

DHC Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 13, 2016

Y. Cui  
T. Li  
C. Liu  
Tsinghua University  
October 11, 2015

**DHCPv6 Prefix Length Hint Issues**  
**draft-cui-dhc-dhcpv6-prefix-length-hint-issue-01**

Abstract

DHCPv6 Prefix Delegation [[RFC3633](#)] allows a client to include a prefix-length hint value in the IA\_PD option to indicate a preference for the size of the prefix to be delegated, but is unclear about how the client and server should act in different situations involving the prefix-length hint. This document provides a summary of the existing problems with the prefix-length hint and guidance on what the client and server could do in different situations.

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 April 13, 2016.

Copyright Notice

Copyright (c) 2015 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

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Requirements Language . . . . .</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Problem Description . . . . .</a>	<a href="#">3</a>
<a href="#">3.1.</a>	<a href="#">Creation of Solicit Message by the Client . . . . .</a>	<a href="#">3</a>
<a href="#">3.2.</a>	<a href="#">Receipt of Solicit message by the Server . . . . .</a>	<a href="#">3</a>
<a href="#">3.3.</a>	<a href="#">Receipt of Advertise Message by the Client . . . . .</a>	<a href="#">4</a>
<a href="#">3.4.</a>	<a href="#">Creation of Renew/Rebind Message by the Client . . . . .</a>	<a href="#">4</a>
<a href="#">3.5.</a>	<a href="#">Receipt of Renew/Rebind Message by the Server . . . . .</a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Proposed Solution . . . . .</a>	<a href="#">5</a>
<a href="#">4.1.</a>	<a href="#">Creation of Solicit Message by the Client . . . . .</a>	<a href="#">5</a>
<a href="#">4.2.</a>	<a href="#">Receipt of Solicit message by the Server . . . . .</a>	<a href="#">5</a>
<a href="#">4.3.</a>	<a href="#">Receipt of Advertise Message by the Client . . . . .</a>	<a href="#">6</a>
<a href="#">4.4.</a>	<a href="#">Creation of Renew/Rebind Message by the Client . . . . .</a>	<a href="#">6</a>
<a href="#">4.5.</a>	<a href="#">Receipt of Renew/Rebind Message by the Server . . . . .</a>	<a href="#">6</a>
<a href="#">5.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">7.</a>	<a href="#">Contributors List . . . . .</a>	<a href="#">7</a>
<a href="#">8.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">7</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">8</a>

## [1.](#) Introduction

DHCPv6 Prefix Delegation [[RFC3633](#)] allows a client to include a prefix-length hint value in the message sent to the server, to indicate a preference for the size of the prefix to be delegated. A prefix-length hint is communicated by a client to the server by including an IA\_PD Prefix Option, encapsulated in an IA\_PD option, with the "IPv6 prefix" field set to zero and the "prefix-length" field set to a non-zero value. The servers are free to ignore the hint values depending on server policy. This would not cause problems for some hint values such as T1 and T2 lifetimes, but it would be an issue for the prefix-length hint. Some clients can't function normally when they're provided with a prefix which length is different from what they requested. E.g. if the client is asking for a /56 and the server returns a /64, the functionality of the client might be limited because it might not be able to split the prefix for all its interfaces. The clients usually have higher preference on the prefix-length hint than the other option hints, and it should be given more consideration.

The current specification is unclear about how the client and server should act in different situations involving the prefix-length hint.



From the client perspective, it should be able to use the prefix-length hint to signal to the server its real time need and it should be able to handle the prefixes which lengths are different from the prefix-length hint. This document provides guidance on what a client should do in different situations, to prevent it from failing. From the server perspective, the server is free to ignore the prefix-length hints depending on server policy, but in cases where the server has a policy for considering the hint, this document provides guidance on how the prefix-length hint should be handled by the server in different situations.

## **2. 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 [[RFC2119](#)].

## **3. Problem Description**

### **3.1. Creation of Solicit Message by the Client**

The Solicit message allows a client to ask servers for addresses and configuration parameters. When the client's configuration changes, it might require a prefix length different from what it had previously gotten. The server usually has a record of the prefix it delegated to the client during previous interactions. How should the client avoid getting the same prefix back from the server?

The server could decide whether to provide the client with the preferred prefix depending on server policy, but the client should be able to signal to the server that it wants a different prefix. The best way to assure a completely new delegated prefix is to send a new IAID in the IA\_PD. However, this would require the client device to have persistent storage, since rebooting the device would cause the client to use the original IAID in the IA\_PD.

### **3.2. Receipt of Solicit message by the Server**

[RFC3633] allows a client to include a prefix-length hint in the Solicit message, to signal its preference to the server. However, it is unclear about how this prefix-length hint should be handled by the server. Some servers will keep a record about prefixes it gave to the client during previous interactions, and give the client the same prefix. When the client includes a prefix-length hint in the Solicit message, the server has to decide whether to honor the newly requested prefix-length hint or give the client the recorded prefix. The client might want a different prefix length due to configuration changes or it might just want the same prefix again after reboot.



The server should interpret these cases differently.

Many servers are configured to provide only prefixes of specific lengths to the client. E.g. If the client requested for a /54, and the server could only provide /30,/48, and /56. How should these servers decide which prefix to give to the client based on the client's prefix-length hint?

### **3.3. Receipt of Advertise Message by the Client**

The server might not be able to honor the prefix-length hint due to server policy. If the prefix length provided by the server in the Advertise message is different from what the client requested in the Solicit message, the question would be whether the client should use the provided prefix length or continue to ask for its preferred prefix length. There are certain situations where the client would fail if it used a prefix which length is different from what it requested in the prefix-length hint. However, if the client ignores the Advertise messages, and continues to solicit for the preferred prefix length, the client might be stuck in the DHCP process.

### **3.4. Creation of Renew/Rebind Message by the Client**

Servers might not be able to provide a prefix matching the prefix-length hint requested by the client. If the client decided to use the prefix provided by the server which doesn't match the prefix-length hint, but would still prefer the prefix-length hint it originally requested in the Solicit message, there should be some way for the client to express this preference during Renew/Rebind. E.g. If the client requested for a /60 but got a /64, the client should be able to signal to the server during Renew/Rebind that it would still prefer a /60. This is to see whether the server has the prefix preferred by the client available in its prefix pool during Renew/Rebind.[\[RFC3633\]](#) is not completely clear on whether the client is allowed to include a prefix-length hint in the Renew/Rebind message.

### **3.5. Receipt of Renew/Rebind Message by the Server**

The prefix preferred by the client might become available in the server's prefix pool during Renew/Rebind, but was unavailable during Solicit. This might be due to server configuration change or because some other client stopped using the prefix.

The question is whether the server should remember the prefix-length hint the client originally included in the Solicit message and check during Renew/Rebind see if it has the prefix length the client preferred. This would require the server to keep extra information about the client. There is also the possibility that the client's



preference for the prefix length might have changed during this time interval, so the prefix-length hint remembered by the server might not be what the client prefers during Renew/Rebind.

Instead of having the server remember the prefix-length hint of the client, another option is for the client to include the prefix-length hint in the Renew/Rebind message. The current specification is unclear about what the server should do if the client also included in the Renew/Rebind message a prefix-length hint value, and whether the server could provide a different prefix to the client during Renew/Rebind.

#### **4. Proposed Solution**

##### **4.1. Creation of Solicit Message by the Client**

When the client prefers a prefix of specific length from the server, the client should send a Solicit message including the preferred prefix-length value in the "prefix-length" field of the IA\_PD Prefix option, and set the "IPv6 prefix" field to zero. This is an indication to the server that the client prefers a prefix of specific length, regardless of what it had gotten before.

When the client wants the same prefix back from the server, it should include the prefix value in the "IPv6 prefix" field of the IA\_PD Prefix option, and the length of the prefix in the "prefix-length" field. This is an indication to the server that the client wants the same prefix back.

##### **4.2. Receipt of Solicit message by the Server**

Upon the receipt of Solicit message, if the client included a prefix-length hint in the message, the server should try to honor the prefix-length hint within bounds of what the server is configured to return, regardless of the prefix record from previous interactions with the client. The server should regard the prefix-length hint in the Solicit message as the prefix length most preferred by the client at the time.

Many servers are configured to provide prefixes of specific lengths to the client. In this situation, the server should provide the shortest prefix length possible which is closest to the prefix-length hint. E.g. If the server could only provide prefixes with lengths /30, /48, and /56, and the client is requesting for a /50 in the prefix-length hint, then the server should provide the /48 to the client.



#### **4.3. Receipt of Advertise Message by the Client**

If none of the prefixes provided by the server in the Advertise messages match the prefix-length hint the client included in the Solicit message, the client could choose to either accept or ignore the prefixes provided by the servers depending on functional need.

If the client could use the prefixes provided by the servers despite being different from the prefix-length hint, the client should choose a prefix length closest to the prefix-length hint.

There are certain situations where the client will fail if it used a prefix which length does not meet its requirement. If the client cannot use the prefixes provided by the servers, it should ignore the Advertise messages and continue to send Solicit messages until it gets the preferred prefix. To avoid traffic congestion, the client should send Solicit messages at defined intervals, as specified in [\[RFC7083\]](#). To prevent the client from not functioning, the client should not ignore other configuration parameters provided by the server such as available IA\_NA addresses.

#### **4.4. Creation of Renew/Rebind Message by the Client**

During the Renew process, if the client prefers a prefix length different from the prefix it is currently using, then the client should send the Renew message with the same IA\_PD, and include two IA\_PD Prefix options, one containing the currently delegated prefix and the other containing the prefix-length hint. This is to extend lifetime of the prefix the client is currently using and also get the prefix the client prefers, and go through a graceful switch over.

If the server is unable to provide the client with the newly requested prefix, the client should continue using the prefix it currently has.

#### **4.5. Receipt of Renew/Rebind Message by the Server**

Upon the receipt of Renew message, if the client included in the IA\_PD both the delegated prefix value and a prefix-length hint value, the server should check to see whether it could extend the lifetime of the original delegated prefix and whether it has any available prefix matching the prefix-length hint, or as close as possible to the requested length, within the server's limit.

The server could do one of the following depending on server policy:

1. Renew just the original delegated prefix.



2. Renew the original delegated prefix and assign a new prefix of the requested length.
3. Mark the original delegated prefix as invalid by giving it 0 lifetimes, and assign a new prefix of requested length. This avoids the complexity of handling multiple delegated prefixes, but may break all the existing connections of the client.
4. Assign the original delegated prefix with 0 preferred-lifetime, a short non-zero valid-lifetime, and assign a new prefix of requested length. This is to provide the original delegated prefix with a short lifetime so the client can go through a graceful switch over.

It's unnecessary for the server to remember the prefix-length hint the client requested during Solicit. It is possible that the client's preference for the prefix length might have changed during this time interval, so the prefix-length hint in the Renew message is reflecting what the client prefers at the time.

## **5. Security Considerations**

TBD.

## **6. IANA Considerations**

This document does not include an IANA request.

## **7. Contributors List**

Many thanks to Qi Sun, Bernie Volz, Ole Troan, Sunil Gandhewar.

## **8. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC7083] Droms, R., "Modification to Default Values of SOL\_MAX\_RT and INF\_MAX\_RT", [RFC 7083](#), DOI 10.17487/RFC7083, November 2013, <<http://www.rfc-editor.org/info/rfc7083>>.



Authors' Addresses

Yong Cui  
Tsinghua University  
Beijing 100084  
P.R.China

Phone: +86-10-6260-3059  
Email: yong@csnet1.cs.tsinghua.edu.cn

Tianxiang Li  
Tsinghua University  
Beijing 100084  
P.R.China

Phone: +86-18301185866  
Email: peter416733@gmail.com

Cong Liu  
Tsinghua University  
Beijing 100084  
P.R.China

Phone: +86-10-6278-5822  
Email: gnocuil@gmail.com

