

HOMENET WG	E. Nordmark
Internet-Draft	Cisco Systems
Intended status: Standards Track	S. Chakrabarti
Expires: May 01, 2012	S. Krishnan
	W. Haddad
	Ericsson
	October 29, 2011

Simple Approach to Prefix Distribution in Basic Home Networks
draft-chakrabarti-homenet-prefix-alloc-01

Abstract

Modern IPv4 home networks are often configured with multi-level of NATs and Residential gateways to separate islands of networks used for different purposes. With the introduction of IPv6 home networks we'd like to be able to maintain the same topological freedom as we have with IPv4 but without requiring any IPv6 NATs. This document specifies the topological restrictions for what we term Basic Home Networks and specifies how DHCPv6 Prefix Delegation can be used to autoconfigure IPv6 address prefixes in such networks.

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 May 01, 2012.

Copyright Notice

Copyright (c) 2011 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](#)
- *2. [Definition Of Terms](#)
- *3. [Conceptual model of a Basic Home Router](#)
- *4. [Topology Assumptions](#)
- *5. [Topology Examples](#)
- *6. [Automatic hierarchical DHCPv6 Prefix Delegation](#)
 - *6.1. [Example](#)
 - *6.2. [Configurability](#)
 - *6.3. [Routing Implications](#)
 - *6.4. [Neighbor Discovery Implications](#)
 - *6.5. [Ensuring stable prefixes](#)
- *7. [Addressing](#)
- *8. [Basic Operations](#)
 - *8.1. [DHCPv6 Prefix Delegation](#)
- *9. [Host Behavior](#)
- *10. [Router Behavior](#)
- *11. [Bootstrapping](#)
- *12. [What if there are loops?](#)
- *13. [Security Considerations](#)
- *14. [IANA Considerations](#)
- *15. [Acknowledgements](#)
- *16. [References](#)
 - *16.1. [Normative References](#)
 - *16.2. [Informative References](#)
- *[Authors' Addresses](#)

1. Introduction

In the past decade due to explosion of IP-enabled devices and home Internet usage, many homes have become testbeds of multi-subnet and often multi-level subnetworks. While the simple case of a single Residential Gateway with NAT functionality is the most common, there is a desire to have separate guest subnets. And the future introduction of new low-power radio technologies will result in additional subnets since such technologies typically can not be bridged to Ethernet networks. Using IPv4 it is possible to get connectivity by connecting several RG's with NAT together to form a tree or a daisy-chain of NATs. That can more or less be performed in a plug and play fashion. It is also possible to manually configure routers with IP subnet numbers, routing protocols, etc, resulting in a home network which does not require any internal NATs. But such configuration requires a fair bit of expertise.

With the introduction of IPv6 in the home networks we would like to avoid assuming IPv6 NATs, yet we want to allow for cases that require separate subnets (for security reasons as in the case of the guest network, or for technology reasons as in the case of new radio networks). We want this without requiring networking experts to manually configure IP subnet numbers and routing protocols.

IPv6 has already taken steps to facilitate some aspects of this configuration through DHCP Prefix delegation [\[RFC3633\]](#) which is used to configure a single Residential Gateway with an IPv6 address prefix that can be used inside the home. However, that does not handle cases where there are multiple routers in the home. The homenet WG desires to solve this more general architecture, with a set of example topologies shown in [\[I-D.chown-homenet-arch\]](#).

In this draft we argue for separating out a subset of those topologies, and focusing on those first. We will call the subset "Basic Homenets". The criteria used for this is the set of topologies that can be implemented using consumer-grade IPv4 Residential Gateways without (significant) manual configuration. As we will see, those topologies end up being constrained to be a single tree rooted in the connection to the ISP. In such a topology we then apply hierarchical DHCPv6 Prefix Delegation in an automated way with sensible defaults. The approach is as robust against misconfiguration and loops as is the use of IPv4 NATs.

Adding the prefix allocation specified in this draft for IPv6 support will have no effect on IPv4; the current use of DHCP, NAT, or even routed IPv4 without NAT in the home routers will function as before. The approach provides stable IPv6 prefixes in the home network by relying on the ability for DHCPv6 servers to keep the assignments in stable storage.

This draft follows the architecture and terminology outlined in the homenet architecture [\[I-D.chown-homenet-arch\]](#).

2. Definition Of Terms

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 [\[RFC2119\]](#).

Some of the following terms are taken from [\[I-D.chown-homenet-arch\]](#):

Customer Edge Router (CER)

The IPv6 router which connects to the ISP network on its uplink interface. Such a router has one or more down-link interfaces which can be used by hosts and routers. This router can be co-located with the CPE.

Interior Router (IR)

The other IPv6 routers in the home network. These routers are not connected to the ISP directly.

Router:

In this document we use "router" to refer to either CERs or IRs. In fact, CER and IR is a topological role which we expect can be played by a router which implements this specification.

Host:

A host in a IPv6 network is a device which does not forward IPv6 packets (that are not addressed to itself). A host can be connected to one or more routers.

CPE:

Customer Premises Equipment aka Home Gateway attached to the DSL or cable modem. The CER and CPE could be co-located in the same device.

Link:

A communication facility or medium over which nodes can communicate at the link layer, i.e., the layer immediately below IPv6.

3. Conceptual model of a Basic Home Router

A key assumption we make is that a Basic Home Router has a designated uplink port. Today such home routers include IPv4 NAT functionality and as IPv6 support is added the goal is to not require IPv6 NAT functionality but instead rely on different prefixes being allocated to different links.

A Basic Home Router can have different configuration and number of downlink ports, whether physical ports (e.g., Ethernet), VLANs, or wireless (e.g., WiFi, 802.15.4). In the case of wireless interfaces it might make sense to think of them as potentially different ports. For example, a WiFi access point with a private and a guest ESSID might be thought of as two separate ports. A device is free to choose which

collections of ports it wants to handle as a single link from IP's perspective. For example, a device with 4 Ethernet downlink ports and a WiFi AP is free to handle that as e.g.:

A single link, by bridging the Ethernet ports and WiFi together.

One link for the 4 Ethernet ports (bridged together), and two links for WiFi - one for the private ESSID and one for the guest ESSID.

Dynamically create a separate link for each station that is authenticated using 802.1X and its WiFi counterparts.

The key point in the conceptual model is the assumption of a single uplink port on a separate IP link, and some set of links covering the set of downlink ports. On typical home router products the uplink port is colored and labeled differently, perhaps with "Internet" or "WAN". While there are IPv6 routers which do not have those limitations, if the device is also operating as a IPv4 home router with NAT, then most likely it would have the single uplink for the purposes of DHCPv4 and NAT.

4. Topology Assumptions


The existence of the single uplink interface naturally drives the topology towards a tree. At first sight one might think that the network can have destructive loops even with a single uplink port. For instance, some set of downlink interfaces on some of the routers could be bridged together using a commonly available Ethernet switch. The key question is whether that would work using IPv4 home routers with NAT. Many IPv4 home routers have a default configuration with 192.168.1.1/24 configured on their LAN interface. Plugging two downlink ports from two different routers into a single switch would cause an IP address conflict; both routers would claim the same above IP address. Such a conflict can be avoided if one of the routers is configured to have a different IP address on its LAN interface such as 10.0.0.1/24. In that case there would still be two uncoordinated DHCP servers on the same LAN. Thus one host might send a DHCP request to one router, and be assigned 192.168.1.5/24 a default router of 192.168.1.1 while some other host happens to use the other router's DHCP server and be assigned 10.0.0.27/24 with 10.0.0.1 as the default router. Thus this doesn't cause any looping problems for IPv4 and NAT, but it isn't useful as a topology since there is no coordinated IPv4 address assignment for the bridged LAN.

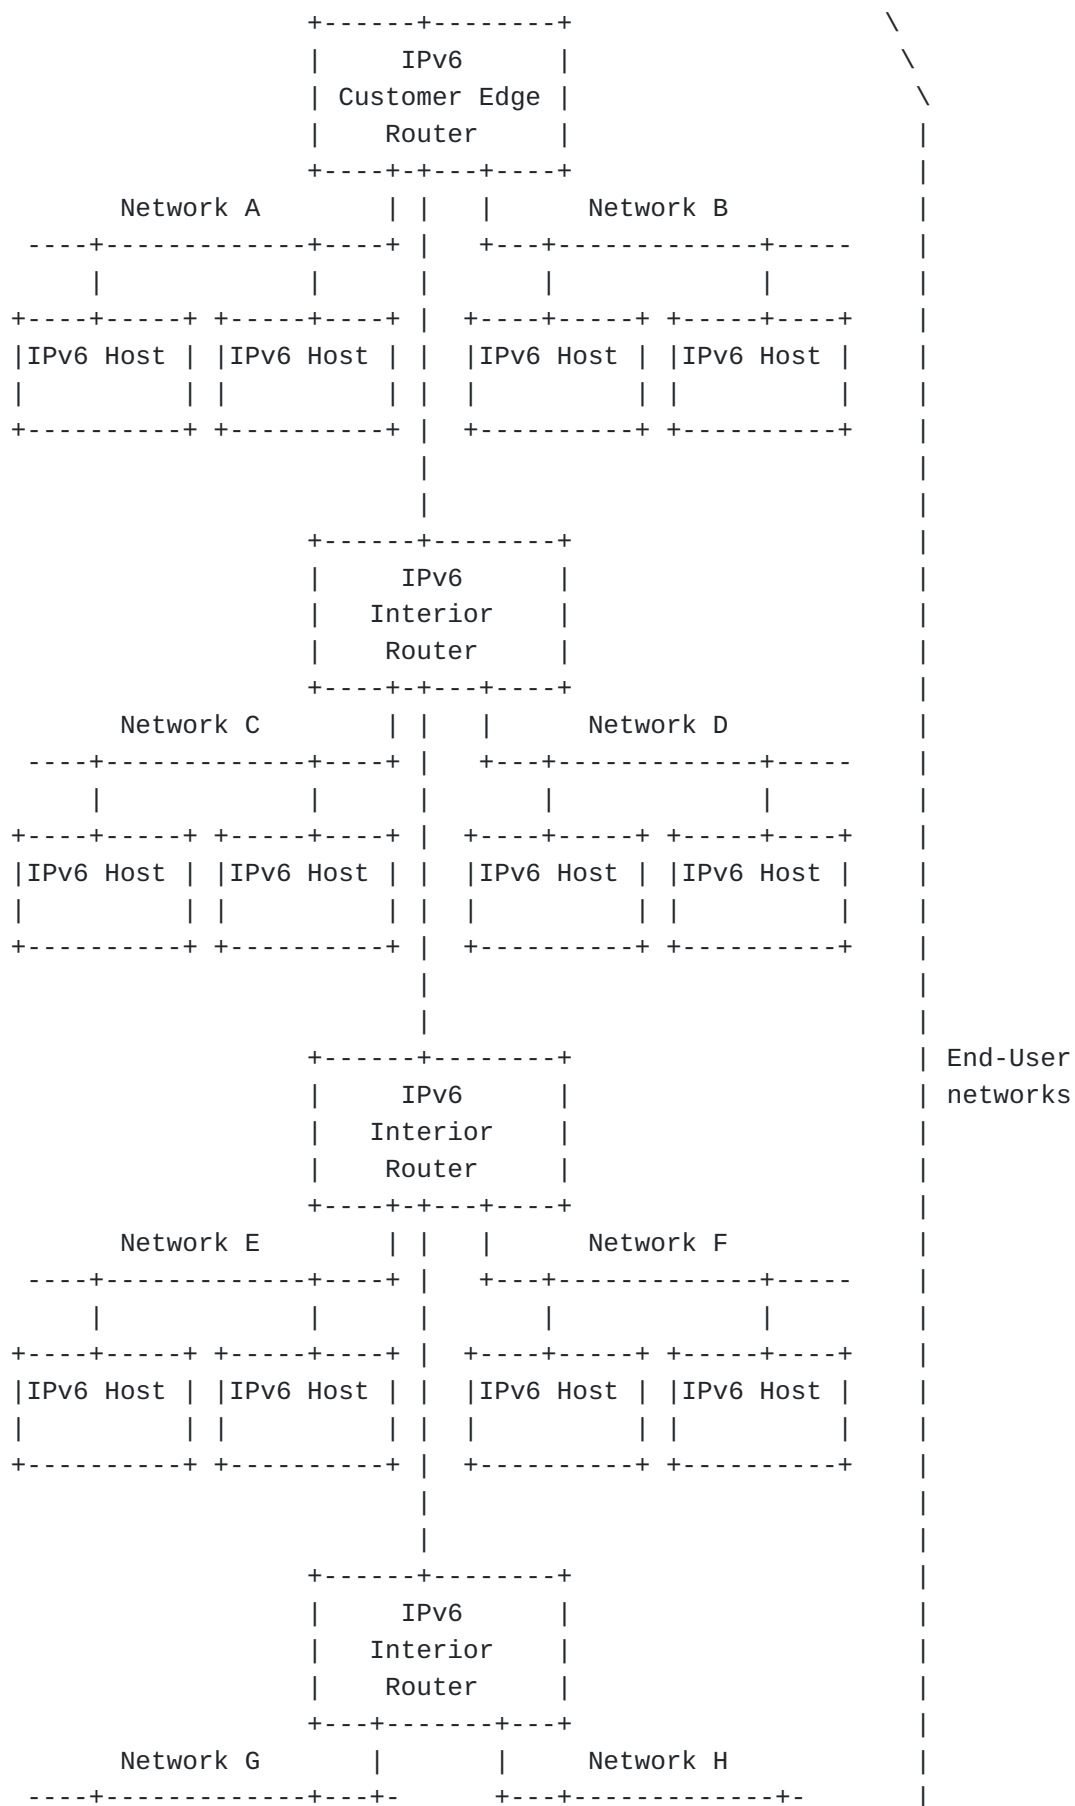
For IPv6 we want to support the same topologies that are useful in the IPv4 NAT case, and ensure that even for non-useful topologies such as the above bridged LAN case IPv6 wouldn't be any worse than the IPv4 NAT case.


5. Topology Examples

The following diagrams show the typical topology scenarios of home network for which the draft is based on. For simplicity the diagram limits the levels of subnets. Figure 1 shows a rather wide tree of routers, and as a result a large number of hosts can be connected using a shallow tree.


Figure 2 shows a daisy-chain of routers, which result in a deeper tree with more levels of routers. But both of those topologies are trees.

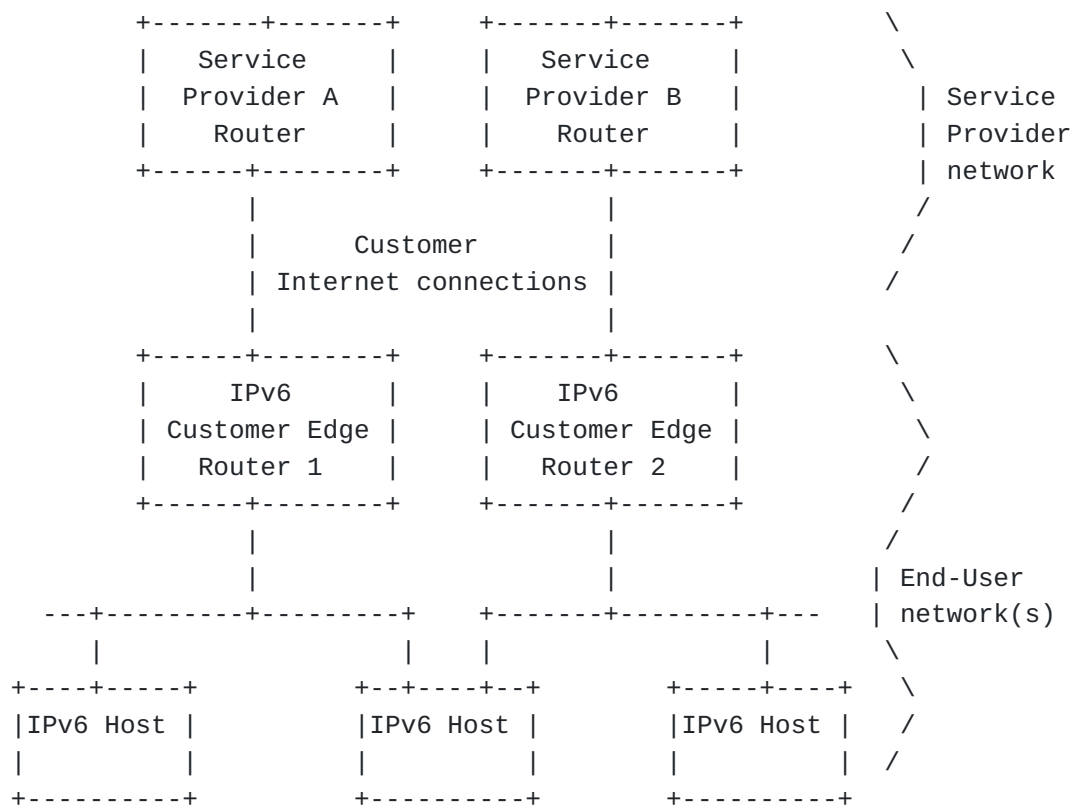






Note that none of the figures about have any multihoming. However, hosts might be multihomed as shown in Figure 3 and Figure 4.





6. Automatic hierarchical DHCPv6 Prefix Delegation

The basic idea is that each router operates a DHCP PD client on their uplink interface, and a DHCP PD server on each of their downlink interfaces. The router can then take the prefix it is delegated on its uplink interface, and carve that up into

- *Prefixes that are advertised in router advertisements on its downlink interfaces for Stateless Address autoconfiguration [\[RFC4862\]](#) of neighboring host and router interfaces.

- *Prefixes that are made available to its DHCP PD server, from which downlink neighboring routers can request allocations.

A router would typically know how many downlink interfaces it has (unless it creates them on the fly based on 802.1X, but that isn't a zero-configuration case). But in general a router does not know how many downlink neighboring routers it might have - whether the topology of routers will look like a wire tree or a narrow daisy-chain. However, we recommend a heuristic approach. If a router has e.g., four wired Ethernet ports and two radio interfaces, it would seem unlikely for it to have more than about six neighboring downlink routers. Based on this we recommend that a router of that size by default reserve seven sub-prefixes for PD allocation. That is the basis for automating the sub-delegations.

6.1. Example

We assume the ISP allocates a /56 prefix to the CER, and that all the routers use the above default of 7 sub-prefixes. Let the prefix be 2001:DB8:0:CD00::/56. The router adds "3" to the prefix length which results in 8 different /59 prefixes:

*2001:DB8:0:CD00::/59

*2001:DB8:0:CD20::/59

*2001:DB8:0:CD40::/59

*2001:DB8:0:CD60::/59

*2001:DB8:0:CD80::/59

*2001:DB8:0:CDA0::/59

*2001:DB8:0:CDC0::/59

The router can use the first /59 to create 32 different /64 prefixes for its downlinks, and has 7 different /59 prefixes it can allocate to downlink neighboring IRs.

When an IR that is directly attached to the CER invokes the DHCP PD client on its uplink interface it might be assigned 2001:DB8:0:CD60::/59. That router operates in exactly the same manner and adds "3" to the prefix length to create 8 different /62 prefixes:

*2001:DB8:0:CD60::/62

*2001:DB8:0:CD64::/62

*2001:DB8:0:CD68::/62

*2001:DB8:0:CD6C::/62

*2001:DB8:0:CD70::/62

*2001:DB8:0:CD74::/62

*2001:DB8:0:CD78::/62

*2001:DB8:0:CD7C::/62

The router can use the first /62 to create 4 different /64 prefixes for its downlink links and has 7 different /62 prefixes to assign to its child IRs should there be any.

Suppose there is a third layer of routers, so that an IR requests a prefix from the above IR. Then it might be assigned

2001:DB8:0:CD78::/62. It carves that into four /64 prefixes (it can't carve into smaller chunks than /64):

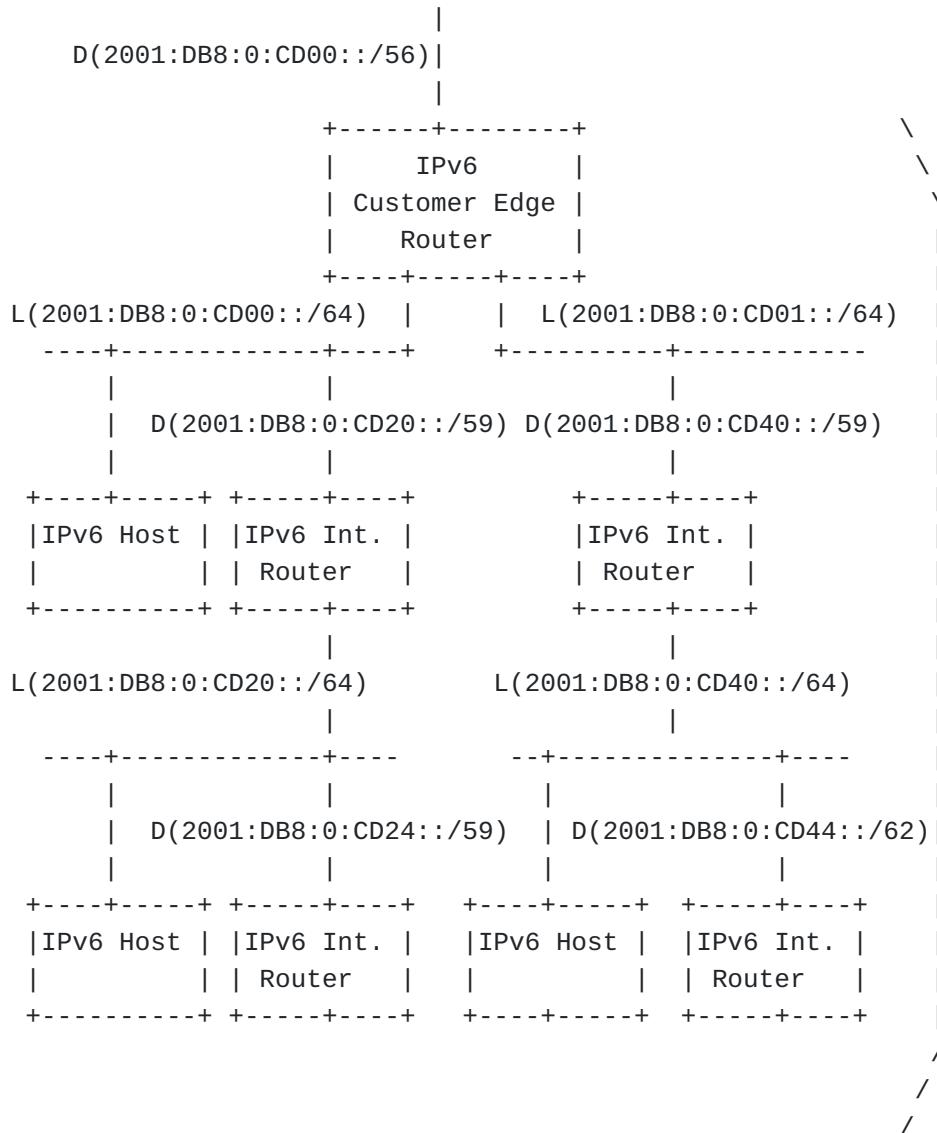
*2001:DB8:0:CD78::/64

*2001:DB8:0:CD79::/64

*2001:DB8:0:CD7A::/64

*2001:DB8:0:CD7B::/64

If the router has less than four downlink interfaces, then it would keep the leftover /64 prefixes in reserve for its DHCP PD client. One such example network is depicted in [Figure 5](#). In this figure L(Prefix) is used to denote that the prefix is being advertised in an RA as an on-link prefix and D(Prefix) is used to denote that the prefix is being delegated from the Delegating Router (DR) to the Requesting Router (RR) in the downward direction.



6.2. Configurability

A router SHOULD provide a configuration interface where that allows both adjusting the added prefix length ("3" in the above example), and also allows manual assignment of prefixes to DHCP PD clients (in the same manner than many IPv4 home routers allow pre-assignment of IPv4 addresses today).

6.3. Routing Implications

If a router (CER or IR) has been assigned a prefix on its uplink interface (e.g., 2001:DB8:0:CD60::/59) then any destination address which is outside of that prefix should be routed to its uplink. The router maintains a default route to its uplink for this purpose using Neighbor Discovery [\[RFC4861\]](#) on its uplink interface. Destination addresses that fall in its delegated prefix will either be routed to

downlink interfaces (if they are assigned as a subnet prefix on an interface, or delegated to a downlink router) or dropped (for unassigned prefixes).

Thus there is no need to run a routing protocol to handle the Basic Homenet topologies.

6.4. Neighbor Discovery Implications

A router (CER or IR) will perform Neighbor Discovery as a host on its uplink interface, thus it will send Router Solicitations and use received Router Advertisements to track its default uplink router. Note that in some suboptimal topologies there might be multiple uplink routers (if some bridge has been inserted) thus a router should handle multiple default routers on its uplink interface.

A CER and IR needs to perform Neighbor Discovery as a router on its downlink interfaces. Thus it will send Router Advertisements periodically and respond to Router Solicitations.

If the prefix delegated by the uplink router changes, this means that the router needs to change both the /64 prefixes it is advertising in RAs and also get the downlink routers to which it has delegated sub-prefixes to get reconfigured. For planned changes that can be handled by ensuring that the lifetime, T0 and T1 [\[RFC3315\]](#) are carried from the PD client in a router to its PD server. But for unplanned changes, for instance when someone manually changes the prefixes on a CER or IR, one would like a way to have that be propagated to downlink PD clients. In theory DHCP Reconfigure messages [\[RFC3315\]](#) could be used, but they require some security configuration. Thus we suggest using prefix changes in received Router Advertisements (on the uplink interface) as a hint that the router's PD client should attempt to renew its DHCP lease and as a result of that discover changes in the delegated prefixes.

6.5. Ensuring stable prefixes

It is highly desirable that the home network maintain the same prefix allocation even if parts or all of the network are powered off and back on, or otherwise fail and come back. That can be handled if the DHCP PD servers in each router (and also in the ISP) maintain the delegated prefixes in stable storage (to guard against the router itself failing) and also retain information about the last holder of a lease even after the lease has expired. That way, as long as the number of downlink routers is less than the size of the pool of prefixes available for delegation (7 in the example above), even if a downlink router is powered off for a long time, when it comes back it will receive the same prefix.

7. Addressing

This document suggests delegating Unique link-local Addresses [\[RFC4193\]](#) and IPv6 Global addresses. The ULA can be only generated or manually configured at the Customer Edge Router (CER) and then delegated down the link the same way IPv6 Global prefix is delegated. A CER SHOULD be capable of delegating a ULA prefix and a IPv6 Global prefix obtained from the ISP.

When the home network is initialized the hosts and routers on the network will start off with only having link local addresses. They will use the link local addresses to bootstrap address acquisition using DHCP PD for the other scopes of addresses.

Depending on whether the CER has working upstream connectivity or not, it is possible that differently scoped addresses/prefixes could be assigned to the home network.

When the home network permanently has no upstream connectivity towards the ISP, it is RECOMMENDED that the CER create an Unique Local Prefix as specified in [\[RFC4193\]](#). We recommend using a /48 ULA prefix as specified in that RFC. Note that it might be difficult to automatically determine whether 1) the home network is permanently disconnected from the ISP and 2) whether a particular router is the CER. Thus it is RECOMMENDED that the generation of the ULA prefix is triggered by manual configuration in the case of a disconnected network.

Even for a connected home network it is RECOMMENDED to trigger the generation of the ULA manually on the CER. The CER will then automatically delegate parts of that prefix concurrently with sub-delegating the global prefix it received from the ISP. Potentially one could do this automatically by leveraging the bootstrapping behavior to determine whether a router is a CER or an IR, with the assumption that the ISP would never delegate a ULA prefix to its customer. In that case, if a router receives a prefix delegation that contains a global prefix but no ULA, then it can assume it is the CER and (if it hasn't already) generate a ULA, store that ULA in its persistent storage, and sub-delegate the global prefix and the ULA in parallel to any downlink PD clients.

In many cases the ISP will select the prefix length it will delegate. Thus it is RECOMMENDED that a router (CER or IR) by default set the prefix-length field [\[RFC3633\]](#) in field of a IA_PD Prefix option (OPTION_IAPREFIX) to zero. A router that has the role of a CER may be manually configured to request a particular prefix-length, but the default allocation scheme in this document assumes that IRs do not set the prefix-length.

8. Basic Operations

It is assumed that CER requests one or more IPv6 prefixes from the ISP Prefix delegating router for IPv6 prefixes for a specified prefix length if the service agreement allows the CER to support multi-level subnets without NAT66 [\[RFC6296\]](#). Currently DHCP-PD [\[RFC3633\]](#) allows a

requesting router to request a specific prefix through the IA prefix option. This document discusses a simple mechanism for assigning and delegating prefixes through the hierarchy in the home network (section 6 and section 6.1). However, the implementation SHOULD also support ability of the requesting router to request a prefix of a specific length by filling-in the Prefix Length field of the IA prefix option while the IPv6-prefix field being the unspecified address.

In addition, this specification requires all IRs to be able to store and delegate prefixes on its downlink interfaces only. The prefix should be stored during reboots and power failure.

The 'Topology' section diagrams are the typical home networking scenarios where the above prefix delegation mechanism is believed to work well.

8.1. DHCPv6 Prefix Delegation

The DHCPv6 prefix delegation at CER follows DHCP-PD [\[RFC3633\]](#) in order to receive the Prefix(es) from the ISP Prefix delegator and it can act as a local prefix delegator for the home network.

DHCP-PD [\[RFC3633\]](#) suggests that in a typical scenario, /48 prefix is assigned to the requesting router. The operational procedures by an ISP might limit this default to a /56. The CER may be configured with a specific prefix length to request from the delegating router.

Thus CER will include IA_PD option(s) as specified in [\[RFC3633\]](#). In the IA_PD Prefix option, the IPv6-Prefix field is set to zero if the requesting router does not have any prior knowledge about its IPv6 Prefix. The prefix length MAY be set between /48 and /64 inclusive when the requesting router likes to specify a prefix len. By default the delegating router (CER and delegating IR) adds bits to the prefix before delegating downwards. The automatic bits calculation and prefix formation is described in section 6 and 7 above.

The IRs also operate as a DHCP PD requesting router on their uplink interface, but unlike the CER there is no need to specify a prefix length that they will request.

Each CER and IR SHOULD act as a default routers on its downlink interfaces by selecting a /64 prefix for each downlink interface and advertising it in Router Advertisements downlink interface. The IRs can use Stateless Address Autoconfiguration to configure the IPv6 addresses on the uplink interface as specified in [\[RFC4862\]](#). If a CER or IR is only delegated a /64 prefix from its delegating router then it can advertise in Router Advertisements for one of its downlink interfaces, but it can not run a DHCP PD server.

There is no need for a dynamic routing protocol since each IR will have a default route towards its delegating router on its uplink interface.

9. Host Behavior

Whether a host uses Stateless Address autoconfiguration or DHCPv6, it does not require any change due to the solutions proposed in this draft.

10. Router Behavior

All home routers (CER and IRs) behave the same as specified in section 6 and section 8, with the exception that a CER might be configured to generate a ULA prefix and delegate sub-prefixes of that ULA.


11. Bootstrapping

It is desirable that the prefix delegation flow in an orderly manner from the ISP to the CER and further down to the IRs, and down to hosts. We do not want any prefix flapping (some IR guessing a prefix to advertise before it has received anything from its uplink), hence it is RECOMMENDED that a router wait until its PD client on the uplink interface has received a prefix allocation, and at that point in time in enable its PD server on its downlink interface and also enable the sending of Router Advertisements on its downlink interfaces. The only exception to this is a CER which has been configured to generate and advertise a ULA prefix even when the ISP connection is down; such a CER would sub-delegate and advertise the ULA prefix in parallel with requesting a prefix delegation from the ISP.

The above behavior implies that when the whole home network is brought up (e.g., after a power failure) it might take a while until a host will start receiving Router Advertisement messages. But once those RAs arrive they will contain at least a ULA prefix and in many cases both a ULA and a global prefix.

12. What if there are loops?

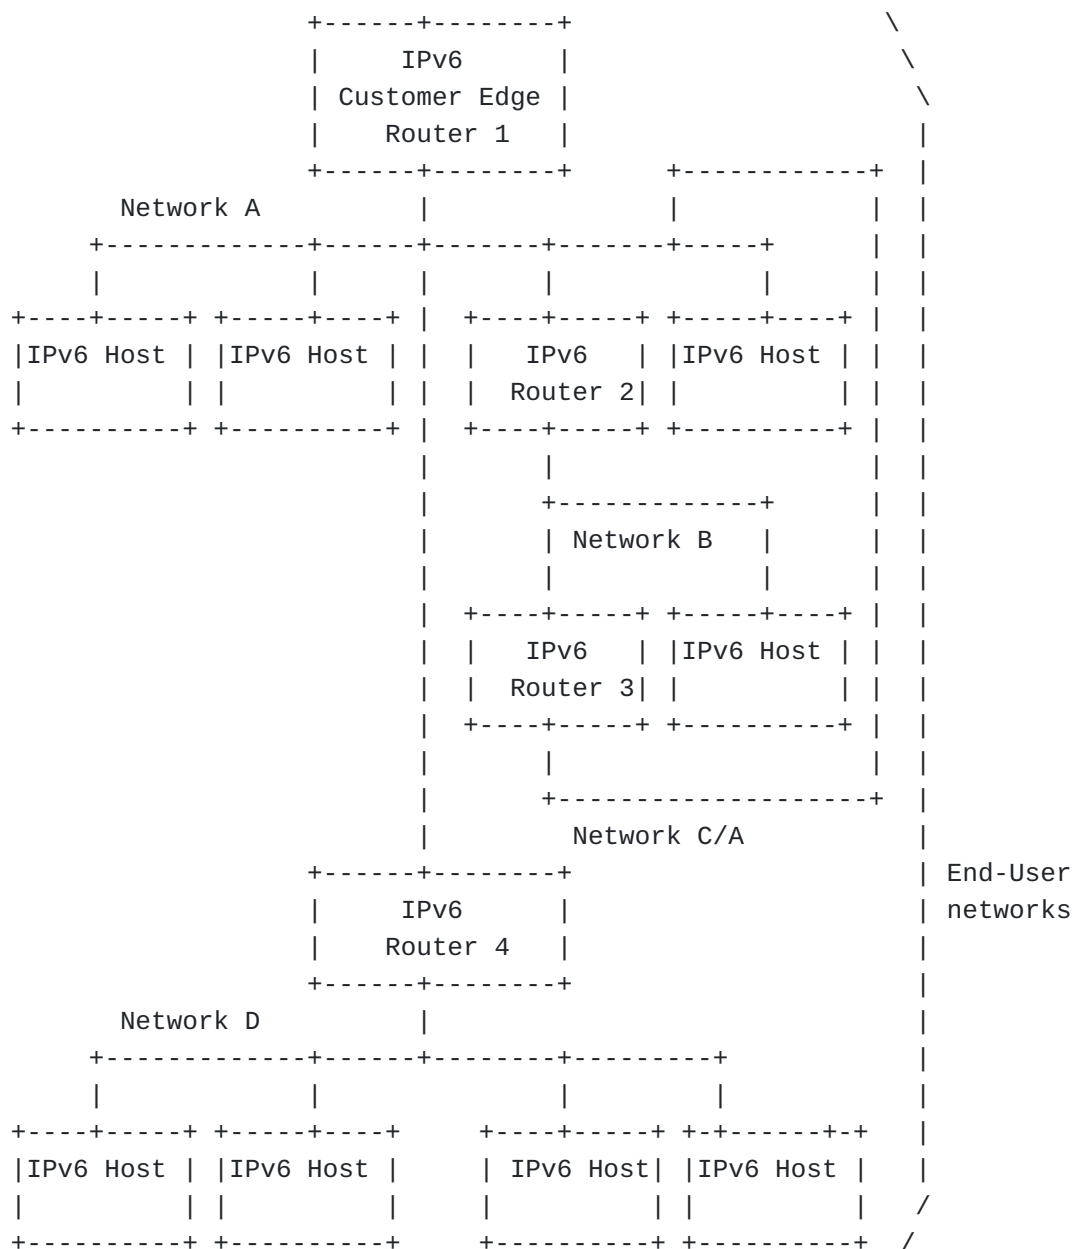
Even if the configuration falls outside of the topology constraints we have specified, we still want the home network to be no worse than the same topology with IPv4 NAT routers. One such topology is when there is a L2 bridge which connects some downlink interfaces on two or more routers, and there are some hosts attached to that bridge and/or there are routers that attach their uplink interface to that bridge. See Figure 6.



In the figure the downlink interfaces of Router X and Router Y have been bridged together. Router X and Y will have received their own prefix delegation from the CER. They will each have pick some /64 prefix from that to advertise in Router Advertisement on Network C. Thus one effect of the bridge is that the hosts that attach to network C will, following [RFC4862](#), configure multiple addresses on their interface. The same might happen for the routers that have an uplink interface to Network C; they might configure multiple addresses on that interface.

A second effect of the bridge is that the PD clients in router Z and W now has two potential DHCP PD servers. Presumably this means that they pick one of them that responds to their DHCP request. Thus router Z and W might end up picking a different uplink router for their PD allocation. That isn't any different than in the DHCPv4 and NAT case. What is different with IPv6 is that the default router assignment is being done using Router Advertisements, thus both router Z and W will end up with two default routers; X and Y. This is independent of which uplink router assigned them a sub-prefix. As long as the home routers do not perform ingress filtering based on the allocated prefixes this

will work, but we might want to consider somehow tying the PD allocation to the choice of default router?



In Figure 7 we see a loop which is caused by having the downlink interface of router 3 be an attached as an uplink of router 2. This means that Router 2 and Router 4 see two different uplink router; router 1 and router 3.

In the IPv4 case, just as above, the default configuration of R1 and R3 might cause IP address conflicts since both might have 192.168.1.1/24 as defaults on their downlink ports in which case the network doesn't work at all. Just as above that can be manually corrected by e.g., configuring R3 to have 10.0.0.1/24 on its downlink interface. In that when case R2 and R4 uses DHCPv4 they might pick the DHCP response from either R1 or R3 and configure themselves to either have a 192.168

address and 192.168.1.1 as their default router, or a 10.0 address and 10.0.0.1 as a default router. If R2 picks picks the latter (R3), then outbound traffic will loop, since it will be sent to R3 which will NAT and send to R2 which will NAT and send to R3. If R2 picks R1 and R4 picks R3 then traffic from R4 to the Internet will merely go through two extra NATs. In general we can't predict which DHCP server R2 and R4 will pick, hence sometimes the network will work and sometimes not. With the proposed prefix delegation scheme and associated bootstrapping for IPv6 things can work a little bit better, since we recommend that a router not enable its PD client on the downlinks until its PD server on the uplink has been delegated a prefix. Thus R1 will be delegated a prefix from the ISP, and then assign a /64 to Network A and enable the PD server. Then R2 and R4 can receive delegations only from R1, since R3 has not yet enabled its PD server. Later when R3 has received a delegation from R2 it will enable the PD server. Note that it isn't much better than for IPv4 since if R4 is powered off and back on or just boots very slowly after a complete power failure it might come up after the R1 -> R2 -> R3 delegation chain has already occurred, in which case R4 might pick R3 as its delegating router. And if R2 crashes and comes back, it might also pick R3 since R2's delegation to R3 will have a non-zero lifetime.

[DISCUSSION: It is possible to improve on the above by having the PD client use the delegated prefix-length to determine which DHCP lease to accept; preferring longer prefixes will make it choose a delegating router which is closer to the ISP. In the above example R1 might delegate /59 prefixes while R3 can delegate only /64 prefixes. But it isn't clear that such added complexity is worth-while. Note that for that to help we'd also need to pick the delegating router as the default router, instead of building a default router list with all the routers which send RAs.

13. Security Considerations

No new threats against Neighbor Discovery beyond what is already documented for IPv6 ND [\[RFC3756\]](#) due to IPv6 Address autoconfiguration and Neighbor Discovery at the last hop of Prefix distribution. The recommendations in this document does not prevent using Secure Neighbor Discovery [\[RFC3971\]](#).

The security threats for this solution is believed to be no worse than DHCPv6 Prefix delegation[\[RFC3633\]](#). See Section 15 of RFC 3633 for further information.

A malicious host inside the end user network can perform a prefix exhaustion attack on the CER or the IRs. It works as follows; the malicious router keeps on requesting prefixes from the delegating router (DR), which could be the CER or another IR, until all the prefixes have been delegated. At this point a legitimate router that attaches to the delegating router will fail to get a prefix delegated as the DR has no more prefixes available to delegate. This means that the subset of the network behind this newly attaching router will not

get any connectivity. This can be avoided by using some form of authorization on the delegating router but the specification of such a mechanism is outside the scope of this document. It might make sense to offer configuration capability so that prefix delegation can be disabled on certain links such as a guest network.

14. IANA Considerations

This document does not require any IANA actions.

15. Acknowledgements

The authors like to thank David Allan, Joel Halpern, Acee Lindem and Jari Arkko for comments on the proposal of the draft. Alan Kavangh suggested using the default prefix len(/56) as per Broadband Forum's recommendation. We thank Tim Chown for the ASCII-art drawings that we reused and in some cases expanded on it this draft.

16. References

16.1. Normative References

[RFC2119]	Bradner, S. , " Key words for use in RFCs to Indicate Requirement Levels ", BCP 14, RFC 2119, March 1997.
[RFC3633]	Troan, O. and R. Droms, " IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6 ", RFC 3633, December 2003.
[RFC4861]	Narten, T., Nordmark, E., Simpson, W. and H. Soliman, " Neighbor Discovery for IP version 6 (IPv6) ", RFC 4861, September 2007.
[RFC4862]	Thomson, S., Narten, T. and T. Jinmei, " IPv6 Stateless Address Autoconfiguration ", RFC 4862, September 2007.
[I-D.chown-homenet-arch]	Arkko, J, Chown, T, Weil, J and O Troan, " Home Networking Architecture for IPv6 ", Internet-Draft draft-chown-homenet-arch-01, October 2011.

16.2. Informative References

[RFC2460]	Deering, S.E. and R.M. Hinden , " Internet Protocol, Version 6 (IPv6) Specification ", RFC 2460, December 1998.
[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.
[RFC3756]	Nikander, P., Kempf, J. and E. Nordmark, " IPv6 Neighbor Discovery (ND) Trust Models and Threats ", RFC 3756, May 2004.
[RFC3769]	

	Miyakawa, S. and R. Droms, " Requirements for IPv6 Prefix Delegation ", RFC 3769, June 2004.
[RFC3971]	Arkko, J., Kempf, J., Zill, B. and P. Nikander, " SEcure Neighbor Discovery (SEND) ", RFC 3971, March 2005.
[RFC4193]	Hinden, R. and B. Haberman, " Unique Local IPv6 Unicast Addresses ", RFC 4193, October 2005.
[RFC6296]	Wasserman, M. and F. Baker, " IPv6-to-IPv6 Network Prefix Translation ", RFC 6296, June 2011.

Authors' Addresses

Erik Nordmark Nordmark Cisco Systems San Jose, CA, USA EMail:
nordmark@cisco.com

Samita Chakrabarti Chakrabarti Ericsson San Jose, CA, USA EMail:
samita.chakrabarti@ericsson.com

Suresh Krishnan Krishnan Ericsson Montreal, , CANADA EMail:
suresh.krishnan@ericsson.com

Wassim Haddad Haddad Ericsson San Jose, CA, USA EMail:
wassim.haddad@ericsson.com