

DHCP Relay Agent Information Option

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 (2001). All Rights Reserved.

Abstract

Newer high-speed public Internet access technologies call for a high-speed modem to have a local area network (LAN) attachment to one or more customer premise hosts. It is advantageous to use the Dynamic Host Configuration Protocol (DHCP) as defined in RFC 2131 to assign customer premise host IP addresses in this environment. However, a number of security and scaling problems arise with such "public" DHCP use. This document describes a new DHCP option to address these issues. This option extends the set of DHCP options as defined in RFC 2132.

The new option is called the Relay Agent Information option and is inserted by the DHCP relay agent when forwarding client-originated DHCP packets to a DHCP server. Servers recognizing the Relay Agent Information option may use the information to implement IP address or other parameter assignment policies. The DHCP Server echoes the option back verbatim to the relay agent in server-to-client replies, and the relay agent strips the option before forwarding the reply to the client.

The "Relay Agent Information" option is organized as a single DHCP option that contains one or more "sub-options" that convey information known by the relay agent. The initial sub-options are defined for a relay agent that is co-located in a public circuit access unit. These include a "circuit ID" for the incoming circuit, and a "remote ID" which provides a trusted identifier for the remote high-speed modem.

Table of Contents

1	Introduction.....	2
1.1	High-Speed Circuit Switched Data Networks.....	2
1.2	DHCP Relay Agent in the Circuit Access Equipment.....	4
2.0	Relay Agent Information Option.....	5
2.1	Agent Operation.....	6
2.1.1	Reforwarded DHCP requests.....	7
2.2	Server Operation.....	7
3.0	Relay Agent Information Suboptions.....	8
3.1	Agent Circuit ID.....	8
3.2	Agent Remote ID.....	9
4.0	Issues Resolved.....	9
5.0	Security Considerations.....	10
6.0	IANA Considerations.....	11
7.0	Intellectual Property Notice.....	12
8.0	References.....	12
9.0	Glossary.....	13
10.0	Author's Address.....	13
11.0	Full Copyright Statement	14

1 Introduction

1.1 High-Speed Circuit Switched Data Networks

Public Access to the Internet is usually via a circuit switched data network. Today, this is primarily implemented with dial-up modems connecting to a Remote Access Server. But higher speed circuit access networks also include ISDN, ATM, Frame Relay, and Cable Data Networks. All of these networks can be characterized as a "star" topology where multiple users connect to a "circuit access unit" via switched or permanent circuits.

With dial-up modems, only a single host PC attempts to connect to the central point. The PPP protocol is widely used to assign IP addresses to be used by the single host PC.

The newer high-speed circuit technologies, however, frequently provide a LAN interface (especially Ethernet) to one or more host PCs. It is desirable to support centralized assignment of the IP addresses of host computers connecting on such circuits via DHCP. The DHCP server can be, but usually is not, co-implemented with the centralized circuit concentration access device. The DHCP server is often connected as a separate server on the "Central LAN" to which the central access device (or devices) attach.

A common physical model for high-speed Internet circuit access is shown in Figure 1, below.

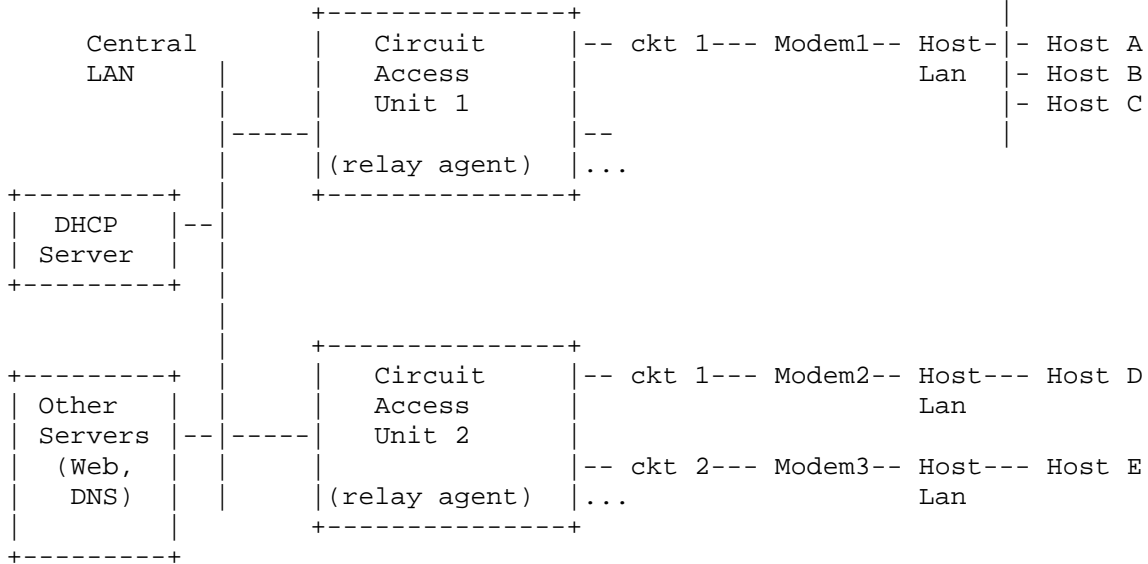


Figure 1: DHCP High Speed Circuit Access Model

Note that in this model, the "modem" connects to a LAN at the user site, rather than to a single host. Multiple hosts are implemented at this site. Although it is certainly possible to implement a full IP router at the user site, this requires a relatively expensive piece of equipment (compared to typical modem costs). Furthermore, a router requires an IP address not only for every host, but for the router itself. Finally, a user-side router requires a dedicated Logical IP Subnet (LIS) for each user. While this model is appropriate for relatively small corporate networking environments, it is not appropriate for large, public accessed networks. In this scenario, it is advantageous to implement an IP networking model that does not allocate an IP address for the modem (or other networking equipment device at the user site), and especially not an entire LIS for the user side LAN.

Note that using this method to obtain IP addresses means that IP addresses can only be obtained while communication to the central site is available. Some host lan installations may use a local DHCP server or other methods to obtain IP addresses for in-house use.

1.2 DHCP Relay Agent in the Circuit Access Unit

It is desirable to use DHCP to assign the IP addresses for public high-speed circuit access. A number of circuit access units (e.g., RAS's, cable modem termination systems, ADSL access units, etc) connect to a LAN (or local internet) to which is attached a DHCP server.

For scaling and security reasons, it is advantageous to implement a "router hop" at the circuit access unit, much like high-capacity RAS's do today. The circuit access equipment acts as both a router to the circuits and as the DHCP relay agent.

The advantages of co-locating the DHCP relay agent with the circuit access equipment are:

DHCP broadcast replies can be routed to only the proper circuit, avoiding, say, the replication of the DHCP reply broadcast onto thousands of access circuits;

The same mechanism used to identify the remote connection of the circuit (e.g., a user ID requested by a Remote Access Server acting as the circuit access equipment) may be used as a host identifier by DHCP, and used for parameter assignment. This includes centralized assignment of IP addresses to hosts. This provides a secure remote ID from a trusted source -- the relay agent.

A number of issues arise when forwarding DHCP requests from hosts connecting publicly accessed high-speed circuits with LAN connections at the host. Many of these are security issues arising from DHCP client requests from untrusted sources. How does the relay agent know to which circuit to forward replies? How does the system prevent DHCP IP exhaustion attacks? This is when an attacker requests all available IP addresses from a DHCP server by sending requests with fabricated client MAC addresses. How can an IP address or LIS be permanently assigned to a particular user or modem? How does one prevent "spoofing" of client identifier fields used to assign IP addresses? How does one prevent denial of service by "spoofing" other client's MAC addresses?

All of these issues may be addressed by having the circuit access equipment, which is a trusted component, add information to DHCP client requests that it forwards to the DHCP server.

2.0 Relay Agent Information Option

This document defines a new DHCP Option called the Relay Agent Information Option. It is a "container" option for specific agent-supplied sub-options. The format of the Relay Agent Information option is:

```

Code   Len   Agent Information Field
+-----+-----+-----+-----+-----+-----+-----+-----+
| 82   |  N   | i1  | i2  | i3  | i4  |   | iN  |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The length N gives the total number of octets in the Agent Information Field. The Agent Information field consists of a sequence of SubOpt/Length/Value tuples for each sub-option, encoded in the following manner:

```

SubOpt Len   Sub-option Value
+-----+-----+-----+-----+-----+-----+-----+-----+
|  1   |  N   | s1  | s2  | s3  | s4  |   | sN  |
+-----+-----+-----+-----+-----+-----+-----+-----+
SubOpt Len   Sub-option Value
+-----+-----+-----+-----+-----+-----+-----+-----+
|  2   |  N   | i1  | i2  | i3  | i4  |   | iN  |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

No "pad" sub-option is defined, and the Information field shall NOT be terminated with a 255 sub-option. The length N of the DHCP Agent Information Option shall include all bytes of the sub-option code/length/value tuples. Since at least one sub-option must be defined, the minimum Relay Agent Information length is two (2). The length N of the sub-options shall be the number of octets in only that sub-option's value field. A sub-option length may be zero. The sub-options need not appear in sub-option code order.

The initial assignment of DHCP Relay Agent Sub-options is as follows:

DHCP Agent Sub-option Code	Sub-Option Description
1	Agent Circuit ID Sub-option
2	Agent Remote ID Sub-option

2.1 Agent Operation

Overall adding of the DHCP relay agent option SHOULD be configurable, and SHOULD be disabled by default. Relay agents SHOULD have separate configurables for each sub-option to control whether it is added to client-to-server packets.

A DHCP relay agent adding a Relay Agent Information field SHALL add it as the last option (but before 'End Option' 255, if present) in the DHCP options field of any recognized BOOTP or DHCP packet forwarded from a client to a server.

Relay agents receiving a DHCP packet from an untrusted circuit with giaddr set to zero (indicating that they are the first-hop router) but with a Relay Agent Information option already present in the packet SHALL discard the packet and increment an error count. A trusted circuit may contain a trusted downstream (closer to client) network element (bridge) between the relay agent and the client that MAY add a relay agent option but not set the giaddr field. In this case, the relay agent does NOT add a "second" relay agent option, but forwards the DHCP packet per normal DHCP relay agent operations, setting the giaddr field as it deems appropriate.

The mechanisms for distinguishing between "trusted" and "untrusted" circuits are specific to the type of circuit termination equipment, and may involve local administration. For example, a Cable Modem Termination System may consider upstream packets from most cable modems as "untrusted", but an ATM switch terminating VCs switched through a DSLAM may consider such VCs as "trusted" and accept a relay agent option added by the DSLAM.

Relay agents MAY have a configurable for the maximum size of the DHCP packet to be created after appending the Agent Information option. Packets which, after appending the Relay Agent Information option, would exceed this configured maximum size shall be forwarded WITHOUT adding the Agent Information option. An error counter SHOULD be incremented in this case. In the absence of this configurable, the agent SHALL NOT increase a forwarded DHCP packet size to exceed the MTU of the interface on which it is forwarded.

The Relay Agent Information option echoed by a server MUST be removed by either the relay agent or the trusted downstream network element which added it when forwarding a server-to-client response back to the client.

The agent SHALL NOT add an "Option Overload" option to the packet or use the "file" or "sname" fields for adding Relay Agent Information option. It SHALL NOT parse or remove Relay Agent Information options that may appear in the sname or file fields of a server-to-client packet forwarded through the agent.

The operation of relay agents for specific sub-options is specified with that sub-option.

Relay agents are NOT required to monitor or modify client-originated DHCP packets addressed to a server unicast address. This includes the DHCP-REQUEST sent when entering the RENEWING state.

Relay agents MUST NOT modify DHCP packets that use the IPSEC Authentication Header or IPSEC Encapsulating Security Payload [6].

2.1.1 Reforwarded DHCP requests

A DHCP relay agent may receive a client DHCP packet forwarded from a BOOTP/DHCP relay agent closer to the client. Such a packet will have giaddr as non-zero, and may or may not already have a DHCP Relay Agent option in it.

Relay agents configured to add a Relay Agent option which receive a client DHCP packet with a nonzero giaddr SHALL discard the packet if the giaddr spoofs a giaddr address implemented by the local agent itself.

Otherwise, the relay agent SHALL forward any received DHCP packet with a valid non-zero giaddr WITHOUT adding any relay agent options. Per RFC 2131, it shall also NOT modify the giaddr value.

2.2 Server Operation

DHCP servers unaware of the Relay Agent Information option will ignore the option upon receive and will not echo it back on responses. This is the specified server behavior for unknown options.

DHCP servers claiming to support the Relay Agent Information option SHALL echo the entire contents of the Relay Agent Information option in all replies. Servers SHOULD copy the Relay Agent Information option as the last DHCP option in the response. Servers SHALL NOT place the echoed Relay Agent Information option in the overloaded sname or file fields. If a server is unable to copy a full Relay Agent Information field into a response, it SHALL send the response without the Relay Information Field, and SHOULD increment an error counter for the situation.

The operation of DHCP servers for specific sub-options is specified with that sub-option.

Note that DHCP relay agents are not required to monitor unicast DHCP messages sent directly between the client and server (i.e., those that aren't sent via a relay agent). However, some relay agents MAY chose to do such monitoring and add relay agent options. Consequently, servers SHOULD be prepared to handle relay agent options in unicast messages, but MUST NOT expect them to always be there.

3.0 Relay Agent Information Sub-options

3.1 Agent Circuit ID Sub-option

This sub-option MAY be added by DHCP relay agents which terminate switched or permanent circuits. It encodes an agent-local identifier of the circuit from which a DHCP client-to-server packet was received. It is intended for use by agents in relaying DHCP responses back to the proper circuit. Possible uses of this field include:

- Router interface number
- Switching Hub port number
- Remote Access Server port number
- Frame Relay DLCI
- ATM virtual circuit number
- Cable Data virtual circuit number

Servers MAY use the Circuit ID for IP and other parameter assignment policies. The Circuit ID SHOULD be considered an opaque value, with policies based on exact string match only; that is, the Circuit ID SHOULD NOT be internally parsed by the server.

The DHCP server SHOULD report the Agent Circuit ID value of current leases in statistical reports (including its MIB) and in logs. Since the Circuit ID is local only to a particular relay agent, a circuit ID should be qualified with the giaddr value that identifies the relay agent.

SubOpt	Len	Circuit ID							
1	n	c1	c2	c3	c4	c5	c6	...	

3.2 Agent Remote ID Sub-option

This sub-option MAY be added by DHCP relay agents which terminate switched or permanent circuits and have mechanisms to identify the remote host end of the circuit. The Remote ID field may be used to encode, for instance:

- a "caller ID" telephone number for dial-up connection
- a "user name" prompted for by a Remote Access Server
- a remote caller ATM address
- a "modem ID" of a cable data modem
- the remote IP address of a point-to-point link
- a remote X.25 address for X.25 connections

The remote ID MUST be globally unique.

DHCP servers MAY use this option to select parameters specific to particular users, hosts, or subscriber modems. The option SHOULD be considered an opaque value, with policies based on exact string match only; that is, the option SHOULD NOT be internally parsed by the server.

The relay agent MAY use this field in addition to or instead of the Agent Circuit ID field to select the circuit on which to forward the DHCP reply (e.g., Offer, Ack, or Nak). DHCP servers SHOULD report this value in any reports or MIBs associated with a particular client.

SubOpt	Len	Agent Remote ID							
2	n	r1	r2	r3	r4	r5	r6	...	

4.0 Issues Resolved

The DHCP relay agent option resolves several issues in an environment in which untrusted hosts access the internet via a circuit based public network. This resolution assumes that all DHCP protocol traffic by the public hosts traverse the DHCP relay agent and that the IP network between the DHCP relay agent and the DHCP server is uncompromised.

Broadcast Forwarding

The circuit access equipment forwards the normally broadcasted DHCP response only on the circuit indicated in the Agent Circuit ID.

DHCP Address Exhaustion

In general, the DHCP server may be extended to maintain a database with the "triplet" of

(client IP address, client MAC address, client remote ID)

The DHCP server SHOULD implement policies that restrict the number of IP addresses to be assigned to a single remote ID.

Static Assignment

The DHCP server may use the remote ID to select the IP address to be assigned. It may permit static assignment of IP addresses to particular remote IDs, and disallow an address request from an unauthorized remote ID.

IP Spoofing

The circuit access device may associate the IP address assigned by a DHCP server in a forwarded DHCP Ack packet with the circuit to which it was forwarded. The circuit access device MAY prevent forwarding of IP packets with source IP addresses -other than- those it has associated with the receiving circuit. This prevents simple IP spoofing attacks on the Central LAN, and IP spoofing of other hosts.

Client Identifier Spoofing

By using the agent-supplied Agent Remote ID option, the untrusted and as-yet unstandardized client identifier field need not be used by the DHCP server.

MAC Address Spoofing

By associating a MAC address with an Agent Remote ID, the DHCP server can prevent offering an IP address to an attacker spoofing the same MAC address on a different remote ID.

5.0 Security Considerations

DHCP as currently defined provides no authentication or security mechanisms. Potential exposures to attack are discussed in section 7 of the DHCP protocol specification in RFC 2131 [1].

This document introduces mechanisms to address several security attacks on the operation of IP address assignment, including IP spoofing, Client ID spoofing, MAC address spoofing, and DHCP server

address exhaustion. It relies on an implied trusted relationship between the DHCP Relay Agent and the DHCP server, with an assumed untrusted DHCP client. It introduces a new identifier, the "Remote ID", that is also assumed to be trusted. The Remote ID is provided by the access network or modem and not by client premise equipment. Cryptographic or other techniques to authenticate the remote ID are certainly possible and encouraged, but are beyond the scope of this document.

This option is targeted towards environments in which the network infrastructure -- the relay agent, the DHCP server, and the entire network in which those two devices reside -- is trusted and secure. As used in this document, the word "trusted" implies that unauthorized DHCP traffic cannot enter the trusted network except through secured and trusted relay agents and that all devices internal to the network are secure and trusted. Potential deployers of this option should give careful consideration to the potential security vulnerabilities that are present in this model before deploying this option in actual networks.

Note that any future mechanisms for authenticating DHCP client to server communications must take care to omit the DHCP Relay Agent option from server authentication calculations. This was the principal reason for organizing the DHCP Relay Agent Option as a single option with sub-options, and for requiring the relay agent to remove the option before forwarding to the client.

While it is beyond the scope of this document to specify the general forwarding algorithm of public data circuit access units, note that automatic reforwarding of IP or ARP broadcast packets back downstream exposes serious IP security risks. For example, if an upstream broadcast DHCP-DISCOVER or DHCP-REQUEST were re-broadcast back downstream, any public host may easily spoof the desired DHCP server.

6.0 IANA Considerations

IANA is required to maintain a new number space of "DHCP Relay Agent Sub-options", located in the BOOTP-DHCP Parameters Registry. The initial sub-options are described in section 2.0 of this document.

IANA assigns future DHCP Relay Agent Sub-options with a "IETF Consensus" policy as described in RFC 2434 [3]. Future proposed sub-options are to be referenced symbolically in the Internet-Drafts that describe them, and shall be assigned numeric codes by IANA when approved for publication as an RFC.

7.0 Intellectual Property Notices

This section contains two notices as required by [5] for standards track documents.

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 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.

8.0 References

- [1] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [2] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extension", RFC 2132, March 1997.
- [3] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [5] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [6] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

9.0 Glossary

DSLAM	Digital Subscriber Link Access Multiplexer
IANA	Internet Assigned Numbers Authority
LIS	Logical IP Subnet
MAC	Message Authentication Code
RAS	Remote Access Server

10.0 Author's Address

Michael Patrick
Motorola Broadband Communications Sector
20 Cabot Blvd., MS M4-30
Mansfield, MA 02048

Phone: (508) 261-5707
EMail: michael.patrick@motorola.com

11.0 Full Copyright Statement

Copyright (C) The Internet Society (2001). 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.

