

Network Working Group
Request for Comments: 3442
Updates: 2132
Category: Standards Track

T. Lemon
Nominum, Inc.
S. Cheshire
Apple Computer, Inc.
B. Volz
Ericsson
December 2002

The Classless Static Route Option for
Dynamic Host Configuration Protocol (DHCP) version 4

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 (2002). All Rights Reserved.

Abstract

This document defines a new Dynamic Host Configuration Protocol (DHCP) option which is passed from the DHCP Server to the DHCP Client to configure a list of static routes in the client. The network destinations in these routes are classless - each routing table entry includes a subnet mask.

Introduction

This option obsoletes the Static Route option (option 33) defined in RFC 2132 [4].

The IP protocol [1] uses routers to transmit packets from hosts connected to one IP subnet to hosts connected to a different IP subnet. When an IP host (the source host) wishes to transmit a packet to another IP host (the destination), it consults its routing table to determine the IP address of the router that should be used to forward the packet to the destination host.

The routing table on an IP host can be maintained in a variety of ways - using a routing information protocol such as RIP [8], ICMP router discovery [6,9] or using the DHCP Router option, defined in RFC 2132 [4].

In a network that already provides DHCP service, using DHCP to update the routing table on a DHCP client has several virtues. It is efficient, since it makes use of messages that would have been sent anyway. It is convenient - the DHCP server configuration is already being maintained, so maintaining routing information, at least on a relatively stable network, requires little extra work. If DHCP service is already in use, no additional infrastructure need be deployed.

The DHCP protocol as defined in RFC 2131 [3] and the options defined in RFC 2132 [4] only provide a mechanism for installing a default route or installing a table of classful routes. Classful routes are routes whose subnet mask is implicit in the subnet number - see section 3.2 of STD 5, RFC 791 [1] for details on classful routing.

Classful routing is no longer in common use, so the DHCP Static Route option is no longer useful. Currently, classless routing [7, 10] is the most commonly-deployed form of routing on the Internet. In classless routing, IP addresses consist of a network number (the combination of the network number and subnet number described in RFC 950 [7]) and a host number.

In classful IP, the network number and host number are derived from the IP address using a bitmask whose value is determined by the first few bits of the IP address. In classless IP, the network number and host number are derived from the IP address using a separate quantity, the subnet mask. In order to determine the network to which a given route applies, an IP host must know both the network number AND the subnet mask for that network.

The Static Routes option (option 33) does not provide a subnet mask for each route - it is assumed that the subnet mask is implicit in whatever network number is specified in each route entry. The Classless Static Routes option does provide a subnet mask for each entry, so that the subnet mask can be other than what would be determined using the algorithm specified in STD 5, RFC 791 [1] and STD 5, RFC 950 [7].

Definitions

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 [2].

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.

"link"

Any set of network attachment points that will all receive a link-layer broadcast sent on any one of the attachment points. This term is used in DHCP because in some cases more than one IP subnet may be configured on a link. DHCP uses a local-network (all-ones) broadcast, which is not subnet-specific, and will therefore reach all nodes connected to the link, regardless of the IP subnet or subnets on which they are configured.

A "link" is sometimes referred to as a broadcast domain or physical network segment.

Classless Route Option Format

The code for this option is 121, and its minimum length is 5 bytes. This option can contain one or more static routes, each of which consists of a destination descriptor and the IP address of the router that should be used to reach that destination.

```

Code Len Destination 1      Router 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| 121 | n | d1 | ... | dN | r1 | r2 | r3 | r4 |
+-----+-----+-----+-----+-----+-----+-----+-----+

Destination 2      Router 2
+-----+-----+-----+-----+-----+-----+
| d1 | ... | dN | r1 | r2 | r3 | r4 |
+-----+-----+-----+-----+-----+-----+

```

In the above example, two static routes are specified.

Destination descriptors describe the IP subnet number and subnet mask of a particular destination using a compact encoding. This encoding consists of one octet describing the width of the subnet mask, followed by all the significant octets of the subnet number.

The width of the subnet mask describes the number of one bits in the mask, so for example a subnet with a subnet number of 10.0.127.0 and a netmask of 255.255.255.0 would have a subnet mask width of 24.

The significant portion of the subnet number is simply all of the octets of the subnet number where the corresponding octet in the subnet mask is non-zero. The number of significant octets is the width of the subnet mask divided by eight, rounding up, as shown in the following table:

Width of subnet mask	Number of significant octets
0	0
1- 8	1
9-16	2
17-24	3
25-32	4

The following table contains some examples of how various subnet number/mask combinations can be encoded:

Subnet number	Subnet mask	Destination descriptor
0	0	0
10.0.0.0	255.0.0.0	8.10
10.0.0.0	255.255.255.0	24.10.0.0
10.17.0.0	255.255.0.0	16.10.17
10.27.129.0	255.255.255.0	24.10.27.129
10.229.0.128	255.255.255.128	25.10.229.0.128
10.198.122.47	255.255.255.255	32.10.198.122.47

Local Subnet Routes

In some cases more than one IP subnet may be configured on a link. In such cases, a host whose IP address is in one IP subnet in the link could communicate directly with a host whose IP address is in a different IP subnet on the same link. In cases where a client is being assigned an IP address on an IP subnet on such a link, for each IP subnet in the link other than the IP subnet on which the client has been assigned the DHCP server MAY be configured to specify a router IP address of 0.0.0.0.

For example, consider the case where there are three IP subnets configured on a link: 10.0.0/24, 192.168.0/24, 10.0.21/24. If the client is assigned an IP address of 10.0.21.17, then the server could include a route with a destination of 10.0.0/24 and a router address of 0.0.0.0, and also a route with a destination of 192.168.0/24 and a router address of 0.0.0.0.

A DHCP client whose underlying TCP/IP stack does not provide this capability MUST ignore routes in the Classless Static Routes option whose router IP address is 0.0.0.0. Please note that the behavior described here only applies to the Classless Static Routes option, not to the Static Routes option nor the Router option.

DHCP Client Behavior

DHCP clients that do not support this option MUST ignore it if it is received from a DHCP server. DHCP clients that support this option MUST install the routes specified in the option, except as specified in the Local Subnet Routes section. DHCP clients that support this option MUST NOT install the routes specified in the Static Routes option (option code 33) if both a Static Routes option and the Classless Static Routes option are provided.

DHCP clients that support this option and that send a DHCP Parameter Request List option MUST request both this option and the Router option [4] in the DHCP Parameter Request List.

DHCP clients that support this option and send a parameter request list MAY also request the Static Routes option, for compatibility with older servers that don't support Classless Static Routes. The Classless Static Routes option code MUST appear in the parameter request list prior to both the Router option code and the Static Routes option code, if present.

If the DHCP server returns both a Classless Static Routes option and a Router option, the DHCP client MUST ignore the Router option.

Similarly, if the DHCP server returns both a Classless Static Routes option and a Static Routes option, the DHCP client MUST ignore the Static Routes option.

After deriving a subnet number and subnet mask from each destination descriptor, the DHCP client MUST zero any bits in the subnet number where the corresponding bit in the mask is zero. In other words, the subnet number installed in the routing table is the logical AND of the subnet number and subnet mask given in the Classless Static Routes option. For example, if the server sends a route with a destination of 129.210.177.132 (hexadecimal 81D4B184) and a subnet

mask of 255.255.255.128 (hexadecimal FFFFFFFF80), the client will install a route with a destination of 129.210.177.128 (hexadecimal 81D4B180).

Requirements to Avoid Sizing Constraints

Because a full routing table can be quite large, the standard 576 octet maximum size for a DHCP message may be too short to contain some legitimate Classless Static Route options. Because of this, clients implementing the Classless Static Route option SHOULD send a Maximum DHCP Message Size [4] option if the DHCP client's TCP/IP stack is capable of receiving larger IP datagrams. In this case, the client SHOULD set the value of this option to at least the MTU of the interface that the client is configuring. The client MAY set the value of this option higher, up to the size of the largest UDP packet it is prepared to accept. (Note that the value specified in the Maximum DHCP Message Size option is the total maximum packet size, including IP and UDP headers.)

DHCP clients requesting this option, and DHCP servers sending this option, MUST implement DHCP option concatenation [5]. In the terminology of RFC 3396 [5], the Classless Static Route Option is a concatenation-requiring option.

DHCP Server Administrator Responsibilities

Many clients may not implement the Classless Static Routes option. DHCP server administrators should therefore configure their DHCP servers to send both a Router option and a Classless Static Routes option, and should specify the default router(s) both in the Router option and in the Classless Static Routes option.

When a DHCP client requests the Classless Static Routes option and also requests either or both of the Router option and the Static Routes option, and the DHCP server is sending Classless Static Routes options to that client, the server SHOULD NOT include the Router or Static Routes options.

Security Considerations

Potential exposures to attack in the DHCP protocol are discussed in section 7 of the DHCP protocol specification [3] and in Authentication for DHCP Messages [11].

The Classless Static Routes option can be used to misdirect network traffic by providing incorrect IP addresses for routers. This can be either a Denial of Service attack, where the router IP address given is simply invalid, or can be used to set up a man-in-the-middle

attack by providing the IP address of a potential snooper. This is not a new problem - the existing Router and Static Routes options defined in RFC 2132 [4] exhibit the same vulnerability.

IANA Considerations

This DHCP option has been allocated the option code 121 in the list of DHCP option codes that the IANA maintains.

Normative References

- [1] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [4] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [5] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", RFC 3396, November 2002.

Informative References

- [6] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.
- [7] Mogul, J. and J. Postel, "Internet Standard Subnetting Procedure", STD 5, RFC 950, August 1985.
- [8] Hedrick, C., "Routing Information Protocol", RFC 1058, June 1988.
- [9] Deering, S., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [10] Pummill, T. and B. Manning, "Variable Length Subnet Table For IPv4", RFC 1878, December 1995.
- [11] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

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.

Authors' Addresses

Ted Lemon
Nominum, Inc.
2385 Bay Road
Redwood City, CA 94063

E-Mail: Ted.Lemon@nominum.com

Stuart Cheshire
Apple Computer, Inc.
1 Infinite Loop
Cupertino
California 95014
USA

Phone: +1 408 974 3207
E-Mail: rfc@stuartcheshire.org

Bernie Volz
Ericsson
959 Concord Street
Framingham, MA, 01701

Phone: +1 508 875 3162
E-Mail: bernie.volz@ericsson.com

Full Copyright Statement

Copyright (C) The Internet Society (2002). 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.