Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 2, 2012

J. Arkko
Ericsson
A. Brandt
Sigma Designs
T. Chown
University of Southampton
J. Weil
Time Warner Cable
O. Troan
Cisco Systems, Inc.
January 30, 2012

Home Networking Architecture for IPv6
draft-ietf-homenet-arch-01

Abstract

This text describes evolving networking technology within small "residential home" networks. The goal of this memo is to define the architecture for IPv6-based home networking and the associated principles, considerations and requirements. The text highlights the impact of IPv6 on home networking, illustrates topology scenarios, and shows how standard IPv6 mechanisms and addressing can be employed in home networking. The architecture describes the need for specific protocol extensions for certain additional functionality. It is assumed that the IPv6 home network is not actively managed, and runs as an IPv6-only or dual-stack network. There are no recommendations in this text for the IPv4 part of the network.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 2, 2012.

Table of Contents

1. Introduction

This document focuses on evolving networking technology within small "residential home" networks and the associated challenges. For example, a trend in home networking is the proliferation of networking technology in an increasingly broad range of devices and media. This evolution in scale and diversity sets requirements on IETF protocols. Some of these requirements relate to the need for multiple subnets, for example for private and guest networks, the introduction of IPv6, and the introduction of specialized networks for home automation and sensors.

While some advanced home networks exist, most operate based on IPv4, employ solutions that we would like to avoid such as (cascaded) network address translation (NAT), or require expert assistance to set up. The assumption of this document is that the homenet is "not actively managed". The architectural constructs in this document are focused on the problems to be solved when introducing IPv6 with an eye towards a better result than what we have today with IPv4, as well as a better result than if the IETF had not given this specific guidance.

This architecture document aims to provide the basis and guiding principles for how standard IPv6 mechanisms and addressing [RFC2460] [RFC4291] can be employed in home networking, while coexisting with existing IPv4 mechanisms. In emerging dual-stack home networks it is vital that introducing IPv6 does not adversely affect IPv4 operation. Future deployments, or specific subnets within an otherwise dual-stack home network, may be IPv6-only.

[RFC6204] defines basic requirements for customer edge routers (CERs). The scope of this text is the homenet, and thus the internal facing interface described in RFC 6204 as well as other components within the home network. While the network may be dual-stack or IPv6-only, the definition of specific transition tools on the CER are out of scope of this text, as is any advice regarding architecture of the IPv4 part of the network. We assume that IPv4 network architecture in home networks is what it is, and can not be affected by new recommendations.

Discussion in the homenet WG has led to a suggestion that there should be a baseline homenet "version 1" architecture, based on protocols and implementations that are as far as possible proven and robust. A future architecture may incorporate more advanced elements. Feedback is sought on what if anything do we want to say about potential homenet versions here.

1.1. Terminology and Abbreviations

In this section we define terminology and abbreviations used throughout the text.

- CER: Customer Edge Router. The border router at the edge of the homenet.
- LLN: Low-power and lossy network.
- NAT: Network Address Translation. Typically referring to Network Address and Port Translation (NAPT).
- NPTv6: Network Prefix Translation for IPv6 [RFC6296].
- PCP: Port Control Protocol [I-D.ietf-pcp-base].
- ULA: Unique Local Addresses [RFC4193].
- uPnP: Universal Plug and Play.
- VM: Virtual machine.

2. Effects of IPv6 on Home Networking

Service providers are deploying IPv6, content is becoming available on IPv6, and support for IPv6 is increasingly available in devices and software used in the home. While IPv6 resembles IPv4 in many ways, it changes address allocation principles, makes multi-addressing the norm, and allows direct IP addressability and routing to devices in the home from the Internet. This section presents an overview of some of the key areas impacted by the introduction of IPv6 into the home network that are both promising and problematic.

2.1. Multiple subnets and routers

Simple layer 3 topologies involving as few subnets as possible are preferred in home networks for a variety of reasons including simpler management and service discovery. However, the incorporation of dedicated (routed) subnets remains necessary for a variety of reasons.

For instance, a common feature in modern home routers is the ability to support both guest and private network subnets. Also, link layer networking technology is poised to become more heterogeneous, as networks begin to employ both traditional Ethernet technology and link layers designed for low-power and lossy networks (LLNs) such as

those used for certain types of sensor devices. Similar needs for subnetting may occur in other cases, such as separating building control or corporate extensions from the Internet access network. Also, different subnets may be associated with parts of the homenet that have different routing and security policies.

Documents that provide some more specific background and depth on this topic include: [I-D.herbst-v6ops-cpeenhancements], [I-D.baker-fun-multi-router], and [I-D.baker-fun-routing-class].

In addition to routing, rather than NATing, between subnets, there are issues of when and how to extend mechanisms such as service discovery which currently rely on link-local addressing to limit scope.

The presence of a multiple subnet, multi-router network implies that there is some kind of automatic routing mechanism in place. In advanced configurations similar to those used in multihomed corporate networks, there may also be a need to discover border router(s) by an appropriate mechanism.

2.2. Multi-Addressing of devices

In an IPv6 network, devices may acquire multiple addresses, typically at least a link-local address and a globally unique address. Thus it should be considered the norm for devices on IPv6 home networks to be multi-addressed, and to also have an IPv4 address where the network is dual-stack. Default address selection mechanisms [I-D.ietf-6man-rfc3484-revise] allow a node to select appropriate src/dst address pairs for communications, though such selection may face problems in the event of multihoming, where nodes will be configured with one address from each upstream ISP prefix, and the presence of upstream ingress filtering thus requires multi-addressed nodes to select the right source address to be used for the corresponding uplink.

2.3. Unique Local Addresses (ULAs)

[RFC4193] defines Unique Local Addresses (ULAs) for IPv6 that may be used to address devices within the scope of a single site. Support for ULAs for IPv6 CERs is described in [RFC6204]. A home network running IPv6 may deploy ULAs for communication between devices within the network. ULAs have the potential to be used for stable addressing in a home network where the externally allocated global prefix changes over time (either due to renumbering within the subscriber's ISP or a change of ISP) or where external connectivity is temporarily unavailable. However, it is undesirable to aggressively deprecate global prefixes for temporary loss of

connectivity, so for this to matter there would have to be a connection breakage longer than the lease period, and even then, deprecating prefixes when there is no connectivity may not be advisable. However, while setting a network up there may be a period with no connectivity.

Another possible reason for using ULAs would be to provide an indication to applications that the traffic is local. This could then be used with security settings to designate where a particular application is allowed to connect to.

ULA addresses will allow constrained LLN devices to create permanent relations between IPv6 addresses, e.g. from a wall controller to a lamp. Symbolic host names would require additional non-volatile memory. Updating global prefixes in sleeping LLN devices might also be problematic.

Address selection mechanisms should ensure a ULA source address is used to communicate with ULA destination addresses. The use of ULAs does not imply use of host-based IPv6 NAT, or NPTv6 prefix-based NAT [RFC6296], rather that external communications should use a node's global IPv6 source address.

2.4. Security, Borders, and the elimination of NAT

Current IPv4 home networks typically receive a single global IPv4 address from their ISP and use NAT with private [RFC1918] addresses for devices within the network. An IPv6 home network removes the need to use NAT given the ISP offers a sufficiently large IPv6 prefix to the homenet, allowing every device on every link to be assigned a globally unique IPv6 address.

The end-to-end communication that is potentially enabled with IPv6 is both an incredible opportunity for innovation and simpler network operation, but it is also a concern as it exposes nodes in the internal networks to receipt of otherwise unwanted traffic from the Internet.

In IPv4 NAT networks, the NAT provides an implicit firewall function. [RFC4864] suggests that IPv6 networks with global addresses utilise "Simple Security" in border firewalls to restrict incoming connections through a default deny policy. Applications or hosts wanting to accept inbound connections then need to signal that desire through a protocol such as uPNP or PCP [I-D.ietf-pcp-base]. In networks with multiple CERs, PCP will need to handle the cases of flows that may use one or both exit routers.

Such an approach would reduce the efficacy of end-to-end connectivity

that IPv6 has the potential to restore, since the need for IPv4 NAT traversal is replaced by a need to use a signalling protocol to request a firewall hole be opened. [RFC6092] provides recommendations for an IPv6 firewall that applies "limitations on end-to-end transparency where security considerations are deemed important to promote local and Internet security." The firewall operation is "simple" in that there is an assumption that traffic which is to be blocked by default is defined in the RFC and not expected to be updated by the user or otherwise. The RFC does however state that CERs should have an option to be put into a "transparent mode" of operation.

It is important to distinguish between addressability and reachability; i.e. while IPv6 offers global addressability through use of globally unique addresses in the home, whether they are globally reachable or not would depend on firewall or filtering configuration, and not the presence or use of NAT.

Advanced Security for IPv6 CPEs [I-D.vyncke-advanced-ipv6-security] takes the approach that in order to provide the greatest end-to-end transparency as well as security, security policies must be updated by a trusted party which can provide intrusion signatures and other "active" information on security threats. This is much like a virus-scanning tool which must receive updates in order to detect and/or neutralize the latest attacks as they arrive. As the name implies "advanced" security requires significantly more resources and infrastructure (including a source for attack signatures) in comparison to "simple" security.

In addition to establishing the security mechanisms themselves, it is important to know where to enable them. If there is some indication as to which router is connected to the "outside" of the home network, this is feasible. Otherwise, it can be difficult to know which security policies to apply where. Further, security policies may be different for various address ranges if ULA addressing is setup to only operate within the homenet itself and not be routed to the Internet at large. Finally, such policies must be able to be applied by typical home users, e.g. to give a visitor in a "guest" network access to media services in the home.

It may be useful to classify the border of the home network as a unique logical interface separating the home network from service provider network/s. This border interface may be a single physical interface to a single service provider, multiple layer 2 sub-interfaces to a single service provider, or multiple connections to a single or multiple providers. This border is useful for describing edge operations and interface requirements across multiple functional areas including security, routing, service discovery, and router

discovery.

2.5. Naming, and manual configuration of IP addresses

In IPv4, a single subnet NATed home network environment is currently the norm. As a result, it is for example common practice for users to be able to connect to a router for configuration via a literal address such as 192.168.1.1 or some other commonly used RFC 1918 address. In IPv6, while ULAs exist and could potentially be used to address internally-reachable services, little deployment experience exists to date. Given a true ULA prefix is effectively a random 48-bit prefix, it is not reasonable to expect users to manually enter such address literals for configuration or other purposes. As such, even for the simplest of functions, naming and the associated discovery of services is imperative for easy administration of the homenet.

In a multi-subnet homenet, naming and service discovery should be expected to operate across the scope of the entire home network, and thus be able to cross subnet boundaries. It should be noted that in IPv4, such services do not generally function across home router NAT boundaries, so this is one area where there is scope for an improvement in IPv6.

3. Architecture

An architecture outlines how to construct home networks involving multiple routers and subnets. In this section, we present a set of typical home network topology models/scenarios, followed by a list of topics that may influence the architecture discussions, and a set of architectural principles that govern how the various nodes should work together. Finally, some guidelines are given for realizing the architecture with the IPv6 addressing, prefix delegation, global and ULA addresses, source address selection rules and other existing components of the IPv6 architecture. The architecture also drives what protocol extensions are necessary, as will be discussed in Section 3.6.

3.1. Network Models

In this section we list six network models.

A) Single ISP, Single CER, Single subnet

B) Single ISP, Single CER, Multiple subnets

C) Single ISP, Single CER, Multiple internal routers

D) Two ISPs, Two CERs, Shared subnets with multiple internal routers

E) Two ISPs, One CER, Isolated subnets with multiple internal routers

F) Two ISPs, One CER, Shared subnets with multiple internal routers

The models are presented to frame the discussion as to which models are in scope for the homenet architecture, and which multi-homing requirements should be met in the architecture.

3.1.1. A: Single ISP, Single CER, Single subnet

Figure 1 shows the simplest possible home network topology, involving just one router, a local area network, and a set of hosts. Setting up such networks is in principle well understood today [RFC6204].

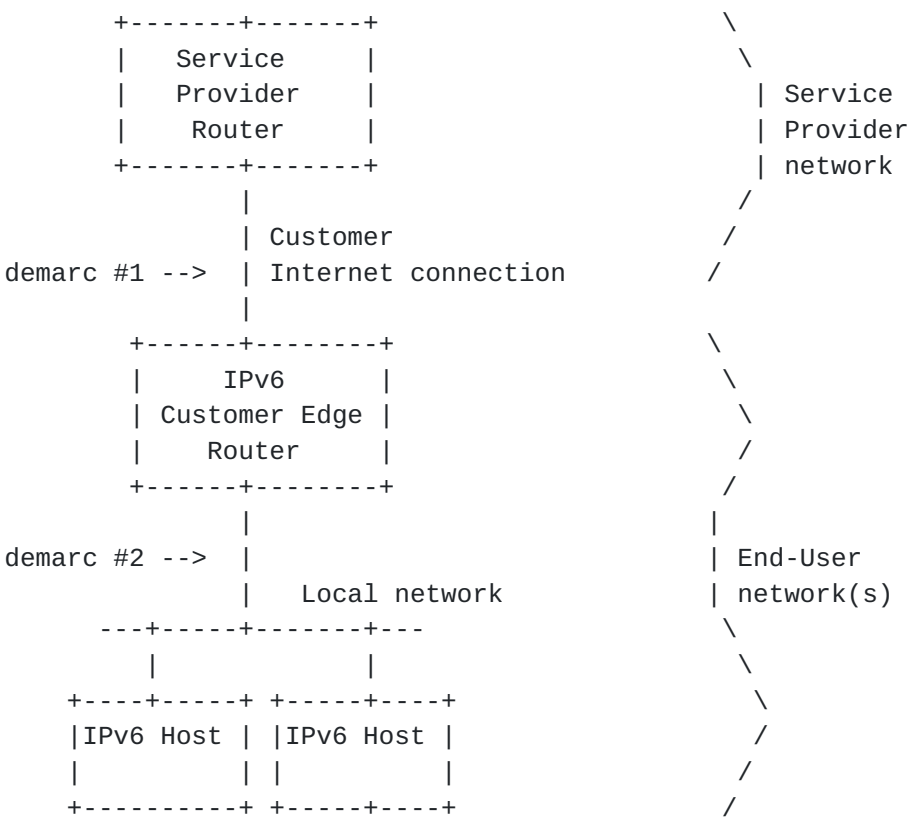

Figure 1

Two possible demarcation points are illustrated in Figure 1, which indicate which party is responsible for configuration or autoconfiguration. Demarcation #1 makes the Customer Edge Router the responsibility of the customer. This is only practical if the Customer Edge Router can function with factory defaults installed. The Customer Edge Router may be pre-configured by the ISP, or by the home user by some suitably simple method. Demarcation #2 makes the Customer Edge Router the responsibility of the provider. Both models of operation must be supported in the homenet architecture, including the scenarios below with multiple ISPs and demarcation points.

3.1.2. B: Single ISP, Single CER, Multiple subnets

Figure 2 shows another network that now introduces multiple local area networks. These may be needed for reasons relating to different link layer technologies in use or for policy reasons. A common arrangement is to have different link types supported on the same router, bridged together. This example however presents two subnets. This could be classic Ethernet in the one subnet and a LLN link layer technology in the other subnet.

This topology is also relatively well understood today [RFC6204],

though it certainly presents additional demands with regards to suitable firewall policies and limits the operation of certain applications and discovery mechanisms (which may typically today only succeed within a single subnet).

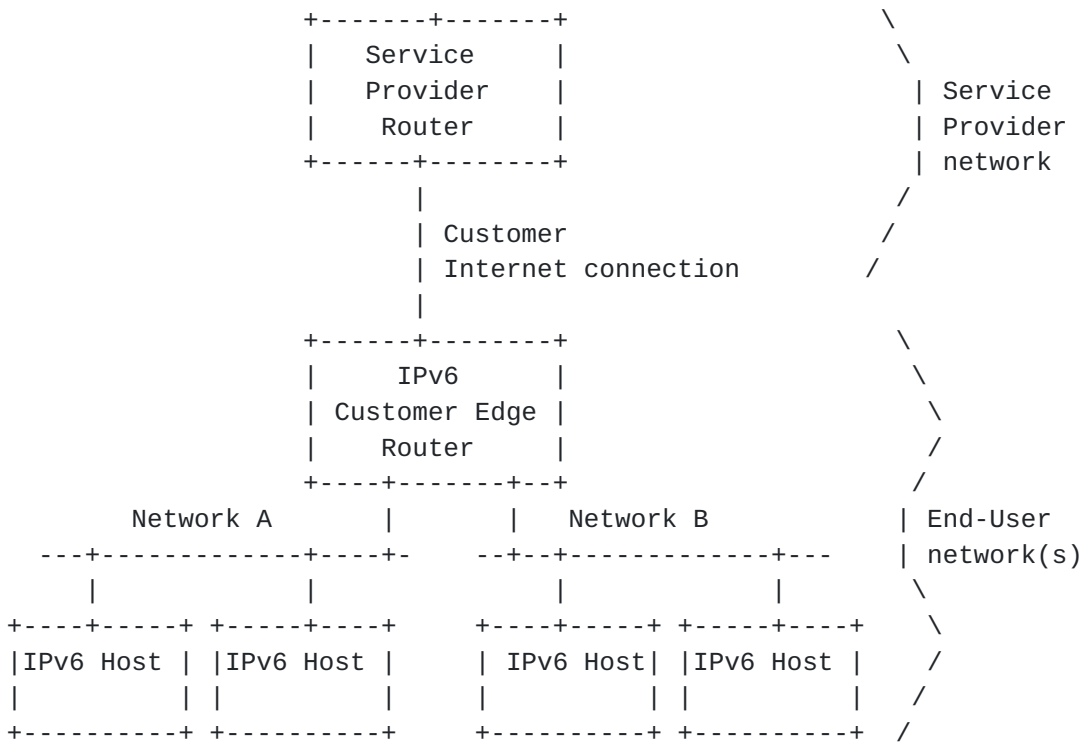

Figure 2

3.1.3. C: Single ISP, Single CER, Multiple internal subnets

Figure 3 shows a little bit more complex network with two routers and eight devices connected to one ISP. This network is similar to the one discussed in [I-D.ietf-v6ops-ipv6-cpe-router-bis]. The main complication in this topology compared to the ones described earlier is that there is no longer a single router that a priori understands the entire topology. The topology itself may also be complex. It may not be possible to assume a pure tree form, for instance. This is a valid consideration as home users may plug routers together to form arbitrary topologies including loops. In the following sections we discuss support for arbitrary topologies.

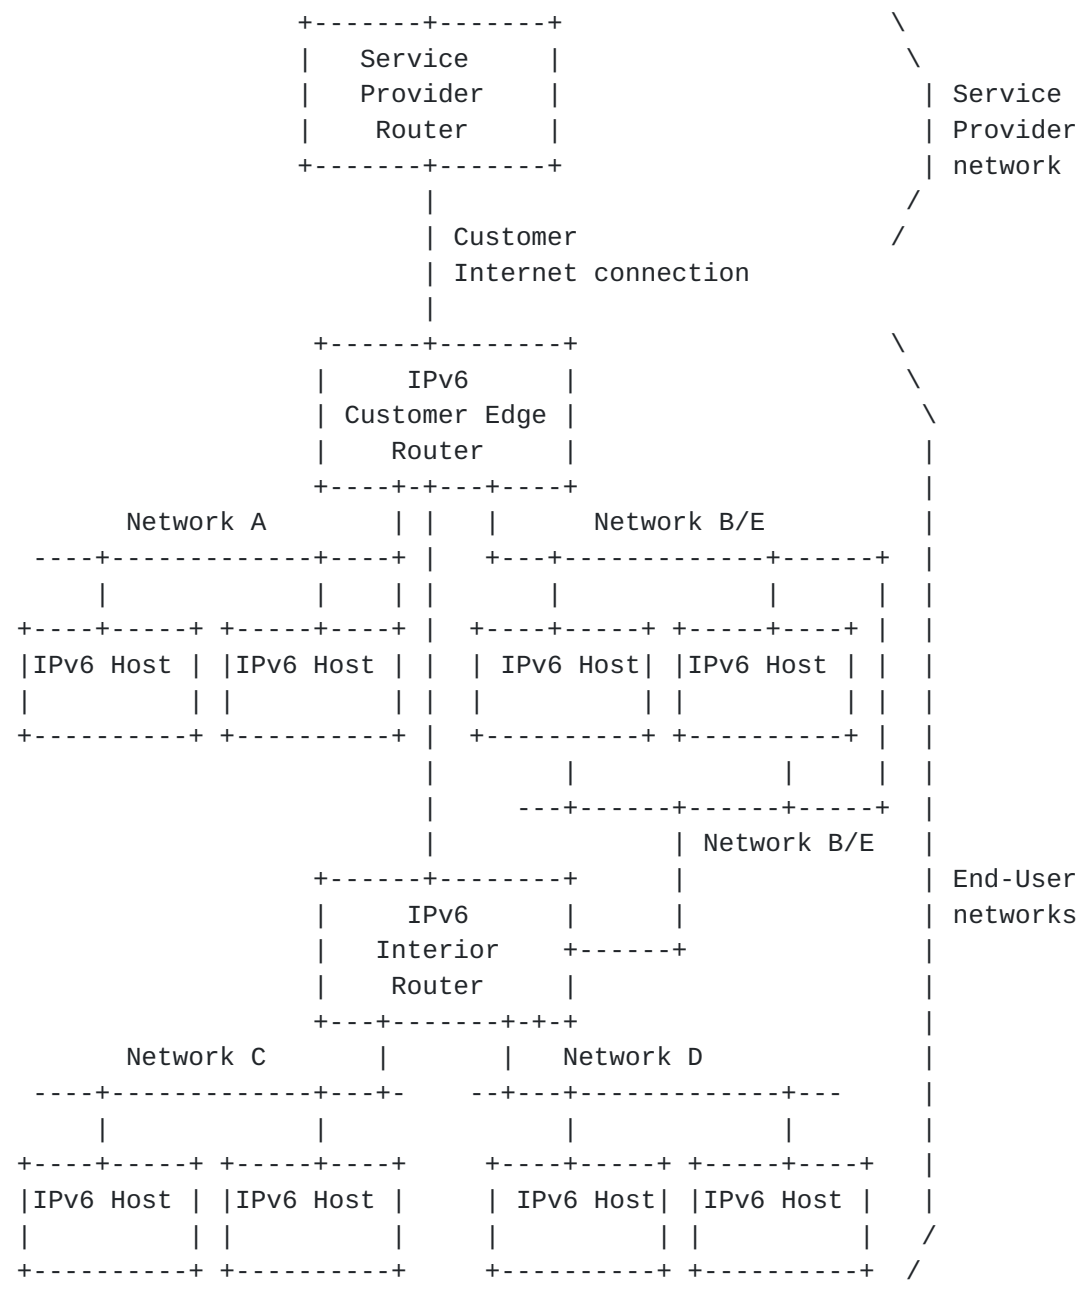
Service
Provider
Router
Service
Provider
network
Customer
Internet connection
IPv6
Customer Edge
Router
Network A
Network B/E
IPv6 Host
IPv6 Host
IPv6 Host
IPv6 Host
Network B/E
IPv6
Interior
Router
Network C
Network D
IPv6 Host
IPv6 Host
IPv6 Host
IPv6 Host
End-User
networks

Figure 3

3.1.4. D: Two ISPs, Two CERs, Shared subnets with multiple internal routers

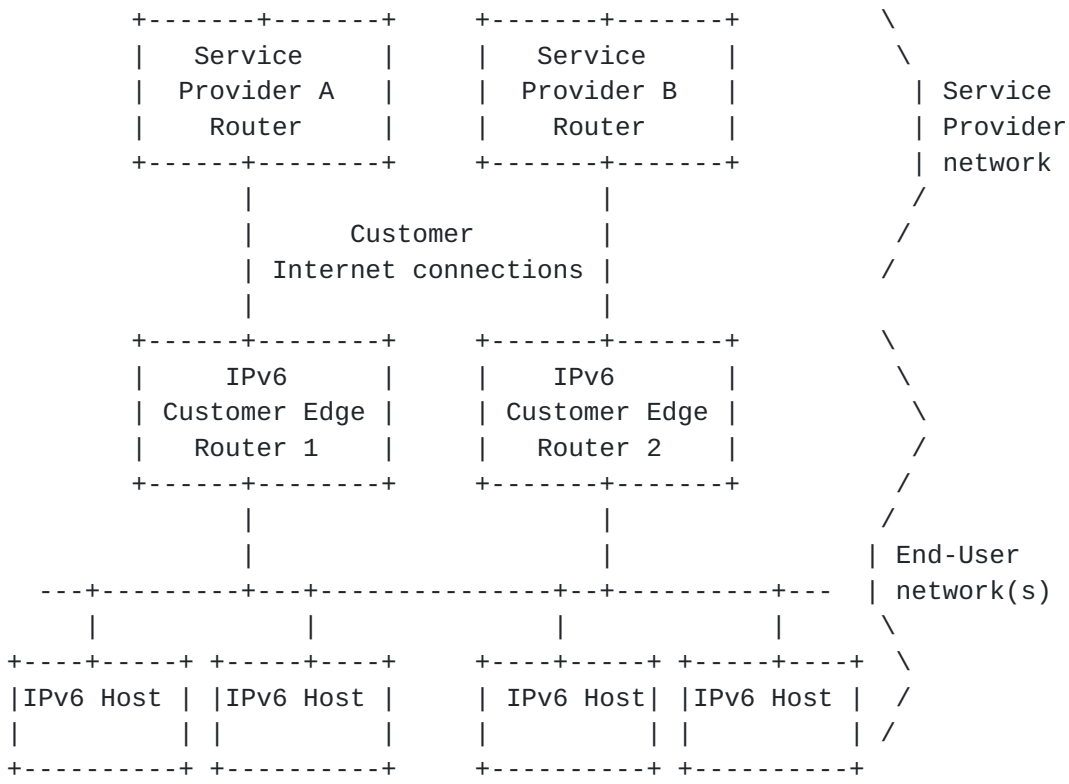

Figure 4

Figure 4 illustrates a multihomed home network model, where the customer has connectivity via CER1 to ISP A and via CER2 to ISP B. This example shows one shared subnet where IPv6 nodes would potentially be multihomed and receive multiple IPv6 global addresses, one per ISP. This model may also be combined with that shown in Figure 3 to create a more complex scenario with subnets that my be behind multiple internal routers.

3.1.5. E: Two ISPs, One CER, Isolated subnets with multiple internal routers

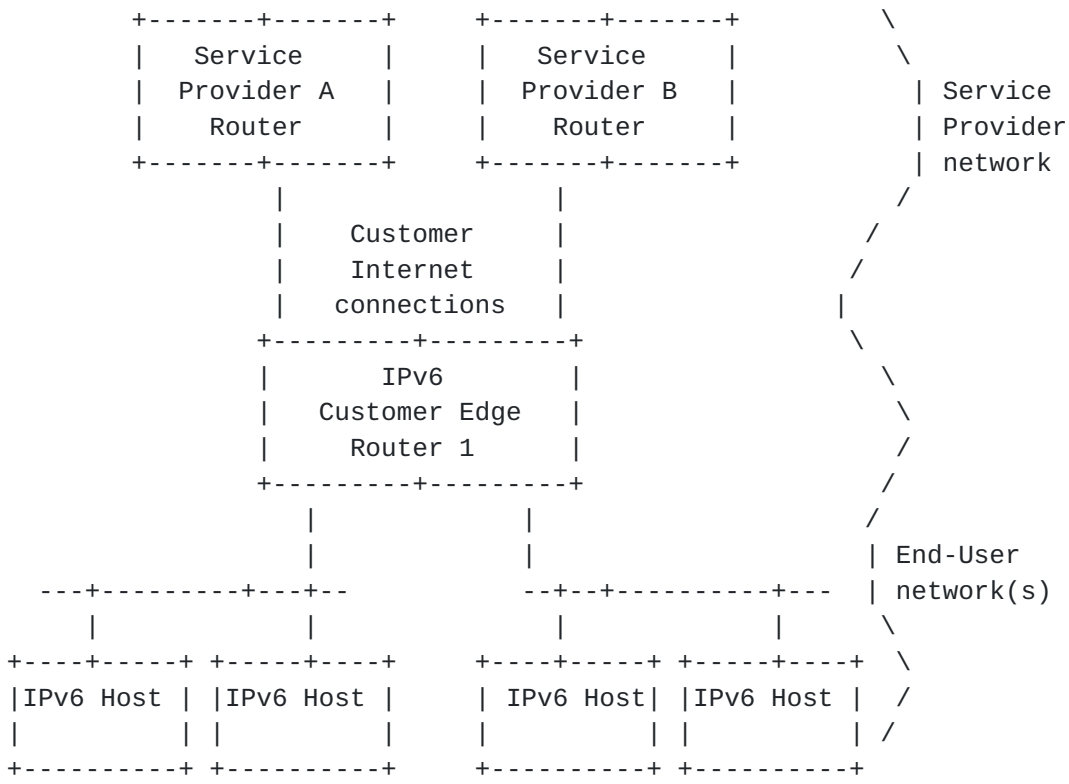

Figure 5

Figure 5 illustrates a model where a home network may have multiple connections to multiple providers or multiple logical connections to the same provider, but the associated subnet(s) are isolated. Some deployment scenarios may require this model.

3.1.6. F: Two ISPs, One CER, Shared subnets with multiple internal routers

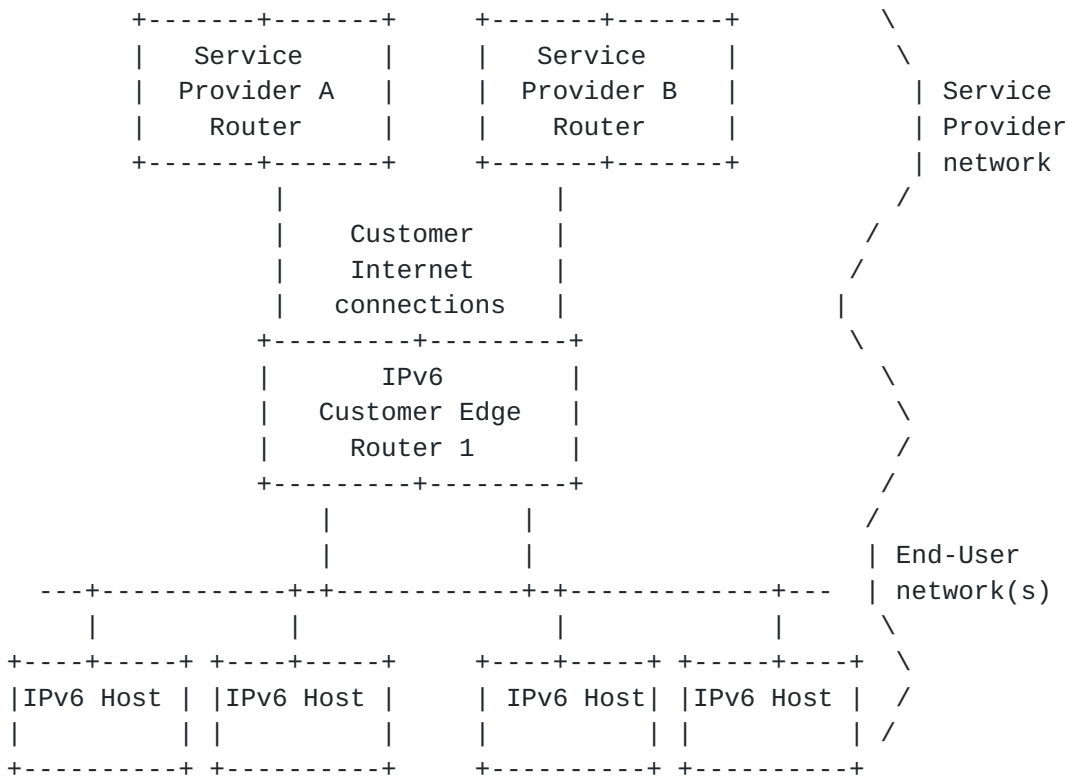

Figure 6

Figure 6 illustrates a model where a home network may have multiple connections to multiple providers or multiple logical connections to the same provider, with shared internal subnets, that may be multiple layers deep.

3.2. Determining the Requirements

[RFC6204] defines "basic" requirements for IPv6 Customer Edge Routers, while [I-D.ietf-v6ops-ipv6-cpe-router-bis] describes "advanced" features. In general, home network equipment needs to cope with the different types of network topologies discussed above. Manual configuration is rarely, if at all, possible, given the knowledge level of typical home users. The equipment needs to be prepared to handle at least

- Prefix configuration for routers

- Managing routing

- Name resolution
- Service discovery
- Network security

The remainder of the architecture document is presented as considerations and principles that lead to more specific requirements for the five general areas listed above.

3.3. Considerations

This section lists some considerations for home networking that may affect the architecture and associated requirements.

3.3.1. Multihoming

A homenet may be multihomed to multiple providers. This may either take a form where there are multiple isolated networks within the home (see Network Model E above) or a more integrated network where the connectivity selection is dynamic (see Network Model D or F above). Current practice is typically of the former kind, but the latter is expected to become more commonplace.

There are some specific multihoming considerations for homenet scenarios. First, it may be the case that multihoming applies due to an ISP migration from a transition method to a native deployment, e.g. a 6rd [RFC5969] sunset scenario. Second, one upstream may be a "walled garden", and thus only appropriate to be used for connectivity to the services of that provider.

In an integrated network, specific appliances or applications may use their own external connectivity, or the entire network may change its connectivity based on the status of the different upstream connections. The complexity of the multihoming solution required will depend on the Network Model deployed. For example, Network Models E and F have a single CER and thus could perform source routing at the single network exit point.

The general approach for IPv6 multihoming is for a hosts to receive multiple addresses from multiple providers, and to select the appropriate source address to communicate via a given provider. An alternative is to deploy ULAs with a site and then use NPTv6 [RFC6296], a prefix translation-based mechanism, at the edge. This obviously comes at some architectural cost, which is why approaches such as [I-D.v6ops-multihoming-without-ipv6nat] have been suggested. There has been much work on multihoming in the IETF, without (yet) widespread deployment of proposed solutions. Host-based methods such

as Shim6 [RFC5533] have also been defined, but of course require support in the hosts.

If multihoming is supported additional requirements apply. The general multihoming problem is broad, and solutions may include complex architectures for monitoring connectivity, traffic engineering, identifier-locator separation, connection survivability across multihoming events, and so on. This implies that if there is any support for multihoming defined in the homenet architecture it should be limited to a very small subset of the overall problem.

The current set of assumptions and requirements proposed by the homenet architecture team is:

MH1) The homenet WG should not try to make another attempt at solving complex multihoming; we should prefer to support scenarios for which solutions exist today.

MH2) Single CER Network Models E and F are in scope, and may be solved by source routing at the CER.

MH3) It is desirable to avoid deployment of NPTv6 at the CER. Hosts should be multi-addressed from each ISP they may communicate with or through.

MH4) Solutions that involve host changes should be avoided.

MH5) Walled garden multihoming is in scope.

MH6) Transition method sunsetting is in scope. The topic of multihoming with specific (6rd) transition coexistence is discussed in [I-D.townsley-troan-ipv6-ce-transitioning].

MH7) "Just" picking the right source address to use to fall foul of ingress filtering on upstream ISP connections (as per Network Model D) is not a trivial task. A solution is highly desirable, but out of scope of homenet.

MH8) Source routing throughout the homenet, a la [I-D.baker-fun-multi-router], requires relatively significant routing changes. The network should "guarantee" routing the packet to the correct exit given the source address, but hosts are responsible for anything extra, e.g. detecting failure, or choosing a new src/dst address combination.

Feedback is sought on the above points.

3.3.2. Quality of Service in multi-service home networks

Support for QoS in a multi-service homenet may be a requirement, e.g. for a critical system (perhaps healthcare related), or for differentiation between different types of traffic (file sharing, cloud storage, live streaming, VoIP, etc). Different media types may have different QoS properties or capabilities.

However, homenet scenarios should require no new QoS protocols. A DiffServ [RFC2475] approach with a small number of predefined traffic classes should generally be sufficient, though at present there is little experience of QoS deployment in home networks.

There may also be complementary mechanisms that could be beneficial in the homenet domain, such as ensuring proper buffering algorithms are used as described in [Gettys11].

3.3.3. Privacy considerations

There are no specific privacy concerns for this text. It should be noted that many ISPs are expected to offer relatively stable IPv6 prefixes to customers, and thus the network prefix associated with the host addresses they use would not generally change over a reasonable period of time, e.g. between restructuring of an ISPs residential network provision.

3.4. Principles

There is little that the Internet standards community can do about the physical topologies or the need for some networks to be separated at the network layer for policy or link layer compatibility reasons. However, there is a lot of flexibility in using IP addressing and inter-networking mechanisms. In this section we discuss how this flexibility should be used to provide the best user experience and ensure that the network can evolve with new applications in the future.

The following principles should be followed when designing homenet solutions. Where requirements are associated with those principles, they are listed here. There is no implied priority by the order in which the principles themselves are listed.

3.4.1. Reuse existing protocols

It is desirable to reuse existing protocols where possible, but at the same time to avoid consciously precluding the introduction of new or emerging protocols.

A generally conservative approach, giving weight to running code, is preferable. Where new protocols are required, evidence of commitment to implementation by appropriate vendors or development communities is highly desirable. Protocols used should be backwardly compatible.

Where possible, changes to hosts should be minimised. Some changes may be unavoidable however, e.g. signalling protocols to punch holes in firewalls where "Simple Security" is deployed in a CER.

Changes to routers should also be minimised, e.g. [I-D.baker-fun-routing-class] suggests introducing a routing protocol that may route on both source and destination addresses, which would be a significant change compared to current practices.

Liaisons with other appropriate standards groups and related organisations is desirable, e.g. the IEEE and Wi-Fi Alliance.

3.4.2. Dual-stack Operation

The homenet architecture targets both IPv6-only and dual-stack networks. While the CER requirements in RFC 6204 are aimed at IPv6-only networks, it is likely that dual-stack homenets will be the norm for some period of time. IPv6-only networking may first be deployed in home networks in "greenfield" scenarios, or perhaps as one element of an otherwise dual-stack network. The homenet architecture must operate in the absence of IPv4, and IPv6 must work in the same scenarios as IPv4 today.

Running IPv6-only may require documentation of additional considerations such as:

- Ensuring there is a way to access content in the IPv4 Internet. This can be arranged through incorporating NAT64 [RFC6144] functionality in the home gateway router, for instance.

- DNS discovery mechanisms are enabled for IPv6. Both stateless DHCPv6 [RFC3736] [RFC3646] and Router Advertisement options [RFC6106] may have to be supported and turned on by default to ensure maximum compatibility with all types of hosts in the network. This requires, however, that a working DNS server is known and addressable via IPv6.

- All nodes in the home network support operations in IPv6-only mode. Some current devices work well with dual-stack but fail to recognize connectivity when IPv4 DHCP fails, for instance.

In dual-stack networks, solutions for IPv6 must not adversely affect IPv4 operation. It is likely that topologies of IPv4 and IPv6

networks would be as congruent as possible.

Note that specific transition tools, particularly those running on the border CER to support transition tools being used inside the homenet, are out of scope. Use of tools, such as 6rd, on the border CER to support ISP access network transition are to be expected, but not within scope of homenet, which focuses on the internal networking.

3.4.3. Largest Possible Subnets

Today's IPv4 home networks generally have a single subnet, and early dual-stack deployments have a single congruent IPv6 subnet, possibly with some bridging functionality.

Future home networks are highly likely to need multiple subnets, for the reasons described earlier. As part of the self-organisation of the network, the network should subdivide itself to the largest possible subnets that can be constructed within the constraints of link layer mechanisms, bridging, physical connectivity, and policy. For instance, separate subnetworks are necessary where two different link layers cannot be bridged, or when a policy requires the separation of a private and visitor parts of the network.

While it may be desirable to maximise the chance of link-local protocols operating across a homenet by maximising the size of a subnet across the homenet, multiple subnet home networks are inevitable, so their support must be included. A general recommendation is to follow the same topology for IPv6 as is used for IPv4, but not to use NAT. Thus there should be routed IPv6 where an IPv4 NAT is used, and where there is no NAT there should be bridging if the link layer allows this.

In some cases IPv4 NAT home networks may feature cascaded NATs, e.g. where NAT routers are included within VMs or Internet connection services are used. IPv6 routed versions of such tools will be required.

3.4.4. Transparent End-to-End Communications

An IPv6-based home network architecture should naturally offer a transparent end-to-end communications model. Each device should be addressable by a unique address. Security perimeters can of course restrict the end-to-end communications, but it is simpler given the availability of globally unique addresses to block certain nodes from communicating by use of an appropriate filtering device than to configure the address translation device to enable appropriate address/port forwarding in the presence of a NAT.

As discussed previously, it is important to note the difference between hosts being addressable and reachable. Thus filtering is to be expected, while host-based IPv6 NAT is not. End-to-end communications are important for their robustness against failure of intermediate systems, where in contrast NAT is dependent on state machines which are not self-healing.

When configuring filters, protocols for securely associating devices are desirable. In the presence of "Simple Security" the use of signalling protocols such as uPnP or PCP may be expected to punch holes in the firewall (and be able to handle cases where there are multiple CERs/firewall(s). Alternatively, RFC 6092 supports the option for a border CER to run in "transparent mode", in which case a protocol like PCP is not required, but the security model is more open.

3.4.5. IP Connectivity between All Nodes

A logical consequence of the end-to-end communications model is that the network should by default attempt to provide IP-layer connectivity between all internal parts as well as between the internal parts and the Internet. This connectivity should be established at the link layer, if possible, and using routing at the IP layer otherwise.

Local addressing (ULAs) may be used within the scope of a home network. It would be expected that ULAs may be used alongside one or more globally unique ISP-provided addresses/prefixes in a homenet. ULAs may be used for all devices, not just those intended to have internal connectivity only. ULAs may then be used for stable internal communications should the ISP-provided prefix (suddenly) change, or external connectivity be temporarily lost. The use of ULAs should be restricted to the homenet scope through filtering at the border(s) of the homenet; thus "end-to-end" for ULAs is limited to the homenet.

In some cases full internal connectivity may not be desirable, e.g. in certain utility networking scenarios, or where filtering is required for policy reasons against guest network subnet(s). Note that certain scenarios may require co-existence of ISP connectivity providing a general Internet service with provider connectivity to a private "walled garden" network.

Some home networking scenarios/models may involve isolated subnet(s) with their own CERs. In such cases connectivity would only be expected within each isolated network (though traffic may potentially pass between them via external providers).

LLNs provide an example of where there may be secure perimeters inside the homenet. Constrained LLN nodes may implement WPA-style network key security but may depend on access policies enforced by the LLN border router.

3.4.6. Routing functionality

Routing functionality is required when there are multiple routers in use. This functionality could be as simple as the current "default route is up" model of IPv4 NAT, or it could involve running an appropriate routing protocol.

The homenet routing environment may include traditional IP networking where existing link-state or distance-vector protocols may be used, but also new LLN or other "constrained" networks where other protocols may be more appropriate. IPv6 VM solutions may also add additional routing requirements. Current home deployments use largely different mechanisms in sensor and basic Internet connectivity networks.

In this section we list the requirements and assumptions for routing functionality within the homenet environment.

RT1) The protocol should preferably be an existing deployed protocol that has been proven to be reliable and robust.

RT2) It is preferable that the protocol is "lightweight".

RT3) The protocol should provide reachability between all nodes in the homement.

RT4) In general, LLN or other networks should be able to attach and participate the same way or map/be gatewayed to the main homenet.

RT5) Multiple interface PHYs must be accounted for in the homenet routed topology. Technologies such as Ethernet, WiFi, MoCA, etc must be capable of coexisting in the same environment and should be tested as part of any routed deployment. The inclusion of the PHY layer characteristics including bandwidth, loss, and latency in path computation should be considered for optimizing communication in the homenet.

RT6) Minimizing convergence time should be a goal in any routed environment, but as a guideline a maximum convergence time of a couple of minutes should be the target.

RT7) It is desirable that the routing protocol has knowledge of the homenet topology, which implies a link-state protocol may be preferable. If so, it is also desirable that the announcements and use of LSAs and RAs are appropriately coordinated.

RT8) Any routed solution will require a means for determining the boundaries of the homenet. Borders may include but are not limited to the interface to the upstream ISP, a gateway device to a separate home network such as a SmartGrid or similar LLN network, and in some cases there may be no border such as before an upstream connection has been established. Devices in the homenet must be able to find the path to the Internet as well as other devices on the home intranet. The border discovery functionality may be integrated into the routing protocol itself, but may also be imported via a separate discovery mechanism.

RT9) The routing environment should be self-configuring, as discussed in the next subsection. An example of how OSPFv3 can be self-configuring in a homenet is described in [I-D.acee-ospf-ospfv3-autoconfig]. The exception is configuration of a "secret" for authentication methods. It is important that self-configuration with "unintended" devices is avoided.

RT10) The protocol should not require upstream ISP connectivity to be established to continue routing within the homenet.

RT11) Multiple upstreams should be supported, as described in the Network Models earlier.

RT12) To support multihoming within a homenet, a routing protocol that can make routing decisions based on source and destination addresses is desirable, to avoid upstream ISP ingress filtering problems. In general the routing protocol should support multiple ISP uplinks and delegated prefixes in concurrent use.

RT13) The routing system should support walled garden environments.

RT14) Load-balancing to multiple providers is not a requirement, but failover from a primary to a backup link when available must be a requirement.

RT15) It is assumed that the typical router designed for residential use does not contain the memory or cpu required to process a full Internet routing table this should not be a requirement for any homenet device.

A new I-D has been published on homenet routing requirements, see [I-D.howard-homenet-routing-comparison] and evaluations of common routing protocols made against those requirements, see [I-D.howard-homenet-routing-requirements]. The requirements from the former document have been worked into this architecture text. Feedback is sought on how these documents move forward.

3.4.7. Self-Organising

A home network architecture should be naturally self-organising and self-configuring under different circumstances relating to the connectivity status to the Internet, number of devices, and physical topology. While the homenet should be self-organising, it should be possible to manually adjust (override) the current configuration.

The most important function in this respect is prefix delegation and management. The requirements and assumptions for the prefix delegation function are summarised as follows:

PD1) From the homenet perspective, a single prefix should be received on the border CER [RFC3633]. The ISP should only see that aggregate, and not single /64 prefixes allocated within the homenet.

PD2) Each link in the homenet should receive a prefix from within the ISP-provided prefix.

PD3) Delegation should be autonomous, and not assume a flat or hierarchical model.

PD4) The assignment mechanism should provide reasonable efficiency, so that typical home network prefix allocation sizes can accommodate all the necessary /64 allocations in most cases. A currently typical /60 allocation gives 16 /64 subnets.

PD5) Duplicate assignment of multiple /64s to the same network should be avoided.

PD6) The network should behave as gracefully as possible in the event of prefix exhaustion.

PD7) Where multiple CERs exist with multiple ISP prefix pools, it is expected that routers within the homenet would assign themselves prefixes from each ISP they communicate with/ through.

PD8) Where ULAs are used, most likely but not necessarily in parallel with global prefixes, one router will need to be elected as the generator of ULA prefixes for the homenet.

PD9) Delegation within the homenet should give each link a prefix that is persistent across reboots, power outages and similar short-term outages.

PD10) Addition of a new routing device should not affect existing persistent prefixes, but persistence may not be expected in the face of significant "replumbing" of the homenet.

PD11) Persistence should not depend on router boot order.

PD12) Persistent prefixes may imply the need for stable storage on routing devices, and also a method for a home user to "reset" the stored prefix should a significant reconfiguration be required (though ideally the home user should not be involved at all).

PD13) The delegation method should support "flash" renumbering.

Several proposals have been made for prefix delegation within a homenet. One group of proposals is based on DHCPv6 PD, as described in [I-D.baker-homenet-prefix-assignment], [I-D.chakrabarti-homenet-prefix-alloc], [RFC3315] and [RFC3633]. The other uses OSPFv3, as described in [I-D.arkko-homenet-prefix-assignment]. More detailed analysis of these approaches needs to be made against the requirements/ assumptions listed above.

Other parameters of the network will need to be self-organising. The network elements will need to be integrated in a way that takes account of the various lifetimes on timers that are used on those different elements, e.g. DHCPv6 PD, router, valid prefix and preferred prefix timers.

The homenet will have one or more borders, with external connectivity providers and potentially parts of the internal network (e.g. for policy-based reasons). It should be possible to automatically perform border discovery at least for the ISP borders. Such borders determine for example the scope of ULAs, site scope multicast boundaries and where firewall policies may be applied.

The network cannot be expected to be completely self-organising, e.g. some security parameters are likely to need manual configuration, e.g. WPA2 configuration for wireless access control. Some existing mechanisms exist to assist home users to associate devices as simply as possible, e.g. "connect" button support.

3.4.8. Fewest Topology Assumptions

There should ideally be no built-in assumptions about the topology in home networks, as users are capable of connecting their devices in ingenious ways. Thus arbitrary topologies will need to be supported.

It is important not to introduce new IPv6 scenarios that would break with IPv4+NAT, given that dual-stack homenets will be commonplace for some time. There may be IPv6-only topologies that work where IPv4 is not used or required.

3.4.9. Naming and Service Discovery

The most natural way to think about naming and service discovery within a homenet is to enable it to work across the entire residence, disregarding technical borders such as subnets but respecting policy borders such as those between visitor and internal networks.

Homenet naming systems will be required that work internally or externally, though the domains used may be different from those different perspectives.

A desirable target may be a fully functional self-configuring secure local DNS service so that all devices are referred to by name, and these FQDNs are resolved locally. This would make clean use of ULAs and multiple ISP-provided prefixes much easier. The local DNS service should be (by default) authoritative for the local name space in both IPv4 and IPv6. A dual-stack residential gateway should include a dual-stack DNS server.

Consideration will also need to be given for existing protocols that may be used within a network, e.g. mDNS, and how these interact with unicast-based DNS services.

With the introduction of new top level domains, there is potential for ambiguity between for example a local host called apple and (if it is registered) an apple gTLD, so some local name space is probably required, which should also be configurable to something else by a home user, e.g. ".home", if desired.

It is also important to note here that there is also potential ambiguity if a mobile device should move between two local name

spaces called ".home", for example.

For service discovery, support may be required for IPv6 multicast across the scope of the home network, and thus at least all routing devices in the network.

3.4.10. Proxy or Extend?

Related to the above, we believe that general existing discovery protocols that are designed to only work within a subnet should be modified/extended to work across subnets, rather than defining proxy capabilities for each of those functions.

Feedback is desirable on which other functions/protocols assume subnet-only operation, in the context of existing home networks. Some experience from enterprises may be relevant here.

3.4.11. Adapt to ISP constraints

The home network may receive an arbitrary length IPv6 prefix from its provider, e.g. /60 or /56. The offered prefix may be stable over time or change frequently. The home network needs to be adaptable to such ISP policies, e.g. on constraints placed by the size of prefix offered by the ISP. The ISP may use [I-D.ietf-dhc-pd-exclude] for example.

The internal operation of the home network should also not depend on the availability of the ISP network at any given time, other than for connectivity to services or systems off the home network. This implies the use of ULAs as supported in RFC6204. If used, ULA addresses should be stable so that they can always be used internally, independent of the link to the ISP.

It is expected that ISPs will deliver a relatively stable home prefix to customers. The norm for residential customers of large ISPs may similar to their single IPv4 address provision; by default it is likely to remain persistent for some time, but changes in the ISP's own provisioning systems may lead to the customer's IP (and in the IPv6 case their prefix pool) changing.

When an ISP needs to restructure and in doing so renumber its customer homenets, "flash" renumbering is likely to be imposed. This implies a need for the homenet to be able to handle a sudden renumbering event which, unlike the process described in [RFC4192], would be without a "flag day". The customer may of course also choose to move to a new ISP, and thus begin using a new prefix. Thus it's desirable that homenet protocols or operational processes don't add unnecessary complexity for renumbering.

The 6renum WG is studying IPv6 renumbering for enterprise networks. It is not currently targetting homenets, but may produce outputs that are relevant.

3.5. Summary of Homenet Architecture Recommendations

Feedback sought on whether a summary section would be useful.

3.6. Implementing the Architecture on IPv6

The necessary mechanisms are largely already part of the IPv6 protocol set and common implementations, though there are some exceptions. For automatic routing, it is expected that existing routing protocols can be used as is. However, a new mechanism may be needed in order to turn a selected protocol on by default. Support for multiple exit routers and multi-homing would also require extensions, even if focused on the problem of multi-addressed hosts selecting the right source address to avoid falling foul of ingress filtering on upstream ISP connections.

For name resolution and service discovery, extensions to existing multicast-based name resolution protocols are needed to enable them to work across subnets, within the scope of the home network.

The hardest problems in developing solutions for home networking IPv6 architectures include discovering the right borders where the domain "home" ends and the service provider domain begins, deciding whether some of necessary discovery mechanism extensions should affect only the network infrastructure or also hosts, and the ability to turn on routing, prefix delegation and other functions in a backwards compatible manner.

4. References

4.1. Normative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, 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.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.

[RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.

[RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.

[RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011.

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.

4.2. Informative References

[RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.

[RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol

(DHCP) Service for IPv6", RFC 3736, April 2004.

[RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.

[I-D.baker-fun-multi-router] Baker, F., "Exploring the multi-router SOHO network", draft-baker-fun-multi-router-00 (work in progress), July 2011.

[I-D.townsley-troan-ipv6-ce-transitioning] Townsley, M. and O. Troan, "Basic Requirements for Customer Edge Routers - multihoming and transition", draft-townsley-troan-ipv6-ce-transitioning-02 (work in progress), December 2011.

[I-D.baker-fun-routing-class] Baker, F., "Routing a Traffic Class", draft-baker-fun-routing-class-00 (work in progress), July 2011.

[I-D.howard-homenet-routing-comparison] Howard, L., "Evaluation of Proposed Homenet Routing Solutions", draft-howard-homenet-routing-comparison-00 (work in progress), December 2011.

[I-D.howard-homenet-routing-requirements] Howard, L., "Homenet Routing Requirements", draft-howard-homenet-routing-requirements-00 (work in progress), December 2011.

[I-D.herbst-v6ops-cpeenhancements] Herbst, T. and D. Sturek, "CPE Considerations in IPv6 Deployments", draft-herbst-v6ops-cpeenhancements-00 (work in progress), October 2010.

[I-D.vyncke-advanced-ipv6-security] Vyncke, E., Yourtchenko, A., and M. Townsley, "Advanced Security for IPv6 CPE", draft-vyncke-advanced-ipv6-security-03 (work in progress), October 2011.

[I-D.ietf-v6ops-ipv6-cpe-router-bis] Singh, H., Beebee, W., Donley, C., Stark, B., and O.

Troan, "Advanced Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-ipv6-cpe-router-bis-01 (work in progress), July 2011.

[I-D.ietf-6man-rfc3484-revise]
Matsumoto, A., Kato, J., Fujisaki, T., and T. Chown, "Update to RFC 3484 Default Address Selection for IPv6", draft-ietf-6man-rfc3484-revise-05 (work in progress), October 2011.

[I-D.ietf-dhc-pd-exclude]
Korhonen, J., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", draft-ietf-dhc-pd-exclude-04 (work in progress), December 2011.

[I-D.v6ops-multihoming-without-ipv6nat]
Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation", draft-v6ops-multihoming-without-ipv6nat-00 (work in progress), March 2011.

[I-D.baker-homenet-prefix-assignment]
Baker, F. and R. Droms, "IPv6 Prefix Assignment in Small Networks", draft-baker-homenet-prefix-assignment-00 (work in progress), October 2011.

[I-D.arkko-homenet-prefix-assignment]
Arkko, J. and A. Lindem, "Prefix Assignment in a Home Network", draft-arkko-homenet-prefix-assignment-01 (work in progress), October 2011.

[I-D.acee-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", draft-acee-ospf-ospfv3-autoconfig-00 (work in progress), October 2011.

[I-D.ietf-pcp-base]
Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", draft-ietf-pcp-base-22 (work in progress), January 2012.

[I-D.chakrabarti-homenet-prefix-alloc]
Nordmark, E., Chakrabarti, S., Krishnan, S., and W. Haddad, "Simple Approach to Prefix Distribution in Basic Home Networks", draft-chakrabarti-homenet-prefix-alloc-01 (work in progress), October 2011.

[Gettys11]
Gettys, J., "Bufferbloat: Dark Buffers in the Internet", March 2011, <http://www.ietf.org/proceedings/80/slides/tsvarea-1.pdf>.

Appendix A. Acknowledgments

The authors would like to thank Brian Carpenter, Mark Andrews, Fred Baker, Ray Bellis, Cameron Byrne, Stuart Cheshire, Lorenzo Colitti, Ralph Droms, Lars Eggert, Jim Gettys, Wassim Haddad, Joel M. Halpern, David Harrington, Lee Howard, Ray Hunter, Joel Jaeggli, Heather Kirksey, Ted Lemon, Erik Nordmark, Michael Richardson, Barbara Stark, Sander Steffann, Dave Thaler, JP Vasseur, Curtis Villamizar, Russ White, and James Woodyatt for their contributions within homenet WG meetings and the mailing list, and Mark Townsley for being an initial editor/author of this text before taking his position as homenet WG co-chair.

Authors' Addresses

Jari Arkko
Ericsson
Jorvas 02420
Finland

Email: jari.arkko@piuha.net

Anders Brandt
Sigma Designs
Emdrupvej 26A, 1
Copenhagen DK-2100
Denmark

Email: abr@sdesigns.dk

Tim Chown
University of Southampton
Highfield
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk

Jason Weil
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
USA

Email: jason.weil@twcable.com

Ole Troan
Cisco Systems, Inc.
Drammensveien 145A
Oslo N-0212
Norway

Email: ot@cisco.com