

DHC Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 14, 2013

L. Yeh
Freelancer Technologies
T. Lemon
Nominum, Inc
M. Boucadair
France Telecom
February 10, 2013

Prefix Pool Option for DHCPv6 Relay Agent on the Provider Edge Routers
draft-ietf-dhc-dhcpv6-prefix-pool-opt-02

Abstract

The DHCPv6 Prefix Pool option provides a mechanism for DHCPv6 Prefix Delegation (DHCPv6-PD), allowing the DHCPv6 server to notify a DHCPv6 relay agent implemented on a Provider Edge (PE) router about active prefix pools allocated by the DHCPv6 server to the PE router. The information of active prefix pools can be used to enforce IPv6 route aggregation on the PE router through adding or removing aggregation routes according to the status of the prefix pools. The advertising of the aggregation routes in the routing protocol enabled on the network-facing interface of PE routers will dramatically decrease the number of the routing table entries in the ISP network.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on August 14, 2013.

Copyright Notice

Copyright (c) 2013 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.

Table of Contents

1.	Introduction	3
2.	Terminology and Conventions	4
3.	Scenario and Network Architecture	5
4.	Prefix Pool Option	6
5.	Relay Agent Behavior	8
6.	Server Behavior	9
7.	Security Considerations	11
8.	IANA Considerations	11
9.	Contributors List	11
10.	Acknowledgements	11
11.	References	12
11.1.	Normative References	12
11.2.	Informative References	12
	Authors' Addresses	12

1. Introduction

The DHCPv6 protocol [[RFC3315](#)] specifies a mechanism for the assignment of IPv6 address and configuration information to IPv6 nodes. The DHCPv6 Prefix Delegation (DHCPv6-PD) [[RFC3633](#)] specifies a mechanism for the delegation of IPv6 prefixes from the Delegating Router (DR) acting as the DHCPv6 server to the Requesting Routers (RR) acting as the DHCPv6 clients. The DHCPv6 servers always maintain authoritative information associated with their operations including, but not limited to, the binding data of the delegated IPv6 prefixes, the lease data of the delegated IPv6 prefixes, the status of their prefix pools, etc. A prefix pool configured and maintained on the server can usually be a short prefix (e.g., a /40 prefix), out of which a longer prefixes (e.g., a /56 prefixes) are delegated to customer networks.

In the scenarios of a centralized DHCPv6 server, the Provider Edge (PE) routers act as DHCPv6 relay agents, when the DHCPv6 server and the Customer Edge (CE) router (a.k.a. Routed-RG or Routed-CPE) acting as RRs and the DHCPv6 clients, are not on the same link. For ensuring reachability, the PE routers always need to add or withdraw the route entries directing to each customer network in their routing table to reflect the status of IPv6 prefixes delegated by the DHCPv6 server to the CE routers (see [Section 6.2](#), [BBF TR-177]).

When a routing protocol is enabled on the network-facing interface of the PE router, all the routes directing to the customer networks are advertised in the ISP network. It will make the number of route entries in the routing table on the ISP routers be unacceptable large. Hence, it is desirable to aggregate the routes directing to the customer networks on the PE router.

Because the prefixes of the customer networks can not be guaranteed to be active and continuous, the routing protocol enabled on the PE router in general can not create one aggregation route automatically to cover all the prefixes delegated within the prefix pool. When the PE router acts as the relay agent, it can not be aware about the status of the prefix pools in general. The way to make the aggregation routes (e.g., black-hole routes) pointing to each of the prefix pools could be to configure them manually and permanently, but it is meant to a large amount of the handwork on each PE router for its operation and maintenance.

This document proposes a new Prefix Pool option for the DHCPv6 relay agent implemented on PE routers, allowing the DHCPv6 server to notify the DHCPv6 relay agent the information about the prefix pools. After the PE router received the prefix pools, the aggregation route entries can be added or withdrawn in the routing table of the PE

router according to the provision status of the prefix pools. The aggregation routes will then be advertised into the ISP network through the routing protocol enabled on the PE's network-facing interface.

DHCPv6 Bulk Leasequery [[RFC5460](#)] specifies a mechanism for bulk transfer of the binding data of each delegated prefix from the server to the requestor, typically a relay agent, to support the replacement or reboot event of a relay agent. In this document, the capability of DHCPv6 Bulk Leasequery will be extended to support the bulk transfer of the prefix and its status of the prefix pools for the route aggregation.

The automatic mechanisms described in this document depend on the existing DHCPv6 protocols and implementations without requiring a new DHCPv6 message or a new interface for the configuration of the aggregation route. The administrator of the ISP network can decide whether to inject the aggregation route or not based on the policies defined on the DHCPv6 server.

2. Terminology and Conventions

This document defines a new DHCPv6 option to communicate the prefix and its status of an IPv6 prefix pool. Definitions for terms and acronyms not specified in this document are defined in [[RFC3315](#)], [[RFC3633](#)], [[RFC5007](#)] and [[RFC5460](#)].

The following terms have been employed in this document:

- o Requesting Router (RR): A router defined in [[RFC3633](#)] that acts as a DHCPv6 client, and is requesting prefix to be delegated.
- o Delegating Router (DR): A router defined in [[RFC3633](#)] that acts as a DHCPv6 server, and is responding to the prefix request.
- o Prefix Pool: An IPv6 address space allocated with a common prefix, out of which the longer prefixes are delegated via prefix delegation.
- o aggregation route: A route entry created on an edge router, is based on the knowledge of a prefix pool of the delegated prefixes.
- o Requestor: A node defined in [[RFC5007](#)] that acts as the leasequery client.

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this

document, are to be interpreted as described in [BCP 14](#), [[RFC2119](#)].

3. Scenario and Network Architecture

Figure 1 and Figure 2 illustrate two typical cases of the targeted network architectures.

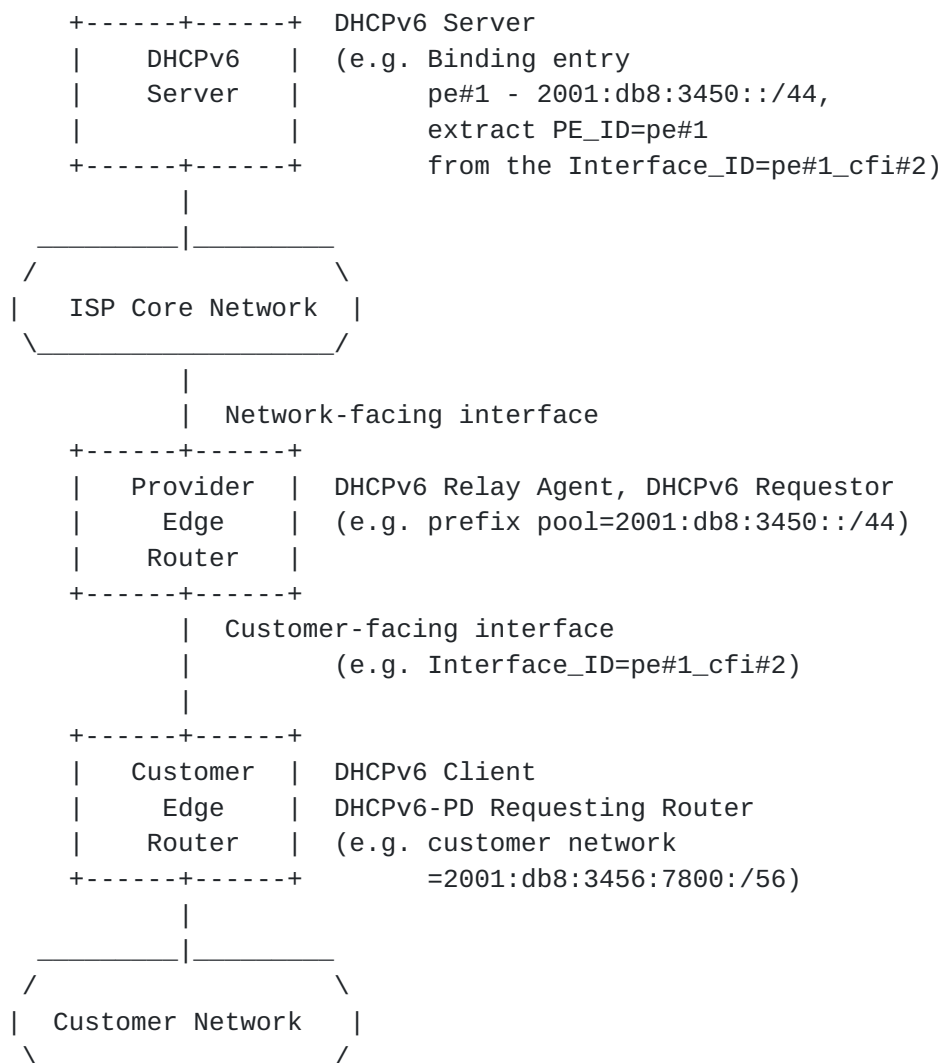


Figure 1: ISP-to-Customer network where CE is directly connected to PE

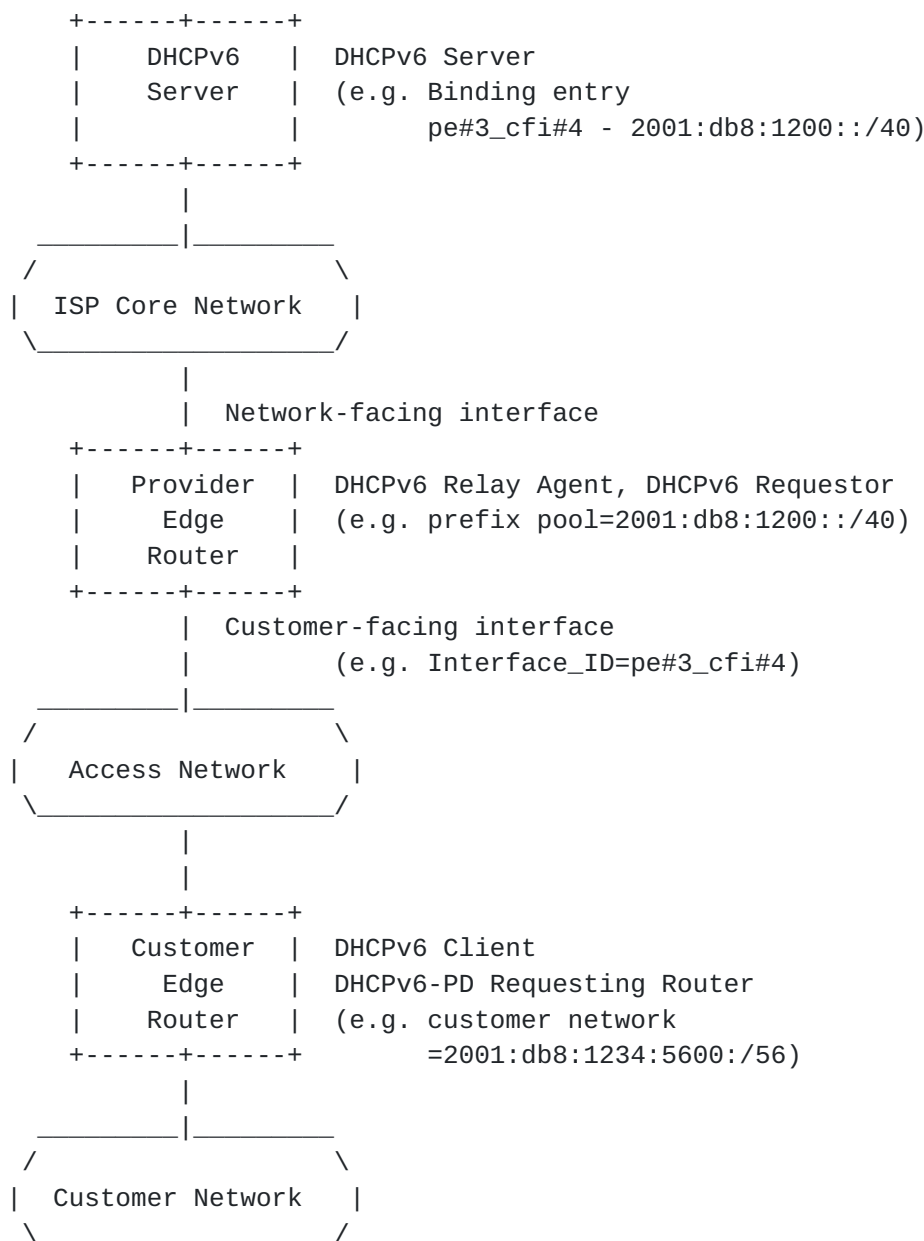
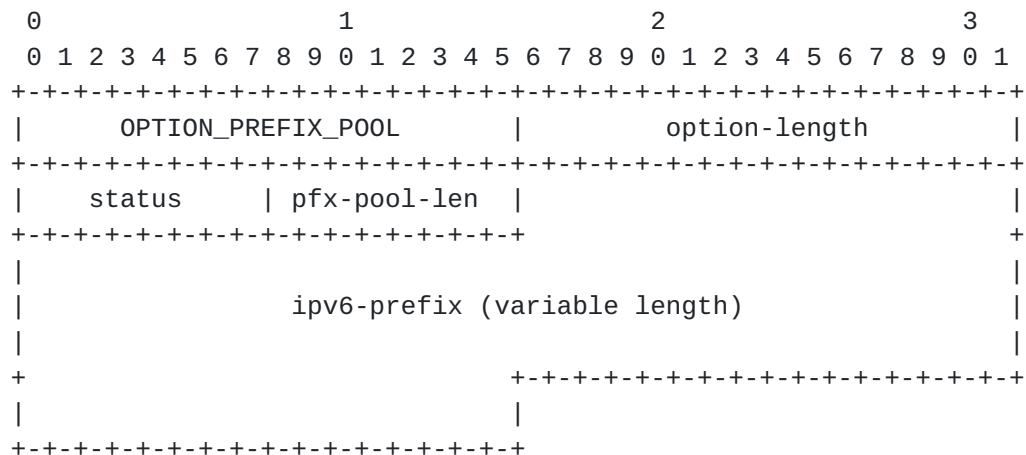


Figure 2: ISP-to-Customer network where CE is connected to PE through access network

4. Prefix Pool Option

The format of the Prefix Pool option is shown in Figure 3.



option-code: OPTION_PREFIX_POOL (TBA-IANA)
 option-length: 2 + length of ipv6-prefix in Octets
 status: Status of the prefix pool, indicating the
 availability of the prefix pool maintained
 on the server.
 pfx-pool-len: Length for the prefix pool in Bits
 ipv6-prefix: IPv6 prefix of the prefix pool, which is up to 16
 octets in length. Bits outside of the
 pfx-pool-len, if included, MUST be zero.

The codes of the status are defined in the following table.

Name	Code
Active	0
Released	1
Reserved	2~255

The 'Active' status of the prefix pool indicated in this option can be used to add the prefix pool and its associated aggregation route on the relay agent; while the 'Released' status of prefix pool indicated in this option can be used to withdraw the prefix pool and its associated aggregation route on the relay agent.

If the administrative policy on the server permits to support route aggregation on the relay agent, the status of prefix pool can be determined by the delegated prefixes within the associated prefix pool: If there is one delegated prefix within the pool that has a valid lease, the status of the prefix pool will be 'Active'; otherwise, the status of the prefix pool is 'Released'. If the administrative policy on the server does not permit to support route aggregation on the relay agent, the status of the prefix pool will always be 'Released'.

Prefix Pool Option MAY be included by the DHCPv6 server in RELAY-REPL

(13), LEASEQUERY-REPLY (15) and LEASEQUERY-DATA (17) message, and MAY be included by the DHCPv6 relay agent in the RELAY-FORW (12).

5. Relay Agent Behavior

The DHCPv6 relay agent who needs the information of prefix pools, SHOULD include the associated requested-option-code in Option Request option (OPTION_ORO, 6) to request the Prefix Pool option (OPTION_PREFIX_POOL, [TBD-IANA]) from the DHCPv6 server, who maintains the status of the prefix pools associated with the relay agent itself (Figure 1) or its particular customer-facing interface (Figure 2), when receiving the DHCPv6-PD message from clients. The relay agent SHOULD include this Option Request option for the Prefix Pool option in the RELAY-FORW (12) message of SOLICIT (1), REQUEST (3), RENEW(5), REBIND (6) and RELEASE (8). The relay agent MAY also include the Prefix Pool option with the values of 'pfx-pool-len' and 'ip6-prefix' to indicate its preference for which prefix pool the relay agent would like the server to return.

The relay agent SHOULD include the Interface ID option (OPTION_INTERFACE_ID, 18) so that the server can identify the relay agent itself or its particular customer-facing interface with which the prefix pool is associated, if the server would not like to use the link-address field specified in the encapsulation of the DHCPv6 RELAY-FORW message to identify the interface of the link on which the clients are located.

The PE router acting as the relay agent MAY set up a table for the lease or status of the prefix pools on it according to the leases of the delegated customer prefixes within the prefix pools. The lease of the prefix pool SHOULD dynamically set to be the maximum lease of the delegated customer prefix within it. If there is no route entry directing to the customer network within the aggregation route associated with the prefix pool or the lease of prefix pool runs out, the PE router acting as the relay agent SHOULD automatically withdraw the aggregation route.

After receiving the Prefix Pool option for the relay agent itself or its particular customer-facing interface in the RELAY-REPL (13) message of REPLY (7) from the server, the PE router acting as the relay agent SHOULD confirm the status of the prefix pool according to the leases of delegated customer prefixes within it. If the status of the prefix pool received and confirmed is 'Active', the PE router acting as the relay agent SHOULD add an aggregation route entry in its routing table, if the same entry has not been added before. If the status of the prefix pool received is 'Released', the PE router acting as the relay agent SHOULD withdraw the associated aggregation

route entry in its routing table, if the same entry has not been withdrawn before.

The PE router acting as the relay agent advertises its routing table including the entries of the aggregation routes based on the information of prefix pools when the routing protocol is enabled on its network-facing interface.

The PE router acting as the relay agent (i.e., Requestor) can use the DHCPv6 Bulk Leasequery [[RFC5460](#)] to query the binding data of prefix pools in the 'Active' status from the server. After established a TCP connection with the server, the relay agent SHOULD include Query option (OPTION_LQ_QUERY, 44) and set the proper query-type (QUERY_BY_RELAY_ID, QUERY_BY_LINK_ADDRESS or QUERY_BY_REMOTE_ID), link-address and query-options in the LEASEQUERY (14) message. The query options SHOULD include Option Request option to request the Prefix Pool option from the server.

6. Server Behavior

According to DHCPv6-PD [[RFC3633](#)], if the prefix of the customer network requested in RELAY-FORW (12) message of SOLICIT (1), REQUEST (3), RENEW(5), REBIND (6) from the DHCPv6 client (i.e., the RR) has a valid lease, the DHCPv6 server (i.e., the DR) will delegate the prefix with the relevant parameters in the RELAY-REPL (13) message of REPLY (7). In order to give a meaningful reply, the server has to maintain the binding data of the delegated IPv6 prefixes with the identification of the client. The Interface ID option (OPTION_INTERFACE_ID, 18) nested in the RELAY-FORW message is usually used to identify the access line of the client.

After receiving the Option Request option (OPTION_ORO, 6) requesting the Prefix Pool option (OPTION_PREFIX_POOL, [TBD]) in the RELAY-FORW messages of the DHCPv6-PD, the server SHOULD include the Prefix Pool option with the status indicated for the associated relay agent itself (Figure 1) or its customer-facing interface (Figure 2) in the RELAY-REPL messages, if the RELAY-FORW messages received are valid.

The server MAY use the link-address specified in RELAY-FORW message to identify the relay agent itself or its particular customer-facing interface where the prefix pool is associated, but the server has to maintain the binding data of prefix pools in association with these link-addresses. To be more readable, the server can alternatively use the Interface ID option included in the RELAY-FORW message by the relay agent to identify the relay agent itself or its particular customer-facing interface where the prefix pool is associated. In order to give a meaningful reply, the server has to maintain the

binding data of prefix pools in association with the information derived from the Interface ID option. According to DHCPv6 [[RFC3315](#)], the server SHOULD copy the Interface ID option from the RELAY-FORW message into the RELAY-REPL message.

If the administrative policy on the server permits to support route aggregation on the relay agents for some particular prefix pool, the status of prefix pool can be determined by the delegated prefixes within the associated prefix pool. If there is at least one delegated prefix within the pool that has a valid lease, the server SHOULD set the status of the associated prefix pool to be 'Active'. After the last prefix released in the associated prefix pool, the server SHOULD set the status of the associated prefix pool to be 'Released'. If the administrative policy on the server does not permit to support route aggregation on the relay agents, the server shall set the status of the associated prefix pools always to be 'Released'.

When the administrator of the server changes the setting to support route aggregation on the relay agent for the particular prefix pool, the status of the prefix pool SHOULD change from 'Released' to be 'Active' if at least one delegated prefix within the prefix pool has the valid lease. When the administrator of the server changes the setting not to support route aggregation on the relay agent for the particular prefix pool, the status of the prefix pool SHOULD change from 'Active' to be 'Released' if at least one delegated prefix within the prefix pool has the valid lease. The server MAY initiate a RELAY-REPL message of RECONFIGURE (10) to immediately trigger RENEW (5) / REPLY (7) prefix delegation message exchange with Prefix Pool option between one active client and the server.

Multiple prefix pools MAY be associated with the same PE router acting as the relay agent, or its customer-facing interface in the binding table on the server. Note that these prefix pools SHOULD not overlay, and the delegated customer prefix is only from one prefix pool.

After receiving the LEASEQUERY (14) message from the relay agent with the OPTION_LQ_QUERY (44) including the OPTION_ORO (6) to request the Prefix Pool option, the server SHOULD include the OPTION_PREFIX_POOL (TBD) in the LEASEQUERY-REPLY (15) and LEASEQUERY-DATA (17) messages to convey the binding data of the associated prefix pools through the established TCP connection according to mechanism defined in the DHCPv6 Bulk Leasequery [[RFC5460](#)]. Each LEASEQUERY-REPLY (15) and LEASEQUERY-DATA (17) message MAY only contain one OPTION_PREFIX_POOL, or and the associated OPTION_INTERFACE_ID, if the status of the prefix pool is 'active'. In order to be able to provide meaningful replies to different query types, the server has to maintain the

relevant association of prefix pools with the Relay_ID, link addresses or Remote_IDs of the relay agent in its binding database.

7. Security Considerations

Security issues related DHCPv6 are described in [Section 23 of \[RFC3315\]](#) and [Section 15 of \[RFC3633\]](#).

8. IANA Considerations

This document requests to assign a new option code for Option_Prefix_Pool in the registry of DHCPv6 Option Codes (<http://www.iana.org/assignments/dhcpv6-parameters/dhcpv6-parameters.xml>).

9. Contributors List

Juergen Schoenwaelder
Jacobs University Bremen
Bremen
Germany

Email: j.schoenwaelder@jacobs-university.de

Jie Hu
China Telecom
Beijing,
P. R. China

Email: huj@ctbri.com.cn

10. Acknowledgements

Thanks to Ralph Droms for the inspiration from his expired [\[I-D.ietf-dhc-dhcpv6-agentopt-delegate-04\]](#), to Tomek Mrugalski, Ole Troan and Alexandru Petrescu for their discussion in the mailing list of DHC, to Acee Lindem for his discussion in the mailing list of routing-discussion, to Christian Jacquenet for pointing out the draft shall cover one more use case of ISP-to-Customer network where CPE is directly connected to PE, to Sven Ooghe for some revisions in the email review, to Shrinivas Ashok Joshi for pointing out the draft shall cover the mechanism against the case of reboot, to Adrian Farrel for the orientation guide on this draft in IETF80 at Prague.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [RFC5007] Brzozowski, J., Kinnear, K., Volz, B., and S. Zeng, "DHCPv6 Leasequery", [RFC 5007](#), September 2007.
- [RFC5460] Stapp, M., "DHCPv6 Bulk Leasequery", [RFC 5460](#), February 2009.

11.2. Informative References

- [BBF TR-177]
Broadband Forum, "IPv6 in the context of TR-101, Issue 1", November 2010.
- [I-D.ietf-dhc-dhcpv6-agentopt-delegate-04]
Droms, R., Volz, B., and O. Troan, "DHCPv6 Relay Agent Assignment Notification (RAAN) Option", July 2009.

Authors' Addresses

Leaf Y. Yeh
Freelancer Technologies
P. R. China

Email: leaf.yeh.sdo@gmail.com

Ted Lemon
Nominum, Inc
USA

Email: Ted.Lemon@nominum.com

Mohamed Boucadair
France Telecom
France

Email: mohamed.boucadair@orange.com