

DHC Working Group
Internet-Draft
Expires: August 11, 2003

O. Troan
R. Droms
Cisco Systems
February 10, 2003

IPv6 Prefix Options for DHCPv6
draft-ietf-dhc-dhcpv6-opt-prefix-delegation-02.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 11, 2003.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

The Prefix Delegation options provide a mechanism for automated delegation of IPv6 prefixes using DHCP. This mechanism is intended for delegating long-lived prefix from a delegating router to a requesting router, across an administrative boundary, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Requirements	3
4.	Model and Applicability	4
5.	Identity Association for Prefix Delegation	6
6.	Overview of DHCP with Prefix Delegation	7
7.	Interface Selection	7
8.	Identity Association for Prefix Delegation Option	8
9.	IA_PD Prefix option	9
10.	Delegating Router Solicitation	11
10.1	Requesting router behaviour	11
10.2	Delegating router behaviour	11
11.	Requesting router initiated prefix delegation	12
11.1	Requesting router behaviour	13
11.2	Delegating Router behaviour	14
12.	Prefix Delegation reconfiguration	15
12.1	Delegating Router behaviour	15
12.2	Requesting Router behaviour	15
13.	Relay agent behaviour	15
14.	Security Considerations	15
15.	IANA Considerations	16
16.	Acknowledgements	16
17.	Changes since revision-01	16
	Normative References	16
	Informative References	17
	Authors' Addresses	17
	Full Copyright Statement	18

1. Introduction

This document describes new options for DHCP, which provide a mechanism for the delegation of IPv6 prefixes. Through these options, a delegating router can delegate prefixes to authorised requesting routers.

The prefix delegation mechanism described in this document is intended for simple delegation of prefixes from a delegating router to requesting routers. It is appropriate for situations in which the delegating router does not have knowledge about the topology of the networks to which the requesting router is attached, and the delegating router does not require other information aside from the identity of the requesting router to choose a prefix for delegation. For example, these options would be used by a service provider to assign a prefix to a CPE device acting as a router between the subscriber's internal network and the service provider's core network.

Many applications expect stable addresses. Even though this mechanism makes automatic renumbering easier, it is expected that prefixes have a long lifespan. During renumbering it is expected that the old and the new prefix co-exist for some time.

2. Terminology

This document uses the terminology defined in [RFC2460](#) [2] and the DHCP specification [6]. In addition, this document uses the following terms:

requesting router The router that acts as a DHCP client and is requesting prefix(es) to be assigned.

delegating router The router that acts as a DHCP server, and is responding to the prefix request.

Identity Association for Prefix Delegation (IA_PD) A collection of prefixes assigned to the requesting router. Each IA_PD has an associated IAID. A requesting router may have more than one IA_PD assigned to it; for example, one for each of its interfaces.

3. Requirements

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 [RFC 2119](#) [1].

4. Model and Applicability

The model of operation for prefix delegation is as follows. A delegating router is provided DHCPv6 prefixes to be delegated to requesting routers. Examples of ways in which the delegating router may be provided these prefixes are given in [Section 11.2](#). A requesting router requests prefix(es) from the delegating router, as described in [Section 11.1](#). The delegating router chooses prefix(es) for delegation, and returns the prefix(es) to the requesting router. The requesting router is then responsible for the delegated prefix(es). For example, the requesting router might assign a subnet from a delegated prefix to one of its interfaces, and begin sending router advertisements for the prefix on that link.

Each prefix has an associated valid and preferred lifetime, which constitutes an agreement about the length of time over which the requesting router is allowed to use the prefix. A requesting router can request an extension of the lifetimes on a delegated prefix and is required to terminate the use of a delegated prefix if the valid lifetime of the prefix expires.

This prefix delegation mechanism would be appropriate for use by an ISP to delegate a prefix to a subscriber, where the delegated prefix would possibly be subnetted and assigned to the links within the subscriber's network.

Figure 1 illustrates a network architecture in which prefix delegation would be used.




Figure 1: An example of prefix delegation.

In this example an AAA server is configured with a prefix assigned to the customer at the time of subscription to the ISP service. The prefix delegation process begins when the requesting router requests configuration information through DHCP. The DHCP messages from the requesting router are received by the delegating router in the aggregation device. When the delegating router receives the request, it consults the AAA server to authenticate and authorise the requesting router. The AAA server returns the subscriber's prefix(es) in a Framed-IPv6-Prefix attribute as described in [RFC 3162 \[7\]](#), and the delegating router returns them to the requesting router.

The requesting router assigns longer prefixes from the delegated prefix for assignment to links in the subscriber's network. In a typical scenario based on the network shown in Figure 1, the requesting router subnets a single delegated /48 prefix into /64 prefixes and assigns one /64 prefix to each of the links in the subscriber network.

The prefix delegation options can be used in conjunction with other DHCP options carrying other configuration information to the requesting router. The requesting router may, in turn, then provide DHCP service to hosts attached to the internal network. For example, the requesting router may obtain the addresses of DNS and NTP servers from the ISP delegating router, and then pass that configuration information on to the subscriber hosts through a DHCP server in the requesting router.

5. Identity Association for Prefix Delegation

An IA_PD is a construct through which a delegating router and a requesting router can identify, group and manage a set of related IPv6 prefixes. Each IA_PD consists of an IAID and associated configuration information. An IA_PD for prefixes is the equivalent of an IA (described in DHCPv6 specification [6]) for addresses.

An IA_PD is different from an IA, in that it does not need to be associated with exactly one interface. One IA_PD can be associated with the requesting router, with a set of interfaces or with exactly one interface. A requesting router must create at least one distinct IA_PD. It may associate a distinct IA_PD with each of its downstream network interfaces and use that IA_PD to obtain a prefix for that interface from the delegating router.

The IAID uniquely identifies the IA_PD and must be chosen to be unique among the IA_PD IDs on the requesting router. The IAID is chosen by the requesting router. For any given use of an IA_PD by the requesting router, the IAID for that IA_PD MUST be consistent across restarts of the requesting router. The requesting router may maintain consistency either by storing the IAID in non-volatile storage or by using an algorithm that will consistently produce the same IAID as long as the configuration of the requesting router has not changed. If the requesting router uses only one IAID, it can use a well-known value, e.g zero.

The configuration information in an IA_PD consists of one or more IPv6 prefixes along with the times T1 and T2 for the IA_PD. See section [Section 8](#) for the representation of an IA_PD in a DHCP message.

6. Overview of DHCP with Prefix Delegation

Prefix delegation with DHCP is independent of address assignment with DHCP. A requesting router can use DHCP for just prefix delegation or for prefix delegation along with address assignment and other configuration information.

A requesting router first creates an IA_PD and assigns it an IAID. The requesting router then transmits a Solicit message containing an IA_PD option describing the IA_PD. Delegating routers that can delegate prefixes to the IA_PD respond to the requesting router with an Advertise message.

The requesting router may include prefixes in the IA_PDs as a hint to the delegating router about specific prefixes for which the requesting router has a preference.

When the requesting router has identified a delegating router, the requesting router uses a Request message to populate the IA_PDs with prefixes. The requesting router includes one or more IA_PD options in the Request message. The delegating router returns prefixes and other information about the IA_PDs to the requesting router in IA_PD options in a Reply message. The requesting router records the lifetimes for the delegated prefix(es) and uses the prefix(es) as described in the previous section.

Before the valid lifetime on each delegated prefix expires, the requesting router includes the prefix in an IA_PD option sent in a Renew message to the delegating router. The delegating router responds by returning the prefix with updated lifetimes to the requesting router.

7. Interface Selection

Delegated prefixes are not associated with a particular interface in the same way as addresses are for address assignment, and the rules described in the section "Client Source Address and Interface Selection" of the DHCP specification [6] do not apply.

When a requesting router sends a DHCP message, it SHOULD be sent on the interface associated with the upstream router (ISP network). The upstream interface is typically determined by configuration. This rule applies even in the case where a separate IA_PD is used for each downstream interface.

When a requesting router sends a DHCP message directly to a delegating router using unicast (after receiving the Server Unicast option from that delegating router), the source address SHOULD be an

T2 The time at which the requesting router contacts any available delegating router to extend the lifetimes of the prefixes assigned to the IA_PD; T2 is a time duration relative to the current time expressed in units of seconds.

IA_PD-options Options associated with this IA_PD.

The IA_PD-options field encapsulates those options that are specific to this IA_PD. For example, all of the IA_PD Prefix Options carrying the prefixes associated with this IA_PD are in the IA_PD-options field.

An IA_PD option may only appear in the options area of a DHCP message. A DHCP message may contain multiple IA_PD options.

The status of any operations involving this IA_PD is indicated in a Status Code option in the IA_PD-options field.

Note that an IA_PD has no explicit "lifetime" or "lease length" of its own. When the valid lifetimes of all of the prefixes in a IA_PD have expired, the IA_PD can be considered as having expired. T1 and T2 are included to give delegating routers explicit control over when a requesting router recontacts the delegating router about a specific IA_PD.

In a message sent by a requesting router to a delegating router, values in the T1 and T2 fields indicate the requesting router's preference for those parameters. The requesting router sets T1 and T2 to 0 if it has no preference for those values. In a message sent by a delegating router to a requesting router, the requesting router MUST use the values in the T1 and T2 fields for the T1 and T2 parameters. The values in the T1 and T2 fields are the number of seconds until T1 and T2.

The delegating router selects the T1 and T2 times to allow the requesting router to extend the lifetimes of any prefixes in the IA_PD before the lifetimes expire, even if the delegating router is unavailable for some short period of time. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the prefixes in the IA_PD, respectively. If the time at which the prefixes in an IA_PD are to be renewed is to be left to the discretion of the requesting router, the delegating router sets T1 and T2 to 0.

9. IA_PD Prefix option

The IA_PD Prefix option is used to specify IPv6 address prefixes associated with an IA_PD. The IA_PD Prefix option must be encapsulated in the IA_PD-options field of an IA_PD option.

In a message sent by a requesting router to a delegating router, the values in the fields can be used to indicate the requesting router's preference for those values. The requesting router may send a value of zero to indicate no preference. A requesting router may set the

IPv6 prefix field to zero and a given value in the prefix-length field to indicate a preference for the size of the prefix to be delegated.

In a message sent by a delegating router the preferred and valid lifetimes should be set to the values of AdvPreferredLifetime and AdvValidLifetime as specified in section "Router Configuration Variables" of [RFC2461](#) [3], unless administratively configured.

The values in the preferred and valid lifetimes are the number of seconds remaining for each lifetime.

An IA_PD Prefix option may appear only in an IA_PD option. More than one IA_PD Prefix Option can appear in a single IA_PD option.

The status of any operations involving this IA_PD Prefix option is indicated in a Status Code option in the IAprefix-options field.

[10. Delegating Router Solicitation](#)

The requesting router locates and selects a delegating router in the same way as described in section "DHCP Server Solicitation" of the DHCP specification [6]. The details of the solicitation process are described in this section.

[10.1 Requesting router behaviour](#)

The requesting router creates and transmits a Solicit message as described in sections "Creation of Solicit Messages" and "Transmission of Solicit Messages" of the DHCP specification [6]. The requesting router creates an IA_PD and assigns it an IAID. The requesting router MUST include the IA_PD option in the Solicit message.

The requesting router processes any received Advertise messages as described in section "Receipt of Advertise Messages" in the DHCP specification [6]. The requesting router MAY choose to consider the presence of advertised prefixes in its decision about which delegating router to respond to.

The requesting router MUST ignore any Advertise message that includes a Status Code option containing the value NoPrefixAvail, with the exception that the requesting router MAY display the associated status message to the user.

[10.2 Delegating router behaviour](#)

The delegating router processes Solicit messages from requesting

routers in the same way as described in section "Receipt of Solicit messages" of the DHCP specification [6]. If the message contains an IA_PD option and the delegating router is configured to delegate prefix(es) to the requesting router, the delegating router selects the prefix(es) to be delegated to the requesting router. The mechanism through which the delegating router selects prefix(es) for delegation is not specified in this document. Examples of ways in which the delegating router might select prefix(es) for a requesting router include: static assignment based on subscription to an ISP; dynamic assignment from a pool of available prefixes; selection based on an external authority such as a RADIUS server using the Framed-IPv6-Prefix option as described in [RFC 3162](#) [7].

If the delegating router cannot delegate any prefixes to an IA_PD in the message from the requesting router, the delegating router MUST include the IA_PD in the Reply message with no prefixes in the IA_PD and a Status Code option in the IA_PD containing status code NoPrefixAvail.

If the requesting router includes an IA_PD Prefix option in the IA_PD option in its Solicit message, the delegating router MAY choose to use the information in that option to select the prefix(es) or prefix size to be delegated to the requesting router.

The delegating router sends an Advertise message to the requesting router in the same way as described in section "Creation and transmission of Advertise messages" in the DHCP specification [6]. The delegating router MUST include an IA_PD option, identifying any prefix(es) that the delegating router will delegate to the requesting router.

If the delegating router will not assign any prefixes to any IA_PDs in a subsequent Request from the requesting router, the delegating router MUST send an Advertise message to the requesting router that includes a Status Code option with code NoPrefixAvail and a status message for the user, a Server Identifier option with the delegating router's DUID and a Client Identifier option with the requesting router's DUID.

11. Requesting router initiated prefix delegation

A requesting router uses the same message exchanges as described in section "DHCP Client-Initiated Configuration Exchange" of the DHCP specification [6] to obtain or update prefix(es) from a delegating router. The requesting router and the delegating router use the IA_PD Prefix option to exchange information about prefix(es) in much the same way IA Address options are used for assigned addresses.

11.1 Requesting router behaviour

The requesting router uses a Request message to populate IA_PDs with prefixes. The requesting router includes one or more IA_PD options in the Request message. The delegating router then returns the prefixes for the IA_PDs to the requesting router in IA_PD options in a Reply message.

The requesting router includes IA_PD options in any Renew, or Rebind messages sent by the requesting router. The IA_PD option include all of the prefixes the requesting router currently has associated with that IA_PD.

In some circumstances the requesting router may need verification that the delegating router still has a valid binding for the requesting router. Examples of times when a requesting router may ask for such verification include:

- o The requesting router reboots.
- o The requesting router's upstream link flaps.
- o The requesting router is physically disconnected from a wired connection.

If such verification is needed the requesting router MUST initiate a Rebind/Reply message exchange as described in the section "Creation and Transmission of Rebind Messages" of the DHCP specification [6], with the exception that the retransmission parameters should be set as for the Confirm message, described in the section "Creation and Transmission of Confirm Messages" of the DHCP specification [6]. The requesting router includes any IA_PDs, along with prefixes associated with those IA_PDs in its Rebind message.

Each prefix has valid and preferred lifetimes whose duration is specified in the IA_PD Prefix option for that prefix. The requesting router uses Renew and Rebind messages to request the extension of the lifetimes of a delegated prefix.

The requesting router uses a Release message to return a delegated prefix to a delegating router. The prefixes to be released MUST be included in the IA_PDs.

The Confirm and Decline message types are not used with Prefix Delegation.

Upon the receipt of a valid Reply message, for each IA_PD the requesting router assigns a subnet from each of the delegated

prefixes to each of the links to which the associated interfaces are attached, with the following exception: the requesting router **MUST NOT** assign any delegated prefixes or subnets from the delegated prefix(es) to the link through which it received the DHCP message from the delegating router.

When a requesting router subnets a delegated prefix, it must assign additional bits to the prefix to generate unique, longer prefixes. For example, if the requesting router in Figure 1 were delegated 3FFE:FFFF:0::/48, it might generate 3FFE:FFFF:0:1::/64 and 3FFE:FFFF:0:2::/64 for assignment to the two links in the subscriber network. If the requesting router were delegated 3FFE:FFFF:0::/48 and 3FFE:FFFF:1::/48, it might assign 3FFE:FFFF:0:1::/64 and 3FFE:FFFF:1:1::/64 to one of the links, and 3FFE:FFFF:0:2::/64 and 3FFE:FFFF:1:2::/64 for assignment to the other link.

If the requesting router assigns a delegated prefix to a link to which the router is attached, and begins to send router advertisements for the prefix on the link, the requesting router **MUST** set the valid lifetime in those advertisements to be no later than the valid lifetime specified in the IA_PD Prefix option. A requesting router **MAY** use the preferred lifetime specified in the IA_PD Prefix option.

11.2 Delegating Router behaviour

When a delegating router receives a Request message from a requesting router that contains an IA_PD option, and the delegating router is authorised to delegate prefix(es) to the requesting router, the delegating router selects the prefix(es) to be delegated to the requesting router. The mechanism through which the delegating router selects prefix(es) for delegation is not specified in this document. [Section 10.2](#) gives examples of ways in which a delegating router might select the prefix(es) to be delegated to a requesting router.

A delegating router examines the prefix(es) identified in IA_PD Prefix options (in an IA_PD option) in Renew and Rebind messages and responds according to the current status of the prefix(es). The delegating router returns IA_PD Prefix options (within an IA_PD option) with updated lifetimes for each valid prefix in the message from the requesting router. If the delegating router cannot find a binding for the requesting router's IA_PD the delegating router returns the IA_PD containing no prefixes with a Status Code option set to NoBinding in the Reply message. If the delegating router finds that any of the prefixes are not in the requesting router's binding entry, the delegating router returns the prefix to the requesting router with lifetimes of 0.

A delegating router may mark any prefix(es) in IA_PD Prefix options in a Release message from a requesting router as "available", dependent on the mechanism used to acquire the prefix, e.g in the case of a dynamic pool.

The delegating router MUST include an IA_PD Prefix option or options (in an IA_PD option) in Reply messages sent to a requesting router.

12. Prefix Delegation reconfiguration

This section describes prefix delegation in Reconfigure message exchanges.

12.1 Delegating Router behaviour

The delegating router initiates a configuration message exchange with a requesting router, as described in the section "DHCP Server-Initiated Configuration Exchange" of the DHCP specification [6]. The delegating router specifies the IA_PD option in the Option Request option to cause the requesting router to include an IA_PD option to obtain new information about delegated prefix(es).

12.2 Requesting Router behaviour

The requesting router responds to a Reconfigure message received from a delegating router as described in the DHCP specification [6]. The requesting router MUST include the IA_PD Prefix option(s) (in an IA_PD option) for prefix(es) that have been delegated to the requesting router by the delegating router from which the Reconfigure message was received.

13. Relay agent behaviour

A relay agent forwards messages containing Prefix Delegation options in the same way as described in section "Relay Behaviour" of the DHCP specification [6].

If a delegating router communicates with a requesting router through a relay agent, the delegating router may need a protocol or other out-of-band communication to add routing information for delegated prefixes into the provider edge router.

14. Security Considerations

Security considerations in DHCP are described in the section "Security Considerations" of the DHCP specification [6].

A rogue delegating router can issue bogus prefixes to a requesting

router. This may cause denial of service due to unreachability.

An intruder requesting router may be able to mount a denial of service attack by repeated requests for delegated prefixes that exhaust the delegating router's available prefixes.

To guard against attacks through prefix delegation, requesting routers and delegating routers SHOULD use DHCP authentication as described in section "Authentication of DHCP messages" in the DHCP specification [6]. For point to point links, where one trusts that there is no man in the middle, or one trusts layer two authentication, DHCP authentication or IPsec may not be necessary. Because a requesting router and delegating routers must each have at least one assigned IPv6 address, the routers may be able to use IPsec for authentication of DHCPv6 messages. The details of using IPsec for DHCPv6 are under development.

15. IANA Considerations

IANA is requested to assign option codes to these options from the option-code space as defined in section "DHCPv6 Options" of the DHCPv6 specification [6].

IANA is requested to assign a status code to the NoPrefixAvail status code from the status-code space as defined in section "Status Codes" of the DHCPv6 specification [6].

16. Acknowledgements

Thanks for the input and review by (in alphabetical order) Steve Deering, Dave Forster, Brian Haberman, Tatuya Jinmei, Shin Miyakawa, Pekka Savola, Bernie Volz, Trevor Warwick and Toshi Yamasaki.

17. Changes since revision-01

- o Clarified the usage of how Preferred/Valid lifetimes should be used in Router Advertisements.
- o Clarified the use of NoPrefixAvail in the case where the delegating router cannot delegate any prefixes.
- o Use Rebind/Reply message exchange for binding confirmation rather than Renew/Reply.

Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [2] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [3] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.
- [4] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 2373](#), July 1998.
- [5] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [6] Droms, R., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [draft-ietf-dhc-dhcpv6-28](#) (work in progress), November 2002.
- [7] Aboba, B., Zorn, G. and D. Mitton, "RADIUS and IPv6", [RFC 3162](#), August 2001.

Informative References

- [8] Miyakawa, S., "Requirements for IPv6 prefix delegation", [draft-ietf-ipv6-prefix-delegation-requirement-00](#) (work in progress), November 2002.

Authors' Addresses

Ole Troan
Cisco Systems
250 Longwater Avenue
Reading RG2 6GB
United Kingdom

Phone: +44 20 8824 8666
EMail: ot@cisco.com

Ralph Droms
Cisco Systems
300 Apollo Drive
Chelmsford, MA 01824
USA

Phone: +1 978 497 4733
EMail: rdroms@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

