DHC Working Group
Internet-Draft
Expires: November 17, 2008

B. Joshi
P. Kurapati
Infosys Technologies Ltd.
May 16, 2008

# Layer 2 Relay Agent Information
# draft-ietf-dhc-l2ra-01.txt

## Status of this Memo



## Copyright Notice



## Abstract

In some networks, DHCP servers rely on Relay Agent Information option appended by Relay Agents for IP address and other parameter assignment policies. This works fine when end hosts are directly connected to Relay Agents. In some network configurations, one or more Layer 2 devices may reside between DHCP clients and Relay agent. In these network scenarios, it is difficult to use the Relay Agent Information option for IP address and other parameter assignment policies effectively. So there is a need for the device that is

closest to the end hosts to append Relay Agent Information option in DHCP messages. These devices are typically known as Layer 2 Relay Agents.

This document aims to describe the network scenarios where Layer 2 Relay Agent is in use and also how it handles DHCP messages.

Table of Contents

## 1. Introduction

DHCP Relay Agents eliminate the necessity of having a DHCP server on each physical network. Relay Agents populate the 'giaddr' field and also append the 'Relay Agent Information' option to the DHCP messages. DHCP servers use this option for IP address and other parameter assignment policies. These DHCP Relay Agents are typically an IP routing aware device and are referred as Layer 3 Relay Agents.

In some network configurations, there is a need for Layer 2 devices to append the Relay Agent Information option as they are closer to the end hosts. These Layer 2 devices are typically operating only as bridges for the network and may not have an IPv4 address on the network in question. Lacking a valid IPv4 source address, they cannot relay packets directly to a DHCP server located on another network. These Layer 2 devices append the Relay Agent Information option and broadcast the DHCP message. A Layer 3 Relay Agent relays it to the DHCP server.

This document provides information about where a Layer 2 Relay Agent fits in and how it is used. This document also looks at various network scenarios with Layer 2 Relay Agent and discusses various issues caused by introduction of Layer 2 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 RFC 2119 [1].

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 "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 [6] and RFC1542[7].

o "DHCP server"

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

o "Unnumbered Interfaces"

An interface with no IP address associated with it. IP packets received on this interface will be processed like any other numbered IP interface. It may use a local IP address while generating IP packets.

## 3. Need of Layer 2 Relay Agent

A Layer 2 device intercepts DHCP messages for following reasons:

1. In some network deployments like xDSL, the subscriber aggregation devices (also known as Access Concentrator or a DSLAM in case of DSL) are configured to act as bridges. As these devices are closest to the subscriber, they are in the best position to provide a unique Relay Agent Information option to enforce policies in DHCP server.

2. In some network deployments, a Layer 2 device can append Relay Agent Information in DHCP messages so that it can use this information to forward the DHCP Replies to the specific port on which request was received.

3. In some networks, the Layer 2 Switch which is closest to the end users, snoops the DHCP messages. These switches extract DHCP Lease Information and use this information to install packet filters. This helps in preventing the Layer 2 and Layer 3 spoofing attempts by the subscribers. A point to note here is that in cases where switches maintain the Lease Information, they have to intercept unicast DHCP messages as well to keep this information up to date.

4. NOTE: Please send an email to the authors if you are aware of any other functionality of Layer 2 Relay Agent. It will be helpful in updating this list. This note will be removed before moving it for IESG review.

## 4. Layer 2 Relay Agent in various network scenarios

This section describes the various network scenarios where a Layer 2 Relay Agent fits in. It also describes how it handles different DHCP messages.

### 4.1. DHCP server and client on same subnet

In certain network configurations, DHCP server may reside on the same subnet as the DHCP clients. A Layer 2 aggregation device resides between the DHCP clients and DHCP server. Following points describe how this Layer 2 device handles various DHCP messages if it acts as a Layer 2 Relay Agent. Figure #1 shows a typical network setup.

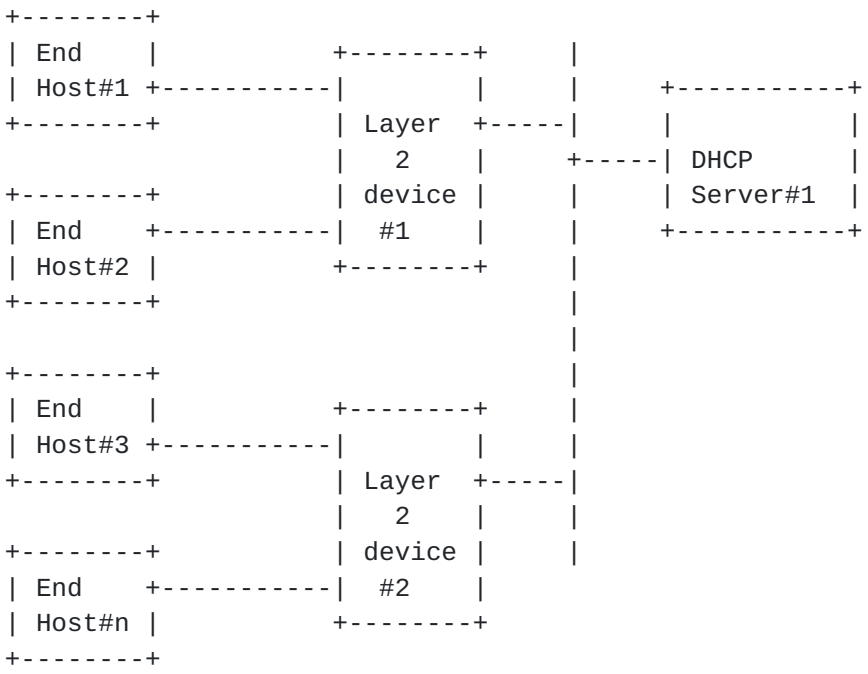


Figure 1

#### 4.1.1. Client-server interaction

The following summary of protocol message exchanges between clients and DHCP servers describes how they are handled in Layer 2 Relay Agent.

1. The client [End Host #1] broadcasts a DHCPDISCOVER message on its local physical subnet. Layer 2 Relay Agent [#1] intercepts this message, appends the Relay Agent Information option and broadcasts it to all the ports except on which it was received.

Relay Agent Information option could be created as suggested in RFC 3046[3]. Layer 2 Relay Agent does not set the 'giaddr' field.

2. Layer 2 device [#2] would also receive this DHCPDISCOVER message from Layer 2 device [#1]. If it is configured as Layer 2 Relay Agent, it intercepts this message but does not append another Relay Agent Information option to the message. It may discard this message if it is coming from an untrusted entity. Otherwise, it will broadcast this on all the ports except on which the message was received.

3. Server responds with a DHCPOFFER message after applying its local policies. DHCP server echoes back the Relay Agent Information option in the DHCPOFFER message. DHCP server can either unicast the reply to MAC address of End Host #1 or broadcast the reply. In both the cases, the Layer 2 Relay Agent intercepts this message and removes the Relay Agent Information option. It identifies the outgoing port using Relay Agent Information option and forwards to the identified interface.

4. Same DHCPOFFER message will be received by the other Layer 2 Device [#2]. If it is configured as Layer 2 Relay Agent, it broadcasts this message normally without removing the Relay Agent option since it had not added the same. A Layer 2 Relay Agent uses the Relay Agent Information option to find out if it had appended it to the request message.

5. The client receives this DHCPOFFER message and it broadcasts a DHCPREQUEST message that contains the 'server identifier' option to indicate the server it has selected. Layer 2 Relay Agent [#1] handles this message similar to how it handles DHCPDISCOVER message.

6. The server receives the DHCPREQUEST message from the client and responds with a DHCPACK message containing the configuration parameters for the requesting client. A DHCP server may unicast the DHCPACK message if the broadcast bit in the DHCPREQUEST message is not set. DHCP server would echo back the Relay Agent Information option in the reply message. A Layer 2 Relay Agent may intercept this unicast message and process it similar to the DHCPOFFER message.

7. A server that is unable to satisfy the DHCPREQUEST message, responds with DHCPNACK. Layer 2 Relay Agent process this similar to DHCPACK message.

8. The client receives the DHCPACK message with configuration parameters. If client detects that the address is already in use, it sends a DHCPDECLINE message to the server. Layer 2 Relay Agent process this message similar to DHCPDISCOVER message.

9. When client knows the address of a DHCP server, it may unicast DHCPDISCOVER, DHCPREQUEST messages to the server. DHCP clients unicast the DHCP messages like DHCPRELEASE and DHCPREQUEST when renewing the lease to the DHCP server. Layer 2 Relay Agent may or may not intercept these messages based on internal configuration. If Layer 2 Relay Agents intercept these messages, they append Relay Agent Information option and forward it towards the DHCP server. They also intercept the reply messages and remove Relay Agent Information option before forwarding them.

### 4.1.2. Issues due to introduction of Layer 2 Relay Agent

1. A DHCP server should be able to handle a DHCP message that contains the Relay Agent Information option without 'giaddr' field set in the message. Some existing DHCP server implementations do not echo back the Relay Agent Information option if giaddr is not set. This may lead to issues at Layer 2 Relay Agents as they will not be able to identify the outgoing port correctly and would broadcast it to all ports. Some Layer 2 Relay Agents discard the reply messages if they do not find a Relay Agent Information option in a DHCP reply.

2. A DHCP server should be able to handle a unicast DHCP message containing Relay Agent Information option. Some existing DHCP server implementations do not echo back the Relay Agent Information option in DHCP reply messages.

## 4.2. Multiple DHCP server and Client on same subnet

In certain network scenarios, there could be multiple DHCP server on the same subnet. Figure #2 shows a typical network setup.

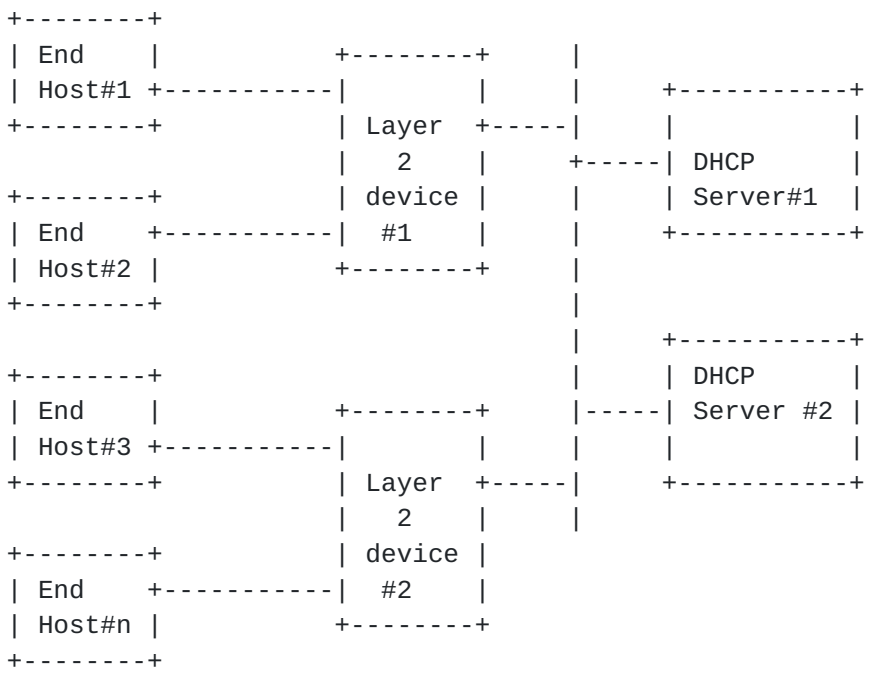


Figure 2

### 4.2.1. Client-server interaction

The message exchange are same as explained in 4.1.1. However, due to introduction of multiple DHCP servers the below additional message exchanges may happen

1. When Host [#1] sends DHCPDISCOVER, it will be received by both the DHCP Servers connected to Layer 2 Relay Agent #1 and both the servers will respond with a DHCPOFFER. So instead of one DHCPOFFER message, Layer 2 Relay Agent would receive two messages. Processing of DHCP messages in the Layer 2 Relay Agent remains the same.

### 4.2.2. Issues due to introduction of Layer 2 Relay Agent

1. Layer 2 relay agents which maintain persistent state, such as updating filters or client registration, must be prepared to handle potentially conflicting responses from different DHCP Servers. Some Layer 2 relay agents may use "the most recent DHCP packet" to update this persistent state but this may not necessarily reflect the actual state of the client. The above is possible when two DHCP servers ack the request of a DHCP client with same address but has different lease times. In this case, if the relay agent selects the server reply which has a shorter lease time, it would expire its state possibly before the client

even has a chance to renew it. Therefore, Layer 2 relay agents SHOULD select the longest lease time of two conflicting but similar replies, by discarding replies that shorten the lease time.

2. Other issues are same as described in section 4.1.2.

## 4.3. DHCP server on another subnet with one Layer 3 Relay Agent

In certain network scenarios, there could be a Layer 3 Relay Agent which relays the DHCP messages from one subnet to DHCP server on another subnet and vice versa. In typical deployments, Access Concentrator acts as Layer 2 Relay Agent and IP edge device (BRAS or IP Services Switch) acts as Layer 3 Relay Agent.

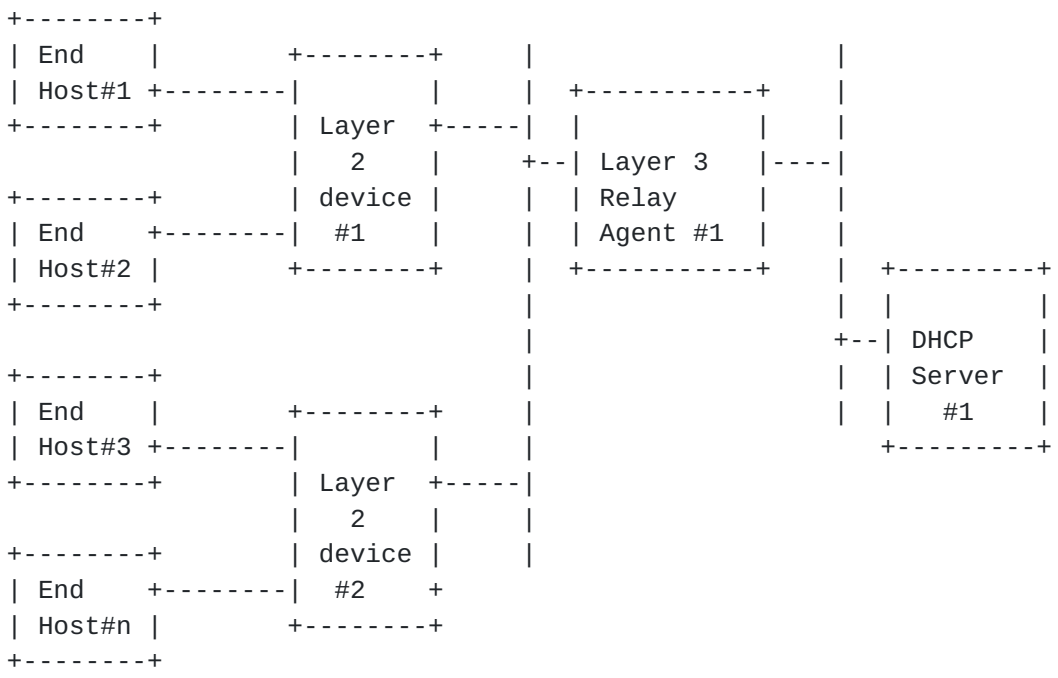


Figure 3

### 4.3.1. Client-server interaction

As far as DHCP message processing is concerned, the presence of Layer 3 Relay Agent is transparent to Layer 2 Relay Agent. So all the messages are handled in the same way as defined in section 4.1.1 by Layer 2 Relay Agent.

Layer 3 Relay Agent are configured to trust/untrust an entity based

on a specific criteria (For example : VLAN/interface on which the message was received). If the DHCP messages coming from the client has a relay agent option present, Layer 3 Relay Agent checks if it is coming on a trusted interface. If it is coming from a trusted interface, it will set the 'giaddr' with one of the local interface address and unicasts it to the configured servers. If the message is coming from an untrusted interface Layer 3 Relay Agent discards the message.

Typical message processing in this scenario is given below.

1. When client sends a DHCPDISCOVER message, Layer 2 Relay Agent forwards it as described in section 4.1.1. Layer 3 Relay Agent receives this message and finds that it contains Relay Agent Information option. It verifies whether the message is from a trusted entity or not. If it is from trusted entity it populates the 'giaddr' as it deems appropriate and relay it to the DHCP server.

2. DHCP Server process the message in the same way as described in section 4.1 and will unicast the DHCPOFFER to Layer 3 Relay Agent on the address specified in 'giaddr' field.

3. Layer 3 Relay Agent process the DHCPOFFER and identifies the outgoing interface. It resets the giaddr and broadcasts the message on the identified outgoing interface

4. Clients receive DHCPOFFER and generate DHCPREQUEST message. Layer 2 Relay Agent process it as described in section 4.1.1. Layer 3 Relay Agent receives the DHCPREQUEST message and process it similar to the DHCPDISCOVER message described in step #1.

5. DHCP Server process the DHCPREQUEST and unicasts the DHCP ACK message to layer 3 Relay Agent if 'broadcast' flag is set or directly to the client if 'broadcast' flag is not set. If Layer 3 Relay Agent receives this message, it will process it similar to DHCPOFFER as described in step #3.

6. In case of unicast messages [For example: DHCPREQUEST in case of DHCPRENEW], a Layer 3 Relay Agent may or may not intercept the message. If it intercepts a unicast DHCP request message, it populates the 'giaddr' and relay it to the DHCP server. When DHCP server sends a reply for this request message, it resets the 'giaddr' field, identifies outgoing interface and forwards the reply on the identified interface.

### 4.3.2. Issues due to introduction of Layer 2 Relay Agent

Though the processing of DHCP messages remain the same in Layer 2 Relay Agent, we see some more issues when a Layer 3 Relay Agent is present to relay the DHCP messages to the DHCP server.

1. When a Layer 2 Relay Agent is configured to intercept unicast messages as well, it appends Relay Agent Information option before forwarding them. A Layer 3 Relay Agent may not intercept these unicast messages. Due to this, a DHCP server may not echo back the Relay Agent Information option because the giaddr is not populated.

2. Existing Layer 3 Relay Agents populate the 'giaddr' with the IP address of the interface on which the request was received. This helps Layer 3 Relay Agent to identify the outgoing interface for the DHCP replies. In some cases, a Layer 3 Relay Agent may use unnumbered interfaces. In this case, it has to use a system wide IP address to populate the 'giaddr' field. Due to this, it becomes difficult to identify the correct outgoing interface for the messages received from the DHCP server. In these cases, some existing Layer 3 Relay Agent implementations maintain an internal state for each DHCP messages and use this state to identify the outgoing interface.

3. DHCP server uses certain parameters to differentiate the RENEW and REBIND state of a client. A DHCP client unicasts a RENEW request to the DHCP server, so DHCP server sees a DHCPREQUEST without 'giaddr' and Relay Agent Information option as RENEW request. While a REBIND request is broadcast and so DHCP server expect it to contain 'giaddr' and Relay Agent Information option. If Layer 2 Relay Agent is configured to intercept unicast messages, it will append Relay Agent Information option to the unicast DHCP messages. Because of this, it could be difficult for DHCP server to differentiate between a RENEWING and REBINDING state.

## 5. Acknowledgments

This document is the result of a discussion on DHC WG mailing list. Thanks to David W. Hankins and Michael Wacker for providing inputs on some of the existing implementations. Thanks to Stefaan.De.Cnodder and Mukund Kamath for reviewing the draft and providing valuable suggestions

## 6. Security Consideration

- A Layer 2 Relay Agent should always be configured to identify a trustable entity so that it appends Relay Agent Information option to DHCP messages coming from a trustable entity and forward it. If a DHCP message is received from a non-trustable entity, it should discard it and may report to the administrator.

- Introduction of Layer 2 Relay Agent does not introduce any new security issue. Security issues pertaining to Relay Agents in general applies to Layer 2 Relay Agents as well.

## 7. IANA Considerations

This document does not introduce any new namespaces for the IANA to manage.

## 8. References

### 8.1. 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] Reynolds, J., "Assigned Numbers", RFC 3232, January 2002.

### 8.2. Informative Reference

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

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

[8] 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/

## Intellectual Property Statement



## Disclaimer of Validity



## Copyright Statement



## Acknowledgment

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