

DHC Working Group
Internet-Draft
Intended status: Informational
Expires: July 16, 2016

Y. Cui
T. Li
C. Liu
Tsinghua University
January 13, 2016

DHCPv6 Prefix Length Hint Issues
draft-ietf-dhc-dhcpv6-prefix-length-hint-issue-00

Abstract

DHCPv6 Prefix Delegation [[RFC3633](#)] allows a requesting router 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 requesting router and delegating router 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 requesting router and delegating router 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 July 16, 2016.

Copyright Notice

Copyright (c) 2016 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

Internet-Draft

DHCPv6 prefix-length hint Issues

January 2016

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	2
2.	Problem Description and Proposed Solutions	3
2.1.	Creation of Solicit Message	3
2.2.	Receipt of Solicit message	4
2.3.	Receipt of Advertise Message	5
2.4.	Creation of Renew/Rebind Message	5
2.5.	Receipt of Renew/Rebind Message	6
3.	Security Considerations	8
4.	IANA Considerations	8
5.	Contributors List	8
6.	Normative References	8
	Authors' Addresses	9

[1.](#) Introduction

DHCPv6 Prefix Delegation [[RFC3633](#)] allows a requesting router to include a prefix-length hint value in the message sent to the delegating router, to indicate a preference for the size of the prefix to be delegated. A prefix-length hint is communicated by a requesting router to the delegating router by including an IA_PD Prefix Option(OPTION_IAPREFIX), 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 delegating routers are free to ignore the prefix-length hint values depending on server policy. However, some requesting routers can't function normally when they're provided with a prefix which length is different from what they requested. E.g. If the requesting router is asking for a /56 and the delegating router returns a /64, the functionality of the requesting router might be limited because it might not be able to split the prefix for all its interfaces.

[RFC3633] is unclear about how the requesting router and delegating router should act in different situations involving the prefix-length hint. From the requesting router perspective, it should be able to use the prefix-length hint to signal to the delegating router 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 requesting router should do in different situations to help it operate properly. From the delegating router perspective, the delegating router is free to ignore the prefix-length hints depending on server policy, but in cases where the

delegating router has a policy for considering the hint, this document provides guidance on how the prefix-length hint should be handled by the delegating router in different situations.

[2.](#) Problem Description and Proposed Solutions

[2.1.](#) Creation of Solicit Message

Problem:

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

The delegating router could decide whether to provide the requesting router with the preferred prefix depending on server policy, but the requesting router should be able to signal to the delegating router whether it wants a different prefix or the same 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 requesting router device to have persistent storage, since rebooting the device would cause the requesting router to use the original IAID in the IA_PD.

Solution:

When the requesting router prefers a prefix of specific length from the delegating router, the requesting router should send a Solicit message including the preferred prefix-length value in the "prefix-length" field of the OPTION_IAPREFIX option, and set the "IPv6 prefix" field to zero. This is an indication to the delegating router that the requesting router prefers a prefix of specific length,

regardless of what it had gotten before.

When the requesting router wants the same prefix back from the delegating router, it should include the prefix value in the "IPv6 prefix" field of the OPTION_IAPREFIX option, and the length of the prefix in the "prefix-length" field. This is an indication to the delegating router that the requesting router wants the same prefix back.

[2.2.](#) Receipt of Solicit message

Problem:

[RFC3633] allows a requesting router to include a prefix-length hint in the Solicit message, to signal its preference to the delegating router. It is unclear about how this prefix-length hint should be handled by the delegating router, whether to honor the prefix-length hint or provide the prefix from previous interactions with the requesting router. The requesting router might want a different prefix length due to configuration changes or it might just want the same prefix again after reboot. The delegating router should interpret these cases differently.

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

Solution:

Upon the receipt of Solicit message, if the requesting router included only a prefix-length hint in the message, the delegating router should try to honor the prefix-length hint within bounds of what the delegating router is configured to return, regardless of the prefix record from previous interactions with the requesting router. The delegating router should regard the prefix-length hint in the

Solicit message as the prefix length most preferred by the requesting router at the time.

If the requesting router included a specific prefix value and the corresponding prefix-length value in the Solicit message, the delegating router should first try to provide the requested prefix to the requesting router. If the requested prefix is not available in the delegating router's prefix pool, then the delegating router should try to provide a prefix matching the prefix-length value.

The delegating router might not have prefixes exactly matching the prefix-length hint. In this situation, the delegating router should provide the shortest prefix length possible which is closest to the prefix-length hint. E.g. If the delegating router could only provide prefixes of lengths /30, /48, and /56, and the requesting router is requesting for a /50 in the prefix-length hint, then the delegating router should provide the /48 to the requesting router.

[2.3.](#) Receipt of Advertise Message

Problem:

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

Solution:

If none of the prefixes provided by the delegating router in the Advertise messages are equal or shorter than the prefix-length hint the requesting router included in the Solicit message, the requesting

router could choose to either accept or ignore the prefixes provided by the delegating routers depending on functional need.

If the requesting router could use the prefixes provided by the delegating routers despite being different from the prefix-length hint, the requesting router should choose the shortest prefix length which is closest to the prefix-length hint.

There are certain situations where the requesting router could not operate if it used a prefix which length does not meet its requirement. If the requesting router cannot use the prefixes provided by the delegating routers, it should ignore the Advertise messages and continue to send Solicit messages until it gets the preferred prefix. To avoid traffic congestion, the requesting router should send Solicit messages at defined intervals, as specified in [\[RFC7083\]](#). If the requesting router also Solicited for IA_NAs, the requesting router should accept the IA_NA addresses and continue to request for the desired IA_PD prefix in subsequent DHCPv6 messages as specified in [\[RFC7550\]](#)..

[2.4.](#) Creation of Renew/Rebind Message

Problem:

Delegating routers might not be able to provide a prefix with length equal or shorter than the prefix-length hint. If the requesting

router decided to use the prefix provided by the delegating router despite being longer than 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 requesting router to express this preference during Renew/Rebind. E.g. If the requesting router requested for a /60 but got a /64, the requesting router should be able to signal to the delegating router during Renew/Rebind that it would still prefer a /60. This is to see whether the delegating router has the prefix preferred by the requesting router available in its prefix pool during Renew/Rebind. [\[RFC3633\]](#) is not completely clear on whether the requesting router is allowed to include a prefix-length hint in the Renew/Rebind message.

Solution:

During Renew/Rebind, if the requesting router prefers a prefix length different from the prefix it is currently using, then the requesting router should send the Renew/Rebind message with the same IA_PD, and include two OPTION_IAPREFIX options, one containing the currently delegated prefix and the other containing the prefix-length hint. This is to extend the lifetime of the prefix the requesting router is currently using and also get the prefix the requesting router prefers, and go through a graceful switch over.

If the delegating router is unable to provide the requesting router with the newly requested prefix, but is able to extend lifetime of the old prefix, the requesting router should continue using the old prefix.

[2.5.](#) Receipt of Renew/Rebind Message

Problem:

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

The question is whether the delegating router should remember the prefix-length hint the requesting router originally included in the Solicit message and check during Renew/Rebind to see if it has the prefix length the requesting router preferred. This would require the delegating router to keep extra information about the requesting router. There is also the possibility that the requesting router's preference for the prefix length might have changed during this time interval, so the prefix-length hint remembered by the delegating router might not be what the requesting router prefers during Renew/

Rebind.

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

could provide a different prefix to the requesting router during Renew/Rebind.

Solution:

Upon the receipt of Renew/Rebind, if the requesting router included in the IA_PD both OPTION_IAPREFIX option with the delegated prefix value and an OPTION_IAPREFIX option with a prefix-length hint value, the delegating router 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 prefix-length hint, within the delegating router's limit.

If the delegating router assigned the prefix included in IA_PD to the requesting router, the delegating router can do one of the following, depending on its policy:

1. Extend lifetime of the original delegated prefix.
2. Extend lifetime of 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 requesting router.
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 allows the requesting router to finish up existing connections with the original prefix, and use the new prefix to establish new connections.
5. Do not include the original delegated prefix in the Reply message, and assign a new prefix of requested length. The original prefix would be valid until it's lifetime expires. This avoids sudden renumbering on the requesting router.

If the delegating router does not know the requesting router's

during Rebind), then the delegating router should ignore the original delegated prefix, and try to assign a new prefix of requested length.

It's unnecessary for the delegating router to remember the prefix-length hint the requesting router requested during Solicit. It is possible that the requesting router'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 requesting router prefers at the time.

3. Security Considerations

This document introduces no new security considerations over those already discussed in [section 15 of RFC3633](#), as this document provides guidance on how the requesting routers and delegating routers interact with regard to the prefix-length hint mechanism introduced in [RFC3633](#).

4. IANA Considerations

This document does not include an IANA request.

5. Contributors List

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

6. Normative References

- [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>>.
- [RFC7550] Troan, O., Volz, B., and M. Siodelski, "Issues and Recommendations with Multiple Stateful DHCPv6 Options", [RFC 7550](#), DOI 10.17487/RFC7550, May 2015, <<http://www.rfc-editor.org/info/rfc7550>>.

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

