

Internet Engineering Task Force (IETF)
Request for Comments: 7368
Category: Informational
ISSN: 2070-1721

T. Chown, Ed.
University of Southampton
J. Arkko
Ericsson
A. Brandt
Sigma Designs
O. Troan
Cisco Systems, Inc.
J. Weil
Time Warner Cable
October 2014

IPv6 Home Networking Architecture Principles

Abstract

This text describes evolving networking technology within residential home networks with increasing numbers of devices and a trend towards increased internal routing. The goal of this document is to define a general architecture for IPv6-based home networking, describing the associated principles, considerations, and requirements. The text briefly highlights specific implications of the introduction of IPv6 for home networking, discusses the elements of the architecture, and suggests 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 document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7368>.

Copyright Notice

Copyright (c) 2014 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	4
1.1.	Terminology and Abbreviations	5
2.	Effects of IPv6 on Home Networking	6
2.1.	Multiple Subnets and Routers	7
2.2.	Global Addressability and Elimination of NAT	8
2.3.	Multi-Addressing of Devices	8
2.4.	Unique Local Addresses (ULAs)	9
2.5.	Avoiding Manual Configuration of IP Addresses	10
2.6.	IPv6-Only Operation	11
3.	Homenet Architecture Principles	11
3.1.	General Principles	12
3.1.1.	Reuse Existing Protocols	12
3.1.2.	Minimise Changes to Hosts and Routers	13
3.2.	Homenet Topology	13
3.2.1.	Supporting Arbitrary Topologies	13
3.2.2.	Network Topology Models	14
3.2.3.	Dual-Stack Topologies	18
3.2.4.	Multihoming	19
3.2.5.	Mobility Support	20
3.3.	A Self-Organising Network	21
3.3.1.	Differentiating Neighbouring Homenets	21
3.3.2.	Largest Practical Subnets	21
3.3.3.	Handling Varying Link Technologies	22
3.3.4.	Homenet Realms and Borders	22
3.3.5.	Configuration Information from the ISP	23
3.4.	Homenet Addressing	24
3.4.1.	Use of ISP-Delegated IPv6 Prefixes	24
3.4.2.	Stable Internal IP Addresses	26
3.4.3.	Internal Prefix Delegation	27
3.4.4.	Coordination of Configuration Information	28
3.4.5.	Privacy	28

3.5.	Routing Functionality	28
3.5.1.	Unicast Routing within the Homenet	30
3.5.2.	Unicast Routing at the Homenet Border	31
3.5.3.	Multicast Support	31
3.6.	Security	32
3.6.1.	Addressability vs. Reachability	32
3.6.2.	Filtering at Borders	33
3.6.3.	Partial Effectiveness of NAT and Firewalls	34
3.6.4.	Exfiltration Concerns	34
3.6.5.	Device Capabilities	34
3.6.6.	ULAs as a Hint of Connection Origin	35
3.7.	Naming and Service Discovery	35
3.7.1.	Discovering Services	35
3.7.2.	Assigning Names to Devices	36
3.7.3.	The Homenet Name Service	37
3.7.4.	Name Spaces	38
3.7.5.	Independent Operation	40
3.7.6.	Considerations for LLNs	40
3.7.7.	DNS Resolver Discovery	41
3.7.8.	Devices Roaming to/from the Homenet	41
3.8.	Other Considerations	41
3.8.1.	Quality of Service	41
3.8.2.	Operations and Management	42
3.9.	Implementing the Architecture on IPv6	43
4.	Conclusions	44
5.	Security Considerations	44
6.	References	44
6.1.	Normative References	44
6.2.	Informative References	44
	Acknowledgments	48
	Authors' Addresses	49

1. Introduction

This document focuses on evolving networking technology within residential home networks with increasing numbers of devices and a trend towards increased internal routing, as well as the associated challenges with their deployment and operation. There is a growing trend in home networking for the proliferation of networking technology through 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 introduction of IPv6, while others relate to the introduction of specialised networks for home automation and sensors.

While at the time of writing some complex home network topologies exist, most are relatively simple single subnet networks and ostensibly operate using just IPv4. While there may be IPv6 traffic within the network, e.g., for service discovery, the homenet is provisioned by the ISP as an IPv4 network. Such networks also typically employ solutions that should be avoided, such as private [RFC1918] addressing with (cascaded) Network Address Translation (NAT) [RFC3022], or they may require expert assistance to set up.

In contrast, emerging IPv6-capable home networks are very likely to have multiple internal subnets, e.g., to facilitate private and guest networks, heterogeneous link layers, and smart grid components, and have enough address space available to allow every device to have a globally unique address. This implies that internal routing functionality is required, and that the homenet's ISP delegates a large enough address block, to allow assignment of a prefix to each subnet in the home network.

It is not practical to expect home users to configure their networks. Thus, the assumption of this document is that the homenet is as far as possible self-organising and self-configuring, i.e., it should function without proactive management by the residential user.

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 aiming for a more consistent solution that addresses as many of the identified requirements as possible. This 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. We assume that the IPv4 network architecture in home networks is what it is and cannot be modified by new recommendations. This document does not discuss how IPv4 home

networks provision or deliver support for multiple subnets. It should not be assumed that any future new functionality created with IPv6 in mind will be backward compatible to include IPv4 support. Further, future deployments, or specific subnets within an otherwise dual-stack home network, may be IPv6-only, in which case considerations for IPv4 impact would not apply.

This document proposes a baseline homenet architecture, using protocols and implementations that are as far as possible proven and robust. The scope of the document is primarily the network-layer technologies that provide the basic functionality to enable addressing, connectivity, routing, naming, and service discovery. While it may, for example, state that homenet components must be simple to deploy and use, it does not discuss specific user interfaces, nor does it discuss specific physical, wireless, or data-link-layer considerations. Likewise, we also do not specify the whole design of a homenet router from top to bottom; rather, we focus on the Layer 3 aspects. This means that Layer 2 is largely out of scope, we're assuming a data-link layer that supports IPv6 is present, and we react accordingly. Any IPv6-over-Foo definitions occur elsewhere.

[RFC7084], which has obsoleted [RFC6204], defines basic requirements for Customer Edge (CE) routers. The update includes the definition of requirements for specific transition tools on the CE router, specifically Dual-Stack Lite (DS-Lite) [RFC6333] and IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) [RFC5969]. Such detailed specification of CE router devices is considered out of scope of this architecture document, and we assume that any required update of the CE router device specification as a result of adopting this architecture will be handled as separate and specific updates to these existing documents. Further, the scope of this text is the internal homenet, and thus specific features on the WAN side of the CE router are out of scope for this text.

1.1. Terminology and Abbreviations

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

- o Border: A point, typically resident on a router, between two networks, e.g., between the main internal homenet and a guest network. This defines a point(s) at which filtering and forwarding policies for different types of traffic may be applied.
- o CE router: Customer Edge router. A border router intended for use in a homenet. A CE router connects the homenet to a service provider network.

- o FQDN: Fully Qualified Domain Name. A globally unique name.
 - o Guest network: A part of the home network intended for use by visitors or guests to the home(net). Devices on the guest network may typically not see or be able to use all services in the home(net).
 - o Homenet: A home network, comprising host and router equipment, with one or more CE routers providing connectivity to a service provider network(s).
 - o ISP: Internet Service Provider. An entity that provides access to the Internet. In this document, a service provider specifically offers Internet access using IPv6 and may also offer IPv4 Internet access. The service provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.
 - o LLN: Low-power and Lossy Network.
 - o LQDN: Locally Qualified Domain Name. A name local to the homenet.
 - o NAT: Network Address Translation. Typically referring to IPv4 Network Address Port Translation (NAPT) [RFC3022].
 - o NPTv6: IPv6-to-IPv6 Network Prefix Translation [RFC6296].
 - o PCP: Port Control Protocol [RFC6887].
 - o Realm: A network delimited by a defined border. A guest network within a homenet may form one realm.
 - o 'Simple Security': Defined in [RFC4864] and expanded further in [RFC6092]; describes recommended perimeter security capabilities for IPv6 networks.
 - o ULA: IPv6 Unique Local Address [RFC4193].
 - o VM: Virtual Machine.
2. Effects of IPv6 on Home Networking

While IPv6 resembles IPv4 in many ways, there are some notable differences in the way it may typically be deployed. It changes address allocation principles, making multi-addressing the norm, and through the vastly increased address space, it allows globally unique IP addresses to be used for all devices in a home network. This section presents an overview of some of the key implications of the

introduction of IPv6 for home networking that are simultaneously both promising and problematic.

2.1. Multiple Subnets and Routers

While simple Layer 3 topologies involving as few subnets as possible are preferred in home networks, the incorporation of dedicated (routed) subnets remains necessary for a variety of reasons. For instance, an increasingly common feature in modern home routers is the ability to support both guest and private network subnets. Likewise, there may be a need to separate home automation or corporate extension LANs (whereby a home worker can have their corporate network extended into the home using a virtual private network, commonly presented as one port on an Ethernet device) from the main Internet access network, or different subnets may in general be associated with parts of the homenet that have different routing and security policies. Further, 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. Constraining the flow of certain traffic from Ethernet links to links of much lower capacity thus becomes an important topic.

The introduction of IPv6 for home networking makes it possible for every home network to be delegated enough address space from its ISP to provision globally unique prefixes for each such subnet in the home. While the number of addresses in a standard /64 IPv6 prefix is practically unlimited, the number of prefixes available for assignment to the home network is not. As a result, the growth inhibitor for the home network shifts from the number of addresses to the number of prefixes offered by the provider; this topic is discussed in BCP 157 [RFC6177], which recommends that "end sites always be able to obtain a reasonable amount of address space for their actual and planned usage."

The addition of routing between subnets raises a number of issues. One is a method by which prefixes can be efficiently allocated to each subnet, without user intervention. Another issue is how to extend mechanisms such as zero-configuration service discovery that currently only operate within a single subnet using link-local traffic. In a typical IPv4 home network, there is only one subnet, so such mechanisms would normally operate as expected. For multi-subnet IPv6 home networks, there are two broad choices to enable such protocols to work across the scope of the entire homenet: extend existing protocols to work across that scope or introduce proxies for existing link-layer protocols. This topic is discussed in Section 3.7.

2.2. Global Addressability and Elimination of NAT

The possibility for direct end-to-end communication on the Internet to be restored by the introduction of IPv6 is, on the one hand, an incredible opportunity for innovation and simpler network operation, but on the other hand, it is also a concern as it potentially exposes nodes in the internal networks to receipt of unwanted and possibly malicious traffic from the Internet.

With devices and applications able to talk directly to each other when they have globally unique addresses, there may be an expectation of improved host security to compensate for this. It should be noted that many devices may (for example) ship with default settings that make them readily vulnerable to compromise by external attackers if globally accessible, or they may simply not be robust by design because it was assumed that either such devices would only be used on private networks or the devices don't have the computing power to apply the necessary security methods. In addition, the upgrade cycle for devices (or their firmware) may be slow and/or lack auto-update mechanisms.

It is thus important to distinguish between addressability and reachability. While IPv6 offers global addressability through the use of globally unique addresses in the home, whether devices are globally reachable or not would depend on any firewall or filtering configuration, and not, as is commonly the case with IPv4, the presence or use of NAT. In this respect, IPv6 networks may or may not have filters applied at their borders to control such traffic, i.e., at the homenet CE router. [RFC4864] and [RFC6092] discuss such filtering and the merits of 'default allow' against 'default deny' policies for external traffic initiated into a homenet. This topic is discussed further in Section 3.6.1.

2.3. Multi-Addressing of Devices

In an IPv6 network, devices will often acquire multiple addresses, typically at least a link-local address and one or more globally unique addresses (GUAs). Where a homenet is multihomed, a device would typically receive a GUA from within the delegated prefix from each upstream ISP. Devices may also have an IPv4 address if the network is dual stack, an IPv6 Unique Local Address (ULA) [RFC4193] (see below), and one or more IPv6 privacy addresses [RFC4941].

It should thus be considered the norm for devices on IPv6 home networks to be multi-addressed and to need to make appropriate address selection decisions for the candidate source and destination address pairs for any given connection. In multihoming scenarios, nodes will be configured with one address from each upstream ISP

prefix. In such cases, the presence of upstream ingress filtering as described in BCP 38 [RFC2827] requires such multi-addressed nodes to select the correct source address to be used for the corresponding uplink. Default address selection for IPv6 [RFC6724] provides a solution for this, but a challenge here is that the node may not have the information it needs to make that decision based on addresses alone. We discuss this challenge in Section 3.2.4.

2.4. Unique Local Addresses (ULAs)

[RFC4193] defines ULAs for IPv6 that may be used to address devices within the scope of a single site. Support for ULAs for IPv6 CE routers is described in [RFC7084]. A home network running IPv6 should deploy ULAs alongside its globally unique prefix(es) to allow stable communication between devices (on different subnets) within the homenet where that externally allocated globally unique prefix may change over time, e.g., due to renumbering within the subscriber's ISP, or where external connectivity may be temporarily unavailable. A homenet using provider-assigned global addresses is exposed to its ISP renumbering the network to a much larger degree than before whereas, for IPv4, NAT isolated the user against ISP renumbering to some extent.

While setting up a network, there may be a period where it has no external connectivity, in which case ULAs would be required for inter-subnet communication. In the case where home automation networks are being set up in a new home/deployment (as early as during construction of the home), such networks will likely need to use their own /48 ULA prefix. Depending upon circumstances beyond the control of the owner of the homenet, it may be impossible to renumber the ULA used by the home automation network so routing between ULA /48s may be required. Also, some devices, particularly constrained devices, may have only a ULA (in addition to a link-local), while others may have both a GUA and a ULA.

Note that unlike private IPv4 space as described in RFC 1918, the use of ULAs does not imply use of an IPv6 equivalent of a traditional IPv4 NAT [RFC3022] or of NPTv6 prefix-based NAT [RFC6296]. When an IPv6 node in a homenet has both a ULA and a globally unique IPv6 address, it should only use its ULA address internally and use its additional globally unique IPv6 address as a source address for external communications. This should be the natural behaviour given support for default address selection for IPv6 [RFC6724]. By using such globally unique addresses between hosts and devices in remote networks, the architectural cost and complexity, particularly to applications, of NAT or NPTv6 translation are avoided. As such, neither IPv6 NAT nor NPTv6 is recommended for use in the homenet architecture. Further, the homenet border router(s) should filter

packets with ULA source/destination addresses as discussed in Section 3.4.2.

Devices in a homenet may be given only a ULA as a means to restrict reachability from outside the homenet. ULAs can be used by default for devices that, without additional configuration (e.g., via a web interface), would only offer services to the internal network. For example, a printer might only accept incoming connections on a ULA until configured to be globally reachable, at which point it acquires a global IPv6 address and may be advertised via a global name space.

Where both a ULA and a global prefix are in use, the ULA source address is used to communicate with ULA destination addresses when appropriate, i.e., when the ULA source and destination lie within the /48 ULA prefix(es) known to be used within the same homenet. In cases where multiple /48 ULA prefixes are in use within a single homenet (perhaps because multiple homenet routers each independently auto-generate a /48 ULA prefix and then share prefix/routing information), utilising a ULA source address and a ULA destination address from two disjoint internal ULA prefixes is preferable to using GUAs.

While a homenet should operate correctly with two or more /48 ULAs enabled, a mechanism for the creation and use of a single /48 ULA prefix is desirable for addressing consistency and policy enforcement.

A counter argument to using ULAs is that it is undesirable to aggressively deprecate global prefixes for temporary loss of connectivity, so for a host to lose its global address, 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, it is assumed in this architecture that homenets should support and use ULAs.

2.5. Avoiding Manual Configuration of IP Addresses

Some IPv4 home networking devices expose IPv4 addresses to users, e.g., the IPv4 address of a home IPv4 CE router that may be configured via a web interface. In potentially complex future IPv6 homenets, users should not be expected to enter IPv6 literal addresses in devices or applications, given their much greater length and the apparent randomness of such addresses to a typical home user. Thus, even for the simplest of functions, simple naming and the associated (minimal, and ideally zero configuration) discovery of services are imperative for the easy deployment and use of homenet devices and applications.

2.6. IPv6-Only Operation

It is likely that IPv6-only networking will be deployed first in new home network deployments, often referred to as 'greenfield' scenarios, where there is no existing IPv4 capability, or perhaps as one element of an otherwise dual-stack network. Running IPv6-only adds additional requirements, e.g., for devices to get configuration information via IPv6 transport (not relying on an IPv4 protocol such as IPv4 DHCP) and for devices to be able to initiate communications to external devices that are IPv4-only.

Some specific transition technologies that may be deployed by the homenet's ISP are discussed in [RFC7084]. In addition, certain other functions may be desirable on the CE router, e.g., to access content in the IPv4 Internet, NAT64 [RFC6144] and DNS64 [RFC6145] may be applicable.

The widespread availability of robust solutions to these types of requirements will help accelerate the uptake of IPv6-only homenets. The specifics of these are, however, beyond the scope of this document, especially those functions that reside on the CE router.

3. Homenet Architecture Principles

The aim of this text is to outline how to construct advanced IPv6-based home networks involving multiple routers and subnets using standard IPv6 addressing and protocols [RFC2460] [RFC4291] as the basis. As described in Section 3.1, solutions should as far as possible reuse existing protocols and minimise changes to hosts and routers, but some new protocols or extensions are likely to be required. In this section, we present the elements of the proposed home networking architecture with discussion of the associated design principles.

In general, home network equipment needs to be able to operate in networks with a range of different properties and topologies, where home users may plug components together in arbitrary ways and expect the resulting network to operate. Significant manual configuration is rarely, if at all, possible or even desirable given the knowledge level of typical home users. Thus, the network should, as far as possible, be self-configuring, though configuration by advanced users should not be precluded.

The homenet needs to be able to handle or provision at least the following:

- o Routing
- o Prefix configuration for routers
- o Name resolution
- o Service discovery
- o Network security

The remainder of this document describes the principles by which the homenet architecture may deliver these properties.

3.1. General 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 internetworking mechanisms. This text discusses how such flexibility should be used to provide the best user experience and ensure that the network can evolve with new applications in the future. The principles described in this text should be followed when designing homenet protocol solutions.

3.1.1. Reuse Existing Protocols

Existing protocols will be used to meet the requirements of home networks. Where necessary, extensions will be made to those protocols. When no existing protocol is found to be suitable, a new or emerging protocol may be used. Therefore, it is important that no design or architectural decisions be made that would preclude the use of new or emerging protocols.

A generally conservative approach, giving weight to running (and available) 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 backward compatible and forward compatible where changes are made.

3.1.2. Minimise Changes to Hosts and Routers

In order to maximise the deployability of new homenets, any requirement for changes to hosts and routers should be minimised where possible; however, solutions that, for example, incrementally improve capability via host or router changes may be acceptable. There may be cases where changes are unavoidable, e.g., to allow a given homenet routing protocol to be self-configuring or to support routing based on source addresses in addition to destination addresses (to improve multihoming support, as discussed in Section 3.2.4).

3.2. Homenet Topology

This section considers homenet topologies and the principles that may be applied in designing an architecture to support as wide a range of such topologies as possible.

3.2.1. Supporting Arbitrary Topologies

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 and arbitrary routing will need to be supported, or at least the failure mode for when the user makes a mistake should be as robust as possible, e.g., deactivating a certain part of the infrastructure to allow the rest to operate. In such cases, the user should ideally have some useful indication of the failure mode encountered.

There should be no topology scenarios that cause a loss of connectivity, except when the user creates a physical island within the topology. Some potentially pathological cases that can be created include bridging ports of a router together; however, this case can be detected and dealt with by the router. Loops within a routed topology are in a sense good in that they offer redundancy. Topologies that include potential bridging loops can be dangerous but are also detectable when a switch learns the Media Access Control (MAC) address of one of its interfaces on another or runs a spanning tree or link-state protocol. It is only topologies with such potential loops using simple repeaters that are truly pathological.

The topology of the homenet may change over time, due to the addition or removal of equipment but also due to temporary failures or connectivity problems. In some cases, this may lead to, for example, a multihomed homenet being split into two isolated homenets or, after such a fault is remedied, two isolated parts reconfiguring back to a single network.

3.2.2. Network Topology Models

As hinted above, while the architecture may focus on likely common topologies, it should not preclude any arbitrary topology from being constructed.

At the time of writing, most IPv4 home network models tend to be relatively simple, typically a single NAT router to the ISP and a single internal subnet but, as discussed earlier, evolution in network architectures is driving more complex topologies, such as the separation of guest and private networks. There may also be some cascaded IPv4 NAT scenarios, which we mention in the next section. For IPv6 homenets, the network architectures described in [RFC7084] should, as a minimum, be supported.

There are a number of properties or attributes of a home network that we can use to describe its topology and operation. The following properties apply to any IPv6 home network:

- o Presence of internal routers. The homenet may have one or more internal routers or may only provide subnetting from interfaces on the CE router.
- o Presence of isolated internal subnets. There may be isolated internal subnets, with no direct connectivity between them within the homenet (with each having its own external connectivity). Isolation may be physical or implemented via IEEE 802.1q VLANs. The latter is, however, not something a typical user would be expected to configure.
- o Demarcation of the CE router. The CE router(s) may or may not be managed by the ISP. If the demarcation point is such that the customer can provide or manage the CE router, its configuration must be simple. Both models must be supported.

Various forms of multihoming are likely to become more prevalent with IPv6 home networks, where the homenet may have two or more external ISP connections, as discussed further below. Thus, the following properties should also be considered for such networks:

- o Number of upstream providers. The majority of home networks today consist of a single upstream ISP, but it may become more common in the future for there to be multiple ISPs, whether for resilience or provision of additional services. Each would offer its own prefix. Some may or may not provide a default route to the public Internet.

- o Number of CE routers. The homenet may have a single CE router, which might be used for one or more providers, or multiple CE routers. The presence of multiple CE routers adds additional complexity for multihoming scenarios and protocols like PCP that may need to manage connection-oriented state mappings on the same CE router as used for subsequent traffic flows.

In the following sections, we give some examples of the types of homenet topologies we may see in the future. This is not intended to be an exhaustive or complete list but rather an indicative one to facilitate the discussion in this text.

3.2.2.1. A: Single ISP, Single CE Router, and Internal Routers

Figure 1 shows a home network with multiple local area networks. These may be needed for reasons relating to different link-layer technologies in use or for policy reasons, e.g., classic Ethernet in one subnet and an LLN link-layer technology in another. In this example, there is no single router that a priori understands the entire topology. The topology itself may also be complex, and it may not be possible to assume a pure tree form, for instance (because home users may plug routers together to form arbitrary topologies, including those with potential loops in them).

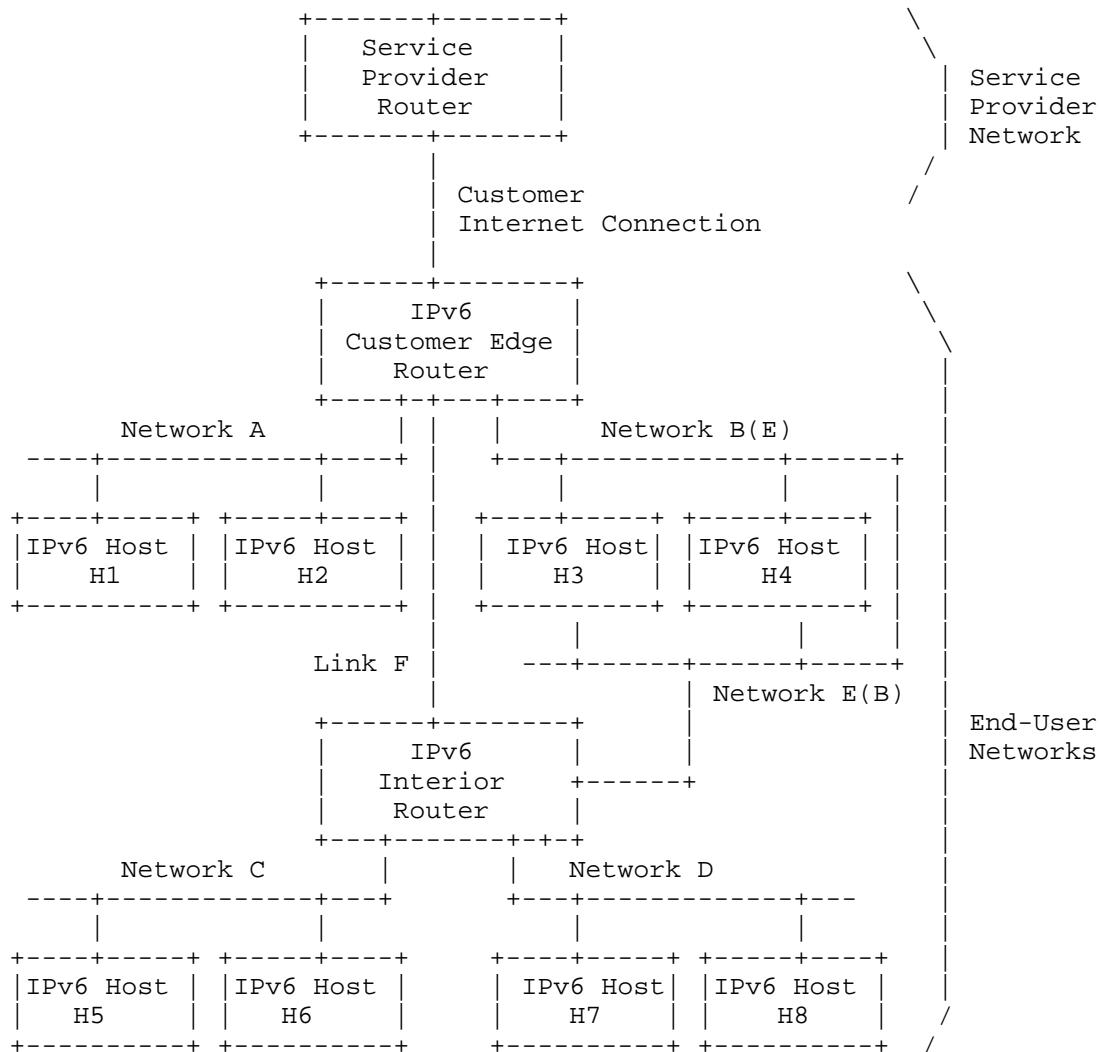


Figure 1

In this diagram, there is one CE router. It has a single uplink interface. It has three additional interfaces connected to Network A, Link F, and Network B. The IPv6 Internal Router (IR) has four interfaces connected to Link F, Network C, Network D, and Network E. Network B and Network E have been bridged, likely inadvertently. This could be as a result of connecting a wire between a switch for Network B and a switch for Network E.

Any of logical Networks A through F might be wired or wireless. Where multiple hosts are shown, this might be through one or more physical ports on the CE router or IPv6 (IR), wireless networks, or through one or more Ethernet switches that are Layer 2 only.

3.2.2.2. B: Two ISPs, Two CE Routers, and Shared Subnet

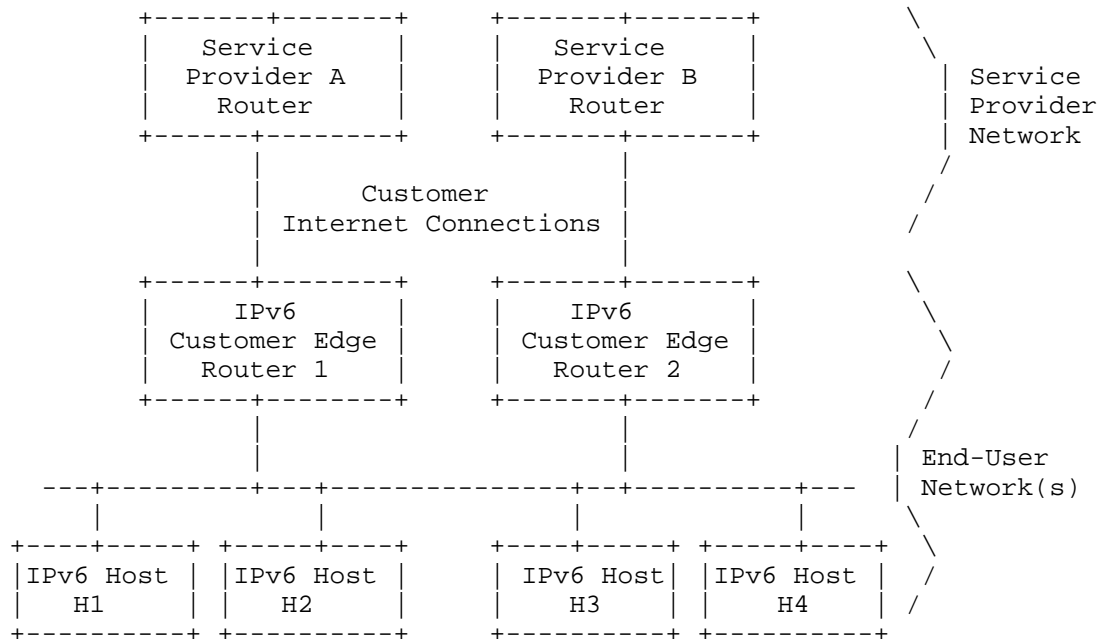


Figure 2

Figure 2 illustrates a multihomed homenet model, where the customer has connectivity via CE router 1 to ISP A and via CE router 2 to ISP B. This example shows one shared subnet where IPv6 nodes would potentially be multihomed and receive multiple IPv6 global prefixes, one per ISP. This model may also be combined with that shown in Figure 1 to create a more complex scenario with multiple internal routers. Or, the above shared subnet may be split in two, such that each CE router serves a separate isolated subnet, which is a scenario seen with some IPv4 networks today.

3.2.2.3. C: Two ISPs, One CE Router, and Shared Subnet

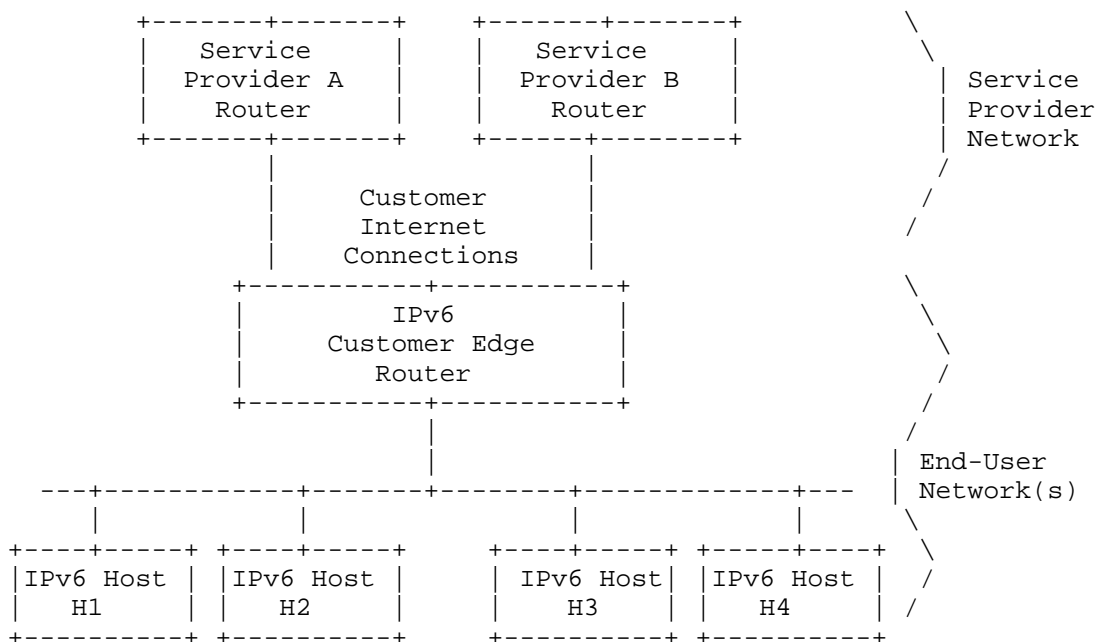


Figure 3

Figure 3 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.

3.2.3. Dual-Stack Topologies

For the immediate future, it is expected that most homenet deployments will be dual-stack IPv4/IPv6. In such networks, it is important not to introduce new IPv6 capabilities that would cause a failure if used alongside IPv4+NAT, given that such dual-stack homenets will be commonplace for some time. That said, it is desirable that IPv6 works better than IPv4 in as many scenarios as possible. Further, the homenet architecture must operate in the absence of IPv4.

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, routing or bridging may be used. Routing may have advantages when compared to bridging together high- and lower-speed shared media, and in

addition, bridging may not be suitable for some networks, such as ad hoc mobile networks.

In some cases, IPv4 home networks may feature cascaded NATs. End users are frequently unaware that they have created such networks, as 'home routers' and 'home switches' are frequently confused. In addition, there are cases where NAT routers are included within Virtual Machine Hypervisors or where Internet connection-sharing services have been enabled. This document applies equally to such hidden NAT 'routers'. IPv6-routed versions of such cases will be required. We should thus also note that routers in the homenet may not be separate physical devices; they may be embedded within other devices.

3.2.4. Multihoming

A homenet may be multihomed to multiple providers, as the network models above illustrate. This may take a form where there are either multiple isolated networks within the home or a more integrated network where the connectivity selection needs to be dynamic. Current practice is typically of the former kind, but the latter is expected to become more commonplace.

In the general homenet architecture, multihomed hosts should be multi-addressed with a global IPv6 address from the global prefix delegated from each ISP they communicate with or through. When such multi-addressing is in use, hosts need some way to pick source and destination address pairs for connections. A host may choose a source address to use by various methods, most commonly [RFC6724]. Applications may of course do different things, and this should not be precluded.

For the single CE Router Network Model C illustrated above, multihoming may be offered by source-based routing at the CE router. With multiple exit routers, as in CE Router Network Model B, the complexity rises. Given a packet with a source address on the home network, the packet must be routed to the proper egress to avoid ingress filtering as described in BCP 38 if exiting through the wrong ISP. It is highly desirable that the packet is routed in the most efficient manner to the correct exit, though as a minimum requirement the packet should not be dropped.

The homenet architecture should support both the above models, i.e., one or more CE routers. However, the general multihoming problem is broad, and solutions suggested to date within the IETF have included complex architectures for monitoring connectivity, traffic engineering, identifier-locator separation, connection survivability across multihoming events, and so on. It is thus important that the

homenet architecture should as far as possible minimise the complexity of any multihoming support.

An example of such a 'simpler' approach has been documented in [RFC7157]. Alternatively, a flooding/routing protocol could potentially be used to pass information through the homenet, such that internal routers and ultimately end hosts could learn per-prefix configuration information, allowing better address selection decisions to be made. However, this would imply router and, most likely, host changes. Another avenue is to introduce support throughout the homenet for routing that is based on the source as well as the destination address of each packet. While greatly improving the 'intelligence' of routing decisions within the homenet, such an approach would require relatively significant router changes but avoid host changes.

As explained previously, while NPTv6 has been proposed for providing multihoming support in networks, its use is not recommended in the homenet architecture.

It should be noted that some multihoming scenarios may see one upstream being a "walled garden" and thus only appropriate for connectivity to the services of that provider; an example may be a VPN service that only routes back to the enterprise business network of a user in the homenet. As per Section 4.2.1 of [RFC3002], we do not specifically target walled-garden multihoming as a goal of this document.

The homenet architecture should also not preclude use of host or application-oriented tools, e.g., Shim6 [RFC5533], Multipath TCP (MPTCP) [RFC6824], or Happy Eyeballs [RFC6555]. In general, any incremental improvements obtained by host changes should give benefit for the hosts introducing them but should not be required.

3.2.5. Mobility Support

Devices may be mobile within the homenet. While resident on the same subnet, their address will remain persistent, but should devices move to a different (wireless) subnet, they will acquire a new address in that subnet. It is desirable that the homenet supports internal device mobility. To do so, the homenet may either extend the reach of specific wireless subnets to enable wireless roaming across the home (availability of a specific subnet across the home) or support mobility protocols to facilitate such roaming where multiple subnets are used.

3.3. A Self-Organising Network

The home network infrastructure 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. At the same time, it should be possible for advanced users to manually adjust (override) the current configuration.

While a goal of the homenet architecture is for the network to be as self-organising as possible, there may be instances where some manual configuration is required, e.g., the entry of a cryptographic key to apply wireless security or to configure a shared routing secret. The latter may be relevant when considering how to bootstrap a routing configuration. It is highly desirable that the number of such configurations is minimised.

3.3.1. Differentiating Neighbouring Homenets

It is important that self-configuration with 'unintended' devices be avoided. There should be a way for a user to administratively assert in a simple way whether or not a device belongs to a given homenet. The goal is to allow the establishment of borders, particularly between two adjacent homenets, and to avoid unauthorised devices from participating in the homenet. Such an authorisation capability may need to operate through multiple hops in the homenet.

The homenet should thus support a way for a homenet owner to claim ownership of their devices in a reasonably secure way. This could be achieved by a pairing mechanism by, for example, pressing buttons simultaneously on an authenticated and a new homenet device or by an enrollment process as part of an autonomic networking environment.

While there may be scenarios where one homenet may wish to intentionally gain access through another, e.g., to share external connectivity costs, such scenarios are not discussed in this document.

3.3.2. Largest Practical 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. More recently, some vendors have started to introduce 'home' and 'guest' functions, which in IPv6 would be implemented as two subnets.

Future home networks are highly likely to have one or more internal routers and thus need multiple subnets for the reasons described earlier. As part of the self-organisation of the network, the

homenet should subdivide itself into the largest practical subnets that can be constructed within the constraints of link-layer mechanisms, bridging, physical connectivity, and policy, and where applicable, performance or other criteria. In such subdivisions, the logical topology may not necessarily match the physical topology. This text does not, however, make recommendations on how such subdivision should occur. It is expected that subsequent documents will address this problem.

While it may be desirable to maximise the chance of link-local protocols operating across a homenet by maximising the size of a subnet, multi-subnet home networks are inevitable, so their support must be included.

3.3.3. Handling Varying Link Technologies

Homenets tend to grow organically over many years, and a homenet will typically be built over link-layer technologies from different generations. Current homenets typically use links ranging from 1 Mbit/s up to 1 Gbit/s -- a throughput discrepancy of three orders of magnitude. We expect this discrepancy to widen further as both high-speed and low-power technologies are deployed.

Homenet protocols should be designed to deal well with interconnecting links of very different throughputs. In particular, flows local to a link should not be flooded throughout the homenet, even when sent over multicast, and, whenever possible, the homenet protocols should be able to choose the faster links and avoid the slower ones.

Links (particularly wireless links) may also have limited numbers of transmit opportunities (txops), and there is a clear trend driven by both power and downward compatibility constraints toward aggregation of packets into these limited txops while increasing throughput. Transmit opportunities may be a system's scarcest resource and, therefore, also strongly limit actual throughput available.

3.3.4. Homenet Realms and Borders

The homenet will need to be aware of the extent of its own 'site', which will, for example, define the borders for ULA and site scope multicast traffic and may require specific security policies to be applied. The homenet will have one or more such borders with external connectivity providers.

A homenet will most likely also have internal borders between internal realms, e.g., a guest realm or a corporate network extension realm. It is desirable that appropriate borders can be configured to

determine, for example, the scope of where network prefixes, routing information, network traffic, service discovery, and naming may be shared. The default mode internally should be to share everything.

It is expected that a realm would span at least an entire subnet, and thus the borders lie at routers that receive delegated prefixes within the homenet. It is also desirable, for a richer security model, that hosts are able to make communication decisions based on available realm and associated prefix information in the same way that routers at realm borders can.

A simple homenet model may just consider three types of realms and the borders between them, namely the internal homenet, the ISP, and a guest network. In this case, the borders will include the border from the homenet to the ISP, the border from the guest network to the ISP, and the border from the homenet to the guest network. Regardless, it should be possible for additional types of realms and borders to be defined, e.g., for some specific LLN-based network, such as Smart Grid, and for these to be detected automatically and for an appropriate default policy to be applied as to what type of traffic/data can flow across such borders.

It is desirable to classify the external border of the home network as a unique logical interface separating the home network from a 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 makes it possible to describe edge operations and interface requirements across multiple functional areas including security, routing, service discovery, and router discovery.

It should be possible for the homenet user to override any automatically determined borders and the default policies applied between them, the exception being that it may not be possible to override policies defined by the ISP at the external border.

3.3.5. Configuration Information from the ISP

In certain cases, it may be useful for the homenet to get certain configuration information from its ISP. For example, the homenet DHCP server may request and forward some options that it gets from its upstream DHCP server, though the specifics of the options may vary across deployments. There is potential complexity here, of course, should the homenet be multihomed.

3.4. Homenet Addressing

The IPv6 addressing scheme used within a homenet must conform to the IPv6 addressing architecture [RFC4291]. In this section, we discuss how the homenet needs to adapt to the prefixes made available to it by its upstream ISP, such that internal subnets, hosts, and devices can obtain and configure the necessary addressing information to operate.

3.4.1. Use of ISP-Delegated IPv6 Prefixes

Discussion of IPv6 prefix allocation policies is included in [RFC6177]. In practice, a homenet may receive an arbitrary length IPv6 prefix from its provider, e.g., /60, /56, or /48. The offered prefix may be stable or change from time to time; it is generally expected that ISPs will offer relatively stable prefixes to their residential customers. Regardless, the home network needs to be adaptable as far as possible to ISP prefix allocation policies and assume nothing about the stability of the prefix received from an ISP or the length of the prefix that may be offered.

However, if, for example, only a /64 is offered by the ISP, the homenet may be severely constrained or even unable to function. BCP 157 [RFC6177] states the following:

A key principle for address management is that end sites always be able to obtain a reasonable amount of address space for their actual and planned usage, and over time ranges specified in years rather than just months. In practice, that means at least one /64, and in most cases significantly more. One particular situation that must be avoided is having an end site feel compelled to use IPv6-to-IPv6 Network Address Translation or other burdensome address conservation techniques because it could not get sufficient address space.

This architecture document assumes that the guidance in the quoted text is being followed by ISPs.

There are many problems that would arise from a homenet not being offered a sufficient prefix size for its needs. Rather than attempt to contrive a method for a homenet to operate in a constrained manner when faced with insufficient prefixes, such as the use of subnet prefixes longer than /64 (which would break stateless address autoconfiguration [RFC4862]), the use of NPTv6, or falling back to bridging across potentially very different media, it is recommended that the receiving router instead enters an error state and issues appropriate warnings. Some consideration may need to be given to how

such a warning or error state should best be presented to a typical home user.

Thus, a homenet CE router should request, for example, via DHCP Prefix Delegation (DHCP PD) [RFC3633], that it would like a /48 prefix from its ISP, i.e., it asks the ISP for the maximum size prefix it might expect to be offered, even if in practice it may only be offered a /56 or /60. For a typical IPv6 homenet, it is not recommended that an ISP offers less than a /60 prefix, and it is highly preferable that the ISP offers at least a /56. It is expected that the allocated prefix to the homenet from any single ISP is a contiguous, aggregated one. While it may be possible for a homenet CE router to issue multiple prefix requests to attempt to obtain multiple delegations, such behaviour is out of scope of this document.

The norm for residential customers of large ISPs may be 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. It is not expected that ISPs will generally support Provider Independent (PI) addressing for residential homenets.

When an ISP does need 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 that, unlike the process described in [RFC4192], would be a 'flag day' event, which means that a graceful renumbering process moving through a state with two active prefixes in use would not be possible. While renumbering can be viewed as an extended version of an initial numbering process, the difference between flash renumbering and an initial 'cold start' is the need to provide service continuity.

There may be cases where local law means some ISPs are required to change IPv6 prefixes (current IPv4 addresses) for privacy reasons for their customers. In such cases, it may be possible to avoid an instant 'flash' renumbering and plan a non-flag day renumbering as per RFC 4192. Similarly, if an ISP has a planned renumbering process, it may be able to adjust lease timers, etc., appropriately.

The customer may of course also choose to move to a new ISP and thus begin using a new prefix. In such cases, the customer should expect a discontinuity, and not only may the prefix change, but potentially also the prefix length if the new ISP offers a different default size prefix. The homenet may also be forced to renumber itself if significant internal 'replumbing' is undertaken by the user.

Regardless, it's desirable that homenet protocols support rapid renumbering and that operational processes don't add unnecessary complexity for the renumbering process. Further, the introduction of any new homenet protocols should not make any form of renumbering any more complex than it already is.

Finally, the internal operation of the home network should also not depend on the availability of the ISP network at any given time, other than, of course, for connectivity to services or systems off the home network. This reinforces the use of ULAs for stable internal communication and the need for a naming and service discovery mechanism that can operate independently within the homenet.

3.4.2. Stable Internal IP Addresses

The network should by default attempt to provide IP-layer connectivity between all internal parts of the homenet as well as to and from the external Internet, subject to the filtering policies or other policy constraints discussed later in the security section.

ULAs should be used within the scope of a homenet to support stable routing and connectivity between subnets and hosts regardless of whether a globally unique ISP-provided prefix is available. In the case of a prolonged external connectivity outage, ULAs allow internal operations across routed subnets to continue. ULA addresses also allow constrained devices to create permanent relationships between IPv6 addresses, e.g., from a wall controller to a lamp, where symbolic host names would require additional non-volatile memory, and updating global prefixes in sleeping devices might also be problematic.

As discussed previously, it would be expected that ULAs would normally be used alongside one or more global prefixes in a homenet, such that hosts become multi-addressed with both globally unique and ULA prefixes. ULAs should be used for all devices, not just those intended to only have internal connectivity. Default address selection would then enable ULAs to be preferred for internal communications between devices that are using ULA prefixes generated within the same homenet.

In cases where ULA prefixes are in use within a homenet but there is no external IPv6 connectivity (and thus no GUAs in use), recommendations ULA-5, L-3, and L-4 in RFC 7084 should be followed to ensure correct operation, in particular where the homenet may be dual stack with IPv4 external connectivity. The use of the Route Information Option described in [RFC4191] provides a mechanism to advertise such more-specific ULA routes.

The use of ULAs should be restricted to the homenet scope through filtering at the border(s) of the homenet, as mandated by RFC 7084 requirement S-2.

Note that in some cases, it is possible that multiple /48 ULA prefixes may be in use within the same homenet, e.g., when the network is being deployed, perhaps also without external connectivity. In cases where multiple ULA /48s are in use, hosts need to know that each /48 is local to the homenet, e.g., by inclusion in their local address selection policy table.

3.4.3. Internal Prefix Delegation

As mentioned above, there are various sources of prefixes. From the homenet perspective, a single global prefix from each ISP should be received on the border CE router [RFC3633]. Where multiple CE routers 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. As discussed above, a ULA prefix should be provisioned for stable internal communications or for use on constrained/LLN networks.

The delegation or availability of a prefix pool to the homenet should allow subsequent internal autonomous assignment of prefixes for use within the homenet. Such internal assignment should not assume a flat or hierarchical model, nor should it make an assumption about whether the assignment of internal prefixes is distributed or centralised. 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, and not waste prefixes. Further, duplicate assignment of multiple /64s to the same network should be avoided, and the network should behave as gracefully as possible in the event of prefix exhaustion (though the options in such cases may be limited).

Where the home network has multiple CE routers and these are delegated prefix pools from their attached ISPs, the internal prefix assignment would be expected to be served by each CE router for each prefix associated with it. Where ULAs are used, it is preferable that only one /48 ULA covers the whole homenet, from which /64s can be assigned to the subnets. In cases where two /48 ULAs are generated within a homenet, the network should still continue to function, meaning that hosts will need to determine that each ULA is local to the homenet.

Prefix assignment within the homenet should result in each link being assigned a stable prefix that is persistent across reboots, power outages, and similar short-term outages. The availability of

persistent prefixes should not depend on the router boot order. The 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. However, assigned ULA prefixes within the homenet should remain persistent through an ISP-driven renumbering event.

Provisioning such 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).

This document makes no specific recommendation towards solutions but notes that it is very likely that all routing devices participating in a homenet must use the same internal prefix delegation method. This implies that only one delegation method should be in use.

3.4.4. Coordination of Configuration Information

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

3.4.5. Privacy

If ISPs offer relatively stable IPv6 prefixes to customers, the network prefix part of addresses associated with the homenet may not change over a reasonably long period of time.

The exposure of which traffic is sourced from the same homenet is thus similar to IPv4; the single IPv4 global address seen through use of IPv4 NAT gives the same hint as the global IPv6 prefix seen for IPv6 traffic.

While IPv4 NAT may obfuscate to an external observer which internal devices traffic is sourced from, IPv6, even with use of privacy addresses [RFC4941], adds additional exposure of which traffic is sourced from the same internal device through use of the same IPv6 source address for a period of time.

3.5. Routing Functionality

Routing functionality is required when there are multiple routers deployed within the internal home network. This functionality could be as simple as the current 'default route is up' model of IPv4 NAT,

or more likely, it would involve running an appropriate routing protocol.

A mechanism is required to discover which router(s) in the homenet is providing the CE router function. Borders may include but are not limited to the interface to the upstream ISP, a gateway device to a separate home network such as an LLN network, or a gateway to a guest or private corporate extension network. In some cases, there may be no border present, which may, for example, occur before an upstream connection has been established.

The routing environment should be self-configuring, as discussed previously. The homenet self-configuration process and the routing protocol must interact in a predictable manner, especially during startup and reconvergence. The border discovery functionality and other self-configuration functionality may be integrated into the routing protocol itself but may also be imported via a separate discovery mechanism.

It is preferable that configuration information is distributed and synchronised within the homenet by a separate configuration protocol.

The homenet routing protocol should be based on a previously deployed protocol that has been shown to be reliable and robust. This does not preclude the selection of a newer protocol for which a high-quality open source implementation becomes available. The resulting code must support lightweight implementations and be suitable for incorporation into consumer devices, where both fixed and temporary storage and processing power are at a premium.

At most, one unicast and one multicast routing protocol should be in use at a given time in a given homenet. In some simple topologies, no routing protocol may be needed. If more than one routing protocol is supported by routers in a given homenet, then a mechanism is required to ensure that all routers in that homenet use the same protocol.

The homenet architecture is IPv6-only. In practice, dual-stack homenets are still likely for the foreseeable future, as described in Section 3.2.3. Whilst support for IPv4 and other address families may therefore be beneficial, it is not an explicit requirement to carry the routing information in the same routing protocol.

Multiple types of physical interfaces must be accounted for in the homenet routing topology. Technologies such as Ethernet, Wi-Fi, Multimedia over Coax Alliance (MoCA), etc., must be capable of coexisting in the same environment and should be treated as part of any routed deployment. The inclusion of physical-layer

characteristics in path computation should be considered for optimising communication in the homenet.

3.5.1. Unicast Routing within the Homenet

The role of the unicast routing protocol is to provide good enough end-to-end connectivity often enough, where good/often enough is defined by user expectations.

Due to the use of a variety of diverse underlying link technologies, path selection in a homenet may benefit from being more refined than minimising hop count. It may also be beneficial for traffic to use multiple paths to a given destination within the homenet where available rather than just a single best path.

Minimising convergence time should be a goal in any routed environment. It is reasonable to assume that convergence time should not be significantly longer than network outages users are accustomed to should their CE router reboot.

The homenet architecture is agnostic as to the choice of underlying routing technology, e.g., link state versus Bellman-Ford.

The routing protocol should support the generic use of multiple customer Internet connections and the concurrent use of multiple delegated prefixes. A routing protocol that can make routing decisions based on source and destination addresses is thus highly desirable, to avoid problems with upstream ISP ingress filtering as described in BCP 38. Multihoming support may also include load balancing to multiple providers and failover from a primary to a backup link when available. The protocol should not require upstream ISP connectivity to be established to continue routing within the homenet.

The homenet architecture is agnostic on a minimum hop count that has to be supported by the routing protocol. The architecture should, however, be scalable to other scenarios where homenet technology may be deployed, which may include small office and small enterprise sites. To allow for such cases, it would be desirable that the architecture is scalable to higher hop counts and to larger numbers of routers than would be typical in a true home network.

At the time of writing, link-layer networking technology is poised to become more heterogeneous, as networks begin to employ both traditional Ethernet technology and link layers designed for LLNs, such as those used for certain types of sensor devices.

Ideally, LLN or other logically separate networks should be able to exchange routes such that IP traffic may be forwarded among the networks via gateway routers that interoperate with both the homenet and any LLNs. Current home deployments use largely different mechanisms in sensor and basic Internet connectivity networks. IPv6 virtual machine (VM) solutions may also add additional routing requirements.

In this homenet architecture, LLNs and other specialised networks are considered stub areas of the homenet and are thus not expected to act as a transit for traffic between more traditional media.

3.5.2. Unicast Routing at the Homenet Border

The current practice defined in [RFC7084] would suggest that routing between the homenet CE router and the service provider router follow the WAN-side requirements model in [RFC7084], Section 4 (WAN-side requirements), at least in initial deployments. However, consideration of whether a routing protocol is used between the homenet CE router and the service provider router is out of scope of this document.

3.5.3. Multicast Support

It is desirable that, subject to the capacities of devices on certain media types, multicast routing is supported across the homenet, including source-specific multicast (SSM) [RFC4607].

[RFC4291] requires that any boundary of scope 4 or higher (i.e., admin-local or higher) be administratively configured. Thus, the boundary at the homenet-ISP border must be administratively configured, though that may be triggered by an administrative function such as DHCP PD. Other multicast forwarding policy borders may also exist within the homenet, e.g., to/from a guest subnet, whilst the use of certain link media types may also affect where specific multicast traffic is forwarded or routed.

There may be different drivers for multicast to be supported across the homenet -- for example,

- o for homenet-wide service discovery, should a multicast service discovery protocol of scope greater than link-local be defined
- o for multicast-based streaming or file-sharing applications

Where multicast is routed across a homenet, an appropriate multicast routing protocol is required, one that as per the unicast routing protocol should be self-configuring. As hinted above, it must be

possible to scope or filter multicast traffic to avoid it being flooded to network media where devices cannot reasonably support it.

A homenet may not only use multicast internally, it may also be a consumer or provider of external multicast traffic, where the homenet's ISP supports such multicast operation. This may be valuable, for example, where live video applications are being sourced to/from the homenet.

The multicast environment should support the ability for applications to pick a unique multicast group to use.

3.6. Security

The security of an IPv6 homenet is an important consideration. The most notable difference to the IPv4 operational model is the removal of NAT, the introduction of global addressability of devices, and thus a need to consider whether devices should have global reachability. Regardless, hosts need to be able to operate securely, end to end where required, and also be robust against malicious traffic directed towards them. However, there are other challenges introduced, e.g., default filtering policies at the borders between various homenet realms.

3.6.1. Addressability vs. Reachability

An IPv6-based home network architecture should embrace the transparent end-to-end communications model as described in [RFC2775]. Each device should be globally addressable, and those addresses must not be altered in transit. However, security perimeters can be applied to restrict end-to-end communications, and thus while a host may be globally addressable, it may not be globally reachable.

[RFC4864] describes a 'Simple Security' model for IPv6 networks, whereby stateful perimeter filtering can be applied to control the reachability of devices in a homenet. RFC 4864 states in Section 4.2 that "the use of firewalls...is recommended for those that want boundary protection in addition to host defences." It should be noted that a 'default deny' filtering approach would effectively replace the need for IPv4 NAT traversal protocols with a need to use a signalling protocol to request a firewall hole be opened, e.g., a protocol such as PCP [RFC6887]. In networks with multiple CE routers, the signalling would need to handle the cases of flows that may use one or more exit routers. CE routers would need to be able to advertise their existence for such protocols.

[RFC6092] expands on RFC 4864, giving a more detailed discussion of IPv6 perimeter security recommendations, without mandating a 'default deny' approach. Indeed, RFC 6092 does not enforce a particular mode of operation, instead stating that CE routers must provide an easily selected configuration option that permits a 'transparent' mode, thus ensuring a 'default allow' model is available.

The topic of whether future home networks as described in this document should have a 'default deny' or 'default allow' position has been discussed at length in various IETF meetings without any consensus being reached on which approach is more appropriate. Further, the choice of which default to apply may be situational, and thus this text makes no recommendation on the default setting beyond what is written on this topic in RFC 6092. We note in Section 3.6.3 below that the implicit firewall function of an IPv4 NAT is commonplace today, and thus future CE routers targeted at home networks should continue to support the option of running in 'default deny mode', whether or not that is the default setting.

3.6.2. Filtering at Borders

It is desirable that there are mechanisms to detect different types of borders within the homenet, as discussed previously, and further mechanisms to then apply different types of filtering policies at those borders, e.g., whether naming and service discovery should pass a given border. Any such policies should be able to be easily applied by typical home users, e.g., to give a user in a guest network access to media services in the home or access to a printer. Simple mechanisms to apply policy changes, or associations between devices, will be required.

There are cases where full internal connectivity may not be desirable, e.g., in certain utility networking scenarios, or where filtering is required for policy reasons against a guest network subnet(s). As a result, some scenarios/models may involve running an isolated subnet(s) with their own CE routers. In such cases, connectivity would only be expected within each isolated network (though traffic may potentially pass between them via external providers).

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

Considerations for differentiating neighbouring homenets are discussed in Section 3.3.1.

3.6.3. Partial Effectiveness of NAT and Firewalls

Security by way of obscurity (address translation) or through firewalls (filtering) is at best only partially effective. The very poor security track record of home computers, home networking, and business PC computers and networking is testimony to this. A security compromise behind the firewall of any device exposes all others, making an entire network that relies on obscurity or a firewall as vulnerable as the most insecure device on the private side of the network.

However, given current evidence of home network products with very poor default device security, putting a firewall in place does provide some level of protection. The use of firewalls today, whether a good practice or not, is common practice, and the capability to afford protection via a 'default deny' setting, even if marginally effective, should not be lost. Thus, while it is highly desirable that all hosts in a homenet be adequately protected by built-in security functions, it should also be assumed that all CE routers will continue to support appropriate perimeter defence functions, as per [RFC7084].

3.6.4. Exfiltration Concerns

As homenets become more complex, with more devices, and with service discovery potentially enabled across the whole home, there are potential concerns over the leakage of information should devices use discovery protocols to gather information and report it to equipment vendors or application service providers.

While it is not clear how such exfiltration could be easily avoided, the threat should be recognised, be it from a new piece of hardware or some 'app' installed on a personal device.

3.6.5. Device Capabilities

In terms of the devices, homenet hosts should implement their own security policies in accordance to their computing capabilities. They should have the means to request transparent communications that can be initiated to them through security filters in the homenet, for either all ports or specific services. Users should have simple methods to associate devices to services that they wish to operate transparently through (CE router) borders.

3.6.6. ULAs as a Hint of Connection Origin

As noted in Section 3.6, if appropriate filtering is in place on the CE router(s), as mandated by requirement S-2 in RFC 7084, a ULA source address may be taken as an indication of locally sourced traffic. This indication could then be used with security settings to designate between which nodes a particular application is allowed to communicate, provided ULA address space is filtered appropriately at the boundary of the realm.

3.7. Naming and Service Discovery

The homenet requires devices to be able to determine and use unique names by which they can be accessed on the network and that are not used by other devices on the network. Users and devices will need to be able to discover devices and services available on the network, e.g., media servers, printers, displays, or specific home automation devices. Thus, naming and service discovery must be supported in the homenet, and given the nature of typical home network users, the service(s) providing this function must as far as possible support unmanaged operation.

The naming system will be required to work internally or externally, whether the user is within or outside of the homenet, i.e., the user should be able to refer to devices by name, and potentially connect to them, wherever they may be. The most natural way to think about such naming and service discovery is to enable it to work across the entire homenet residence (site), disregarding technical borders such as subnets but respecting policy borders such as those between guest and other internal network realms. Remote access may be desired by the homenet residents while travelling but also potentially by manufacturers or other 'benevolent' third parties.

3.7.1. Discovering Services

Users will typically perform service discovery through graphical user interfaces (GUIs) that allow them to browse services on their network in an appropriate and intuitive way. Devices may also need to discover other devices, without any user intervention or choice. Either way, such interfaces are beyond the scope of this document, but the interface should have an appropriate application programming interface (API) for the discovery to be performed.

Such interfaces may also typically hide the local domain name element from users, especially where only one name space is available. However, as we discuss below, in some cases the ability to discover available domains may be useful.

We note that current zero-configuration service discovery protocols are generally aimed at single subnets. There is thus a choice to make for multi-subnet homenet as to whether such protocols should be proxied or extended to operate across a whole homenet. In this context, that may mean bridging a link-local method, taking care to avoid packets entering looping paths, or extending the scope of multicast traffic used for the purpose. It may mean that some proxy or hybrid service is utilised, perhaps co-resident on the CE router. Or, it may be that a new approach is preferable, e.g., flooding information around the homenet as attributes within the routing protocol (which could allow per-prefix configuration). However, we should prefer approaches that are backward compatible and allow current implementations to continue to be used. Note that this document does not mandate a particular solution; rather, it expresses the principles that should be used for a homenet naming and service discovery environment.

One of the primary challenges facing service discovery today is lack of interoperability due to the ever increasing number of service discovery protocols available. While it is conceivable for consumer devices to support multiple discovery protocols, this is clearly not the most efficient use of network and computational resources. One goal of the homenet architecture should be a path to service discovery protocol interoperability through either a standards-based translation scheme, hooks into current protocols to allow some form of communication among discovery protocols, extensions to support a central service repository in the homenet, or simply convergence towards a unified protocol suite.

3.7.2. Assigning Names to Devices

Given the large number of devices that may be networked in the future, devices should have a means to generate their own unique names within a homenet and to detect clashes should they arise, e.g., where a second device of the same type/vendor as an existing device with the same default name is deployed or where a new subnet is added to the homenet that already has a device of the same name. It is expected that a device should have a fixed name while within the scope of the homenet.

Users will also want simple ways to (re)name devices, again most likely through an appropriate and intuitive interface that is beyond the scope of this document. Note that the name a user assigns to a device may be a label that is stored on the device as an attribute of the device, and it may be distinct from the name used in a name service, e.g., 'Study Laser Printer' as opposed to printer2.<somedomain>.

3.7.3. The Homenet Name Service

The homenet name service should support both lookups and discovery. A lookup would operate via a direct query to a known service, while discovery may use multicast messages or a service where applications register in order to be found.

It is highly desirable that the homenet name service must at the very least coexist with the Internet name service. There should also be a bias towards proven, existing solutions. The strong implication is thus that the homenet service is DNS based, or DNS compatible. There are naming protocols that are designed to be configured and operate Internet-wide, like unicast-based DNS, but also protocols that are designed for zero-configuration local environments, like Multicast DNS (mDNS) [RFC6762].

When DNS is used as the homenet name service, it typically includes both a resolving service and an authoritative service. The authoritative service hosts the homenet-related zone. One approach when provisioning such a name service, which is designed to facilitate name resolution from the global Internet, is to run an authoritative name service on the CE router and a secondary authoritative name service provided by the ISP or perhaps an external third party.

Where zero-configuration name services are used, it is desirable that these can also coexist with the Internet name service. In particular, where the homenet is using a global name space, it is desirable that devices have the ability, where desired, to add entries to that name space. There should also be a mechanism for such entries to be removed or expired from the global name space.

To protect against attacks such as cache poisoning, where an attacker is able to insert a bogus DNS entry in the local cache, it is desirable to support appropriate name service security methods, including DNS Security Extensions (DNSSEC) [RFC4033], on both the authoritative server and the resolver sides. Where DNS is used, the homenet router or naming service must not prevent DNSSEC from operating.

While this document does not specify hardware requirements, it is worth noting briefly here that, e.g., in support of DNSSEC, appropriate homenet devices should have good random number generation capability, and future homenet specifications should indicate where high-quality random number generators, i.e., with decent entropy, are needed.

Finally, the impact of a change in the CE router must be considered. It would be desirable to retain any relevant state (configuration) that was held in the old CE router. This might imply that state information should be distributed in the homenet, to be recoverable by/to the new CE router, or to the homenet's ISP or a third-party externally provided service by some means.

3.7.4. Name Spaces

If access to homenet devices is required remotely from anywhere on the Internet, then at least one globally unique name space is required, though the use of multiple name spaces should not be precluded. One approach is that the name space(s) used for the homenet would be served authoritatively by the homenet, most likely by a server resident on the CE router. Such name spaces may be acquired by the user or provided/generated by their ISP or an alternative externally provided service. It is likely that the default case is that a homenet will use a global domain provided by the ISP, but advanced users wishing to use a name space that is independent of their provider in the longer term should be able to acquire and use their own domain name. For users wanting to use their own independent domain names, such services are already available.

Devices may also be assigned different names in different name spaces, e.g., by third parties who may manage systems or devices in the homenet on behalf of the resident(s). Remote management of the homenet is out of scope of this document.

If, however, a global name space is not available, the homenet will need to pick and use a local name space, which would only have meaning within the local homenet (i.e., it would not be used for remote access to the homenet). The .local name space currently has a special meaning for certain existing protocols that have link-local scope and is thus not appropriate for multi-subnet home networks. A different name space is thus required for the homenet.

One approach for picking a local name space is to use an Ambiguous Local Qualified Domain Name (ALQDN) space, such as .sitelocal (or an appropriate name reserved for the purpose). While this is a simple approach, there is the potential in principle for devices that are bookmarked somehow by name by an application in one homenet to be confused with a device with the same name in another homenet. In practice, however, the underlying service discovery protocols should be capable of handling moving to a network where a new device is using the same name as a device used previously in another homenet.

An alternative approach for a local name space would be to use a Unique Locally Qualified Domain Name (ULQDN) space such as `.<UniqueString>.sitelocal`. The `<UniqueString>` could be generated in a variety of ways, one potentially being based on the local /48 ULA prefix being used across the homenet. Such a `<UniqueString>` should survive a cold restart, i.e., be consistent after a network power-down, or if a value is not set on startup, the CE router or device running the name service should generate a default value. It would be desirable for the homenet user to be able to override the `<UniqueString>` with a value of their choice, but that would increase the likelihood of a name conflict. Any generated `<UniqueString>` should not be predictable; thus, adding a salt/hash function would be desirable.

In the (likely) event that the homenet is accessible from outside the homenet (using the global name space), it is vital that the homenet name space follow the rules and conventions of the global name space. In this mode of operation, names in the homenet (including those automatically generated by devices) must be usable as labels in the global name space. [RFC5890] describes considerations for Internationalizing Domain Names in Applications (IDNA).

Also, with the introduction of new 'dotless' top-level domains, there is also potential for ambiguity between, for example, a local host called 'computer' and (if it is registered) a .computer Generic Top Level Domain (gTLD). Thus, qualified names should always be used, whether these are exposed to the user or not. The IAB has issued a statement that explains why dotless domains should be considered harmful [IABdotless].

There may be use cases where different name spaces may be desired for either different realms in the homenet or segmentation of a single name space within the homenet. Thus, hierarchical name space management is likely to be required. There should also be nothing to prevent an individual device(s) from being independently registered in external name spaces.

It may be the case that if there are two or more CE routers serving the home network, if each has a name space delegated from a different ISP, there is the potential for devices in the home to have multiple fully qualified names under multiple domains.

Where a user is in a remote network wishing to access devices in their home network, there may be a requirement to consider the domain search order presented where multiple associated name spaces exist. This also implies that a domain discovery function is desirable.

It may be the case that not all devices in the homenet are made available by name via an Internet name space, and that a 'split view' (as described in [RFC6950], Section 4) is preferred for certain devices, whereby devices inside the homenet see different DNS responses to those outside.

Finally, this document makes no assumption about the presence or omission of a reverse lookup service. There is an argument that it may be useful for presenting logging information to users with meaningful device names rather than literal addresses. There are also some services, most notably email mail exchangers, where some operators have chosen to require a valid reverse lookup before accepting connections.

3.7.5. Independent Operation

Name resolution and service discovery for reachable devices must continue to function if the local network is disconnected from the global Internet, e.g., a local media server should still be available even if the Internet link is down for an extended period. This implies that the local network should also be able to perform a complete restart in the absence of external connectivity and have local naming and service discovery operate correctly.

As described above, the approach of a local authoritative name service with a cache would allow local operation for sustained ISP outages.

Having an independent local trust anchor is desirable, to support secure exchanges should external connectivity be unavailable.

A change in ISP should not affect local naming and service discovery. However, if the homenet uses a global name space provided by the ISP, then this will obviously have an impact if the user changes their network provider.

3.7.6. Considerations for LLNs

In some parts of the homenet, in particular LLNs or any devices where battery power is used, devices may be sleeping, in which case a proxy for such nodes may be required that could respond (for example) to multicast service discovery requests. Those same devices or parts of the network may have less capacity for multicast traffic that may be flooded from other parts of the network. In general, message utilisation should be efficient considering the network technologies and constrained devices that the service may need to operate over.

There are efforts underway to determine naming and discovery solutions for use by the Constrained Application Protocol (CoAP) [RFC7252] in LLN networks. These are outside the scope of this document.

3.7.7. DNS Resolver Discovery

Automatic discovery of a name service to allow client devices in the homenet to resolve external domains on the Internet is required, and such discovery must support clients that may be a number of router hops away from the name service. Similarly, it may be desirable to convey any DNS domain search list that may be in effect for the homenet.

3.7.8. Devices Roaming to/from the Homenet

It is likely that some devices that have registered names within the homenet Internet name space and that are mobile will attach to the Internet at other locations and acquire an IP address at those locations. Devices may move between different homenets. In such cases, it is desirable that devices may be accessed by the same name as is used in their home network.

Solutions to this problem are not discussed in this document. They may include the use of Mobile IPv6 or Dynamic DNS -- either of which would put additional requirements on the homenet -- or establishment of a (VPN) tunnel to a server in the home network.

3.8. Other Considerations

This section discusses two other considerations for home networking that the architecture should not preclude but that this text is neutral towards.

3.8.1. Quality of Service

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

However, homenet scenarios should require no new QoS protocols. A Diffserv [RFC2475] approach with a small number of predefined traffic classes may generally be sufficient, though at present there is little experience of QoS deployment in home networks. It is likely that QoS, or traffic prioritisation, methods will be required at the

CE router and potentially around boundaries between different link media types (where, for example, some traffic may simply not be appropriate for some media and need to be dropped to avoid overloading the constrained media).

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

3.8.2. Operations and Management

In this section, we briefly review some initial considerations for operations and management in the type of homenet described in this document. It is expected that a separate document will define an appropriate operations and management framework for such homenets.

As described in this document, the homenet should have the general goal of being self-organising and self-configuring from the network-layer perspective, e.g., prefixes should be able to be assigned to router interfaces. Further, applications running on devices should be able to use zero-configuration service discovery protocols to discover services of interest to the home user. In contrast, a home user would not be expected, for example, to have to assign prefixes to links or manage the DNS entries for the home network. Such expert operation should not be precluded, but it is not the norm.

The user may still be required to, or wish to, perform some configuration of the network and the devices on it. Examples might include entering a security key to enable access to their wireless network or choosing to give a 'friendly name' to a device presented to them through service discovery. Configuration of link- and application-layer services is out of scope of this architectural principles document but is likely to be required in an operational homenet.

While not being expected to actively configure the networking elements of their homenet, users may be interested in being able to view the status of their networks and the devices connected to it, in which case appropriate network monitoring protocols will be required to allow them to view their network, and its status, e.g., via a web interface or equivalent. While the user may not understand how the network operates, it is reasonable to assume they are interested in understanding what faults or problems may exist on it. Such monitoring may extend to other devices on the network, e.g., storage devices or web cameras, but such devices are beyond the scope of this document.

It may also be the case that an ISP, or a third party, might wish to offer a remote management service for the homenet on behalf of the user, or to be able to assist the user in the event of some problem they are experiencing, in which case appropriate management and monitoring protocols would be required.

Specifying the required protocols to facilitate homenet management and monitoring is out of scope of this document. As stated above, it is expected that a separate document will be produced to describe the operations and management framework for the types of home networks presented in this document.

As a final point, we note that it is desirable that all network management and monitoring functions should be available over IPv6 transport, even where the homenet is dual stack.

3.9. Implementing the Architecture on IPv6

This architecture text encourages reuse of existing protocols. Thus, 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 solutions can be found based on existing protocols. Some relatively smaller updates are likely to be required, e.g., a new mechanism may be needed in order to turn a selected protocol on by default, or a mechanism may be required to automatically assign prefixes to links within the homenet.

Some functionality, if required by the architecture, may need more significant changes or require development of new protocols, e.g., support for multihoming with multiple exit routers would likely require extensions to support source and destination address-based routing within the homenet.

Some protocol changes are, however, required in the architecture, e.g., for name resolution and service discovery, extensions to existing zero-configuration link-local name resolution protocols are needed to enable them to work across subnets, within the scope of the home network site.

Some of the hardest problems in developing solutions for home networking IPv6 architectures include discovering the right borders where the 'home' domain ends and the service provider domain begins, deciding whether some of the 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. Conclusions

This text defines principles and requirements for a homenet architecture. The principles and requirements documented here should be observed by any future texts describing homenet protocols for routing, prefix management, security, naming, or service discovery.

5. Security Considerations

Security considerations for the homenet architecture are discussed in Section 3.6 above.

6. References

6.1. Normative References

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005, <<http://www.rfc-editor.org/info/rfc4193>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.

6.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996, <<http://www.rfc-editor.org/info/rfc1918>>.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998, <<http://www.rfc-editor.org/info/rfc2475>>.

- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000, <<http://www.rfc-editor.org/info/rfc2775>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3002] Mitzel, D., "Overview of 2000 IAB Wireless Internetworking Workshop", RFC 3002, December 2000, <<http://www.rfc-editor.org/info/rfc3002>>.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001, <<http://www.rfc-editor.org/info/rfc3022>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005, <<http://www.rfc-editor.org/info/rfc4033>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005, <<http://www.rfc-editor.org/info/rfc4191>>.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005, <<http://www.rfc-editor.org/info/rfc4192>>.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006, <<http://www.rfc-editor.org/info/rfc4607>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007, <<http://www.rfc-editor.org/info/rfc4864>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.

- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009, <<http://www.rfc-editor.org/info/rfc5533>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, August 2010, <<http://www.rfc-editor.org/info/rfc5890>>.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010, <<http://www.rfc-editor.org/info/rfc5969>>.
- [RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011, <<http://www.rfc-editor.org/info/rfc6092>>.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011, <<http://www.rfc-editor.org/info/rfc6144>>.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011, <<http://www.rfc-editor.org/info/rfc6145>>.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, March 2011, <<http://www.rfc-editor.org/info/rfc6177>>.
- [RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011, <<http://www.rfc-editor.org/info/rfc6204>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011, <<http://www.rfc-editor.org/info/rfc6296>>.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011, <<http://www.rfc-editor.org/info/rfc6333>>.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012, <<http://www.rfc-editor.org/info/rfc6555>>.

- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013, <<http://www.rfc-editor.org/info/rfc6762>>.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013, <<http://www.rfc-editor.org/info/rfc6824>>.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013, <<http://www.rfc-editor.org/info/rfc6887>>.
- [RFC6950] Peterson, J., Kolkman, O., Tschofenig, H., and B. Aboba, "Architectural Considerations on Application Features in the DNS", RFC 6950, October 2013, <<http://www.rfc-editor.org/info/rfc6950>>.
- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, November 2013, <<http://www.rfc-editor.org/info/rfc7084>>.
- [RFC7157] Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation", RFC 7157, March 2014, <<http://www.rfc-editor.org/info/rfc7157>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.
- [IABdotless] IAB, "IAB Statement: Dotless Domains Considered Harmful", February 2013, <<http://www.iab.org/documents/correspondence-reports-documents/2013-2/iab-statement-dotless-domains-considered-harmful>>.
- [Gettys11] Gettys, J., "Bufferbloat: Dark Buffers in the Internet", March 2011, <<http://www.ietf.org/proceedings/80/slides/tsvarea-1.pdf>>.

Acknowledgments

The authors would like to thank Mikael Abrahamsson, Aamer Akhter, Mark Andrews, Dmitry Anipko, Ran Atkinson, Fred Baker, Ray Bellis, Teco Boot, John Brzozowski, Cameron Byrne, Brian Carpenter, Stuart Cheshire, Julius Chroboczek, Lorenzo Colitti, Robert Cragie, Elwyn Davies, Ralph Droms, Lars Eggert, Jim Gettys, Olafur Gudmundsson, Wassim Haddad, Joel M. Halpern, David Harrington, Lee Howard, Ray Hunter, Joel Jaeggli, Heather Kirksey, Ted Lemon, Acee Lindem, Kerry Lynn, Daniel Migault, Erik Nordmark, Michael Richardson, Mattia Rossi, Barbara Stark, Sander Steffann, Markus Stenberg, Don Sturek, Andrew Sullivan, Dave Taht, Dave Thaler, Michael Thomas, Mark Townsley, JP Vasseur, Curtis Villamizar, Russ White, Dan Wing, and James Woodyatt for their comments and contributions within homenet WG meetings and on the WG mailing list. An acknowledgment generally means that a person's text made it into the document or was helpful in clarifying or reinforcing an aspect of the document. It does not imply that each contributor agrees with every point in the document.

Authors' Addresses

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

EEmail: tjc@ecs.soton.ac.uk

Jari Arkko
Ericsson
Jorvas 02420
Finland

EEmail: jari.arkko@piuha.net

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

EEmail: anders_brandt@sigmadesigns.com

Ole Troan
Cisco Systems, Inc.
Philip Pedersensvei 1
Lysaker, N-1325
Norway

EEmail: ot@cisco.com

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

EEmail: jason.weil@twcable.com