

DHC Working Group
Internet-Draft
Expires: October 23, 2009

B. Joshi
P. Kurapati
Infosys Technologies Ltd.
April 21, 2009

Layer 2 Relay Agent Information
draft-ietf-dhc-l2ra-04.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on October 23, 2009.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

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 a 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 a Layer 2 Relay Agent is in use and also how it handles DHCP messages.

Table of Contents

| | | |
|------------------------|--|--------------------|
| 1. | Introduction | 4 |
| 2. | Terminology | 5 |
| 3. | Need of Layer 2 Relay Agent | 6 |
| 4. | Layer 2 Relay Agent in various network scenarios | 7 |
| 4.1. | DHCP server and client on same subnet | 7 |
| 4.1.1. | Client-server interaction | 7 |
| 4.1.2. | Issues due to introduction of Layer 2 Relay Agent | 9 |
| 4.2. | Multiple DHCP server and Client on same subnet | 9 |
| 4.2.1. | Client-server interaction | 10 |
| 4.2.2. | Issues due to introduction of Layer 2 Relay Agent | 10 |
| 4.3. | DHCP server on another subnet with one Layer 3 Relay Agent | 11 |
| 4.3.1. | Client-server interaction | 11 |
| 4.3.2. | Issues due to introduction of Layer 2 Relay Agent | 13 |
| 5. | Acknowledgments | 14 |
| 6. | Security Considerations | 15 |
| 7. | IANA Considerations | 16 |
| 8. | References | 17 |
| 8.1. | Normative Reference | 17 |
| 8.2. | Informative Reference | 17 |
| | Authors' Addresses | 18 |

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 Agents and discusses various issues caused by the introduction of Layer 2 Relay Agents.

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](#) [[RFC2119](#)].

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](#)] and [[RFC1542](#)].

- o "BRAS"

BRAS or Broadband Remote Access Server is a network element which acts as an aggregation device terminating end user sessions. BRAS is usually the first IP edge device in a Layer 2 Access Network architecture.

- 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 generated from this interface may use a local loopback address which may be shared with other unnumbered interfaces.

3. Need of Layer 2 Relay Agent

A Layer 2 device intercepts DHCP messages for the 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 the 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 the 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 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 this draft 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, a 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. The 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 a 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 the one on which it was received. The Relay Agent Information option could be created as suggested in [RFC 3046](#) [[RFC3046](#)]. The 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 the one on which the message was received.
3. The DHCP server responds with a DHCPOFFER message after applying its local policies. It echoes back the Relay Agent Information option in the DHCPOFFER message. The DHCP server can either unicast the reply to the MAC address of End Host #1 or broadcast the reply. If the reply is broadcast, the Layer 2 Relay Agent intercepts this message and removes the Relay Agent Information option. It identifies the outgoing port using the Relay Agent Information option and forwards the message to the identified interface. A Layer 2 Relay Agent may be configured to intercept unicast DHCP messages. In such a case, the Layer 2 Relay Agent intercepts unicast DHCP messages and handles them similar to broadcast messages.
4. The same DHCPOFFER message will be received by 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. Layer 2 Relay Agent #1 handles this message similar to how it handles a DHCPDISCOVER message.
6. The server receives the DHCPREQUEST message from the client and responds with a DHCPACK/DHCPNACK message. A DHCP server may unicast the DHCPACK message. The Layer 2 Relay Agent processes the DHCPACK message similar to a DHCPOFFER message.
7. The Layer 2 Relay Agent processes a DHCPNAK messages similar to a DHCPACK message.
8. The Layer 2 Relay Agent processes a DHCPDECLINE message similar to a DHCPDISCOVER message.

9. The DHCP client can unicast some of the DHCP messages. The 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 a Relay Agent Information option and forward the message towards the DHCP server. They also intercept the reply messages and remove the 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. There is a case when the DHCP client receives a unicast reply message like DHCPACK with a Relay Agent Information option. This may happen when the DHCP server unicasts the DHCPACK message and the Layer 2 Relay Agent is configured not to intercept unicast messages. In such a case, the DHCP client can ignore the Relay Agent Information option.
3. A DHCP server should be able to handle a unicast DHCP message containing a Relay Agent Information option. Some existing DHCP server implementations do not echo back the Relay Agent Information option in responses to unicast messages.

4.2. Multiple DHCP server and Client on same subnet

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




Figure 2

4.2.1. Client-server interaction

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

1. When Host #1 sends DHCPDISCOVER, it will be received by both DHCP Servers connected to Layer 2 Relay Agent #1 and both servers will respond with a DHCPOFFER. So instead of one DHCPOFFER message, the Layer 2 Relay Agent would receive two messages. The processing of DHCP messages in the Layer 2 Relay Agents 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 acknowledge the request of a DHCP client with the same address but different lease times. In this case, if the relay agent selects the server reply with the shorter lease time, it would expire its state possibly before the

The Layer 3 Relay Agents are configured to trust/untrust an entity

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

Typical message processing in this scenario is given below.

1. When the client sends a DHCPDISCOVER message, the Layer 2 Relay Agent forwards it as described in [section 4.1.1](#). The Layer 3 Relay Agent receives this message and finds that it contains a Relay Agent Information option. It verifies whether the message is from a trusted entity or not. If it is from a trusted entity, the Layer 2 Relay Agent populates the 'giaddr' field as it deems appropriate and relays the message to the DHCP server.
2. The DHCP Server processes the message in the same way as described in [section 4.1](#) and unicasts the DHCPOFFER to the Layer 3 Relay Agent on the address specified in the 'giaddr' field.
3. The Layer 3 Relay Agent processes the DHCPOFFER and identifies the outgoing interface. It resets the 'giaddr' field and broadcasts the message on the identified outgoing interface.
4. The client receives the DHCPOFFER and generates a DHCPREQUEST message. The Layer 2 Relay Agent processes it as described in [section 4.1.1](#). The Layer 3 Relay Agent receives the DHCPREQUEST message and processes it similar to the DHCPDISCOVER message described in step #1.
5. The DHCP Server processes the DHCPREQUEST and unicasts the DHCP ACK message to the layer 3 Relay Agent if the 'broadcast' flag is set, or directly to the client if the 'broadcast' flag is not set. If the Layer 3 Relay Agent receives this message, it processes it similar to the DHCPOFFER as described in step #3.
6. In the 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' field and relays the message to the DHCP server. When the DHCP server sends a reply for this request message, it resets the 'giaddr' field, identifies the 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 remains the same in Layer 2 Relay Agents, 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 a Relay Agent Information option before forwarding the request message. 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' field is not populated.
2. Existing Layer 3 Relay Agents populate the 'giaddr' field with the IP address of the interface on which the request was received. This helps the 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 message and use this state to identify the outgoing interface.
3. A 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 the DHCP server sees a DHCPREQUEST without 'giaddr' and Relay Agent Information option as a RENEW request. On the other hand, a REBIND request is broadcast and so the DHCP server expects it to contain 'giaddr' field and a Relay Agent Information option. If the Layer 2 Relay Agent is configured to intercept unicast messages, it will append a Relay Agent Information option to the unicast DHCP message. Because of this, it could be difficult for the 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 Ted Lemon, Mukund Kamath, Alfred Hoenes and Stefaan De Cnodder for reviewing the draft and providing valuable suggestions.

6. Security Considerations

- o A Layer 2 Relay Agent should always be configured to identify a trustable entity so that it appends a Relay Agent Information option to a DHCP message coming from a trustable entity and forwards it. If a DHCP message is received from a non-trustable entity, the Layer 2 Relay Agent should discard it and may report to the administrator.
- o The introduction of Layer 2 Relay Agents does not introduce any new security issues. Security issues pertaining to Relay Agents in general apply to Layer 2 Relay Agents as well.

7. IANA Considerations

This document does not introduce any new namespaces for the IANA to manage and does not request any new code point assignments.

8. References

8.1. Normative Reference

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", [RFC 3046](#), January 2001.
- [RFC3118] Droms, R. and B. Arbaugh, "Authentication for DHCP Messages", [RFC 3118](#), June 2001.
- [RFC3232] Reynolds, J., "Assigned Numbers", [RFC 3232](#), January 2002.

8.2. Informative Reference

- [RFC951] Croft, B. and J. Gilmore, "Bootstrap Protocol (BOOTP)", [RFC 951](#), September 1985.
- [RFC1542] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", [RFC 1542](#), October 1993.
- [RFC2132] 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/>

