

Internet Engineering Task Force (IETF)
Request for Comments: 6148
Updates: 4388
Category: Standards Track
ISSN: 2070-1721

P. Kurapati
Juniper Networks
R. Desetti
B. Joshi
Infosys Technologies Ltd.
February 2011

DHCPv4 Lease Query by Relay Agent Remote ID

Abstract

Some relay agents extract lease information from the DHCP messages exchanged between the client and DHCP server. This lease information is used by relay agents for various purposes like antispoofing and prevention of flooding. RFC 4388 defines a mechanism for relay agents to retrieve the lease information from the DHCP server when this information is lost. The existing lease query mechanism is data-driven, which means that a relay agent can initiate the lease query only when it starts receiving data to and from the clients. In certain scenarios, this model is not scalable. This document first looks at issues in the existing mechanism and then proposes a new query type, query by Remote ID, to address these issues.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6148>.

Copyright Notice

Copyright (c) 2011 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Motivation	6
4. Protocol Details	7
4.1. Sending the DHCPLEASEQUERY Message	7
4.2. Responding to the DHCPLEASEQUERY Message	8
4.3. Building a DHCPLEASEACTIVE or DHCPLEASEUNKNOWN Message	8
4.4. Determining the IP Address to Be Used in Response	9
4.5. Sending a DHCPLEASEACTIVE or DHCPLEASEUNKNOWN Message	9
4.6. Receiving a DHCPLEASEACTIVE or DHCPLEASEUNKNOWN Message	9
4.7. Receiving No Response to the DHCPLEASEQUERY Message	10
4.8. Lease-Binding Data Storage Requirements	10
4.9. Using the DHCPLEASEQUERY Message with Multiple DHCP Servers	10
5. RFC 4388 Considerations	11
6. Security Considerations	11
7. Acknowledgments	11
8. References	12
8.1. Normative References	12
8.2. Informative References	12

1. Introduction

DHCP relay agents snoop DHCP messages and append a Relay Agent Information option before relaying them to the configured DHCP server. In this process, some relay agents also glean the lease information sent by the server and maintain this locally. This information is used to prevent spoofing attempts from clients and also sometimes to install routing information. When a relay agent reboots, this information is lost. RFC 4388 [RFC4388] has defined a mechanism to retrieve this lease information from the DHCP server. The existing query types defined by RFC 4388 [RFC4388] are data-driven. When a client sends data upstream, the relay agent can query the server about the related lease information, based on the source MAC/IP address. These mechanisms do not scale well when there are thousands of clients connected to the relay agent. In the data-driven model, lease query does not provide the full and consolidated active lease information associated with a given connection/circuit, which will result in inefficient anti-spoofing. The relay agent also has to contend with considerable resources for negative caching, especially under spoofing attacks.

We need a mechanism for a relay agent to retrieve the consolidated lease information for a given connection/circuit before upstream traffic is sent by the clients.




Figure 1

For example, when a DSLAM (Digital Subscriber Line Access Multiplexer) acting as a relay agent is rebooted, it should query the server for the lease information for all the connections/circuits. Also, as shown in the above figure, there could be multiple clients on one DSL circuit. The relay agent should get the lease information of all the clients connected to a DSL circuit. This is possible by introducing a new query type based on the Remote ID sub-option of the Relay Agent Information option. This document talks about the motivation for the new query type and the method to perform it.

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 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 includes the DHCP relay agent functionality.

o DHCP client

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

- o DHCP relay agent

A DHCP relay agent is a third-party agent that transfers Bootstrap Protocol (BOOTP) and DHCP messages between clients and servers residing on different subnets, per RFC 951 [RFC951] and RFC 1542 [RFC1542].

- o DHCP server

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

- o Fast path

Fast path refers to data transfer that happens through a network processor or an Application Specific Integrated Circuit (ASIC) programmed to forward the data at very high speeds.

- o Gleaning

Gleaning is the extraction of location information from DHCP messages as the messages are forwarded by the DHCP relay agent function.

- o Location information

Location information is information needed by the access concentrator to forward traffic to a broadband-accessible node. This information includes knowledge of the node's hardware address, the port or virtual circuit that leads to the node, and/or the hardware address of the intervening subscriber modem.

- o MAC address

In the context of a DHCP packet, a MAC address consists of the following fields: hardware type ("htype"), hardware length ("hlen"), and client hardware address ("chaddr").

- o Slow path

Slow path refers to data transfer that happens through the control plane. This has very limited buffers to store data, and the speeds are very low compared to the fast path data transfer.

- o Upstream

Upstream is the direction from the broadband subscriber towards the access concentrator.

3. Motivation

Consider an access concentrator (e.g., DSLAM) working also as a DHCP relay agent. A "fast path" and a "slow path" generally exist in most networking boxes. Fast path processing is done in a network processor or an ASIC. Slow path processing is done in a normal processor. As much as possible, regular data forwarding should be done in the fast path. Slow path processing should be reduced, as it may become a bottleneck.

For an access concentrator having multiple access ports, multiple IP addresses may be assigned to a single port using DHCP, and the number of clients on a port may be unknown. The access concentrator may also not know the network portions of the IP addresses that are assigned to its DHCP clients.

The access concentrator gleans IP address or other information from DHCP negotiations for antispoofing and other purposes. The antispoofing itself is done in the fast path. The access concentrator keeps track of only one list of IP addresses: the list of IP addresses that are assigned by the DHCP servers; upstream traffic from all other IP addresses is dropped. If a client starts its data transfer after its DHCP negotiations have been gleaned by the access concentrator, no legitimate packets will be dropped because of antispoofing. In other words, antispoofing is effective (no legitimate packets are dropped, and all spoofed packets are dropped) and efficient (antispoofing is done in the fast path). The intention is to achieve similar effective and efficient antispoofing in the lease query scenario also, when an access concentrator loses its gleaned information (for example, because of a reboot).

After a deep analysis, we found that the three existing query types supported by RFC 4388 [RFC4388] do not provide effective and efficient antispoofing for the above scenario, and a new mechanism is required.

The existing query types necessitate a data-driven approach: the lease queries can only be performed when the access concentrator receives data. This results in

- o increased outage time for clients
- o excessive negative caching, consuming a lot of resources under a spoofing attack
- o antispoofing being done in the slow path instead of the fast path

4. Protocol Details

This section talks about the protocol details for query by Remote ID. Most of the message handling is similar to RFC 4388 [RFC4388], and this section highlights only the differences. Readers are advised to go through RFC 4388 [RFC4388] before going through this section for complete understanding of the protocol.

When used in this document, the unqualified term "DHCPLEASEQUERY" indicates a lease query by Remote ID, unless otherwise specified.

RFC 3046 [RFC3046] defines two sub-options for the Relay Agent Information option. Sub-option 1 corresponds to the Circuit ID that identifies the local circuit of the access concentrator. This sub-option is unique to the relay agent. Sub-option 2 corresponds to the Remote ID that identifies the remote node connected to the access concentrator. The Remote ID is globally unique in the network and is configured per circuit/connection in the relay agent.

This document defines a new query type based on the Remote ID sub-option. Suppose that the access concentrator (e.g., DSLAM) lost the lease information when it was rebooted. When the access concentrator comes up, it initiates (for each connection/circuit) a DHCP lease query by Remote ID as defined in this section. For this query, the requester supplies an option 82 that includes only a Remote ID sub-option in the DHCPLEASEQUERY message. The Remote ID is normally pre-provisioned in the access concentrator per circuit/connection and hence will remain available to the access concentrator after reboot.

The DHCP server MUST reply with a DHCPLEASEACTIVE message if there is an active lease corresponding to the Remote ID that is present in the DHCPLEASEQUERY message. Otherwise, the server MUST reply with a DHCPLEASEUNKNOWN message. Servers that do not implement DHCP lease query based on Remote ID SHOULD simply not respond.

4.1. Sending the DHCPLEASEQUERY Message

The lease query defined in this document will mostly be used by access concentrators, but it may also be used by other authorized elements in the network. The DHCPLEASEQUERY message uses the DHCP message format as described in RFC 2131 [RFC2131], and uses message number 10 in the DHCP Message Type option (option 53). The DHCPLEASEQUERY message has the following pertinent message contents:

- o There MUST be a Relay Agent Information option (option 82) with only a Remote ID sub-option (sub-option 2) in the DHCPLEASEQUERY message.
- o The Parameter Request List option [RFC2132] MUST be populated by the access concentrator with the Associated-IP option code. The giaddr field and other option codes listed in the Parameter Request List option are set as explained in Section 6.2 of RFC 4388 [RFC4388].
- o The ciaddr field MUST be set to zero.
- o The values of htype, hlen, and chaddr MUST be set to zero.
- o The Client Identifier option (option 61) MUST NOT appear in the packet.

The DHCPLEASEQUERY message SHOULD be sent to a DHCP server that is known to possess authoritative information concerning the Remote ID. The DHCPLEASEQUERY message MAY be sent to more than one DHCP server, and in the absence of information concerning which DHCP server might possess authoritative information concerning the Remote ID, it SHOULD be sent to all DHCP servers configured for the associated relay agent (if any are known).

4.2. Responding to the DHCPLEASEQUERY Message

There are two possible responses to a DHCPLEASEQUERY message:

- o DHCPLEASEUNKNOWN

The DHCPLEASEUNKNOWN message indicates that the client associated with the Remote ID sub-option of the DHCPLEASEQUERY message is not allocated any lease or it is not managed by the server.

- o DHCPLEASEACTIVE

The DHCPLEASEACTIVE message indicates that the server not only knows the client specified in the DHCPLEASEQUERY message, but also knows that there is an active lease for that client.

4.3. Building a DHCPLEASEACTIVE or DHCPLEASEUNKNOWN Message

A DHCPLEASEACTIVE message is built by populating information pertaining to the client associated with the IP address specified in the ciaddr field.

In the case where more than one IP address has been involved in a DHCP message exchange with the client specified by the Remote ID, then the list of all those IP addresses MUST be returned in the Associated-IP option, whether or not that option was requested as part of the Parameter Request List option. This is intended for maintaining backwards compatibility with RFC 4388 [RFC4388].

All other options specified in the Parameter Request List [RFC2132] are processed as mentioned in Section 6.4.2 of RFC 4388 [RFC4388].

In a DHCPLEASEUNKNOWN response message, the DHCP server MUST echo the option 82 received in the DHCPLEASEQUERY message. No other option is included in the message.

4.4. Determining the IP Address to Be Used in Response

The IP address placed in the ciaddr field of a DHCPLEASEACTIVE message MUST be the IP address with the latest client-last-transaction-time associated with the client described by the Remote ID specified in the DHCPLEASEQUERY message.

If there is only a single IP address that fulfills this criteria, then it MUST be placed in the ciaddr field of the DHCPLEASEACTIVE message.

In the case where more than one IP address has been accessed by the client specified by the Remote ID, then the DHCP server MUST return the IP address returned to the client in the most recent transaction with the client, unless the DHCP server has been configured by the server administrator to use some other preference mechanism.

4.5. Sending a DHCPLEASEACTIVE or DHCPLEASEUNKNOWN Message

The server unicasts the DHCPLEASEACTIVE or DHCPLEASEUNKNOWN message to the address specified in the giaddr field of the DHCPLEASEQUERY message.

4.6. Receiving a DHCPLEASEACTIVE or DHCPLEASEUNKNOWN Message

When a DHCPLEASEACTIVE message is received in response to the DHCPLEASEQUERY message, it means that there is currently an active lease associated with the Remote ID in the DHCP server. The access concentrator SHOULD use the information in the htype, hlen, and chaddr fields of the DHCPLEASEACTIVE message as well as the Relay

Agent Information option included in the packet to refresh its location information for this IP address. An access concentrator is likely to query by IP address for all the IP addresses specified in the Associated-IP option in the response, if any, at this point in time.

When a DHCPLEASEUNKNOWN message is received by an access concentrator that had sent out a DHCPLEASEQUERY message, it means that the DHCP server does not have definitive information concerning the DHCP client specified in the Remote ID sub-option of the DHCPLEASEQUERY message. The access concentrator MAY store this information for future use. However, another DHCPLEASEQUERY message to the same DHCP server SHOULD NOT be attempted with the same Remote ID sub-option.

For lease query by Remote ID, the impact of negative caching is greatly reduced, as the response leads to "definitive" information on all the nodes connected behind the connection. Note that in the case of the data-driven approach [RFC4388], a node spoofing several IP addresses can lead to negative caching of greater magnitude. Another important change that this document brings is the removal of periodic lease queries generated from negative caching caused by DHCPLEASEUNKNOWN messages. Since the information obtained through query by Remote ID is complete, there is no need to attempt lease query again for the same connection.

4.7. Receiving No Response to the DHCPLEASEQUERY Message

The condition of an access concentrator receiving no response to a DHCPLEASEQUERY message is handled in the same manner as suggested in RFC 4388 [RFC4388].

4.8. Lease-Binding Data Storage Requirements

Implementation Note:

To generate replies for a lease query by Remote ID efficiently, a DHCP server should index the lease-binding data structures using Remote ID.

4.9. Using the DHCPLEASEQUERY Message with Multiple DHCP Servers

This scenario is handled in the same way it is done in RFC 4388 [RFC4388].

5. RFC 4388 Considerations

This document is compatible with RFC 4388-based [RFC4388] implementations, which means that a client that supports this extension can work with a server not supporting this document, provided it uses RFC 4388-defined query types. Also, a server supporting this document can work with a client not supporting this query type. However, there are some changes that this document proposes with respect to RFC 4388 [RFC4388]. Implementers extending RFC 4388 [RFC4388] implementations to support this document should take note of the following points:

- o There may be cases where a query by IP address/MAC address/Client Identifier has an option 82 containing a Remote ID. In that case, the query will still be recognized as a query by IP address/MAC address/Client Identifier as specified by RFC 4388 [RFC4388].
- o Section 6.4 of RFC 4388 [RFC4388] suggests that a DHCPLEASEUNKNOWN message MUST NOT have any other option present. But for a query by Remote ID, option 82 MUST be present in the reply.

6. Security Considerations

This document inherits the security concerns present in the original lease query protocol specification (RFC 4388 [RFC4388]).

This specification introduces one additional issue, beyond those described in RFC 4388 [RFC4388]. A query by Remote ID will result in the server replying with consolidated lease-binding information. Such a query, if done from an unauthorized source, may lead to a leak of lease-binding information. It is critical to deploy authentication techniques mentioned in RFC 3118 [RFC3118] to prevent such unauthorized lease queries.

7. Acknowledgments

Copious amounts of text in this document are derived from RFC 4388 [RFC4388]. Kim Kinnear, Damien Neil, Stephen Jacob, Ted Lemon, Ralph Droms, and Alfred Hoenes provided valuable feedback on this document.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4388] Woundy, R. and K. Kinnear, "Dynamic Host Configuration Protocol (DHCP) Leasequery", RFC 4388, February 2006.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.
- [RFC3118] Droms, R., Ed. and W. Arbaugh, Ed., "Authentication for DHCP Messages", RFC 3118, June 2001.

8.2. Informative References

- [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.

Authors' Addresses

Pavan Kurapati
Juniper Networks
Embassy Prime Buildings, C.V. Raman Nagar
Bangalore 560 093
India

EEmail: kurapati@juniper.net
URI: <http://www.juniper.net/>

D.T.V Ramakrishna Rao
Infosys Technologies Ltd.
44 Electronics City, Hosur Road
Bangalore 560 100
India

EEmail: ramakrishnadtv@infosys.com
URI: <http://www.infosys.com/>

Bharat Joshi
Infosys Technologies Ltd.
44 Electronics City, Hosur Road
Bangalore 560 100
India

EEmail: bharat_joshi@infosys.com
URI: <http://www.infosys.com/>