available address space to allocate to clients. Now suppose server A responds to clients' DHCPDISCOVERs significantly faster than server B. Without some other means for expressing preference, the clients will likely repeatedly choose DHCPOFFERS from A, thus depleting its pool at a much faster rate than B. Now suppose server B encounters a failure or is taken out of operation for maintenance. Server A, with perhaps very few addresses remaining for allocation, is left to service the entire subnet.

DHCP Server Selection based on address availability can be used to maintain a balance of addresses across aggregated servers and thus optimize availability of services in the event of a server becomes inoperative (e.g., for maintenance or due to failure).

Appendix D: Profile 3 [Rank, with Address Balancing, Binding]

The DHCP Server Selection Profiles 0, 1, and 2 introduce mechanisms to allow DHCP clients to differentiate between DHCPOFFERS from servers based on a "rank", "previous address binding", and "address availability". In Profile 3, these mechanisms are all combined into a single priority system with the given preference of:

1. rank
2. address availability
3. previous address binding

D.1 DHCP Server Selection Option Format

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

Code	Len	Priority
TBD	2	r v flags

where:

r = rank, or relative priority, of server
(higher order 4 bits)

```

v =      percentage of available leases remaining in pool
      (lower order 4 bits)
      calculated as:
      v = (#remaining addresses / #total addresses) / 6
flags =  b`xAxP' where:
      x = reserved, must be x'0'
      A = offer contains currently active lease for client
      P = offer contains previously active
      lease available to client

```

The priority field is actually a concatenation of the "rank", "address availability, and "flags" field. The "flags" field, represented in the lower order byte, can be used to instruct the client to give preference to offers from servers which have a currently active or recently expired lease association. The priority field, represented in the high order byte, can be used to instruct the client to accept a lease from a server according to some other criteria (see "DHCP Server Behavior" below).

DISCUSSION:

Can we cram all fields into 8 bits if we limit rank values to 0,1,2,3 (2 bits)?

D.2 DHCP Server Behavior

A DHCP server offering support for the DHCP Server Selection option MUST support and set all three fields as described in Profiles 0, 1, and 2.

D.3 Application Notes

Using the given combination and preference of DHCP Server Selection mechanisms, a client will choose the DHCP OFFER with the highest designated rank. In the event the highest rank servers are aggregated, the aggregated server with best address availability status will be selected. Given equal rank and address availability status, preference for an active or (else) previous address binding may determine the selection.

Appendix E: Profile 4 [Rank with Binding, Address Balancing]

Profile 4 is identical to Profile 3, with the exception that the bindings flags field, "flags" is swapped with the "address availability" field, "v", thus yielding client preference to an active or previous binding over a server's address depletion status (assuming identical server rank "r" values].

E.1 DHCP Server Selection Option Format

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

Code	Len	Priority
+-----+	+-----+	+-----+

	TBD		2		r		flags		v	
+	-----	+	-----	+	-----	+	-----	+	-----	+

where:

r = rank, or relative priority, of server
 (higher order 4 bits)
 v = percentage of available leases remaining in pool
 (lower order 4 bits)
 calculated as:
 $v = (\text{\#remaining addresses} / \text{\#total addresses}) / 6$
 $\text{flags} = \text{b`xAxP'}$ where:
 x = reserved, must be $x'0'$
 A = offer contains currently active lease for client
 P = offer contain previously active lease available
 to client

E.2 Application Notes

Using the given combination and preference of DHCP Server Selection mechanisms, a client will choose a DHCP OFFER with the highest designated rank. In the event the highest rank servers are aggregated, the aggregated server with an active or (else) previous address binding will be selected. Given equal rank and binding values, preference for the server best address availability status may determine the selection.

Full Copyright Statement

Copyright (C) The Internet Society (1998,1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF
MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.