Network Working Group   T. Chown, Ed.
Internet-Draft   University of Southampton
Intended status: Informational   J. Arkko
Expires: December 31, 2012   Ericsson
                                                              A. Brandt
                                                          Sigma Designs
                                                               O. Troan
                                                    Cisco Systems, Inc.
                                                                J. Weil
                                                      Time Warner Cable
                                                          June 29, 2012

                     Home Networking Architecture for IPv6
                          draft-ietf-homenet-arch-03

Abstract

   This text describes evolving networking technology within
   increasingly large residential home networks. The goal of this
   document is to define the architecture for IPv6-based home networking
   through the associated principles, considerations and requirements.
   The text briefly highlights the implications of the introduction of
   IPv6 for home networking, discusses topology scenarios, 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.

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 in an increasingly broad range of devices and media. This
   evolution in scale and diversity sets requirements on IETF protocols.
   Some of these requirements relate to the introduction of IPv6, others
   to the introduction of specialised networks for home automation and
   sensors. There are likely to be scenarios where internal routing is
   required, for example to support private and guest networks, in which
   case home networks may use increasing numbers of subnets.

   While at the time of writing some complex home network topologies
   exist, most operate based on IPv4, employ solutions that we would
   like to avoid such as (cascaded) network address translation (NAT),
   or require expert assistance to set up. The assumption of this
   document is that the homenet is as far as possible self-organising
   and self-configuring, and is thus not 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. 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. We assume that the IPv4 network architecture
   in home networks is what it is, and can not be affected by new
   recommendations.

   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 consider specific physical, wireless or data-
   link layer considerations.

   [RFC6204] defines basic requirements for customer edge routers

(CERs). The scope of this text is the homenet, and thus the relevant part of RFC 6204 is the internal facing interface as well as any other components within the home network. While the network may be dual-stack or IPv6-only, the definition of specific transition tools on the CER, as introduced in RFC 6204-bis [I-D.ietf-v6ops-6204bis] with DS-Lite [RFC6333] and 6rd [RFC5969], are considered issues for that RFC, and are thus out of scope of this text.

1.1. Terminology and Abbreviations

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

o "Advanced Security". Describes advanced security functions for a CER, as defined in [I-D.vyncke-advanced-ipv6-security], where the default inbound connection policy is generally "default allow".

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

o LLN: Low-power and lossy network.

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.

o ULA: IPv6 Unique Local Addresses [RFC4193].

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

o VM: Virtual machine.

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

2. Effects of IPv6 on Home Networking

Service providers are deploying IPv6, content is becoming available on IPv6, and support for IPv6 is increasingly available in devices

and software used in the home. While IPv6 resembles IPv4 in many
ways, it changes address allocation principles, making multi-
addressing the norm, and allowing direct IP addressability home
networking devices from the Internet. This section presents an
overview of some of the key implications of the introduction of IPv6
for home networking, that are 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.

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

The addition of routing between subnets raises the issue of how to
extend mechanisms such as service discovery which currently rely on
link-local addressing to limit scope. There are two broad choices;
extend existing protocols to work across the scope of the homenet, or
introduce proxies for existing link-layer protocols. This 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. These include an appropriate service
discovery protocol, or the use of a well-known name, resolved by some
local name service. Both might have to deal with handling more than
one router responding in multihomed environments.

2.2. Global addressability and elimination of NAT

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

link to be assigned a globally unique IPv6 address.

The end-to-end communication that is potentially enabled with IPv6 is
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 otherwise unwanted traffic from
the Internet. There may thus be an expectation of improved host
security to compensate for this, at least in general networked
devices, but it must be noted that many devices may also (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.

IPv6 networks may or may not have filters applied at their borders,
i.e. at the homenet CER. [RFC4864], [RFC6092] and
[I-D.vyncke-advanced-ipv6-security] discuss such filtering, and the
merits of "default allow" against "default deny" policies for
external traffic initiated into a homenet. It is important to
distinguish between addressability and reachability. While IPv6
offers global addressability through use of globally unique addresses
in the home, whether they 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.

2.3. Multi-Addressing of devices

In an IPv6 network, devices may acquire multiple addresses, typically
at least a link-local address and a globally unique address. They
may also have an IPv4 address if the network is dual-stack, a Unique
Local Address (ULA) [RFC4193] (see below), and one or more IPv6
Privacy Addresses [RFC4941].

Thus it should 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. Default Address Selection for IPv6
[I-D.ietf-6man-rfc3484bis] provides a solution for this, though it
may face problems in the event of multihoming, where nodes will be
configured with one address from each upstream ISP prefix. In such
cases the presence of upstream ingress filtering requires multi-
addressed nodes to select the right source address to be used for the
corresponding uplink, to avoid ISP BCP 38 ingress filtering, but the
node may not have the information it needs to make that decision
based on addresses alone. We discuss such challenges in multihoming
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 may deploy ULAs for stable communication between devices
   (on different subnets) within the network where the externally
   allocated global prefix changes over time (e.g. due to renumbering
   within the subscriber's ISP) or where external connectivity is
   temporarily unavailable.

   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. It should also be noted that there may be timers on
   the prefix lease to the homenet, on the internal prefix delegations,
   and on the Router Advertisements to the hosts. Despite this counter-
   argument, while setting a network up there may 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.

   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.

   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 where a particular
   application is allowed to connect to or receive traffic from.

   Default address selection mechanisms should ensure a ULA source
   address is used to communicate with ULA destination addresses when
   appropriate, in particular when the ULA destination lies within a /48
   ULA prefix known to be used within the same homenet. Unlike the 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 external
   communications should use a node's additional globally unique IPv6
   source address.

2.5. Security and borders

   The filtering policy to/from the homenet is an important
   consideration, but the homenet/ISP border may not be the only border
   in a homenet. It is desirable that there are mechanisms to detect
   other types of borders, and then the means to 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 visitor in a "guest" network access to media services in the
   home, or access to a printer in the residence. Simple mechanisms to
   apply policy changes, or associations between devices, will be
   required.

   A simple homenet model may just consider three types of realm and the
   borders between them. For example if the realms are the homenet, the
   ISP and the visitor network, then the borders will include that from
   the homenet to the ISP, and that from the homenet to a 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 or configured, 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.

   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 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 only one such key is needed for any set
   of functions, to increase usability for the homenet user.

   Advanced Security for IPv6 CPEs [I-D.vyncke-advanced-ipv6-security]
   takes the approach that in order to provide the greatest end-to-end
   transparency as well as security, security policies must be updated
   by a trusted party which can provide intrusion signatures and other
   "active" information on security threats. This might for example
   allow different malware detection profiles to be configured on a CER.

Such methods should be able to be automatically updating.

There is no defined "threat model" as such for the type of IPv6
homenet described in this text. Such a document may be very useful.
It may include a variety of perspectives, from probing for specific
types of home appliance being present, to potential denial of service
attacks. Hosts need to be able to operate securely, end-to-end where
required, but also be robust against malicious traffic direct towards
them. We simply note at this point that software on home devices
will have an increase in security if it allows its software to be
updated regularly.

2.6. Naming, and 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. Users should not be expected to enter IPv6 literal
addresses in homenet devices or applications, given their much
greater length and apparent randomness to a typical home user. While
shorter addresses, perhaps ones registered with IANA from ULA-C space
[I-D.hain-ipv6-ulac], could be used for specific devices/services, in
general it is better to not expose users to real IPv6 addresses.
Thus, even for the simplest of functions, simple naming and the
associated (ideally zero configuration) discovery of services is
imperative for the easy deployment and use of homenet devices and
applications.

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

3. Architecture

The aim of this architecture text is to outline how to construct home
networks involving multiple routers and subnets. In this section, we
present a set of typical home network topology models/scenarios,
followed by a list of topics that may influence the architectural
discussions, and a set of architectural principles and requirements
that govern how the various nodes should work together. The
architecture also drives what protocol extensions are necessary, as
will be discussed briefly in Section 3.5.

3.1. Network 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 visitor and private networks.

   In general, the models described in [RFC6204] and its successor RFC
   6204-bis [I-D.ietf-v6ops-6204bis] should be supported by an IPv6 home
   networking architecture.

   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. Isolation may be physical, or implemented via IEEE
      802.1q VLANs.

   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.

   It has also been suggested that various forms of multihoming are more
   prevalent with IPv6 home networks. Thus the following properties may
   also apply to such networks:

   o Number of upstream providers. A typical homenet might just have a
      single upstream ISP, but it may become more common 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
      be walled gardens.

   o Number of CERs. The homenet may have a single CER, which might be
      used for one or more providers, or multiple CERs. Multiple CERs
      adds additional complexity for multihoming scenarios, and
      protocols like PCP that need to manage connection-oriented state
      mappings.

   Some separate discussion of physical infrastructures for homenets is
   included in and [I-D.arkko-homenet-physical-standard].

In principle, we might argue that an architecture for IPv6 homenets
should support any arbitrary topology. We discuss this topic later
in the text. 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 discussion in this text.

3.1.1. A: Single ISP, Single CER, Internal routers

Figure 1 shows a 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. This is a
valid consideration as home users may plug routers together to form
arbitrary topologies including loops (we discuss support for
arbitrary topologies in layer sections).

```
                +-------+-------+                        \
                |    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 | | |
     |          | |          | | |          | |          | | |
     +----------+ +----------+ | +----------+ +----------+ | |
                               |     |             |     | |
                               |   ---+------+------+-----+ |
                               |     | Network B/E    |
                +------+--------+     |             | End-User
                |     IPv6     |      |             | networks
                |   Interior   +------+             |
                |    Router    |                    |
                +---+-------+-+-+                    |
         Network C   |       |   Network D           |
     ----+-----------+---+-  --+---+------------+---   |
         |           |         |             |        |
     +----+-----+ +-----+----+   +----+-----+ +-----+----+   |
     |IPv6 Host | |IPv6 Host |   | IPv6 Host| |IPv6 Host |   |
     |          | |          |   |          | |          |  /
     +----------+ +----------+   +----------+ +----------+  /
```
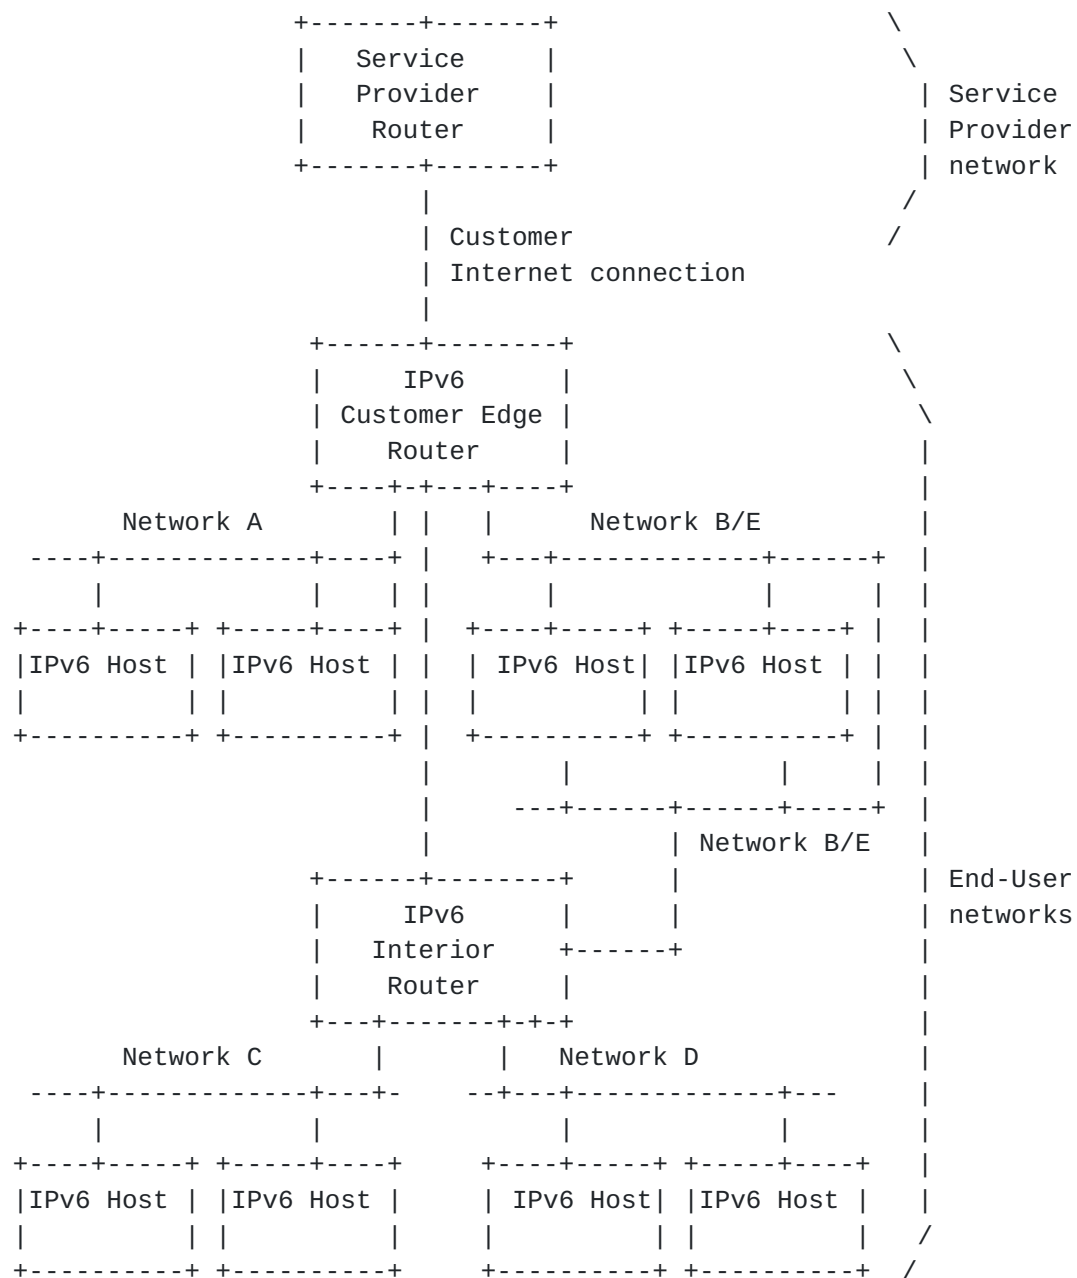
                              Figure 1

3.1.2. **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 |  /
  |          | |          |   |          | |          | /
  +----------+ +----------+   +----------+ +----------+
```
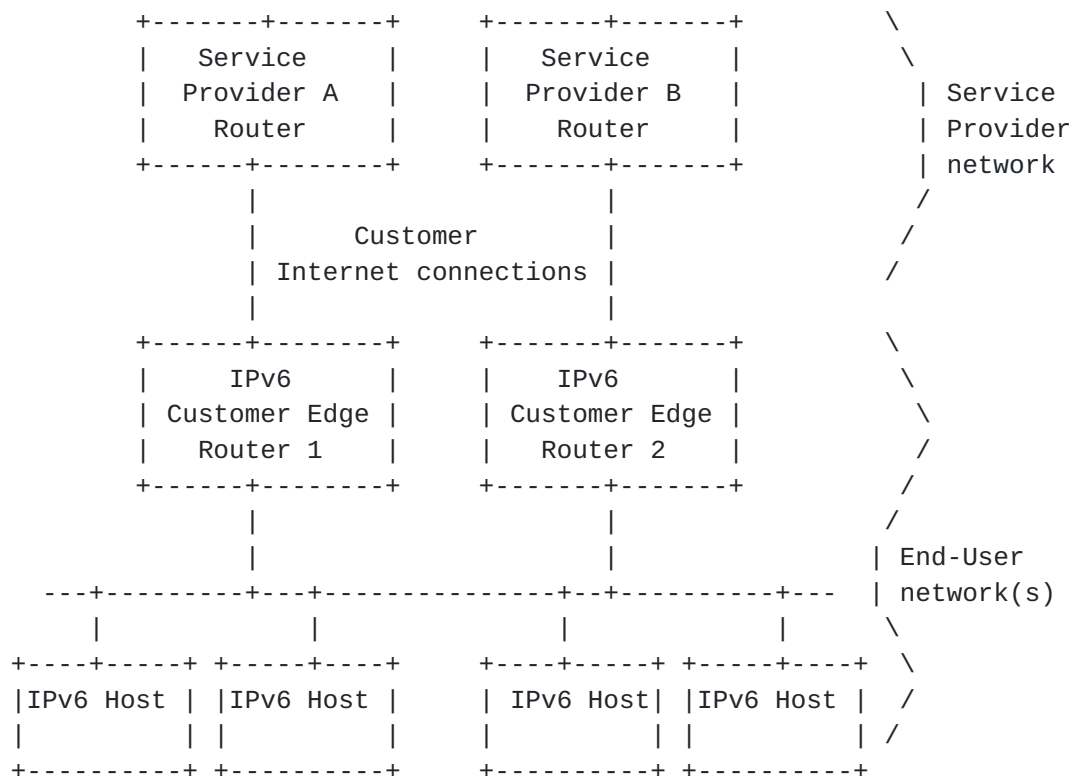
                            Figure 2

   Figure 2 illustrates a multihomed home network model, where the
   customer has connectivity via CER1 to ISP A and via CER2 to ISP B.
   This example shows one shared subnet where IPv6 nodes would
   potentially be multihomed and receive multiple IPv6 global addresses,
   one per ISP. This model may also be combined with that shown in
   Figure 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.

3.1.3. 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 |  /
   |          | |          |   |          | |          | /
   +----------+ +----------+   +----------+ +----------+
```

   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. Determining the Requirements

   [RFC6204] defines "basic" requirements for IPv6 Customer Edge
   Routers, while [I-D.ietf-v6ops-6204bis] extends RFC 6204 to describe
   additional features. In general, home network equipment needs to
   cope with the different types of network properties and topologies as
   exemplified above. Significant manual configuration is rarely, if at
   all, possible, given the knowledge level of typical home users. The
   network should as far as possible be self-configuring. The equipment
   needs to be prepared to handle at least

   o Routing

   o Prefix configuration for routers

o Name resolution

o Service discovery

o Network security

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

3.3. Considerations

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

3.3.1. Multihoming

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

The general multihoming problem is broad, and solutions suggested to
date within the IETF may include 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. So we
should limit the support to the smallest subset of the overall
problem to meet the requirements of the topologies described above.
This means that the homenet architecture should not try to make
another attempt at solving complex multihoming, and we should prefer
to support scenarios for which solutions exist today.

In the general homenet architecture, hosts should be multi-addressed
with globally unique prefixes from each ISP they may communicate with
or through. The alternative for a homenet would be to deploy NPTv6
[RFC6296] at the CER, with ULAs then typically used internally.
NPTv6 has some architectural cost, due to the prefix translation
used, but the internal part of the homenet (which is the scope of
this text) sees only the one prefix in use.

When 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, which would typically
include [I-D.ietf-6man-rfc3484bis]. Applications may of course do
different things, and this should not be precluded.

For the single CER Network Model C, multihoming may be offered by
source routing at the CER. With multiple exit routers, the
complexity rises. Given a packet with a source address on the
network, the packet must be routed to the proper egress to avoid
ingress filtering at a wrong ISP. While the packet might not take an
optimal path to the correct exit CER, the minimum requirement is that
the packet is not dropped, and it is highly desirable that the packet
is routed in the most efficient manner to the correct exit.

There are various potential approaches to this problem, one example
being described in [I-D.v6ops-multihoming-without-ipv6nat]. Another
is discussed in [I-D.baker-fun-multi-router], which explores support
for source routing throughout the homenet. This approach would
however likely require relatively significant routing changes to
route the packet to the correct exit given the source address. Such
changes should preferably be minimised.

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] sunset scenario, as discussed in
[I-D.townsley-troan-ipv6-ce-transitioning]. 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.

Host-based methods such as Shim6 [RFC5533] have been defined, but of
course require support in the hosts. There are also application-
oriented approaches such as Happy Eyeballs
[I-D.ietf-v6ops-happy-eyeballs]; simplified versions of this are for
example implemented in some commonly-used web browsers. The homenet
architecture should not preclude use of such tools. Solutions that
require host changes should be avoided, but solutions which
incrementally improve with host changes may be acceptable.

3.3.2. 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 should 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].

3.3.3. 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.

3.3.4. Privacy considerations

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

Hosts inside an IPv6 homenet may get new IPv6 addresses over time
regardless, e.g. through Privacy Addresses [RFC4941].

3.4. Design Principles and Requirements

There is little that the Internet standards community can do about
the physical topologies or the need for some networks to be separated
at the network layer for policy or link layer compatibility reasons.
However, there is a lot of flexibility in using IP addressing and

inter-networking mechanisms. In this section we discuss how this
flexibility should be used to provide the best user experience and
ensure that the network can evolve with new applications in the
future.

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

3.4.1. Reuse existing protocols

It is desirable to reuse existing protocols where possible, but at
the same time to avoid consciously precluding the introduction of new
or emerging protocols. A generally conservative approach, giving
weight to running code, is preferable. Where new protocols are
required, evidence of commitment to implementation by appropriate
vendors or development communities is highly desirable. Protocols
used should be backwardly compatible, and forward compatible where
changes are made.

Where possible, any requirement for changes to hosts and routers
should be minimised.

3.4.2. Dual-stack Operation

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

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

o Ensuring there is a way to access content in the IPv4 Internet.
   This can be arranged through incorporating NAT64 [RFC6144] and
   DNS64 [RFC6145] functionality in the home gateway router, for
   instance. Such features are outside the scope of this document
   however, being CER functions.

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 such discovery options
can operate 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.

In dual-stack networks, solutions for IPv6 should not adversely
affect IPv4 operation. It is desirable that topologies of IPv4 and
IPv6 networks would be as congruent as possible.

3.4.3. Largest Possible Subnets

Today's IPv4 home networks generally have a single subnet, and early
dual-stack deployments have a single congruent IPv6 subnet, possibly
with some bridging functionality. 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 possible subnets that
can be constructed within the constraints of link layer mechanisms,
bridging, physical connectivity, and policy.

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. A general recommendation is to follow the same
topology for IPv6 as is used for IPv4, but not to use NAT. Thus
there should be routed IPv6 where an IPv4 NAT is used, and where
there is no NAT there should be bridging if the link layer allows
this.

In some cases IPv4 NAT home networks may feature cascaded NATs, which
may include cases where NAT routers are included within VMs or
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 devices.

3.4.4. Security vs Transparent, End-to-End Communications

An IPv6-based home network architecture should embrace and naturally
offer a transparent end-to-end communications model as described in

[RFC2775]. Each device should be addressable by a globally unique
address, and those addresses must not be altered in transit.
Security perimeters can (via policy) 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 proscribe a particular
mode of operation, instead stating that CERs must provide an easily
selected configuration option that permits a "transparent" mode of
operation, 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, which in turn simply
states that a CER should provide functionality sufficient to support
the recommendations in that RFC.

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.

3.4.5. Internal IP Connectivity

A logical consequence of the end-to-end communications model is that
the network should by default attempt to provide IP-layer
connectivity between all internal parts of the homenet as well as to
and from the external Internet.

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. However, it would be expected that ULAs would
also 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. 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.

ULAs may be used for all devices, not just those intended to only
have internal connectivity. ULAs used in this way provide stable
internal communications should the ISP-provided prefix (suddenly)
change, or external connectivity be temporarily lost. The use of
ULAs should be restricted to the homenet scope through filtering at
the border(s) of the homenet, as described in RFC 6092; thus "end-to-
end" for ULAs is limited to the homenet.

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

LLNs provide an example of where there may be secure perimeters
inside the homenet. Constrained LLN nodes may implement WPA2-style
network key security but may depend on access policies enforced by
the LLN border router.

3.4.6. Routing functionality

Routing functionality is required when there are multiple routers
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 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 in
the next subsection. An example of how OSPFv3 can be self-
configuring in a homenet is described in
[I-D.acee-ospf-ospfv3-autoconfig]. 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 SmartGrid or similar LLN network. In
some cases there may be no border such as occurs 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.

[I-D.howard-homenet-routing-comparison] contains evaluations of
common routing protocols made against the type of requirements
described above.

3.4.7. 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. While the homenet should be self-organising, it should be
   possible to manually adjust (override) the current configuration.

   The homenet will need to be aware of the extent of its own "site", as
   discussed in the previous section. The homenet "site" defines the
   borders for ULAs, site scope multicast, service discovery and
   security policies. The homenet will thus have one or more borders
   with external connectivity providers and potentially also have
   borders within the internal network (e.g. for policy-based reasons).
   It should be possible to automatically perform border discovery for
   the different borders. Such borders determine for example the scope
   of where prefixes, routing information, network traffic, service
   discovery and naming may be shared. The default internally should be
   to share everything.

   The most important function in this respect is prefix delegation and
   management. There are various sources of prefixes, e.g. they may be
   globally unique prefixes originating from ISP(s), they may be
   globally unique or ULA prefixes allocated by "master" router(s) in
   the homenet, or they may be ULAs allocated by LLN gateways. There
   may also be a prefix associated with NAT64, if in use in the homenet.

   From the homenet perspective, a single prefix from each ISP should be
   received on the border CER [RFC3633]. Then each subnet in the
   homenet should receive a prefix from within the ISP-provided
   prefix(es). The ISP should only see the aggregate from the homenet,
   and not single /64 prefixes allocated within the homenet.

   Delegation should be autonomous, and not assume a flat or
   hierarchical model. This text makes no assumption about whether the
   delegation of 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. A currently
   typical /60 allocation gives 16 /64 subnets. Duplicate assignment of
   multiple /64s to the same network should be avoided. The network
   should behave as gracefully as possible in the event of prefix
   exhaustion, though the options in such cases may be limited.

   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.

Where ULAs are used, most likely but not necessarily in parallel with
global prefixes, one router should be elected to offer ULA prefixes
for the homenet. The router should generate a /48 ULA for the site,
and then delegate /64's from that ULA prefix to subnets. In the
normal state, a single /48 ULA should be used within the homenet. In
cases where two /48 ULAs are generated within a homenet, the network
should still continue to function.

Delegation within the homenet should give each link 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. 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],
[I-D.chakrabarti-homenet-prefix-alloc], [RFC3315] and [RFC3633]. The
other uses OSPFv3, as described in
[I-D.arkko-homenet-prefix-assignment]. More detailed analysis of
these approaches needs to be made against the requirements/principles
described above.

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

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

   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.

[3.4.8](#). Fewest Topology Assumptions

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

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

[3.4.9](#). Naming and Service Discovery

   Naming and service discovery must be supported in the homenet. The
   most natural way to think about such naming and service discovery is
   to enable it to work across the entire residence (site), disregarding
   technical borders such as subnets but potentially respecting policy
   borders such as those between visitor and internal networks.
   Discovery of a DNS service for access to external Internet resources
   is also a fundamental requirement in a multi-subnet homenet; the
   problem is not just name and service discovery within the homent
   itself.

   Users will need simple ways to name devices, or be provided with
   appropriate ways for devices to generate unique names within the
   homenet. The naming system will be required to work internally or
   externally, be the user within the homenet or outside it, and there
   may be multiple naming domains, e.g. Internet, home or guest
   domains. It is highly likely that a home user will want access to
   many of the devices and services in their home while "roaming"
   elsewhere. It may be the case that not all devices in the homenet
   are made available by name via any Internet-facing domain, and that a
   "split-view" naming system is preferred for certain devices. Also,
   name resolution for reachable devices must continue to function if
   the local network is disconnected from the global Internet.

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

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 environments, like mDNS.
Consideration should be made for how these interact with each other
in a homenet scenario.

With the introduction of new "dotless" top level domains, there is
potential for ambiguity between for example a local host called
"computer" and (if it is registered) a .computer gTLD. This suggests
some implicit local name space is probably required. Such a name
space should also be configurable to something else by the user.

The use of standard local domain name across adjacent homenets
potentially introduces some ambiguity if, for example, a mobile
device should move between two such networks.

Current service discovery protocols are generally aimed at single
subnets. If service discovery is to operate across the an entire
homenet, by adopting an approach like that proposed as Extended mDNS
(xmDNS) [I-D.lynn-homenet-site-mdns], then support may be required
for IPv6 multicast across the scope of the whole homenet.

In some parts of the homenet, e.g. LLNs, devices may be sleeping, in
which case a proxy for such nodes may be required, that can respond
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.

3.4.10. Proxy or Extend?

There are two broad choices for allowing services that would
otherwise be link-local to work across a homenet site. In the
example of service discovery, one is to take protocols like mDNS and
have them run over site multicast within the homenet. This is fine
if all hosts support the extension, and the scope within any internal
borders is well-understood. But it's not backwards-compatible with
existing link-local protocols. The alternative is to proxy service
discovery across each link, to propagate it. This is more complex,

but is backwards-compatible. It would need to work with IPv6, and
dual-stack.

The homenet architecture proposes that any existing protocols that
are designed to only work within a subnet should be extended to work
across subnets, rather than defining proxy capabilities for each of
those functions. However, while it is desirable to extend protocols
to site scope operation rather than providing proxy functions on
subnet boundaries, the reality is that until all hosts can use site-
scope discovery protocols, existing link-local protocols would need
to be proxied anyway.

Some protocols already have proxy functions defined and in use, e.g.
DHCPv6 relays, in which case those protocols would be expected to
continue to operate that way.

3.4.11. Adapt to ISP constraints

Different homenets may be subject to different behaviour by their
ISP(s). 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 home network needs to
be adaptable to such ISP 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.

The internal operation of the home network should also not depend on
the availability of the ISP network at any given time, other than for
connectivity to services or systems off the home network. This
implies the use of ULAs as supported in RFC 6204. If used, ULA
addresses should be stable so that they can always be used
internally, independent of the link to the ISP.

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 in general residential homenets.

When an ISP needs to restructure and in doing so renumber its
customer homenets, "flash" renumbering is likely to be imposed. This
implies a need for the homenet to be able to handle a sudden
renumbering event which, unlike the process described in [RFC4192],
would be 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 is 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. In such cases, not only may the prefix change, but
potentially the prefix length, if the new ISP offers a different
default size prefix, e.g. a /60 rather than a /56. Regardless, it's
desirable that homenet protocols support rapid renumbering and that
operational processes don't add unnecessary complexity for the
renumbering process.

The 6renum WG is studying IPv6 renumbering for enterprise networks.
It is not currently targetting homenets, but may produce outputs that
are relevant. The introduction of any new homenet protocols should
not make any form of renumbering any more complex than it already is.

3.5. 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. There are though 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 domain "home" 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.

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

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

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

   [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-09 (work in progress), May 2012.

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.

   [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.

   [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.

[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.

[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.

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

[I-D.lynn-homenet-site-mdns]
            Lynn, K. and D. Sturek, "Extended Multicast DNS",
            draft-lynn-homenet-site-mdns-00 (work in progress),
            March 2012.

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

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

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

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

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

   [I-D.ietf-6man-rfc3484bis]
              Thaler, D., Draves, R., Matsumoto, A., and T. Chown,
              "Default Address Selection for Internet Protocol version 6
              (IPv6)", draft-ietf-6man-rfc3484bis-06 (work in progress),
              June 2012.

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

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

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

   [I-D.acee-ospf-ospfv3-autoconfig]
              Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration",
              draft-acee-ospf-ospfv3-autoconfig-02 (work in progress),
              May 2012.

   [I-D.ietf-pcp-base]
              Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P.
              Selkirk, "Port Control Protocol (PCP)",
              draft-ietf-pcp-base-26 (work in progress), June 2012.

   [I-D.hain-ipv6-ulac]
              Hain, T., Hinden, R., and G. Huston, "Centrally Assigned
              IPv6 Unicast Unique Local Address Prefixes",
              draft-hain-ipv6-ulac-02 (work in progress), July 2010.

   [I-D.ietf-v6ops-happy-eyeballs]
              Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with
              Dual-Stack Hosts", draft-ietf-v6ops-happy-eyeballs-07
              (work in progress), December 2011.

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

   [I-D.arkko-homenet-physical-standard]
              Arkko, J. and A. Keranen, "Minimum Requirements for
              Physical Layout of Home Networks",
              draft-arkko-homenet-physical-standard-00 (work in
              progress), March 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, Fred Baker, Ray Bellis, Cameron Byrne, Brian Carpenter,
   Stuart Cheshire, Lorenzo Colitti, Robert Cragie, Ralph Droms, Lars
   Eggert, Jim Gettys, Wassim Haddad, Joel M. Halpern, David Harrington,
   Lee Howard, Ray Hunter, Joel Jaeggli, Heather Kirksey, Ted Lemon,
   Kerry Lynn, Erik Nordmark, Michael Richardson, Barbara Stark, Sander
   Steffann, Dave Thaler, JP Vasseur, Curtis Villamizar, Dan Wing, Russ
   White, and James Woodyatt for their contributions within homenet WG
   meetings and the mailing list, and Mark Townsley for being an initial
   editor/author of this text before taking his position as homenet WG
   co-chair.

Appendix B. Changes

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

B.1. **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 targeting 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.

   o Noted routers may not by separate devices, they may be embedded in
      devices.

   o Clarified simple and advanced security some more, and RFC 4864 and
      6092.

   o Stated that there should be just one secret key, if any are used
      at all.

   o For multihoming, support multiple CERs but note that routing to
      the correct CER to avoid ISP filtering may not be optimal within
      the homenet.

   o Added some ISPs renumber due to privacy laws.

   o Removed extra repeated references to Simple Security.

   o Removed some solution creep on RIOs/RAs.

   o Load-balancing scenario added as to be supported.

B.2. Version 02

   Changes made include:

   o Made the IPv6 implications section briefer.

   o 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.

   o Text to state multihoming support focused on single CER model.
      Multiple CER support is desirable, but not required.

   o Stated that NPTv6 not supported.

   o Added considerations section for operations and management.

   o Added bullet point principles/requirements to Section 3.4.

   o Changed IPv6 solutions must not adversely affect IPv4 to should
      not.

   o End-to-end section expanded to talk about "Simple Security" and
      borders.

   o 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