

Network Working Group
Request for Comments: 3527
Category: Standards Track

K. Kinnear
M. Stapp
R. Johnson
J. Kumarasamy
Cisco Systems
April 2003

Link Selection sub-option
for the Relay Agent Information Option for DHCPv4

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract


This document describes the link selection sub-option of the relay-agent-information option for the Dynamic Host Configuration Protocol (DHCPv4). The giaddr specifies an IP address which determines both a subnet, and thereby a link on which a Dynamic Host Configuration Protocol (DHCP) client resides as well as an IP address that can be used to communicate with the relay agent. The subnet-selection option allows the functions of the giaddr to be split so that when one entity is performing as a DHCP proxy, it can specify the subnet/link from which to allocate an IP address, which is different from the IP address with which it desires to communicate with the DHCP server. Analogous situations exist where the relay agent needs to specify the subnet/link on which a DHCP client resides, which is different from an IP address that can be used to communicate with the relay agent.

1. Introduction

In RFC 2131, the giaddr specifies an IP address which determines a subnet (and from there a link) on which a DHCP client resides as well as an IP address which can be used to communicate with the relay agent. The subnet-selection option [RFC 3011] allows these functions of the giaddr to be split, so that when one entity is performing as a

DHCP proxy, it can specify the subnet/link from which to allocate an IP address that is different from the IP address with which it desires to communicate with the DHCP server.

Analogous situations exist where the relay agent needs to specify the subnet/link on which a DHCP client resides, which is different from an IP address that can be used to communicate with the relay agent. Consider the following architecture:



In the usual approach, the relay agent would put IP address Y into the giaddr of any packets that it forwarded to the DHCP server. However, if for any reason, IP address Y is not accessible from the DHCP server, this approach will fail. There are several reasons why IP y might be inaccessible from the DHCP server:

- o There might be some firewall capability in the network element in which the relay agent resides that does not allow the DHCP server to access the relay agent via IP y.
- o There might not be an IP y. An example would be the case where there was only one host and this was a point to point link.

In any of these or other cases, the relay agent needs to be able to communicate to the DHCP server the subnet/link from which to allocate an IP address. The IP address, which will communicate to the DHCP server the subnet/link information, cannot be used as a way to communicate between the DHCP server and the relay agent.

Since the relay agent can modify the client's DHCP DHCPREQUEST in only two ways, the giaddr and the relay-agent-info option, there is a need to extend the relay-agent-info option with a new sub-option, the link-selection sub-option, to allow separation of the specification of the subnet/link from the IP address to use when communicating with the relay agent.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [RFC 2119].

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 relay agent"

A DHCP relay agent is a third-party agent that transfers BOOTP and DHCP messages between clients and servers residing on different subnets, per [RFC 951] and [RFC 1542].

- o "DHCP server"

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

- o "link"

A link is a communications facility or medium over which nodes can communicate at the link layer, i.e., the layer immediately below IPv4. Examples are Ethernets (simple or bridged); PPP links; X.25, Frame Relay, or ATM networks; and internet (or higher) layer "tunnels", such as tunnels over IPv4 or IPv6 itself.

- o "subnet"

A subnet (for the purposes of this document) consists of a routable address range. It may be one of several that exist on a link at the same time.

3. Link selection sub-option definition

The link-selection sub-option is used by any DHCP relay agent that desires to specify a subnet/link for a DHCP client request that it is relaying but needs the subnet/link specification to be different from the IP address the DHCP server should use when communicating with the relay agent.

The sub-option contains a single IP address that is an address contained in a subnet. The value for the subnet address is determined by taking any IP address on the subnet and ANDing that address with the subnet mask (i.e., the network and subnet bits are left alone and the remaining (address) bits are set to zero). This determines a single subnet, and when allocating an IP address, all of the other related subnets on the same link will also be considered in the same way as currently specified for the processing of the giaddr in [RFC 2131, Section 4.3.1, first group of bullets, bullet 4].

In scenarios where this sub-option is needed, the relay agent adds it whenever it sets the giaddr value (i.e., on all messages relayed to the DHCP server).

When the DHCP server is allocating an address and this sub-option is present, then the DHCP server MUST allocate the address on either:

- o the subnet specified in the link-selection sub-option, or;
- o a subnet on the same link (also known as a network segment) as the subnet specified by the link-selection sub-option.

The format of the sub-option is:

SubOpt	Len	subnet IP address			
-----+	-----+	-----+	-----+	-----+	-----+
5	4	a1	a2	a3	a4
-----+	-----+	-----+	-----+	-----+	-----+

A relay agent which uses this sub-option MUST assume that the server receiving the sub-option supports the sub-option and uses the information available in the sub-option to correctly allocate an IP address. A relay agent which uses this sub-option MUST NOT take different actions based on whether this sub-option appears or does not appear in the response packet from the server.

It is important to ensure, using administrative techniques, that any relay agent employing this sub-option is directed to only send packets to a server that supports this sub-option.

Support for this sub-option does not require changes to operations or features of the DHCP server other than to select the subnet (and link) on which to allocate an address. For example, the handling of DHCPDISCOVER for an unknown subnet should continue to operate unchanged.

In the event that a DHCP server receives a packet that contains both a subnet-selection option [RFC 3011], as well as a link-selection sub-option, the information contained in the link-selection sub-option MUST be used to control the allocation of an IP address in preference to the information contained in the subnet-selection option.

When this sub-option is present and the server supports this sub-option, the server MUST NOT offer an address that is not on the requested subnet or the link (network segment) with which that subnet is associated.

The IP address to which a DHCP server sends a reply MUST be the same as it would choose when this sub-option is not present.

4. Security Considerations

Potential attacks on DHCP are discussed in section 7 of the DHCP protocol specification [RFC 2131], as well as in the DHCP authentication specification [RFC 3118].

The link-selection sub-option allows a relay agent to specify the subnet/link on which to allocate an address for a DHCP client. Given that the subnet-selection option already exists [RFC 3011], no fundamental new security issues are raised by the existence of the link-selection sub-option specified in this document beyond those implied by the subnet-selection option [RFC 3011].

The existence of either the subnet-selection option or link-selection sub-option documented here would allow a malicious DHCP client to perform a more complete address-pool exhaustion attack than could be performed without the use of these options, since the client would no longer be restricted to attacking address-pools on just its local subnet.

There is some minor protection against this form of attack using this sub-option that is not present for the subnet-selection option, in that a trusted relay agent that supports the relay-agent-info option MUST discard a packet it receives with a zero giaddr and a relay-agent-info option when that packet arrives on an "untrusted" circuit [RFC 3046, section 2.1].

5. IANA Considerations

IANA has assigned a value of 5 from the DHCP Relay Agent sub-options space [RFC 3046] for the link-selection sub-option defined in Section 3.

6. Acknowledgments

Eric Rosen helped the authors to understand the need for this sub-option. Much of this document was borrowed, with only minimal modifications, from the document describing the subnet-selection option [RFC 3011].

7. References

7.1. Normative References

- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC 2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC 3011] Waters, G. "The IPv4 Subnet Selection Option for DHCP", RFC 3011, November 2000.
- [RFC 3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

7.2. Informative References

- [RFC 951] Croft, W. and J. Gilmore, "Bootstrap Protocol", RFC 951, September 1985.
- [RFC 1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, October 1993.

8. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

9. Authors' Addresses

Kim Kinnear
Cisco Systems
1414 Massachusetts Ave
Boxborough, Ma. 01719

Phone: (978) 936-0000
EMail: kkinnear@cisco.com

Mark Stapp
Cisco Systems
1414 Massachusetts Ave
Boxborough, Ma. 01719

Phone: (978) 936-0000
EMail: mjs@cisco.com

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

Phone: (408) 526-4000
EMail: jayk@cisco.com

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

Phone: (408) 526-4000
EMail: raj@cisco.com

10. Full Copyright Statement

Copyright (C) The Internet Society (2003). 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

