DHC Working Group
Internet-Draft
Expires: December 16, 2008

B. Joshi
P. Kurapati
M. Kamath
Infosys Technologies Ltd.
S. De Cnodder
Alcatel-Lucent
June 14, 2008

Extensions to Layer 2 Relay Agent
draft-ietf-dhc-l2ra-extensions-00.txt

Status of this Memo

Copyright Notice

Abstract

As per industry trends, Access Networks have been migrating from traditional ATM based networks to Ethernet networks. In Ethernet based access networks, Access Concentrators are typically configured to act as a transparent bridge in Layer 2 mode. These Access Concentrators also act as Layer 2 relay agents. Layer 2 Relay Agent

functionality does not provide means to avoid flooding of DHCP messages and also needs to be extended to support DHCP LeaseQuery This draft discusses these issues and provides solutions for the same.

Table of Contents

1. Introduction

DHCP Relay Agents eliminate the necessity of having a DHCP server on each physical network. RFC 3046 [3] defines a new option 'Relay Agent Information' which is added to DHCP messages by Relay Agents. DHCP servers may use this option for IP address and other parameter assignment policies.

In case of Layer 2 Access Networks, Access Concentrators typically act as Layer 2 Relay Agents [7].

This document proposes enhancements in Layer 2 Relay Agent [7] which addresses issues like flooding between Layer 3 Relay Agent and Layer 2 Relay Agent and retrieving lease information from server using DHCP leasequery mechanism.

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 RFC 2119 [1].

This document uses the following terms:

o "Access Concentrator"

An Access Concentrator is a router or switch at the broadband access provider's edge of a public broadband access network. This document assumes that the Access Concentrator acts as a Transparent Bridge and includes the DHCP relay agent functionality. For example: In DSL environment, this is typically known as DSLAM.(Digital Subscriber Line Access Multiplexer)

o "DHCP client"

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

o "Layer 3 Relay Agent"

A Layer 3 Relay Agent is a third-party agent that transfers Bootstrap Protocol (BOOTP) and DHCP messages between clients and servers residing on different subnets, per RFC951 [8] and RFC1542[9].

o "DHCP server"

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

o "downstream"

Downstream is the direction from the edge network towards the DHCP Clients.

o "Transparent Bridge"

A device which does bridging based on MAC learning principles. Bridge learns the Source MAC of the incoming frames and updates a table with MAC/Interface information. While forwarding data packets, bridge looks at this table to find the outgoing interface.

o "upstream"

Upstream is the direction from the DHCP Clients towards the edge

network.

3. Enhancements in Layer 2 Relay Agent

This section looks at various enhancements possible in Layer 2 Relay Agents. Following issues are seen in a typical Layer 2 Relay Agent[7] deployments

- Broadcasting DHCP requests on all interfaces

A normal Layer 2 Relay Agent[7] would broadcast a DHCP request message to all its interfaces except on which the message was received. Because of this, a DHCP request message is received by those devices which would not be interested in it. Configuring an uplink port that leads to a Layer 3 Relay Agent or DHCP server can solve this issue. Some of the existing implementations [Mostly in xDSL Access Concentrators] already supports this.

o Recovering Lease Information from Server

A Layer 2 Relay Agent[7] may snoop DHCP messages and maintain the lease information. This information is lost if the Layer 2 Relay Agent reboots. RFC 4388 suggests Leasequery mechanism to get the lease information from the server. This document extends the same for Layer 2 Relay Agent.

o Layer 3 Relay Agent broadcasting DHCP replies

Layer 3 Relay Agents generally broadcast DHCP replies towards Layer 2 Relay Agents. This will be received by those devices which would not be interested in it. In general, broadcasts should be avoided in Layer 2 networks. A new sub-option in Relay Agent Information option can be used to solve this issue. To avoid broadcasts in case of replies to Leasequery, a new option is defined.

3.1. Reference Network

Following network configuration is used as a reference network to explain the various issues and solutions in Layer 2 Networks. This network configuration is a typical Ethernet Aggregated Access Network.

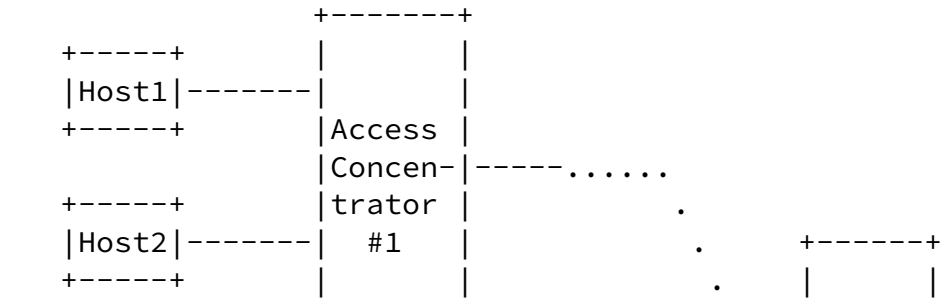

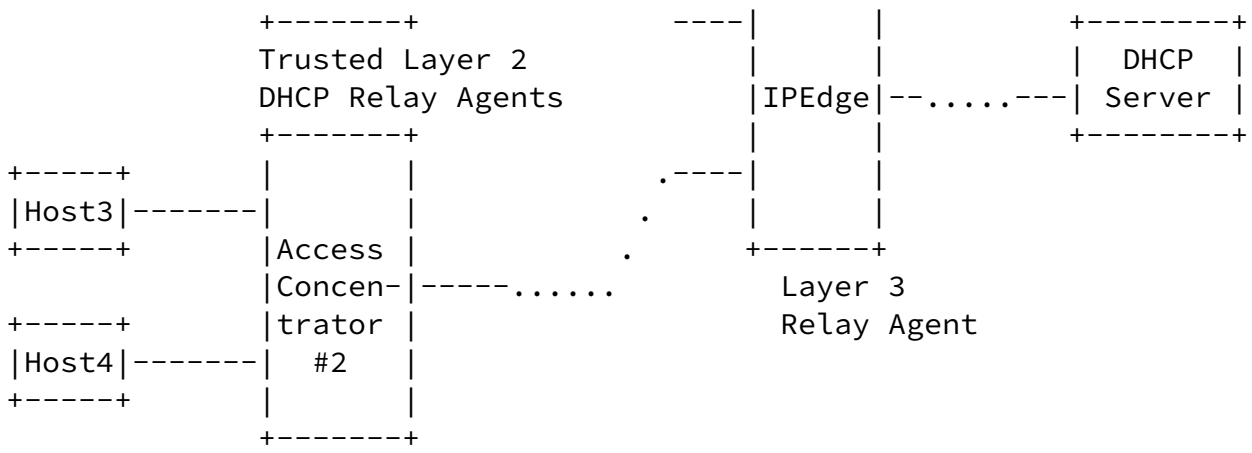

Figure 1

4. Uplink port

A Layer 2 Relay Agent broadcasts the DHCP request messages [Messages which are broadcast by Clients] to all the interfaces within the same broadcast domain except the interface on which it was received. This leads to flooding of DHCP messages which is unnecessary. Hence there is a need to identify an "Uplink Port", through which the DHCP request messages could be relayed towards the DHCP server. The uplink port SHOULD be a configurable parameter.

5. Extension of DHCPLEASEQUERY for Layer 2 Relay Agent

5.1. Protocol Extension Overview

A Layer 2 Relay Agent [7] may want to maintain the information of outgoing interface, MAC Address, IP address and Lease information for each DHCP Client. This information [MAC-IP-Interface Binding] could be used to prevent MAC/IP Spoofing attacks. It could also be used for bridging frames. Maintain this information makes a Layer 2 Relay Agent vulnerable to the same issue [location/lease information lost when Layer 2 Relay Agent gets rebooted] which has been addressed in RFC 4388 [5] for Layer 3 networks. This document extends mechanism proposed in [5] to address this issue for layer 2 networks.

When Layer 2 Relay Agent needs to bridge a frame, it MAY refer to location/lease information to verify the IP address or MAC address. If the location/lease information is not available, it can query DHCP server to obtain the lease/location information using DHCPLEASEQUERY message.

A Layer 2 Relay Agent can generate a DHCPLEASEQUERY [Query by IP address, MAC address, client identifier [10]] with all the fields properly populated as defined in RFC 4388 [5].

5.2. Protocol Extension Details

5.2.1. Generating DHCPLEASEQUERY Message

When a data packet is received from a host, Layer 2 Relay Agent [7] may verify if it has location/lease information for the source IP address or source MAC address of data packet received. Similarly when Layer 2 Relay Agent receives a data packet from the uplink port, it may verify location/lease information for the destination IP address or destination MAC address of the data packet. A Layer 2 Relay Agent would typically generate DHCPLEASEQUERY message if the location/lease information is not available for the corresponding IP address or MAC address assuming that it has lost the location/lease information during last reboot. The DHCPLEASEQUERY message uses the DHCP message format as described in RFC 2131 [2], and uses message number 10 in the DHCP Message Type option (option 53). The DHCPLEASEQUERY message has the following pertinent message contents:

- "giaddr" field MUST NOT be set. Though RFC 4388 [5] mandates that an Access Concentrator [in Layer 3 mode] 'MUST' set the "giaddr" field, this document suggest that a Layer 2 Relay Agent acting as Transparent Bridge must not set the "giaddr" field.

- The Parameter Request List option (option 55) MUST include the Relay Agent Information option (option 82).
- All the other options in Parameter Request List option (option 55) SHOULD be set as per the interest of the requester. The options of interest are likely to be the IP Address Lease Time option (option 51) and possibly the Vendor class identifier option (option 60).
- Source IP address of the DHCPLEASEQUERY message MUST be set to 0.0.0.0.
- Destination IP address of the DHCPLEASEQUERY message MUST be set to broadcast address 255.255.255.255.
- Destination MAC address of the DHCPLEASEQUERY message MUST be set to FF:FF:FF:FF:FF:FF.
- Source MAC address of the DHCPLEASEQUERY message MUST be set to the hardware address of the interface on which this request is sent out.

All other fields in MAC header, IP header and DHCP header SHOULD be set as per RFC 2131 [2]. Additional details concerning different query types are same as defined in RFC 4388 [5].

5.2.2. Handling DHCPLEASEQUERY Message in Layer 3 Relay Agent

A Layer 3 Relay Agent conforming to this document, MUST process the DHCP LEASEQUERY message received on its downstream interface similar to the other DHCP messages.

5.2.3. Handling DHCPLEASEQUERY Message in DHCP Server

While generating a DHCP reply for a DHCPLEASEQUERY message, if the message type is DHCPLEASEUNASSIGNED or DHCPLEASEUNKNOWN, it MUST echo back the Relay Agent Information received in the DHCPLEASEQUERY message. If the message type is DHCPLEASEACTIVE, DHCP server prepares the message as described in RFC 4388 and ignores the Relay

Agent Information option received in the DHCPLEASEQUERY message.

This document does not propose any other changes to RFC 4388 [5] for handling DHCPLEASEQUERY message in DHCP server.

5.2.4. Handling DHCP Reply Message in Layer 3 Relay Agent

When Layer 3 Relay Agent receives a DHCP Reply message with message type as DHCPLEASEUNASSIGNED, DHCPLEASEACTIVE or DHCPLEASEUNKNOWN, it

must have a way to identify if it had generated the leasequery message or it had relayed it for a Layer 2 Relay Agent.

When the DHCP Reply message is received, a Layer 3 Relay Agent may use 'giaddr' or 'state information' to identify the outgoing interface.

5.2.5. Handling DHCP Reply Message in Layer 2 Relay Agent

5.2.5.1. Handling DHCPLEASEUNASSIGNED Reply Message

When a DHCPLEASEUNASSIGNED message is received by a Layer 2 Relay Agent, it means that there is no active lease for the IP address present in the DHCP server, but that a server does in fact manage that IP address. Layer 2 Relay Agent SHOULD cache this information for later use.

5.2.5.2. Handling DHCPLEASEUNKNOWN Reply Message

When a DHCPLEASEUNKNOWN message is received by Layer 2 Relay Agent, it SHOULD cache this information but only for a short lifetime, approximately for 5 minutes as suggested in RFC 4388 [5].

5.2.5.3. Handling DHCPLEASEACTIVE Reply Message

When Layer 2 Relay Agent receives a DHCPLEASEACTIVE message, it MUST update its location/lease information.

5.2.5.4. Handling multiple responses for DHCPLEASEQUERY Message

A Layer 3 Relay Agent can forward a DHCPLEASEQUERY request to more than one DHCP server and so a Layer 2 Relay Agent may receive more

than one reply for a DHCPLEASEQUERY message.

A Layer 2 Relay Agent MUST be able to process multiple responses for a DHCPLEASEQUERY message. For example:

- It should be able to ignore all other responses once it receives DHCPLEASEACTIVE response from one of the DHCP server.

5.2.5.5. Handling No Response to the DHCPLEASEQUERY Message

This has been discussed in detail in RFC 4388 [5] and the same holds good for this document as well.

5.2.5.6. Handling DHCPLEASEQUERY messages not belonging to Layer 2 Relay Agent

- Since Layer 3 Relay Agent can broadcast the reply of DHCPLEASEQUERY message, it will be processed by all the Layer 2 Relay Agents connected to the same LAN. Using either Transaction Id or Relay Agent Information Option, a Layer 2 Relay Agent should be able to correctly identify if the DHCPLEASEQUERY response is meant for itself. Responses which do not belong to an Access Concentrator MUST be silently discarded.

- In a typical bridged network, multiple Layer 2 Relay Agents may share the same LAN. As a DHCPLEASEQUERY message generated by a Layer 2 Relay Agent is broadcast, it will be received by other Layer 2 Relay Agents also. Layer 2 Relay Agents MUST silently discard any DHCPLEASEQUERY message received from the uplink port.

5.3. DHCPLEASEQUERY using Management IP address of Layer 2 Relay Agent

Though rare, but if a Layer 2 Relay Agent allows the use of Management IP address for communication with DHCP server, it can generate DHCPLEASEQUERY message as described in RFC 4388 instead of using the extension of DHCPLEASEQUERY message described in this document.

6. Prevention of flooding of DHCP replies from Layer 3 Relay Agent

Figure 1 shows an example where each access concentrator adds the relay agent information option containing the port information of the host sending the DHCP messages. IP edge router relays these DHCP messages to the server.

RFC 2131[2] defines the meaning of the broadcast flag in the flags field: it indicates whether the client wishes to receive the DHCPOFFER and DHCPACK message as a broadcast or a unicast from the DHCP server or the DHCP relay agent. In the scenario of Figure 1, this means that the IP edge router will broadcast the DHCPOFFER and DHCPACK messages to all access concentrators if the broadcast flag is set. Whether or not broadcast is used between the Layer 3 Relay Agent and the trusted Layer 2 Relay Agents depends on the behavior of the DHCP clients. However broadcasts in the aggregation network are to be avoided. So it is preferred to always use unicast from the Layer 3 DHCP relay agent to the trusted layer 2 DHCP relay agent.

Between the trusted layer 2 DHCP relay agent and the host, broadcast flag has to be honored.

Even though the DHCP clients are not setting the broadcast flag, it is still possible that the DHCPOFFER and DHCPACK messages from the DHCP server are sent to all access concentrators. This is when the access concentrator implements a MAC concentration or MAC translation function. When such a MAC operation is performed, the access concentrator replaces the source MAC address of all upstream frames by another MAC address, for instance with its own MAC address. In this case, the MAC addresses of the hosts will remain unknown in the network between the trusted layer 2 DHCP relay agent and the Layer 3 DHCP relay agent. Hence all unicast messages sent by the Layer 3 DHCP relay agent using this MAC address will be flooded to all access concentrators.

6.1. Flooding of DHCP reply messages from Layer 3 Relay Agent

To overcome these two previously mentioned problems, a new sub-option 'unicast-address' is defined for the Relay Agent Information option. With this sub-option, the Layer 3 Relay Agent will always unicast the messages towards the trusted Layer 2 Relay Agent with a hardware address that is known in the network.

6.1.1. Unicast-Address Sub-Option

6.1.1.1. Unicast-Address Sub-Option Definition

The unicast-address sub-option of the relay-agent-information option MAY be used by any trusted layer 2 DHCP relay agent such that the

Layer 3 relay agent unicasts the messages from the DHCP server with a hardware address known in the network. The hardware address in the unicast-address sub-option MUST be an address that can be used to send unicast packets towards the client.

The format of the option is as follows:

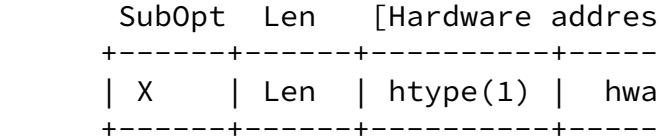

Figure 2

- 'X' is the sub-option code which needs to be allocated by IANA.
- 'Len' represents the length of the 'value' which includes both htype and hwaddr fields
- "htype" represents Hardware type. See the 'ARP parameters' maintained in the database referenced by Assigned numbers RFC 3232 [6].
- "hwaddr" is the unicast hardware address.

6.1.1.2. Layer 3 Relay Agent Behavior

When Layer 3 DHCP Relay Agent receives a DHCP packet with unicast-address sub-option added, it SHOULD unicast that message towards the layer 2 DHCP relay agent with destination address set to the value contained in the hwaddr field of the sub-option. A Layer 3 relay agent that supports this option SHOULD ignore the broadcast flag if this sub-option is present in the DHCP message. In the absence of this sub-option a Layer 3 relay agent SHOULD behave as earlier and forward the message as per the broadcast bit set in the message.

6.1.1.3. Layer 2 Relay Agent Behavior

The Layer 2 Relay Agent may add this sub-option only in the case when the intermediate network elements do MAC learning ensuring that when the Layer 3 relay agent unicasts the messages to this hardware address, the messages will arrive at the same layer 2 DHCP relay agent. The Layer 2 DHCP relay agent SHOULD still be able to receive broadcast messages from the Layer 3 DHCP relay agent in order to remain compatible with relay agents that do not support the unicast-address sub-option.

Layer 2 DHCP relay agent MUST always process the broadcast flag as

described in [RFC2131]. This means that it is possible that the layer 2 DHCP relay agents receive a unicast message from the Layer 3 DHCP relay agent, and that it has to forward it as a broadcast. It is also possible that the unicast message stays unicast and that only

the destination MAC address has to be changed to the content of the chaddr field.

If the layer 2 DHCP relay agent performs a MAC address concentration function, it SHOULD add the unicast-address sub-option to all upstream DHCP messages in order to avoid flooding of unknown destination MAC addresses. On the other hand, if the layer 2 DHCP relay agent acts as a bridge, it MAY add the unicast-address sub-option only to the DHCPDISCOVER and DHCPREQUEST messages as these are the only messages which may result in a downstream broadcast.

6.1.1.4. DHCP Server Behavior

Although rather unlikely, it is also possible that no Layer 3 DHCP relay agent is configured in the network and that the DHCP server has layer 2 connectivity with the trusted layer 2 DHCP relay agent. In this case the DHCP server, supporting the unicast address option, SHOULD act as a Layer 3 DHCP relay agent would do.

So if the DHCP server receives DHCP messages with giaddr set to zero and a valid unicast-address sub-option, the DHCP server SHOULD ignore the broadcast flag and unicast the DHCP messages to the hardware address in the unicast-address sub-option. The DHCP Server SHOULD also include this sub-option in the option 82 of its reply.

6.1.1.5. Example Scenarios

- The trusted layer 2 DHCP relay agent acts as a bridge. In such a case, the layer 2 DHCP relay agent puts the MAC address in the chaddr field of DHCP messages in the unicast-address sub-option. The Layer 3 DHCP relay agent will then send the DHCPOFFER and DHCPACK messages from the DHCP server as unicast to the layer 2 DHCP relay agent, which converts the message to broadcast if the broadcast flag is set.

- The Layer 2 Relay Agent does MAC translation/concentration function. In this case layer 2 DHCP relay agent adds unicast-address sub-option which contains the MAC address that the Layer 2 DHCP Relay Agent is using for upstream frames.

6.2. Flooding of DHCPLEASEQUERY reply messages from Layer 3 Relay Agent

The above suboption would not work for reply message for a LEASEQUERY request because the reply message type other than LEASEACTIVE for a

LEASEQUERY message will not have Relay Agent Information option. This can be resolved by creating a new option which is echoed back by the DHCP server in DHCP reply messages for a LEASEQUERY message.

This document need the definition of following new option for DHCP packet beyond those defined by [RFC2131] and [RFC2132]. See also Section 9, IANA Considerations.

6.2.1. Relay Agent Hardware Address option

"relay-agent-hwaddr" option allows a Layer 3 Relay agent to unicast a DHCP reply for a DHCPLEASEQUERY message to the Layer 2 Relay Agent which had generated the DHCPLEASEQUERY message. The code for this option need to be allocated by IANA.

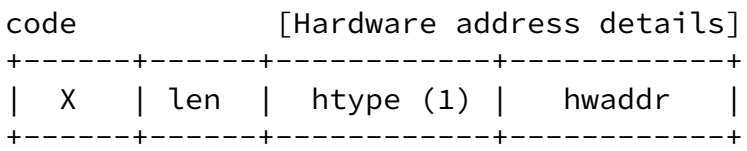

Figure 3

In the above option:

- 'X' need to be allocated by IANA.
- "len" field contains the length of the "Hardware address details" and can be used to deduce length of "hwaddr" field.
- "htype" represents Hardware type. See the 'ARP parameters' maintained in the database referenced by Assigned numbers RFC 3232[4].
- "hwaddr" is Relay Agent hardware address.

6.2.1.1. Layer 2 Relay Agent Behavior

Layer 2 Relay agents which has the capability to receive a unicast reply for DHCPLEASEQUERY message SHOULD add option "relay-agent-hwaddr" in DHCPLEASEQUERY message. Option "relay-agent-hwaddr" SHOULD be populated based on the interface on which this request is sent out.

6.2.1.2. Layer 3 Relay Agent Behavior

While forwarding a reply for Lease Query request, a Layer 3 Relay

Agent MUST look for "relay-agent-hwaddr" option [code 'X'] in the

DHCP reply and if it finds this option, it SHOULD extract the hardware address and use it to unicast the reply to the Layer 2 Relay Agent.

DHCP reply message with message type 'DHCPLEASEACTIVE' can have Relay Agent Information option which may have 'unicast-address' sub-option. In such a case, both 'relay-agent-hwaddr' option and 'unicast-address' sub-option MAY be present. A Layer 3 Relay Agent conforming to this document MUST always prefer hardware address extracted from 'unicast-address' sub-option of Relay Agent Information option over 'relay-agent-hwaddr' option.

6.2.1.3. DHCP server Behavior

DHCP servers conforming to this document MUST echo the entire contents of the "relay-agent-hwaddr" option [code 'X'] in the reply for a DHCPLEASEQUERY request. DHCP servers SHALL NOT place the echoed "relay-agent-hwaddr" option in the overloaded sname or file fields. If a server is unable to copy a full "relay-agent-hwaddr" option into a response, it SHALL send the response without the "relay-agent-hwaddr" option, and SHOULD increment an error counter for the situation.

DHCP Server MUST NOT add or echo back this option in any other DHCP reply messages it generates.

7. Acknowledgments

Stig Venaas, Wojciech Dec, Richard Pruss and Andre Kostur provided good feedback on this memo. A detailed discussion with Ted Lemon, Andre Kostur on how a Layer 3 Relay Agent can unicast the various DHCP replies to a Layer 2 Relay Agent was very helpful.

The authors would like to acknowledge Ludwig Pauwels and Paul Reynders for their feedback on 'unicast-address' sub-option. Thanks to Patrick Mensch who contributed for the initial version of the document which had defined 'unicast-address' sub-option.

Description of authentication for DHCPLEASEQUERY messages in security section are taken from RFC 4388.

8. Security Consideration

- Layer 3 Relay Agent that relays the DHCP message are essentially DHCP clients for the purposes of the DHCP messages relayed by Layer 2 Relay Agent. Layer 3 Relay Agent MUST relay a DHCP message only when it comes from a trusted circuit. Thus, RFC3118[4] is an appropriate mechanism for DHCP messages relayed by Layer 2 Relay Agent.

- This document suggest new option which MAY be added by Layer 2 Relay Agents in DHCP message. If a server finds this new option included in a received message, the server MUST compute any hash function as if the option were NOT included in the message without changing the order of options. Whenever the server sends back this option to a relay agent, the server MUST not include this option in the computation of any hash function over the message.

9. IANA Considerations

This document needs IANA to provide a unique number for the new option to carry Hardware address of a Relay Agent. Please refer to section 6.1 for more details.

This document also needs IANA to provide a unique number for the following new suboptions in Relay Agent Information option [Option 82]:

- To carry the hardware address of a Relay Agent. Please refer to section 6.2 for more details.

10. References

10.1. Normative Reference

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[3] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

[4] Droms, R. and B. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

[5] Woundy, R. and K. Kinnear, "Dynamic Host Configuration Protocol (DHCP) Leasequery", RFC 4388, February 2006.

[6] Reynolds, J., "Assigned Numbers", RFC 3232, January 2002.

[7] Joshi, B. and P. Kurapati, "Layer 2 Relay Agent Information", draft draft-ietf-dhc-l2ra-01.txt, May 2008.

10.2. Informative Reference

[8] Croft, B. and J. Gilmore, "Bootstrap Protocol (BOOTP)", RFC 951, September 1985.

[9] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1542, October 1993.

[10] Droms, R. and S. Alexander, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

Authors' Addresses

Bharat Joshi
Infosys Technologies Ltd.
44 Electronics City, Hosur Road

Bangalore 560 100
India

Email: bharat_joshi@infosys.com
URI: http://www.infosys.com/

Pavan Kurapati
Infosys Technologies Ltd.
44 Electronics City, Hosur Road
Bangalore 560 100
India

Email: pavan_kurapati@infosys.com
URI: http://www.infosys.com/

Mukund Kamath
Infosys Technologies Ltd.
44 Electronics City, Hosur Road
Bangalore 560 100
India

Email: mukund_kamath@infosys.com
URI: http://www.infosys.com/

Stefaan De Cnodder
Alcatel-Lucent
Francis Wellesplein 1,
B-2018 Antwerp
Belgium

Email: stefaan.de_cnodder@alcatel-lucent.be
URI: http://www.alcatel-lucent.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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