

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 12, 2019

T. Lemon
Nibbhaya Consulting
March 11, 2019

Homenet vs. the Market: Gap Analysis
draft-lemon-homenet-review-00

Abstract

This document discusses the homenet problem space and tries to compare what we have both with what the market is now providing, and also with what we need.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on September 12, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Internet-Draft

Homenet Review

March 2019

Table of Contents

1.	Introduction	2
2.	Finding the Gaps	4
2.1.	Connectivity between hosts on the homenet	4
2.2.	Connectivity from hosts on the homenet to hosts on the internet (single egress)	5
2.3.	Connectivity from hosts on the internet to hosts on the homenet	5
2.4.	Support for multi-homing (more than one egress)	5
2.5.	Service discovery	6
2.6.	Roaming between APs	6
2.7.	IPv4 Connectivity within the home	6
2.8.	Connectivity from hosts on IoT leaf networks to hosts on the internet	6
2.9.	Connectivity from hosts on the Internet to hosts on IoT leaf networks	7
2.10.	Connectivity between hosts on the same IoT leaf network	7
2.11.	Connectivity between hosts on different IoT leaf networks within the same home	7
2.12.	Connectivity from hosts on the homenet to hosts on the IoT network	7
2.12.1.	Connectivity from hosts on the IoT network to hosts on the homenet	7
2.13.	Isolation between hosts that shouldn't be communicating on the homenet	8
3.	Themes	8
	Author's Address	9

[1.](#) Introduction

The Homenet working group has been developing a set of specifications for some years with the goal of providing a self-configuring home network with potentially multiple routers. Since Homenet began, the market has changed significantly, and it is worth spending some time looking at how it has changed and how that changes what, if anything, the Homenet working group should be doing.

Homenet originally set out to provide for a multi-homed network with a layer 3 routing topology that would isolate individual subnets in the hope of better performance, while preserving end-to-end service and allowing for service discovery throughout the home.

At the time, a typical home network either had a single router, or several routers connected together with one or more layers of network address translation. Each router provided an isolated "LAN" link, connected to a "WAN" upstream. Service discovery was restricted to

individual LAN links. Some routers provided "air-to-air bridging" or "Wi-Fi range extension" service.

In recent years, several competing technologies have shown up. First, there are access points that provide a hub-and-spoke infrastructure service. Each access point provides service, and there is a single layer two with no layer three routing topology. Either a single wired switch is used as an edge router, or one of the access points acts as an edge router, bridging air-to-air to the others.

A second technology that has gained market share is wifi mesh. IEEE's 802.11s was not a success when it was initially introduced, but private changes to 802.11s have allowed for deployment of layer two mesh networks using Wi-Fi access points. This functions similarly to infrastructure network, except that all traffic is sent over the Wi-Fi mesh. As with Wi-Fi infrastructure routers, the network appears to the host as a single layer two.

There are some significant advantages to the layer 2 approach. First, it means that a host can roam from AP to AP without needing to renumber. Second, at least in principle, service discovery can be accomplished using multicast. Third, if there are multiple egress routes, these can be presented to hosts using router advertisements, allowing the hosts to choose between routes without any new protocol infrastructure as is required by homenet.

What homenet provides is a layer three topology with a lot of new protocol infrastructure. Roaming from one AP to another will always result in renumbering, which means that the user experience will be less than ideal. Service discovery is a solved problem at this point, and in fact probably works better than multicast service discovery on a large network.

However, by and large, homenets are a lot of work for not much return. We don't have any field experience with them, so we don't

know what they look like in terms of customer support load, but it's easy to conjecture that they will be orders of magnitude more expensive to support than layer two mesh networks.

Looking at this from the perspective of the IETF, an SDO that deals mostly with layer 3 and above, the question is, what value do we add, or can we add? What should a home network look like? Is layer 2 actually the right solution? What gaps exist?

[2.](#) Finding the Gaps

If we look at the anatomy of a home network, there are some clear problems that any home network needs to solve (or fails to solve):

- o Connectivity between hosts on the home network
- o Connectivity from hosts on the home network to hosts on the Internet (single egress)
- o Connectivity from hosts on the internet to hosts on the homenet network
- o Support for multi-homing (more than one egress)
- o Service discovery
- o Roaming between APs
- o IPv4 Connectivity within the home
- o IoT connectivity, specifically:
 - * Connectivity from hosts on IoT leaf networks to hosts on the internet
 - * Connectivity from hosts on the internet to hosts on IoT leaf networks
 - * Connectivity between hosts on the same IoT leaf network
 - * Connectivity between hosts on different IoT leaf networks within the same home
 - * Connectivity from hosts on the homenet to hosts on the IoT network
 - * Connectivity from hosts on the IoT network to hosts on the homenet
- o Isolation between hosts that shouldn't be communicating on the homenet

If we consider each type of network, we can see how each of these applies. In the sections below, we analyze each case.

[2.1.](#) Connectivity between hosts on the homenet

The problem here is twofold: there needs to be numbering, and there needs to be routing. That is, every host on the homenet needs to have a unique IP address, and every subnet has to be reachable--there has to be a routing in both directions. Additionally, the numbering has to be stable; if when the upstream ISP goes away, the prefix goes away, that doesn't count. For the unique IP address:

Single NAT specified
Multi-NATted not specified
Flat layer 2 specified
Homenet specified, but optional for v4

For routing:

Lemon

Expires September 12, 2019

[Page 4]

Internet-Draft

Homenet Review

March 2019

Single NAT not needed
Multi-NATted not specified, no workaround
Flat layer 2 not needed
Homenet specified

[2.2.](#) Connectivity from hosts on the homenet to hosts on the internet (single egress)

In order for this to work, there just has to be a default route to the internet. Since this is the most-needed use case, it's not surprising that it works in all cases. "Specified" here is a bit wrong for NAT, but the problem is sufficiently well-understood that we can say there is a clear spec.

Single NAT specified
Multi-NATted specified
Flat layer 2 specified
Homenet specified

[2.3.](#) Connectivity from hosts on the internet to hosts on the homenet

This works in a variety of ways. If the host has a global IP address

and there is no firewall, or there is a hole in the firewall that was set up manually or using PCP, then it's reachable. NATted hosts can be reached if they are able to set up a port forward in the NAT either manually or using PCP. Most home gateways will let you set up a manual forward, but this is technically challenging. Support for PCP is nearly nonexistent; in principle it's specified, and works well, but in practice it's not available to most users. In all non-NAT cases, IPv6 support works if it is supported on the router, but only if there is no firewall, PCP is supported, or the user sets it up manually.

Single NAT PCP or manual
Multi-NATted PCP or manual
Flat layer 2 PCP or manual
Homenet Specified; optional for v4

[2.4.](#) Support for multi-homing (more than one egress)

With NAT, multi-egress support is possible, but there is no standard way of doing it, and this is not generally supported. Flat Layer 2 networks can support multi-homing with IPv6 simply by connecting more than one egress router up to the layer 2 and having each one advertise routes. For IPv4 it's just like the Single NAT case.

Single NAT not likely
Multi-NATted not likely

Flat layer 2 v6 only
Homenet specified, optional for v4

[2.5.](#) Service discovery

Service Discovery generally works on single links using multicast, so whether it works depends on whether multicast works. Some routers disable multicast because of its performance characteristics, particularly in the flat layer 2 case. Ad hoc workarounds exist for some use cases, and solutions are specified for homenet-like environments. Service discovery protocols like UPnP work when multicast works. While specifications do exist for multi-link UPnP, it's a safe assumption that no home network routers implement them.

Single NAT DNS-SD, UPnP

Multi-NATted Not specified
Flat layer 2 DNS-SD, UPnP
Homenet DNS-SD is specified, but UPnP isn't

[2.6.](#) Roaming between APs

Single NAT Not specified, doesn't work
Multi-NATted Not specified, doesn't work
Flat layer 2 Specified
Homenet Specified, bad user experience

[2.7.](#) IPv4 Connectivity within the home

Single NAT Specified
Multi-NATted Specified
Flat layer 2 Specified
Homenet Specified, optional

[2.8.](#) Connectivity from hosts on IoT leaf networks to hosts on the internet

In all of these cases, there is a specification for how an IoT network can get routing using HNCP on a homenet, but as far as I know there is no specification for an IoT gateway to translate from compressed IP headers to regular IP headers.

Single NAT Double NAT
Multi-NATted Triple NAT
Flat layer 2 Double NAT
Homenet Specified (HNCP)

[2.9.](#) Connectivity from hosts on the Internet to hosts on IoT leaf networks

Single NAT PCP
Multi-NATted PCP
Flat layer 2 PCP
Homenet specified (HNCP)

[2.10.](#) Connectivity between hosts on the same IoT leaf network

Single NAT works
Multi-NATted works
Flat layer 2 works
Homenet works

[2.11.](#) Connectivity between hosts on different IoT leaf networks within the same home

Single NAT not specified
Multi-NATted not specified
Flat layer 2 not specified
Homenet specified (HNCP)

[2.12.](#) Connectivity from hosts on the homenet to hosts on the IoT network

Single NAT not specified
Multi-NATted not specified
Flat layer 2 not specified
Homenet specified (HNCP)

[2.12.1.](#) Connectivity from hosts on the IoT network to hosts on the homenet

In this case, it's possible for an IoT host to connect to a NAT host if the IoT edge router does NAT, but in that case there is no guarantee that there won't be an address conflict.

Single NAT not specified
Multi-NATted not specified
Flat layer 2 not specified
Homenet specified (HNCP)

[2.13.](#) Isolation between hosts that shouldn't be communicating on the

homenet

There are some devices that really shouldn't be able to connect to everything on the network, and to which everything on the network shouldn't be able to connect. This can be managed using MUD profiles, at least in principle, but only if there is a way to isolate the devices. It's common nowadays for people to set up isolated IoT-only networks in the home, but this is of limited value, and requires manual configuration. There is no reason in principle why this type of isolation couldn't be done for homenets or for Flat Layer 2 solutions. It can also be done in principle on legacy home networks. But at present, how to do this automatically is an unsolved or at best partially-solved problem.

Single NAT not specified
Multi-NATted not specified
Flat layer 2 not specified
Homenet not specified

[3.](#) Themes

One of the common themes here is that it's no surprise that Flat Layer 2 networks are gaining in popularity. Things work better with a flat layer 2 than with a multi-layer NAT, and even a single NAT is too limited for a lot of use cases.

It seems to be the case that a lot of the work we have done in Homenet is applicable to FL2 networks. HNCP and routing are useful because they provide end-to-end connectivity to and between IoT leaf networks. Even though an FL2 network can in principle support multicast, the more L2 segments there are, the worse this scales. So in practice, the solution's we've been working on for Homenet Naming are likely applicable in the FL2 use case.

The bit about isolation of hosts may seem like a non-sequitur in this analysis but I bring it up because I think it's applicable in two ways. First, it's applicable to the multicast scaling problem. With our DNS-SD solutions for homenet, we can specify a multicast-like service discovery framework that works reliably, but doesn't spray multicasts everywhere. And the problem of isolation of IoT nodes is likely to be something that needs to be addressed; while it's not specifically a homenet problem, my suspicion is that there is interest here.

Author's Address

Ted Lemon
Nibbhaya Consulting
P.O. Box 958
Brattleboro, Vermont 05301
United States of America

Email: mellon@fugue.com

Lemon

Expires September 12, 2019

[Page 9]