Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 14, 2013

T. Chown, Ed.
University of Southampton
J. Arkko
Ericsson
A. Brandt
Sigma Designs
O. Troan
Cisco Systems, Inc.
J. Weil
Time Warner Cable
February 10, 2013

# Home Networking Architecture for IPv6
## draft-ietf-homenet-arch-07

Abstract

This text describes evolving networking technology within increasingly large residential home networks. 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

Table of Contents

## 1. Introduction

This document focuses on evolving networking technology within increasingly large residential home networks and 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, others to the introduction of specialised networks for home automation and sensors.

While at the time of writing some complex home network topologies exist, but most are relatively simple single subnet networks, and ostensibly operate using just IPv4 (there may be IPv6 traffic within the network, e.g. for service discovery, but the homenet is provisioned by the ISP as an IPv4 network). However, they also typically employ solutions that we would like to avoid 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 support private and guest networks, and have enough address space to allow every device to have a globally unique address. Thus there are likely to be scenarios where internal routing is required, in which case such networks require methods for IPv6 prefixes to be delegated to those subnets. It is not practical to expect home users to configure such prefixes, thus the assumption of this document is that the homenet is as far as possible self-organising and self-configuring, i.e. it need not be pro-actively managed 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 a better result than if the IETF had not given this specific guidance. The 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 can not be affected by new recommendations. 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 architecture document proposes a baseline homenet architecture, based on 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.

[RFC6204] defines basic requirements for customer edge routers (CERs). The scope of this text is the internal homenet, and thus specific features on the CER are out of scope for this text. While the network may be dual-stack or IPv6-only, the definition of specific transition tools on the CER, as introduced in RFC 6204-bis [I-D.ietf-v6ops-6204bis] with DS-Lite [RFC6333] and 6rd [RFC5969], are also considered out of scope of this text.

### 1.1. Terminology and Abbreviations

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

o ALQDN: Ambiguous Locally Qualified Domain Name. An example would be .sitelocal.

o CER: Customer Edge Router. A border router at the edge of the homenet.

o FQDN: Fully Qualified Domain Name. A globally unique name space.

o LLN: Low-power and lossy network.

o LQDN: Locally Qualified Domain Name. A name space local to the homenet.

o NAT: Network Address Translation. Typically referring to IPv4 Network Address and Port Translation (NAPT) [RFC3022].

o NPTv6: Network Prefix Translation for IPv6 [RFC6296].

o PCP: Port Control Protocol [I-D.ietf-pcp-base].

o 'Simple Security'. Defined in [RFC4864] and expanded further in [RFC6092]; describes recommended perimeter security capabilities for IPv6 networks.

- ULA: IPv6 Unique Local Addresses [RFC4193].

- ULQDN: Unique Locally Qualified Domain Name. An example might be .<UniqueString>.sitelocal.

- UPnP: Universal Plug and Play. Includes the Internet Gateway Device (IGD) function, which for IPv6 is UPnP IGD Version 2 [IGD-2].

- VM: Virtual machine.

- WPA2: Wi-Fi Protected Access, as defined by the Wi-Fi Alliance.

## 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, 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 building control or corporate extensions 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 much lower capacity links thus becomes an important topic.

The introduction of IPv6 for home networking enables the potential for every home network to be delegated enough address space to provision globally unique prefixes for each such subnet in the home. As discussed later, this assumes the customer's ISP delegates enough address space to the home. While the number of addresses in a

standard /64 IPv6 prefix is practically infinite, 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.

The addition of routing between subnets raises the issue of how to extend mechanisms such as service discovery which 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 later in the document.

There will also be the need to discover which routers in the homenet are the border router(s) by an appropriate mechanism. Here, there are a number of choices, including the use of an appropriate service discovery protocol. Whatever method is chosen would likely have to deal with handling more than one router responding in multihomed environments.

## 2.2. Global addressability and elimination of NAT

The end-to-end communication that is potentially enabled with IPv6 is on the one hand an incredible opportunity for innovation and simpler network operation, but it is also a concern as it exposes nodes in the internal networks to receipt of potentially unwanted 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 may simply not have robustness designed-in because it was either assumed such devices would only be used on private networks or the device itself doesn't have the computing power to apply the necessary security methods.

It is important to distinguish between addressability and reachability. While IPv6 offers global addressability through use of globally unique addresses in the home, whether devices are globally reachable or not would depend on the 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 CER. [RFC4864] and [RFC6092] discuss such filtering, and the merits of 'default allow' against 'default deny' policies for external traffic initiated into a homenet. This document takes no position on which mode is the default, but assumes the choice to use either would be made available.

## 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. Where a homenet is multihomed, a device would typically receive a globally unique address 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. Default Address Selection for IPv6 [RFC6724] provides a solution for this, though it may face problems in the event of multihoming where, as described above, nodes will be configured with one address from each upstream ISP prefix. In such cases the presence of upstream BCP 38 [RFC2827] ingress filtering requires multi-addressed nodes to select the correct source address to be used for the corresponding uplink, but the node may not have the information it needs to make that decision based on addresses alone. We discuss such challenges in the multihoming section later in this document.

## 2.4. Unique Local Addresses (ULAs)

[RFC4193] defines Unique Local Addresses (ULAs) for IPv6 that may be used to address devices within the scope of a single site. Support for ULAs for IPv6 CERs is described in [RFC6204]. A home network running IPv6 should deploy ULAs alongside its globally unique prefix(es) to allow stable communication between devices (on different subnets) within the hoemnet 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. While setting up a network there may also be a period with no connectivity, in which case ULAs would be required for inter-subnet communication. In the case where LLNs are being set up in a new home/deployment, individual LLNs may, at least initially, each use their own /48 ULA prefix.

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. It may thus be expected that one router in the homenet be elected a 'master' to delegate ULA prefixes to subnets from a single /48 ULA prefix.

Where both a ULA and a global prefix are in use, the default address selection mechanisms described above should ensure that a ULA source address is used to communicate with ULA destination addresses when appropriate, i.e. when the ULA destination lies within the /48 ULA prefix(es) known to be used within the same homenet. Note that unlike private IPv4 RFC 1918 space, the use of ULAs does not imply use of host-based IPv6 NAT, or NPTv6 prefix-based NAT [RFC6296], rather that in an IPv6 homenet a node should use its ULA address internally, and its additional globally unique IPv6 address as the source address for external communications. By using such globally unique addresses between networks, the architectural cost and complexity, particulrly to applications, of NAT or NPTv6 translation is avoided. As such, neither IPv6 NAT or NPTv6 is recommended for use in the homenet architecture.

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 will need to support and use ULAs.

As noted later in this text, if appropriate filtering is in place on the CER(s), a ULA source address may be taken as an indication of locally sourced traffic.

### 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 CER 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 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 is imperative for the easy deployment and use of homenet devices and applications.

As mentioned previously, this means that zeroconf naming and service discovery protocols must be capable of operating across subnet

boundaries.

### 2.6. IPv6-only operation

It is likely that IPv6-only networking will be deployed first in 'greenfield' homenet scenarios, 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. Thus, for example, the following requirements are amongst those that should be considered in IPv6-only environments:

o Ensuring there is a way to access content in the IPv4 Internet. This can be arranged through appropriate use of NAT64 [RFC6144] and DNS64 [RFC6145], for example, or via a node-based DS-Lite [RFC6333] approach.

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

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

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 CER.

## 3. Homenet Architecture

The aim of this architecture text is to outline how to construct advanced IPv6-based home networks involving multiple routers and subnets using standard IPv6 protocols and addressing [RFC2460] [RFC4291]. In this section, we present the elements of such a home networking architecture, with discussion of the associated design principles.

Existing IETF work [RFC6204] defines the 'basic' requirements for

CERs, while [I-D.ietf-v6ops-6204bis] updates the current requirements based on operator feedback and adds new requirements for IP transition technologies and transition technology coexistence. This document describes a homenet architecture which is focused on the internal homenet, rather than the CER(s).

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

- 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 a 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 inter-networking mechanisms. This architecture text discusses how this flexibility should be used to provide the best user experience and ensure that the network can evolve with new applications in the future. The principles described in this text should be followed when designing homenet solutions.

### 3.1.1. Reuse existing protocols

It is desirable to reuse existing protocols where possible, but at the same time to avoid consciously precluding the introduction of new or emerging protocols. A generally conservative approach, giving weight to running code, is preferable. Where new protocols are

required, evidence of commitment to implementation by appropriate vendors or development communities is highly desirable. Protocols used should be backwardly compatible, and forward compatible where changes are made.

#### 3.1.2. Minimise changes to hosts and routers

Where possible, any requirement for changes to hosts and routers should be minimised, though solutions which, for example, incrementally improve with host or router changes may be acceptable.

### 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. de-activating 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 which cause 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. Bridging loops can be dangerous but are also detectable when a switch learns the MAC of one of its interfaces on another or runs a spanning tree or link state protocol. It is only loops using simple repeaters that are truly pathological.

#### 3.2.2. Network topology models

Most IPv4 home network models at the time of writing 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.

In general, the models described in [RFC6204] and its successor RFC 6204-bis [I-D.ietf-v6ops-6204bis] should be supported by the IPv6 home networking architecture. The functions resident on the CER itself are, as stated previously, out of scope of this text.

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 CER.

- 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 CER. The CER(s) may or may not be managed by the ISP. If the demarcation point is such that the customer can provide or manage the CER, its configuration must be simple. Both models must be supported.

Various forms of multihoming are likely to become more prevalent with IPv6 home networks, 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 CERs. The homenet may have a single CER, which might be used for one or more providers, or multiple CERs. The presence of multiple CERs adds additional complexity for multihoming scenarios, and protocols like PCP that need to manage connection-oriented state mappings.

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, rather an indicative one to facilitate the discussion in this text.

#### 3.2.2.1. A: Single ISP, Single CER, 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 a 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 loops).

```
                     +-------+-------+                     \
                     |   Service     |                      \
                     |   Provider    |                       | Service
                     |    Router     |                       | Provider
                     +-------+-------+                       | network
                             |                              /
                             | Customer                    /
                             | Internet connection
                             |
                      +------+--------+                    \
                      |     IPv6      |                     \
                      | Customer Edge |                      \
                      |    Router     |                      |
                      +----+-+---+----+                      |
          Network A        | |   |      Network B(E)         |
    ----+-------------+----+ |   +---+-------------+------+  |
        |             |      |       |             |      |  |
   +----+-----+ +-----+----+ |  +----+-----+ +-----+----+ |  |
   |IPv6 Host | |IPv6 Host | |  | IPv6 Host| |IPv6 Host | |  |
   |    H1    | |    H2    | |  |    H3    | |    H4    | |  |
   +----------+ +----------+ |  +----------+ +----------+ |  |
                             |        |             |     |  |
                      Link F |     ---+------+------+-----+  |
                             |               | Network E(B)  |
                      +------+--------+      |               | End-User
                      |     IPv6      |      |               | networks
                      |   Interior    +------+               |
                      |    Router     |                      |
                      +---+-------+-+-+                      |
          Network C       |       |   Network D              |
    ----+-------------+---+       +---+-------------+---     |
        |             |               |             |        |
   +----+-----+ +-----+----+     +----+-----+ +-----+----+   |
   |IPv6 Host | |IPv6 Host |     | IPv6 Host| |IPv6 Host |   |
   |   H5     | |   H6     |     |    H7    | |    H8    |   /
   +----------+ +----------+     +----------+ +----------+  /
```

Figure 1

In this diagram there is one CER. It has a single uplink interface. It has three additional interfaces connected to Network A, Link F, and Network B. 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 CER or IPv6 (IR), wireless networks, or through one or more layer-2 only Ethernet switches.

#### 3.2.2.2. B: Two ISPs, Two CERs, Shared subnet

```
           +-------+-------+     +-------+-------+         \
           |   Service     |     |   Service     |          \
           |  Provider A   |     |  Provider B   |           | Service
           |    Router     |     |    Router     |           | Provider
           +------+--------+     +-------+-------+           | network
                  |                      |                   /
                  |      Customer        |                  /
                  | Internet connections |                 /
                  |                      |
           +------+--------+     +-------+-------+         \
           |     IPv6      |     |    IPv6       |          \
           | Customer Edge |     | Customer Edge |           \
           |   Router 1    |     |   Router 2    |           /
           +------+--------+     +-------+-------+          /
                  |                      |                 /
                  |                      |                | End-User
     ---+---------+---+---------------+--+----------+---  | network(s)
        |             |               |             |      \
   +----+-----+ +-----+----+     +----+-----+ +-----+----+  \
   |IPv6 Host | |IPv6 Host |     | IPv6 Host| |IPv6 Host |  /
   |   H1     | |   H2     |     |    H3    | |    H4    | /
   +----------+ +----------+     +----------+ +----------+
```

Figure 2

Figure 2 illustrates a multihomed homenet model, where the customer has connectivity via CER1 to ISP A and via CER2 to ISP B. This example shows one shared subnet where IPv6 nodes would potentially be multihomed and receive multiple IPv6 global addresses, one per ISP. This model may also be combined with that shown in Figure 1 to create a more complex scenario with multiple internal routers.  Or the above shared subnet may be split in two, such that each CER serves a separate isolated subnet, which is a scenario seen with some IPv4 networks today.

#### 3.2.2.3. C: Two ISPs, One CER, Shared subnet

```
           +-------+-------+     +-------+-------+         \
           |   Service     |     |   Service     |          \
           |  Provider A   |     |  Provider B   |           | Service
           |    Router     |     |    Router     |           | Provider
           +-------+-------+     +-------+-------+           | network
                    |                 |                     /
                    |    Customer     |                   /
                    |    Internet     |                  /
                    |   connections   |                 |
                   +---------+---------+                 \
                   |       IPv6        |                   \
                   |   Customer Edge   |                    \
                   |      Router       |                    /
                   +---------+---------+                   /
                             |                            /
                             |                            | End-User
     ---+------------+-------+--------+-------------+---  | network(s)
        |            |                |             |      \
   +----+-----+ +----+-----+     +----+-----+ +-----+----+  \
   |IPv6 Host | |IPv6 Host |     | IPv6 Host| |IPv6 Host |  /
   |   H1     | |   H2     |     |    H3    | |   H4     | /
   +----------+ +----------+     +----------+ +----------+
```

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.

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

### 3.2.3. Dual-stack topologies

It is expected that most homenet deployments will for the immediate future 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 speed and lower speed shared media, and in addition bridging may not be suitable for some media, such as ad-hoc mobile networks.

In some cases IPv4 home networks may feature cascaded NATs, which could include cases where NAT routers are included within VMs, or where Internet connection sharing services are used. IPv6 routed versions of such cases will be required. We should thus 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 either take a form where there are 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, 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 CER Network Model C illistrated above, multihoming may be offered by source routing at the CER. With multiple exit routers, as in CER 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 BCP 38 filtering at an 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 archiecture should support both the above models, i.e. one or more CERs. 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 [I-D.ietf-v6ops-ipv6-multihoming-without-ipv6nat]. 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 probably host and certainly router changes. Or another avenue is to introduce support for source routing throughout the homenet; while greatly improving the 'intelligence' of routing decisions within the homenet, such an approach would require relatively significant router changes.

As explained previously, NPTv6 is not recommended in the homenet architecture.

There are some other multihoming considerations for homenet scenarios. First, it may be the case that multihoming applies due to an ISP migration from a transition method to a native deployment, e.g. a 6rd [RFC5969] sunsetting scenario. Second, one upstream may be a "walled garden", and thus only appropriate to be used 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. While we should not specifically target walled garden multihoming as a principal goal, it should not be precluded.

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

### 3.3. A Self-Organising Network

A home network architecture should be naturally self-organising and self-configuring under different circumstances relating to the connectivity status to the Internet, number of devices, and physical topology. 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 is avoided. Methods are needed for devices to know whether they are intended to be part of the same homenet site or not. Thus methods to ensure separation between neighbouring homenets are required. This may require use of some unique 'secret' for devices/protocols in each homenet. Some existing mechanisms exist to assist home users to associate devices as simply as possible, e.g. 'connect' button support.

### 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 to 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. For example, bridging a busy Gigabit Ethernet subnet and a wireless subnet together may impact wireless performance.

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. 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 should be possible to automatically discover these borders, which will 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 which receive delegated prefixes within the homenet. It is also desirable for a richer security model that hosts, which may be running in a transparent communication mode, 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 realm and the borders between them, namely the internal homenet, the ISP and a guest network. In this case the borders will include that from the homenet to the ISP, that from the guest network to the ISP, and that 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 Grid or LLN-based network, 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 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.

Some initial proposals towards border discovery are presented in [I-D.kline-default-perimeter].

### 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 the and configure the necessary addressing information to operate.

### 3.4.1. Use of ISP-delegated IPv6 prefixes

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. Some ISPs may offer relatively stable prefixes, while others may change the prefix whenever the CER is reset. Some discussion of IPv6 prefix allocation policies is included in [RFC6177] which discusses why, for example, a one-size-fits-all /48 allocation is not desirable.

The homenet architecture expects internal host subnets to be /64 in size. While it may be possible to operate a DHCPv6-only network with prefixes longer than /64, doing so would break SLAAC, and is thus not recommended.

The home network needs to be adaptable to ISP prefix allocation policies, and thus make no assumptions about the stability of the prefix received from an ISP, or the length of the prefix that may be offered. However, if only a /64 is offered by the ISP, the homenet may be severely constrained or even unable to function. As stated above, attempting to use internal subnet prefixes longer than /64 would break SLAAC, and is thus not recommended. Using ULA prefixes internally with NPTv6 at the boundary is not recommended for reasons given elsewhere. Reverting to bridging would destroy subnetting, breaks multicast if bridged onto 802.11 wireless networks and has serious limitations with regard to heterogeneous link layer technologies and LLNs. For those reasons it is recommended that DHCP-PD or OSPFv3 capable routers have the ability to issue a warning upon receipt of a /64 if required to assign further prefixes within the home network. Though some consideration needs to be given to how that should be presented to a typical home user.

Thus the border CER router should 'hint', most likely via DHCP-PD, 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, but in practice it may only be offered a /56 or /60. For a typical IPv6 homenet, it is not recommended that an ISP offer less than a /60 prefix, and should preferably offer at least a /56.

In practice, it is expected that ISPs will deliver a relatively stable home prefix to customers. 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 support Provider Independent (PI) addressing for general 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 which, 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.

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. 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 routing between subnets regardless of whether a globally unique ISP-provided prefix is available. As discussed previously, it would be expected that ULAs would be used alongside one or more such 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.

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

The use of ULAs should be restricted to the homenet scope through filtering at the border(s) of the homenet, as described in RFC 6092.

Note that it is possible that in some cases 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. It is expect that routers in the homenet would somehow elect a 'master' that would be responsible for delegating /64 prefixes to internal requesting routers, much as routers obtain /64 global prefixes from the prefix pool delegated by the ISP to the CER. In cases where multiple ULA /48's 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 CER [RFC3633]. Where multiple CERs exist with multiple ISP prefix pools, it is expected that routers within the homenet would assign themselves prefixes from each ISP they communicate with/through. As discussed above, a ULA prefix can be made available for stable internal communications, or for use on constrained/LLN networks. There may also be a prefix associated with NAT64, if in use in the homenet.

The delegation or availability of a prefix pool to the homenet should allow subsequent internal autonomous delegation of prefixes for use within the homenet. Such internal delegation should not assume a flat or hierarchical model, nor should it make an assumption about whether the delegation 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 CERs and these are delegated prefix pools from their attached ISPs, the internal prefix delegation would be expected to be served by each CER for each prefix associated with it. However, where ULAs are used, most likely but not necessarily in parallel with global prefixes, one router should be elected as 'master' for delegation of ULA prefixes for the homenet, such that only one /48 ULA covers the whole homenet where possible. That router should generate a /48 ULA for the site, and then delegate /64's from that ULA prefix to 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.

Delegation within the homenet should give each subnet a prefix that is persistent across reboots, power outages and similar short-term outages. 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. Persistent prefixes should not depend on router boot order. However, 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).

The delegation method should support renumbering, which would typically be 'flash' renumbering in that the homenet would not have advance notice of the event or thus be able to apply the types of approach described in [RFC4192]. As a minimum, delegated ULA prefixes within the homenet should remain persistent through an ISP-driven renumbering event.

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

### 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

There are no specific privacy concerns discussed in this text. It should be noted that, in general, ISPs are expected to offer relatively stable IPv6 prefixes to customers, and thus the network prefix associated with the host addresses they use may not change over a reasonably long period of time. This exposure is similar to IPv4 networks that expose the same IPv4 global address via use of NAT, where the IPv4 address received from the ISP may change over time, but not necessarily that frequently.

Hosts inside an IPv6 homenet may get new IPv6 addresses over time regardless, e.g. through Privacy Addresses [RFC4941]. This may benefit mutual privacy of users within a home network, but not mask which home network traffic is sourced from.

### 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.

The homenet unicast routing protocol should preferably be an existing deployed protocol that has been shown to be reliable and robust, and it is preferable that the protocol is 'lightweight'. It is desirable that the routing protocol has knowledge of the homenet topology, which implies a link-state protocol is preferable. If so, it is also desirable that the announcements and use of LSAs and RAs are appropriately coordinated. This would mean the routing protocol gives a consistent view of the network, and that it can pass around more than just routing information.

Multiple interface PHYs must be accounted for in the homenet routed topology. Technologies such as Ethernet, WiFi, MoCA, etc must be capable of coexisting in the same environment and should be treated as part of any routed deployment. The inclusion of the PHY layer characteristics including bandwidth, loss, and latency in path computation should be considered for optimising communication in the homenet. Multiple upstreams should be supported, as described in the multihoming section earlier. This should include load-balancing to multiple providers, and failover from a primary to a backup link when available. The protocol however should not require upstream ISP connectivity to be established to continue routing within the homenet.

To support multihoming within a homenet, a routing protocol that can

make routing decisions based on source and destination addresses is desirable, to avoid upstream ISP ingress filtering problems. In general the routing protocol should support multiple ISP uplinks and delegated prefixes in concurrent use.

The routing environment should be self-configuring, as discussed previously. An example of how OSPFv3 can be self-configuring in a homenet is described in [I-D.acee-ospf-ospfv3-autoconfig]. Minimising convergence time should be a goal in any routed environment, but as a guideline a maximum convergence time of around 30 seconds should be the target.

Any routed solution will require a means for determining the boundaries of the homenet. Borders may include but are not limited to the interface to the upstream ISP, or a gateway device to a separate home network such as a LLN network. In some cases there may be no border present, which may for example occur before an upstream connection has been established. The border discovery functionality may be integrated into the routing protocol itself, but may also be imported via a separate discovery mechanism.

In general, LLN or other networks should be able to attach and participate the same way as the main homenet, or alternatively map/be gatewayed to the main homenet. Current home deployments use largely different mechanisms in sensor and basic Internet connectivity networks. IPv6 VM solutions may also add additional routing requirements.

### 3.5.1. Multicast support

It is desirable that, subject to the capacities of devices on certain media types, multicast routing is supported across the homenet. The natural scopes for multicast would be link-local or site-local, with the latter constrained within the homenet, but other policy borders, e.g. to a guest subnet, or to certain media types, may also affect where specific multicast traffic is routed.

There may be different drivers for multicast to be supported across the homenet, e.g. for service discovery should a proposal such as xmDNS [I-D.lynn-homenet-site-mdns] be deployed, or potentially for novel streaming or filesharing 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. It must be possible to scope or filter multicast traffic to avoid it being flooded to network media where devices cannot reasonably support it.

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 direct towards them. However, there are other challenges introduced, e.g. default filtering policies at the borders between other 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.

In IPv4 NAT networks, the NAT provides an implicit firewall function. [RFC4864] describes a 'Simple Security' model for IPv6 networks, whereby stateful perimeter filtering can be applied instead where global addresses are used. RFC 4864 implies an IPv6 'default deny' policy for inbound connections be used for similar functionality to IPv4 NAT. It should be noted that such a 'default deny' 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. Thus to support applications wanting to accept connections initiated into home networks where a 'default deny' policy is in place support for a signalling protocol such as UPnP or PCP [I-D.ietf-pcp-base] is required. In networks with multiple CERs, the signalling would need to handle the cases of flows that may use one or more exit routers. CERs 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 CERs must provide an easily selected configuration option that permits a 'transparent' mode, thus ensuring a 'default allow' model is available. The homenet architecture text makes no recommendation on the default setting, and refers the reader to RFC 6092.

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

LLNs provide an 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.

### 3.6.3. Marginal Effectiveness of NAT and Firewalls

Security by way of obscurity (address translation) or through firewalls (filtering) is at best marginally effective. The very poor security track record of home computer, home networking and business PC computers and networking is testimony to its ineffectiveness. A 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 home network products with very poor security, putting a firewall in place does provide some protection, even if only marginally effective. The use of firewalls today, whether a good practice or not, is common practice and whatever protection afforded, even if marginally effective, must not be lost.

### 3.6.4. 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 to be initiated to them, either for all ports or for specific services.

Users should have simple methods to associate devices to services that they wish to operate transparently through (CER) borders.

#### 3.6.5. ULAs as a hint of connection origin

It has been suggested that using ULAs would provide an indication to applications that received traffic is locally sourced. This 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

Naming and service discovery must be supported in the homenet, and 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, be the user within the homenet or outside it, 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.

#### 3.7.1. Discovering services

Users will typically perform service discovery through GUI interfaces that allow them to browse services on their network in an appropriate and intuitive way. Such interfaces are beyond the scope of this document, but the interface should have an appropriate 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 service discovery protocols are generally aimed at single subnets. There is thus a choice to make for multi-subnet homenets 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 loops, or extending the scope of multicast traffic used for the purpose. This document does not mandate either solution, rather it expresses the principles that should be used for a homenet naming and service discovery environment. 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). In general we should prefer approaches that are backwardly compatible, and allow current implementations to continue to be used.

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 either through a standards based translation scheme, hooks into current protocols to allow some for 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 two running network elements with such devices are suddenly joined. For example, mDNS [I-D.cheshire-dnsext-multicastdns] section 8 describes such a mechanism for a single subnetwork and the '.local' zone. Before assigning a name to the device and the .local naming space, the device checks whether the name already belongs to another device by sending a multicast DNS query.

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 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 may be distinct from the name used in a name service, e.g. 'Study Laser Printer' as opposed to printer2.<somedomain>.

### 3.7.3. Name spaces

It is desirable that only one name space is in use in the homenet, and that this name space is served authoritatively by a server in the homenet, most likely resident on the CER.

If a user wishes to access their home devices remotely from elsewhere on the Internet a globally unique name space is required. This may be acquired by the user or provided/generated by their ISP. It is

expected 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. Examples of provider name space delegation approaches are described in [I-D.mglt-homenet-naming-delegation] and [I-D.mglt-homenet-front-end-naming-delegation]. For users wanting to use their own independent domain names, such services are already available.

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 which 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 an application in one homenet to be confused with a device with the same name 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 CER or device running the name service should generate a default value. It could 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.

Whichever approach is used, the intent of using a ULQDN is to disambiguate the name space across different homenets, not to create a new IANA name space for such networks. However, in practice an ALQDN may typically suffice, because 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. And regardless, if remote access to a homenet is required, a global domain is required, which implictly disambiguates devices.

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 gTLD. Thus qualified names should always be used, whether these are exposed to the user or not.

There may be use cases where segmentation of the name space is desirable, e.g. for use in different realms within the homenet. Thus hierarchical name space management is likely to be required.

Where a user may be 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 two accompanying name spaces exist. In such cases, a GUI may present the user a choice of domains to use, where the name of their devices is thus relative to that domain. This 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' is preferred for certain devices.

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.

#### 3.7.4. 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 co-exist 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 mDNS [I-D.cheshire-dnsext-multicastdns]. Note that when DNS is used as the homenet name service, it includes both a resolving service and an authoritative service. The authoritative service hosts the homenet related zone, that may be requested by the resolving service.

As described in [I-D.mglt-homenet-naming-delegation], one approach is to run an authoritative name service in the homenet as well as a resolving name service, most likely on the CER. The homenet resolving name service relies both on the homenet authoritative

service as well as on a secondary resolving name service provided by the ISP, for global Internet naming resolution.

For a service such as mDNS to coexist with an Internet name service, where the homenet is preferably using a global domain name, it is desirable that the zeroconf devices have a way to add their names to the global name space in use. One solution could be for zeroconf protocols to be used to indicate global FQDNs, e.g. an mDNS service could return a FQDN in a SRV record.

Regardless, a method for local name service entries to be populated automatically by devices is desirable. Interfaces to devices might choose to give users the option as to whether the device should register itself in the global name space. There should also be a defined mechanism for device entries to be removed or expired from the global name space.

It has been suggested that Dynamic DNS could be made to operate in a zero-configuration mode using a locally significant root domain and with minimal configuration or, using a DHCPv6 based means of automated delegation, populate a global DNS zone.

To protect against attacks such as cache poisoning, it is desirable to support appropriate name service security methods, including DNSSEC.

The CER is an appropriate location to host the naming service. However, it introduces an additional load due to the name service management, e.g. signing the zone, or resolving naming queries. This additional load must be balanced with the CER capabilities, else the function(s) may need to be offloaded elsewhere, e.g. with the ISP, though this may impact on the independent operation principle.

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

#### 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 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.

The approach described above 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, 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 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 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) 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 the search domains for local FQDN-derived zones should be included.

### 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 QoS in a multi-service homenet may be a requirement, e.g. for a critical system (perhaps healthcare related), or for differentiation between different types of traffic (file sharing, cloud storage, live streaming, VoIP, etc). Different media types may have different 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 CER, and potentially around boundaries between different media types (where for example some traffic may simply not be appropriate for some media, and need to be dropped to avoid drowning 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

The homenet should be self-organising and configuring as far as possible, and thus not be pro-actively managed by the home user. Thus protocols to manage the network are not discussed in this architecture text.

However, users may be interested in the status of their networks and devices on the network, in which case simplified monitoring mechanisms may be desirable. It may also be the case that an ISP, or a third party, might offer management of the homenet on behalf of a user, in which case management protocols would be required. How such management is done is out of scope of this document; many solutions exist.

## 3.9. Implementing the Architecture on IPv6

This architecture text encourages re-use 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 existing routing protocols can be used as is. However, a new mechanism may be needed in order to turn a selected protocol on by default.

Some functionality, if required by the architecture, would add 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

multicast-based 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. References

### 5.1. Normative References

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

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

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

[RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.

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

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

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

## 5.2. 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.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.

[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.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

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

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

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.

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

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

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

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

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

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.

[RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, March 2011.

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

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

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012.

[RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

[I-D.mglt-homenet-front-end-naming-delegation]
Migault, D., Cloetens, W., Lemordant, P., and C. Griffiths, "IPv6 Home Network Front End Naming Delegation", draft-mglt-homenet-front-end-naming-delegation-01 (work in progress), November 2012.

[I-D.mglt-homenet-naming-delegation]
Cloetens, W., Lemordant, P., and D. Migault, "IPv6 Home Network Naming Delegation Architecture", draft-mglt-homenet-naming-delegation-00 (work in progress), July 2012.

[I-D.lynn-homenet-site-mdns]
Lynn, K. and D. Sturek, "Extended Multicast DNS", draft-lynn-homenet-site-mdns-01 (work in progress), September 2012.

[I-D.ietf-v6ops-ipv6-multihoming-without-ipv6nat]
Matsushima, S., Okimoto, T., Troan, O., Miles, D., and D. Wing, "IPv6 Multihoming without Network Address

Translation",
draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat-04 (work in progress), February 2012.

[I-D.baker-homenet-prefix-assignment]
Baker, F. and R. Droms, "IPv6 Prefix Assignment in Small Networks", draft-baker-homenet-prefix-assignment-01 (work in progress), March 2012.

[I-D.arkko-homenet-prefix-assignment]
Arkko, J., Lindem, A., and B. Paterson, "Prefix Assignment in a Home Network", draft-arkko-homenet-prefix-assignment-03 (work in progress), October 2012.

[I-D.acee-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", draft-acee-ospf-ospfv3-autoconfig-03 (work in progress), July 2012.

[I-D.cheshire-dnsext-multicastdns]
Cheshire, S. and M. Krochmal, "Multicast DNS", draft-cheshire-dnsext-multicastdns-15 (work in progress), December 2011.

[I-D.ietf-pcp-base]
Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", draft-ietf-pcp-base-29 (work in progress), November 2012.

[I-D.kline-default-perimeter]
Kline, E., "Default Border Definition", draft-kline-default-perimeter-01 (work in progress), November 2012.

[I-D.ietf-v6ops-6204bis]
Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-6204bis-12 (work in progress), October 2012.

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

[IGD-2] UPnP Gateway Committee, "Internet Gateway Device (IGD) V 2.0", September 2010, <http://upnp.org/specs/gw/

UPnP-gw-WANIPConnection-v2-Service.pdf>.

## Appendix A. Acknowledgments

The authors would like to thank Aamer Akhter, Mark Andrews, Dmitry Anipko, Ran Atkinson, Fred Baker, Ray Bellis, Cameron Byrne, Brian Carpenter, Stuart Cheshire, Lorenzo Colitti, Robert Cragie, 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, Don Sturek, Dave Taht, Dave Thaler, Michael Thomas, Mark Townsley, JP Vasseur, Curtis Villamizar, Dan Wing, Russ White, and James Woodyatt for their comments and contributions within homenet WG meetings and on the WG mailing list. An acknowledgement generally means that person's text made it in to the document, or was helpful in clarifying or reinforcing an aspect of the document.

## Appendix B. Changes

This section will be removed in the final version of the text.

### B.1. Version 07

Changes made include:

- Removed reference to NPTv6 in section 3.2.4. Instead now say it has an architectural cost to use in the earlier section, and thus it is not recommended for use in the homenet architecture.

- Removed 'proxy or extend?' section. Included shorter text in main body, without mandating either approach for service discovery.

- Made it clearer that ULAs are expected to be used alongside globals.

- Removed reference to 'advanced security' as described in draft-vyncke-advanced-ipv6-security.

- Balanced the text between ULQDN and ALQDN.

- Clarify text does not assume default deny or allow on CER, but that either mode may be enabled.

- Removed ULA-C reference for 'simple' addresses. Instead only suggested service discovery to find such devices.

- Reiterated that single/multiple CER models to be supported for multihoming.

- Reordered section 3.3 to improve flow.

- Added recommendation that homenet is not allocated less than /60, and a /56 is preferable.

- Tidied up first few intro sections.

- Other minor edits from list feedback.

## B.2. Version 06

Changes made include:

- Stated that unmanaged goal is 'as far as possible'.

- Added note about multiple /48 ULAs potentially being in use.

- Minor edits from list feedback.

## B.3. Version 05

Changes made include:

- Some significant changes to naming and SD section.

- Removed some expired drafts.

- Added notes about issues caused by ISP only delegating a /64.

- Recommended against using prefixes longer than /64.

- Suggested CER asks for /48 by DHCP-PD, even if it only receives less.

- Added note about DS-Lite but emphasised transition is out of scope.

- Added text about multicast routing.

## B.4. Version 04

Changes made include:

- o Moved border section from IPv6 differences to principles section.

- o Restructured principles into areas.

- o Added summary of naming and service discovery discussion from WG list.

### B.5. Version 03

Changes made include:

- o Various improvements to the readability.

- o Removed bullet lists of requirements, as requested by chair.

- o Noted 6204bis has replaced advanced-cpe draft.

- o Clarified the topology examples are just that.

- o Emphasised we are not targetting walled gardens, but they should not be precluded.

- o Also changed text about requiring support for walled gardens.

- o Noted that avoiding falling foul of ingress filtering when multihomed is desirable.

- o Improved text about realms, detecting borders and policies at borders.

- o Stated this text makes no recommendation about default security model.

- o Added some text about failure modes for users plugging things arbitrarily.

- o Expanded naming and service discovery text.

- o Added more text about ULAs.

- o Removed reference to version 1 on chair feedback.

- o Stated that NPTv6 adds architectural cost but is not a homenet matter if deployed at the CER. This text only considers the internal homenet.

- o Noted multihoming is supported.

- Noted routers may not by separate devices, they may be embedded in devices.
- Clarified simple and advanced security some more, and RFC 4864 and 6092.
- Stated that there should be just one secret key, if any are used at all.
- For multihoming, support multiple CERs but note that routing to the correct CER to avoid ISP filtering may not be optimal within the homenet.
- Added some ISPs renumber due to privacy laws.
- Removed extra repeated references to Simple Security.
- Removed some solution creep on RIOs/RAs.
- Load-balancing scenario added as to be supported.

## B.6. Version 02

Changes made include:

- Made the IPv6 implications section briefer.
- Changed Network Models section to describe properties of the homenet with illustrative examples, rather than implying the number of models was fixed to the six shown in 01.
- Text to state multihoming support focused on single CER model. Multiple CER support is desirable, but not required.
- Stated that NPTv6 not supported.
- Added considerations section for operations and management.
- Added bullet point principles/requirements to Section 3.4.
- Changed IPv6 solutions must not adversely affect IPv4 to should not.
- End-to-end section expanded to talk about "Simple Security" and borders.
- Extended text on naming and service discovery.

o Added reference to RFC 2775, RFC 6177.

o Added reference to the new xmDNS draft.

o Added naming/SD requirements from Ralph Droms.

Authors' Addresses

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

Email: tjc@ecs.soton.ac.uk

Jari Arkko
Ericsson
Jorvas 02420
Finland

Email: jari.arkko@piuha.net

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

Email: abr@sdesigns.dk

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

Email: ot@cisco.com

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

Email: jason.weil@twcable.com