DHC Working Group
Internet-Draft
Intended status: Informational
Expires: May 22, 2009

D. Miles, Ed.
S. Ooghe
Alcatel-Lucent
W. Dec
Cisco Systems
S. Krishnan
A. Kavanagh
Ericsson

November 18, 2008

Lightweight DHCPv6 Relay Agent (LDRA)
draft-miles-dhc-dhcpv6-ldra-02

Status of this Memo

Abstract

This document proposes a Lightweight DHCPv6 Relay Agent (LDRA) that is used to insert relay agent options in DHCPv6 message exchanges identifying client-facing interfaces. The LDRA can be implemented in existing access nodes (such as DSLAMs and Ethernet switches) that do not support IPv6 control or routing functions.

Table of Contents

1. Introduction

DHCPv6 Relay-Agents [RFC3315] are deployed to forward DHCPv6 messages between clients and servers when they are not on the same IPv6 link and are often implemented alongside a routing function in a common node. A Lightweight DHCPv6 Relay Agent (LDRA) allows Relay Agent Information to be inserted by an access node that performs a link-layer bridging (i.e. non-routing) function. A LDRA resides on the same IPv6 link as the client and a DHCPv6 Relay Agent or Server and is functionally the equivalent of the Layer 2 DHCP Relay draft[L2RA] proposed for DHCPv4 operation.

Unlike a DHCPv6 Relay-Agent specified in [RFC3315], a LDRA does not implement any IPv6 control functions (e.g. ICMPv6) or have any routing capability in the node.

1.1. Requirements Language

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].

2. Background

A variety of different link-layer network topologies exist for the aggregation of IPv6 nodes into one or more routers. In Layer 2 aggregation networks (IEEE 802.1D bridging or similar) that have many nodes on a single link, a DHCPv6 server or DHCP relay agent would normally be unaware of how a DHCP client is attached to the network. The LDRA allows Relay Agent Information, including the Interface-ID option [RFC3315], to be inserted by the access node so that it may be used by the DHCPv6 server for client identification. A typical application in a broadband service provider may be as an equivalent to the Broadband Forum TR-101 Layer 2 DHCP Relay Agent[TR-101] described in [L2RA]

3. Terminology

Address	An IP layer identifier for an interface or set of interfaces
Host	A non-routing IPv6 node that is participating in DHCPv6 message exchange

IP	Internet Protocol Version 6 (IPv6)
LDRA	Lightweight DHCPv6 Relay Agent
Link	A communication facility or medium over which nodes can communicate at the link layer
Link-local address	An IP address having only local-scope, indicated by having the address prefix FE80::/10, that can be used to reach neighbouring nodes attached to the same link. Every interface has a link-local address.
Node	A device that implements IPv6
Router	A node that forwards packets not directly addressed to itself
Access Node	A device that combines many interfaces onto one link. An access node is not IP-aware in the data path
Relay Agent	A node that acts as an intermediary to deliver DHCP messages between clients and servers and being on the same link as the client
Lightweight Relay Agent	A function on the access node that intercepts DHCP messages between clients and servers. The function exists as a bump in the wire on the IP link.
Unspecified address	An IP address that is comprised entirely of zeros

4. Message Format

The Lightweight DHCPv6 Relay Agent (LDRA) exchanges DHCP messages between clients and servers using the message formats established in [RFC3315].

To maintain interoperability with existing DHCP relays and servers the message format is unchanged from [RFC3315]. The LDRA implements the same message types as a normal DHCPv6 Relay Agent. They are:

- Relay-Forward Messages
- Relay-Reply Messages

4.1. Relay-Forward Message

The Relay-Forward message is created by any DHCPv6 Relay Agent, including an LDRA, to forward messages between clients and servers or other relay agents. These messages are built as specified in [RFC3315]

The Relay-Forward message contains relay agent parameters that identify the client-facing interface on which any reply messages should be forwarded. These parameters are link-address, peer-address and Interface-ID. The link-address parameter MUST be set to the unspecified address. The Interface-ID Relay Agent Option MUST be included in the Relay-Forward message. The LDRA MAY insert additional relay agent options.

4.2. Relay-Reply Message

The Relay-Reply message is constructed by a DHCPv6 server to send parameters to a DHCP client when a relay agent is present between the server and the client. The Relay-Reply message may be sent after an initial Relay-Forward message as the parameters link-address, peer-address, Interface-ID and the relay agent's IP address are learnt from the Relay-Forward message.

The server MUST include the Interface-ID option in the Relay-Reply Message to indicate to the LDRA the interface on which the de-capsulated message should be forwarded.

4.3. Mandatory DHCP Options

Parameters are exchanged between DHCP client, relay-agent and server through the use of DHCP options. There are a set of mandatory DHCP options that MUST be included by the LDRA in all Relay-Forward and Relay-Reply messages. These are the:

- Relay-Message Option
- Interface-ID Option

4.3.1. Relay-Message Option

A DHCPv6 Relay Agent relays messages between clients and servers or other relay agents through Relay-Forward and Relay-Reply message types. The original client DHCP message (i.e. the packet payload,

excluding UDP and IP headers) is encapsulated in a Relay Message option [RFC3315].

As an LDRA does not implement ICMPv6, fragmentation of Relay-Messages is not supported. If a Relay-Message would exceed the MTU of the outgoing interface it MUST be discarded and an error condition SHOULD be logged.

4.3.2. Interface-ID Option

The LDRA MUST include the Interface-ID option [RFC3315] in all Relay-Forward messages. When a LDRA receives a Relay-reply message with an Interface-ID option present and link-address unspecified, the LDRA MUST relay the decapsulated message to the client on the interface identified in the Interface-ID option.

Servers MAY use the Interface-ID for parameter assignment policies. The format of the Interface-ID is outside the scope of this contribution. The Interface-ID SHOULD be considered an opaque value, i.e. the server SHOULD NOT try to parse the contents of the Interface-ID option. The LDRA SHOULD use the same Interface-ID value for a given interface, and this value SHOULD be retained across restarts. This is because, if the Interface-ID changes, a server will not be able to use it reliably in parameter assignment policies.

5. Agent Behaviour

The LDRA MUST have each of its interfaces configured as either client-facing or network (DHCPv6 server) facing. The LDRA uses the notion of client-facing and network-facing interfaces to process the DHCPv6 messages.

5.1. Relaying a Client Message

When a DHCPv6 message (defined in [RFC3315]) is received on any client-facing interface, the LDRA MUST intercept and process the message. The LDRA MUST also prevent the original message from being forwarded on the network facing interface.

The lightweight relay agent adds any other options it is configured or required to include in the Relay-Forward message. The LDRA MUST set the link-address field of the Relay-forward message to the Unspecified Address (::) and MUST include the Interface-ID option in all DHCP Relay-Forward messages.

If the message received on the client-facing interface is a Relay-Forward message, the LDRA MUST set the Hop-Count field in the newly

created Relay-Forward message to the value of the hop-count field in the received message incremented by 1 as specified in [RFC3315].

The LDRA MUST copy the IP destination and link-layer destination addresses from the client-originated message into the IP destination address and link-layer destination address of the Relay-forward message.

The LDRA MUST copy the IP source address from the client-originated message into the peer-address field of the Relay-forward message. The LDRA MUST copy the link-layer source address from the client-originated message into the link-layer source address of the Relay-forward message.

5.1.1. Client Message Validation

On receipt of a DHCP message on the client facing interface, the LDRA MUST discard a message if it is of one of following message types:

- ADVERTISE (2)
- REPLY (7)
- RECONFIGURE (10)
- RELAY-REPLY (13)

Options contained in the DHCPv6 message MUST NOT be validated by the LDRA, making it the responsibility of the DHCP server to check message option validity and allow new options to be introduced without changes on the LDRA.

5.1.2. Trusted and Untrusted Interfaces

In [RFC3046] DHCPv4 relay-agents had their client-facing interfaces set to trusted and untrusted. An LDRA MUST implement a trusted/untrusted definition for all client-facing interfaces that SHOULD be configurable per interface. When a client-facing interface is deemed untrusted the LDRA MUST discard any message received from the client-facing interface of type RELAY-FORWARD (12).

In DHCPv4 it was not possible for a DHCP server to determine whether the client or an intermediate relay agent had added relay agent options and thus trusted interfaces (relay-agent interfaces that would allow any DHCP options to be present on incoming messages) and untrusted interfaces (relay-agent interfaces that would ensure there are no relay-agent options on incoming messages) were defined. In DHCPv6, relay agents encapsulate the received message into the Relay-

Message Option in addition to adding any relay-agent options. This nested message behaviour allows a server to identify the options each relay-agent has inserted along the path, whenever the data path between LDRA and server falls within a protected or operator controlled environment.

When an LDRA is deployed, DHCPv6 servers MAY be configured with the Relay-Forward hop-count of the LDRA to instruct at which level of nesting the relay-agent options should be parsed. This removes the need for an interface to be configured as trusted or untrusted by providing the DHCPv6 server with an awareness of the LDRA logical location in the DHCP relay path. This behaviour is dependent on the interception of all DHCP messages by the LDRA and the incrementing of the Relay-Forward hop-count if a Relay-Forward message is received from the client-facing LDRA interface.

5.2. Relaying a Relay-Reply message from the network

When a valid Relay-Reply is received on any network-facing access node interface, it MUST be intercepted by the LDRA. The LDRA MUST listen to all IP traffic that has a link-local scoped source address, link-local scoped destination address, protocol type UDP and destination port 547. The LDRA SHOULD ignore any message that does not meet this criteria and MUST allow it to be forwarded like any other packet. The LDRA MAY be configured to listen only to a specific destination address if it has been configured as a node (implementing a full IP stack).

The LDRA MUST intercept and process all DHCP Relay-Reply messages and MUST silently discard all other DHCP message types.

In addition to the validity checks performed by a relay agent in [RFC3315] , the Relay-Reply message is considered valid by the LDRA if:

- The Interface-ID Option is present and the value corresponds to a valid interface in the access node,

- the Relay-Reply peer-address and the destination IP address MUST be identical and MUST be a link-local scoped address when no IP address is configured on the LDRA, and

- the link-address is the Unspecified Address when no IP address is configured on the LDRA

The LDRA copies the peer-address into the destination IP address field, and MAY use the destination link-layer address (MAC address) or Interface-ID to determine which interface to send the message to

the client. The contents of the Relay Message Option is put into an IP/UDP packet and then forwarded to the client.

The LDRA MUST copy the link-layer and IP source address from the Relay-Reply message into the IP/UDP packet that is forwarded to the client.

6. Network Topology

The LDRA intercepts any DHCPv6 message received on client-facing interfaces with a destination IP address of All_DHCP_Relay_Agents_and_Servers (FF02::1:2). The LDRA MUST NOT forward the original client message to a network-facing interface, it MUST process the message and add the appropriate Relay-Forward options as described in previous sections.

6.1. Client and Server on Same Link

The access node acts as a bridge; it has no information about any IP prefixes that are valid on the link, thus a server should consider address and parameter assignment as if the client DHCP message was not relayed.

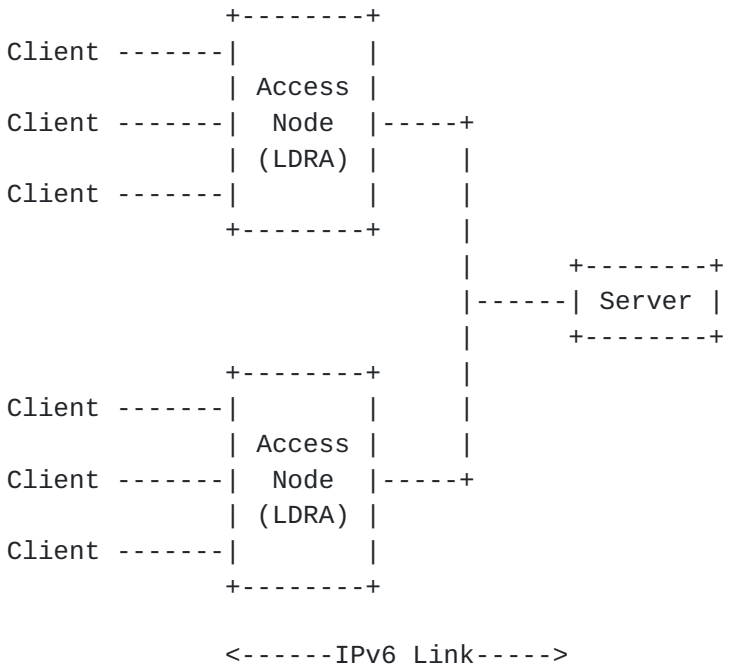

For example, if a client sent a DHCP solicit message that was relayed by the LDRA and then to the server, the server would receive the following Relay-Forward message from the LDRA:

```
src-ip:              client link-local address
dst-ip:              All_DHCP_Relay_Agents_and_Servers
  msg-type:          RELAY-FORWARD
  hop-count:         0
  link-address:      Unspecified_Address
  peer-address:      client link-local address
  Interface-ID Option:
    interface-id:    LDRA-inserted interface-id
  Relay-Message Option, which contains:
    msg-type:        SOLICIT
    Solicit Message Options: <from client>
```

6.2. Client and Server behind Relay Agent

The client and server are on different IPv6 links, separated by one or more relay agents that will typically act as a router. The LDRA will send Relay-Forward messages upstream towards the second relay agent which in turn will process the messages.

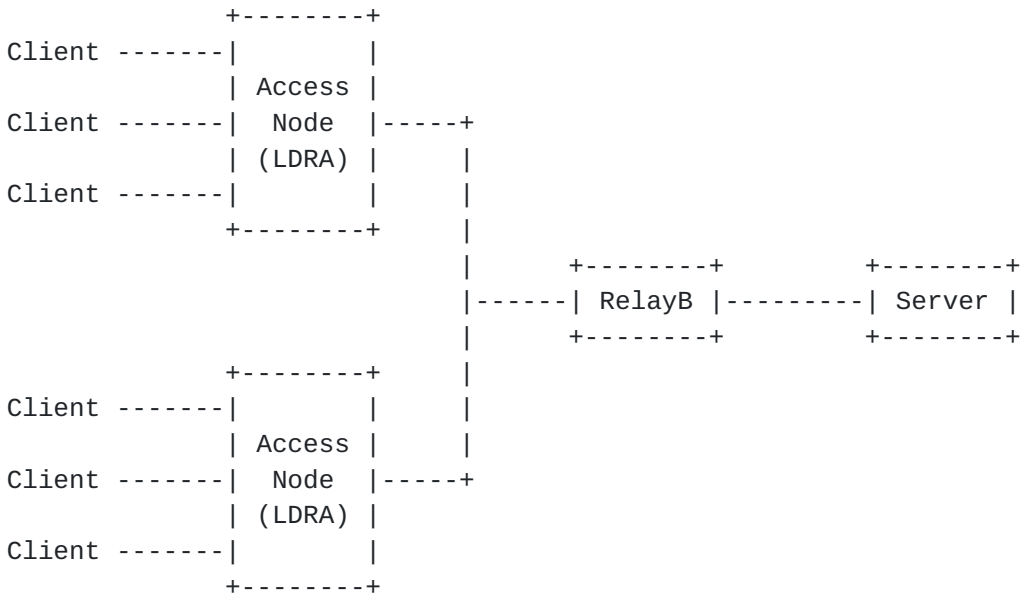

For example, if a client sent a DHCP solicit message that was relayed by the LDRA to another relay agent and then to the server, the server would receive the following Relay-Forward message from the LDRA:

```
src-ip:              relayB
dst-ip:              server
  msg-type:          RELAY-FORWARD
  hop-count:         1
  link-address:      relayB address from link A
  peer-address:      client link-local address
  Relay-Message Option, which contains:
    msg-type:        RELAY-FORWARD
    hop-count:       0
    link-address:    Unspecified_Address
    peer-address:    client link-local address
    Interface-ID Option:
      interface-id:  LDRA-inserted interface-id
    Relay-Message Option, which contains:
      msg-type:      SOLICIT
      Solicit Message Options: <from client>
```

6.3. Relay Agent in Front of LDRA

The client and server are on different IPv6 links, separated by one or more relay agents that will typically act as a router and there is an [RFC3315] Relay Agent on the client-facing Interface of the LDRA. The LDRA will send Relay-Forward messages upstream towards the second relay agent which in turn will process the messages.

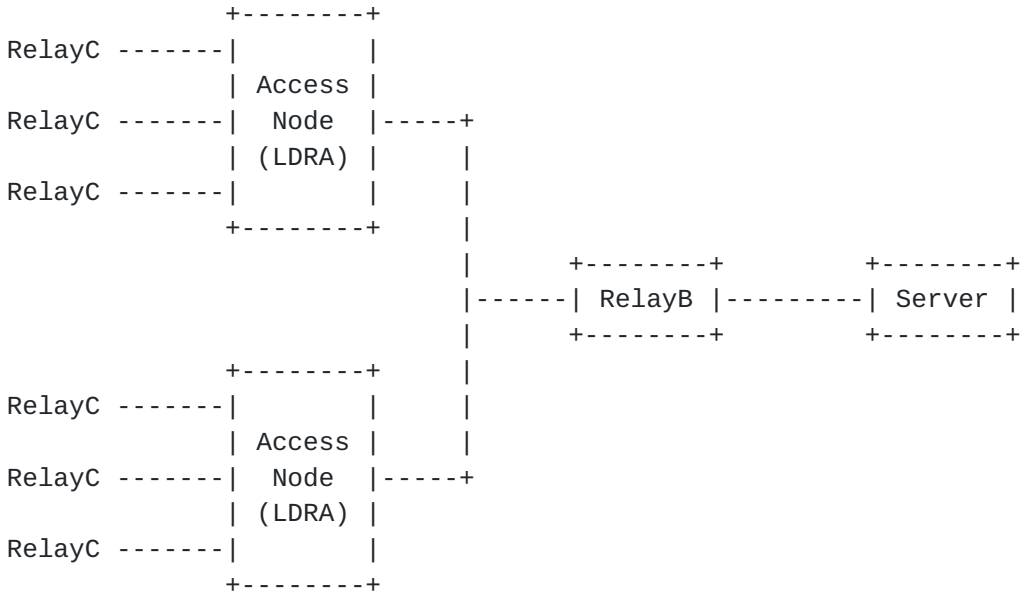

For example, if a client sent a DHCP solicit message that was relayed by the LDRA to another relay agent and then to the server, the server would receive the following Relay-Forward message from the LDRA:

```
src-ip:               relayB
dst-ip:               server
  msg-type:           RELAY-FORWARD
  hop-count:          2
  link-address:       relayB address from link A
  peer-address:       relayC
  Relay-Message Option, which contains:
    msg-type:         RELAY-FORWARD
    hop-count:        1
    link-address:     Unspecified_Address
    peer-address:     relayC
    Interface-ID Option:
      interface-id:   LDRA-inserted interface-id
    Relay-Message Option, which contains:
      msg-type:       RELAY-FORWARD
      hop-count:      0
      link-address:   global or unspecified address
      peer-address:   end client address
      Interface-ID Option: (if required)
        interface-id: relayC inserted Interface-ID
      Relay-Message Option, which contains:
        msg-type:      SOLICIT
        Solicit Message Options: <from end client>
```

7. Server Considerations

Although permitted in [RFC3315] the LDRA makes specific use of Relay-Forward link-address fields with a zero value.

- A DHCPv6 server MUST ignore any Relay-Forward link-addresses field with a zero value in Relay-Forward messages when searching for the inner-most link-address field. This allows the DHCPv6 server to select an address appropriate to the L3 link and supports a combination of L3 DHCPv6 relay agents and LDRA.

- If no non-zero Relay-Forward link-address is found, the DHCPv6 server should act as though the message were directly received. This is the case where no LDRA is present.

8. Contributors

The authors would like to thank the following for their support, Lieven Levrau, Alastair Johnson, Robert Haylock, Mickey Vucic, Ludwig Pauwels, Fernando Cuervo, John Kaippallimalil, Fredrik Garneij and Alfred Hoenes.

Comments are solicited and should be addressed to the DHC WG mailing list (dhcwg@ietf.org) and/or the author.

9. IANA Considerations

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

10. Security Considerations

Although the LDRA only listens to client-originated IPv6 traffic sent to the All_DHCPv6_Servers_and_Relay_Agents address on UDP port 547, the LDRA SHOULD implement some form of rate-limiting on received messages to prevent excessive process utilisation. As DHCP is session-oriented, messages in excess of the rate-limit may be silently discarded.

The hop count based determination of the trustworthiness of the LDRA can be easily defeated by a rogue relay agent on the network-facing interface of the LDRA.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[RFC3315] Droms, R., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

11.2. Informative References

[L2RA] Joshi, B., "Layer 2 Relay Agent Information", IETF Draft draft-ietf-dhc-l2ra-02, May 2008.

[RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

[TR-101] The Broadband Forum, "Migration to Ethernet-Based DSL Aggregation", Technical Report TR-101, April 2006.

Authors' Addresses

David Miles (editor)
Alcatel-Lucent
L3 / 215 Spring St
Melbourne, Victoria 3000
Australia

Phone: +61 3 9664 3308
Email: david.miles@alcatel-lucent.com

Sven Ooghe
Alcatel-Lucent
Copernicuslaan 50
2018 Antwerp,
Belgium

Phone:
Email: sven.ooghe@alcatel-lucent.com

Wojciech Dec
Cisco Systems
Haarlerberdweg 13-19
1101 CH Amsterdam,
The Netherlands

Phone:
Email: wdec@cisco.com

Suresh Krishnan
Ericsson
8400 Blvd Decarie
Town of Mount Royal, Quebec
Canada

Email: suresh.krishnan@ericsson.com

Alan Kavanagh
Ericsson
8400 Blvd Decarie
Town of Mount Royal, Quebec
Canada

Email: alan.kavanagh@ericsson.com

Full Copyright Statement

Intellectual Property