DHC Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 29, 2013

L. Yeh, Ed.
Huawei Technologies
M. Boucadair
France Telecom
T. Lemon
Nominum, Inc
J. Hu
China Telecom
July 28, 2012

Prefix Pool Option for DHCPv6 Relay Agents on Provider Edge Routers
draft-yeh-dhc-dhcpv6-prefix-pool-opt-08

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 by adding or removing aggregated routes according to the status of the prefix pools. The advertising of the aggregated routes in the routing protocol enabled on the network-facing interface of PE routers will dramatically decreases 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 January 29, 2013.

Table of Contents

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. DHCPv6 servers always maintain authoritative information associated to their operations including, but not limited to, the binding data of the delegated IPv6 prefixes, the lease data of the delegated IPv6 prefixes, and the status of their prefix pools. A prefix pool configured and maintained on the server can usually be a short prefix (e.g., a /40 prefix) out of which the longer prefixes (e.g., /56 prefixes) are delegated to customer networks.

In the scenario 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 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 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. This will make the number of route entries in the routing table on the ISP router 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 aggregated route automatically to cover all the prefixes delegated within the prefix pool. One way to make the aggregated routes (e.g., black-hole routes) pointing to each of the prefix pools is to configure them manually and permanently, but the PE router is not really aware about the status of the prefix pools, especially when it acts as the relay agent.

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 about the prefix of pools. After the PE router received information about the prefix pools, the aggregated route entries per the provision status of the prefix pools can be added or withdrawn in the routing table of the PE router. The aggregated 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 (i.e., a DHCPv6 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 status of the prefix pools for route aggregation.

The automatic mechanisms described in this document depends on the existing DHCPv6 protocols and implementations without requiring a new DHCPv6 message or a new interface for the configuration of the aggregated route. The administrator of the ISP network can decide whether to inject the aggregated 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 of an IPv6 prefix pool. This document should be read in conjunction with the DHCPv6 specifications, [RFC3315], [RFC3633], [RFC5007] and [RFC5460], for understanding the complete mechanism. Definitions for terms and acronyms not specified in this document are defined in [RFC3315], [RFC3633], [RFC3769], [RFC5007] and [RFC5460].

The following terms can be found in this document:

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

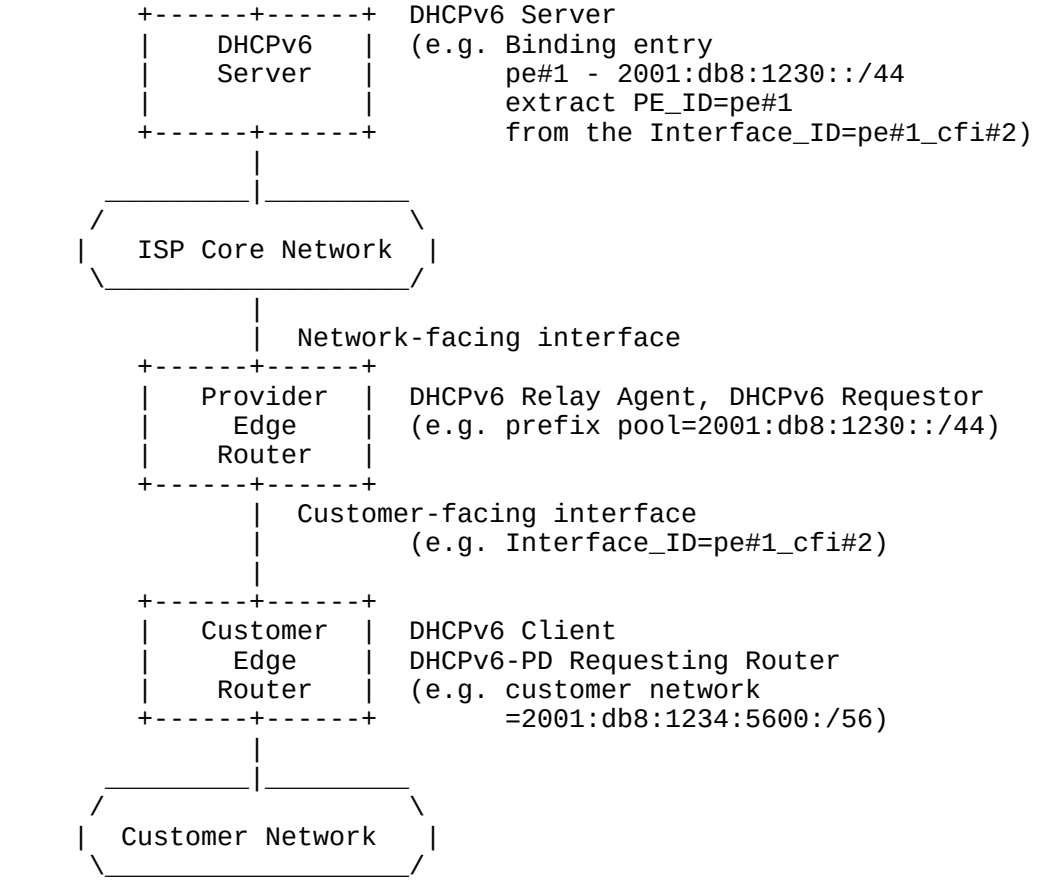

Figure 1: Use case of ISP-Customer network where CPE is directly connected to PE

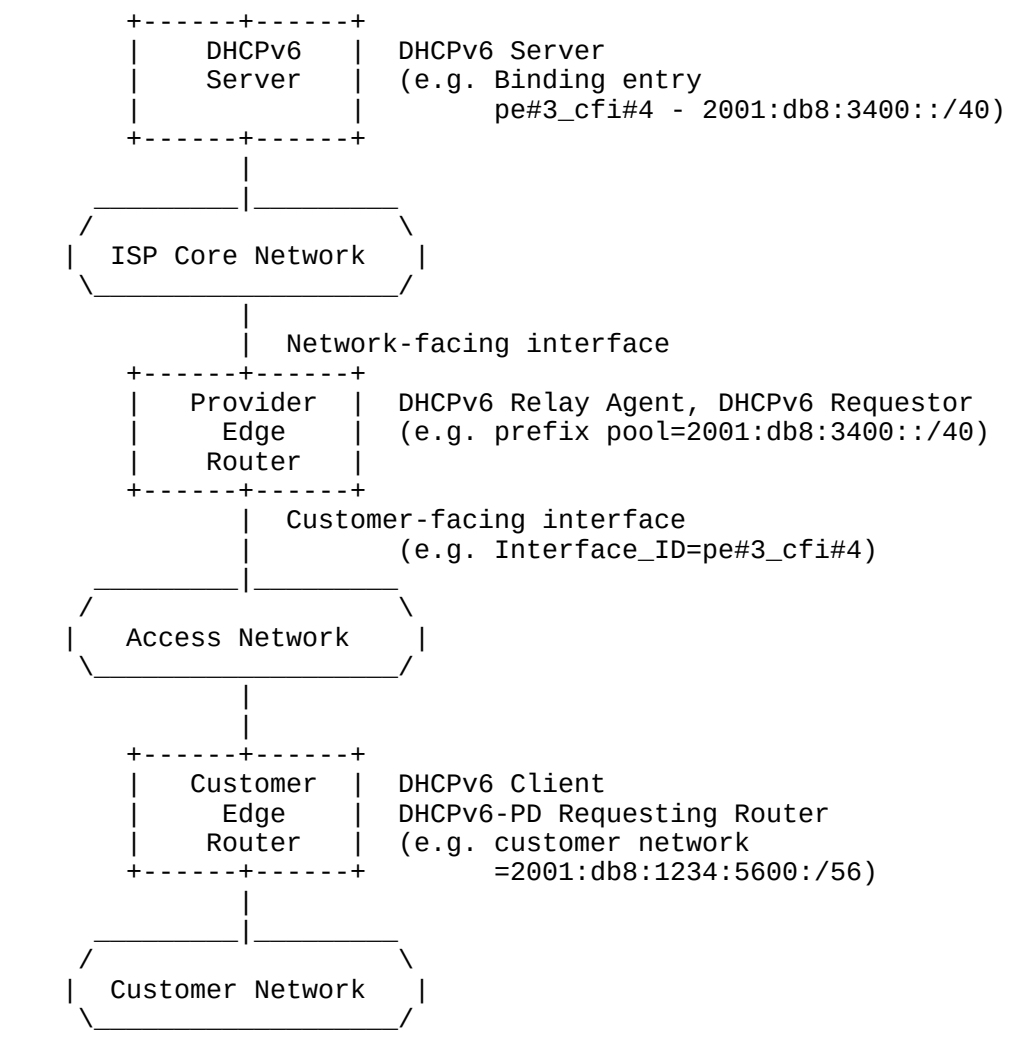

Figure 2: Use case of ISP-Customer network where CPE is connected to PE through access network

4. Prefix Pool Option

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

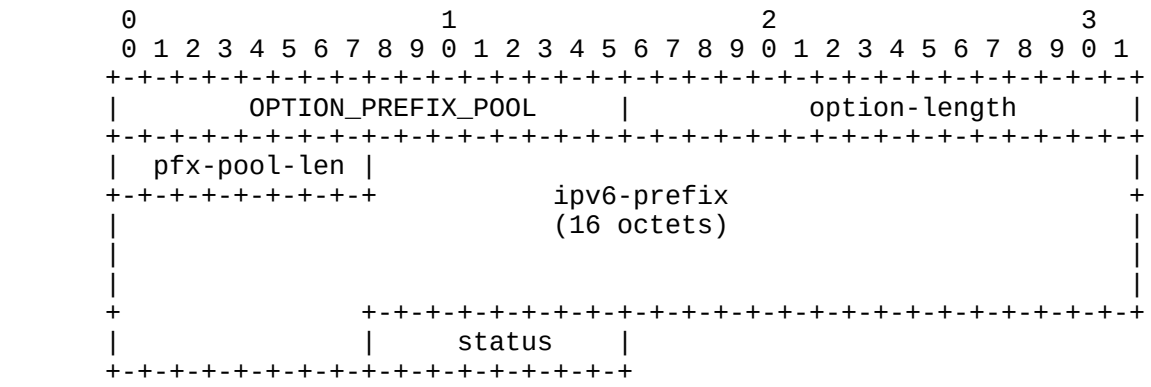

option-code: OPTION_PREFIX_POOL (TBD)
option-length: 18
pfx-pool-len: Length for the prefix pool in bits
ipv6-prefix: IPv6 prefix of the prefix pool
status: Status of the prefix pool, indicating the availability of the prefix pool maintained on the server.

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 aggregated 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 aggregated route on the relay agent.

If the administrative policy on the DHCPv6 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 DHCPv6 relay agent, the status of the prefix pool will always be 'Released'.

Discussion:

The alternative option might include the lease information in the prefix pool, then populate it to relay agent, make the state machine on the relay agent keep synchronizing the lease and status of the associated prefix pool with the server. But the solution proposed in this draft is to let relay agent confirm the received status of the prefix pool by itself as per the leases of delegated customer prefixes within it, and build its own lease for the prefix pool.

5. Relay Agent Behavior

The relay agent who needs the information of prefix pools, must include the associated requested-option-code in Option Request option (OPTION_ORO, 6) to request the Prefix Pool option (OPTION_PREFIX_POOL, [TBD]) from the DHCPv6 server, who maintains the status of the prefix pools associated to the relay agent itself (Figure 1) or its particular customer-facing interface (Figure 2), when receiving the DHCPv6-PD message from clients. The DHCPv6 relay agent can include this Option Request option for the Prefix Pool option in the relay-forward (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 IPv6-prefix to indicate its preference, which the 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 DHCPv6 server can identify the relay agent itself or its particular customer-facing interface to 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-forward message to identify the interface of the link on which the clients are located.

The relay agent may set up a table for the leases or status of the prefix pools on it as per the delegated customer prefixes within it. The lease of the prefix pools must dynamically set to be the maximum lease of the delegated customer prefixes. If there is no route entry directing to the customer network within the aggregated route associated with the prefix pool, the relay agent shall automatically withdraw the aggregated route.

After receiving the Prefix Pool option for the relay agent itself or its particular customer-facing interface in the relay-reply message (13) of REPLY (7) from the DHCPv6 server, the relay agent acting as the PE router shall confirm the status of the prefix pool as per the leases of delegated customer prefixes within it, then add or withdraw the aggregated route entry per the status of the prefix pool. If the

status of the prefix pool received and confirmed is 'Active', the relay agent shall add an aggregated route entry in its routing table, if the same entry has not been added in before. If the status of the prefix pool received is 'Released', the relay agent shall withdraw the associated aggregated route entry in its routing table, if the same entry has not been withdrawn before.

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

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 DHCPv6 server, the relay agent must include Query option (OPTION_LQ_QUERY, 44) and set the proper query-type (QUERY_BY_RELAY_ID, QUERY_BY_LINK_ADDRESS, QUERY_BY_REMOTE_ID), link-address and query-options in the LEASEQUERY (14) message. The query options must include Option Request option (OPTION_ORO, 6) to request the Prefix Pool option (OPTION_PREFIX_POOL, [TBD]) from the server.

6. Server Behavior

Per DHCPv6-PD [RFC3633], if the prefix of the customer network requested in relay-forward (12) message of SOLICIT, REQUEST, RENEW, REBIND 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-reply (13) message of REPLY. In order to give a meaningful reply, the server has to be able 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-forward 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-forward messages of the PD, the server must 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-reply messages if the relay-forward messages received are valid.

The server may use the link-address specified in relay-forward 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 (OPTION_INTERFACE_ID, 18) included in the relay-forward message by the relay agent to identify the relay agent itself (Figure 1) or its particular customer-facing interface (Figure 2) 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.

Per DHCPv6 [RFC3315], the server shall copy the same Interface ID option received via the relay-forward message into the relay-reply message.

If the administrative policy on the DHCPv6 server permits to support route aggregation on the relay agent for some particular prefix, 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 shall set the status of the associated prefix pool to be 'Active'. After the last prefix releasing in the associated prefix pool, the server shall 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 DHCPv6 relay agent, the server shall set the status of the 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 may 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 may change from 'Active' to be 'Released' if at least one delegated prefix within the prefix pool has the valid lease. Then the server may send a relay-reply message of RECONFIGURE (10) to initiate immediately a Renew (5) / Reply (7) PD message exchange with Prefix Pool option between one active client and the server.

Multiple prefix pools may be associated with the same PE router implementing a DHCPv6 relay agent (Figure 1) or its customer-facing interface (Figure 2) in the binding table on the server. Note that the delegated prefix is only from one prefix pool.

After receiving the LEASEQUERY (14) message from the relay agent with the Query option (OPTION_LQ_QUERY, 44) including the Option Request option (OPTION_ORO, 6) to request the Prefix Pool option (OPTION_PREFIX_POOL, [TBD]), the server must include the Client Data options (OPTION_CLIENT_DATA, 45) in the LEASEQUERY-REPLY (15) and

LEASEQUERY-DATA (16) message to convey the binding data of the associated prefix pools with the 'Active' status through the established TCP connection per [RFC5460]. Each Client Data option shall contain a Prefix Pool option, and may contain the Interface ID option (OPTION_INTERFACE_ID, 18). In order to be able to provide meaningful replies to different query types, the server has to be able to maintain the relevant association of prefix pools with the RELAY_ID, link addresses or Remote_ID of the relay agent in its binding database.

7. Security Considerations

Security issues related DHCPv6 are described in section 23 of [RFC3315].

8. IANA Considerations

IANA is requested to assign an option code to Option_Prefix_Pool from the "DHCPv6 and DHCPv6 options" registry (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

Tina Tsou
Huawei Technologies
Santa Clara, CA
USA

Email: tina.tsou.zouting@huawei.com

10. Acknowledgements

Thanks to Ralph Droms for the inspiration from his expired draft (RAAN option), to the DHC working group members, Bernie Volz, Ole Troan and Roberta Maglione for the discussion in the mailing list, to

Christian Jacquenet for pointing out the draft should cover one more use case of ISP-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 should cover the robust 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.

[RFC3769] Miyakawa, S. and R. Droms, "Requirements for IPv6 Prefix Delegation", RFC 3769, June 2004.

Authors' Addresses

Leaf Y. Yeh (editor)
Huawei Technologies
Shenzhen,
P. R. China

Email: leaf.y.yeh@huawei.com

Mohamed Boucadair
France Telecom
Rennes,
France

Email: mohamed.boucadair@orange.com

Ted Lemon
Nominum, Inc
USA

Email: Ted.Lemon@nominum.com

Jie Hu
China Telecom
Beijing,
P. R. China

Email: huj@ctbri.com.cn