

homenet architecture draft

IETF82, Taipei
15th November 2011

Jari Arkko, Tim Chown, Ole Troan, Jason Weil

Where we're at

Draft discussed heavily at Philadelphia interim

Documented:

- Standardization approach

- New IPv6 considerations

- Topology considerations

- Requirements

- Principles

Now need to agree from that work what the homenet architecture looks like

- Have some initial views from the interim

- Would like more feedback

Potential Conclusions Preview

Discussions in Philadelphia seemed to lead to the following:

- Support multiple subnets and routers

- Route where you had an IPv4 NAT

- Use link-state routing protocols (e.g., OSPF) for routing

- LLNs, VMs, etc. can attach to home networks

- For multihoming, we only deal with ingress filtering

- Prefix delegation from the ISP

- Stable & efficient prefix assignment within home

- Simple Security + PCP

Standardization approach

Many perspectives

Operational - “works well for me”

Experience - “enough experience to recommend”

Implementations - “available in most devices”

Functionality - “we need this feature”

Specification - “IETF develops new mechanisms”

Authors in the operational-experience-implementation camp

Start with what works and push envelope from that

Might imply Homenet versions

Practical Example

Making a useful HOMENET recommendation

Make a recommendation to turn on the things that already exist: DHCPv6 PD, RIP/OSPF, ...

Add small enhancements where needed to ensure automatic self-configuration

What can HOMENET do for Jari's home network?

Already does routing, subnets, local DNS servers, etc.

HOMENET can help with turning routing automatically on, automatic prefix assignment, zero config naming services

New IPv6 considerations

Multi-addressed devices by default

ULAs available - open question whether we should use them

For stable internal addressing, not for NAT

Support included in RFC 6204 (Basic Requirements for IPv6 Customer Edge Routers)

One way to indicate traffic is sourced within homenet

The opposing opinion is that they complicate things

Global addressability (removal of NAT)

Though perhaps not reachability

- Depends on security borders and policies

Topologies

Basic network architectures:

RFC 6204

draft-ietf-v6ops-ipv6-cpe-router-bis

draft-baker-fun-multi-router

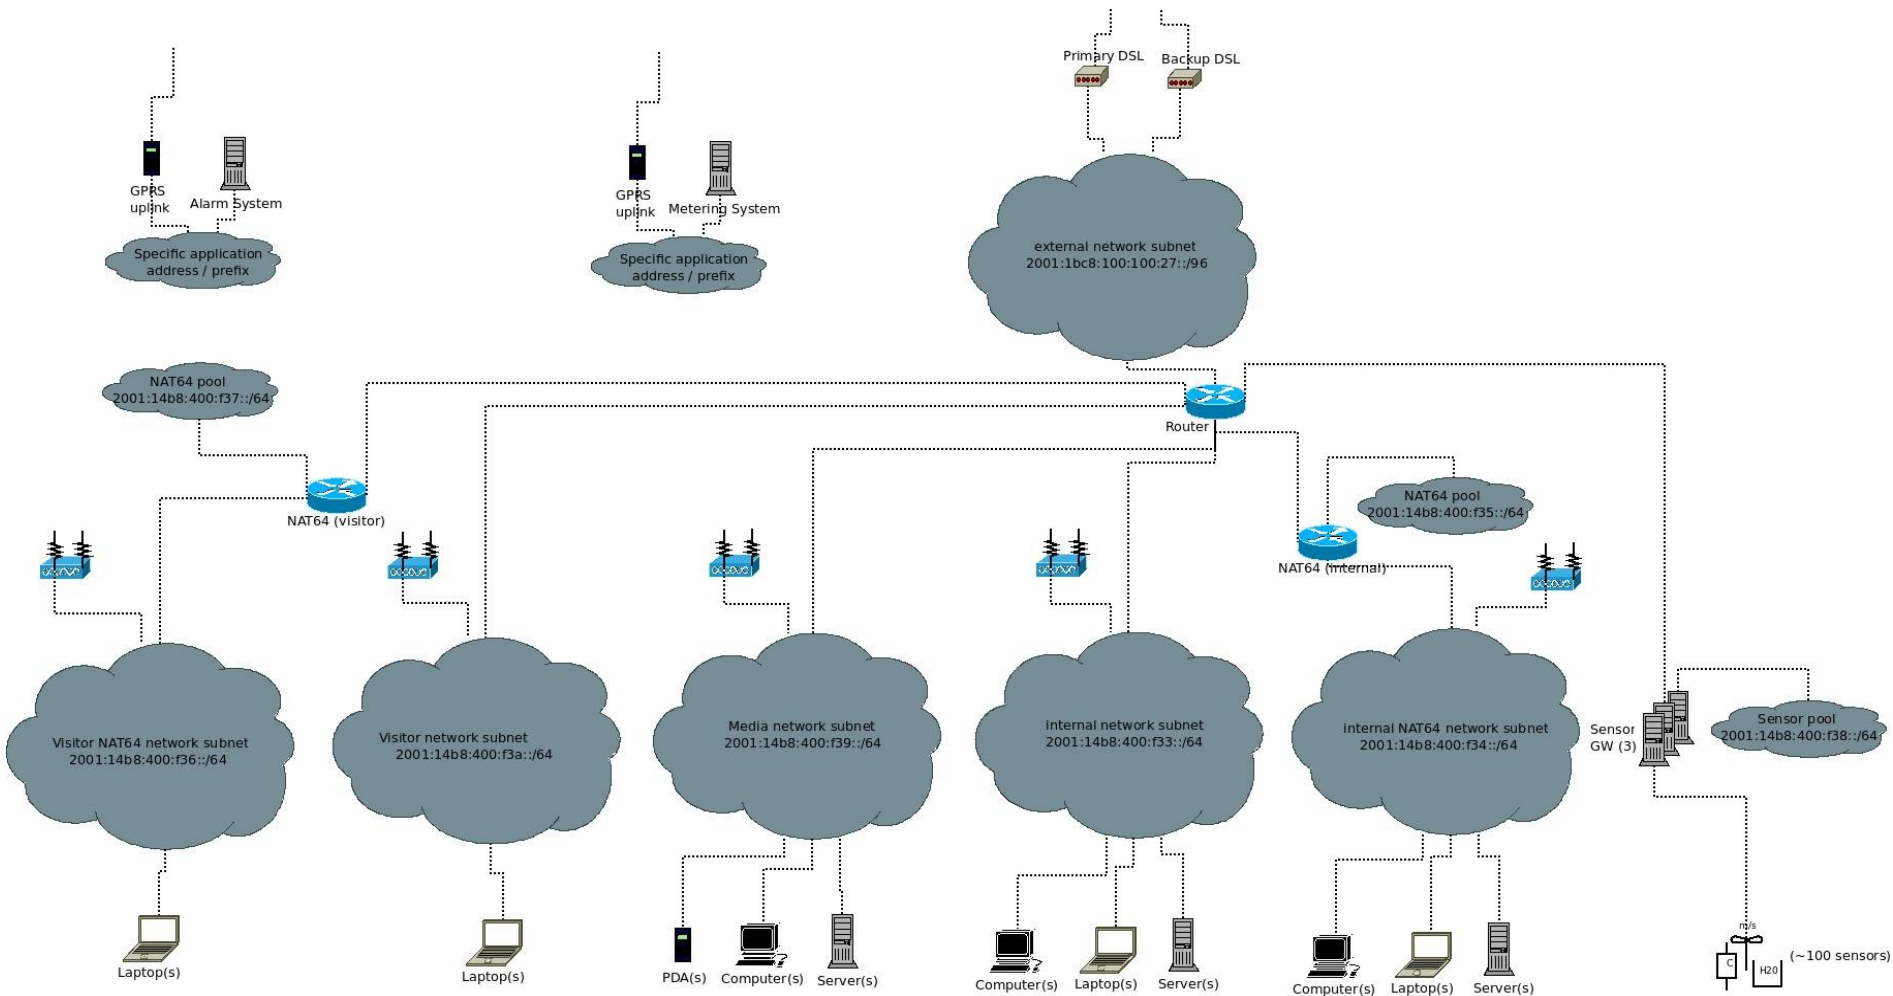
Potentially multiple subnets and routers

Five examples described in the draft

Multihoming in some form likely; must be considered

Heterogeneous link layer technology, mixture of old and new devices, routers, servers, and hosts

A non-trivial example



Topology considerations

May be affected by practical issues
e.g. chaining of devices

Or by policy issues
For private net, guest net, utility net, etc

Or by technical issues
Separating wired from wireless

Can we assume an arbitrary topology?

If so, may influence solution space

e.g. for internal prefix delegation arbitrary topologies work better with routing-protocol-like designs than with delegation designs

Requirements

Five areas set in Charter

- Prefix configuration for routers

- Managing routing

- Name resolution

- Service discovery

- Network security

Most deep dive work so far is on routing and prefix delegation approaches

Principles

The draft discusses architecture principles

Some implied by the homenet charter

e.g. dual-stack/IPv6-only, self-organizing

Others by consensus

Goal is to draw statement on homenet architecture from applying principles

Dual-stack homenet

The most likely deployment model today
IPv6-only in the future

Additional considerations need to be documented for v6-only: DNS discovery, need for NAT64, etc. – early text in the arch draft, may need a more full-blown description in its own document

Do nothing in IPv6 to break IPv4

But the IPv6 part may work in cases where IPv4 would fail

Route in IPv6 where IPv4 NAT is used today

Should include VM and ICS scenarios

Benefit from IPv6 addressability

Subject to reachability based on security borders

Transition tools out of scope

Limited to CPE, see RFC6204-bis

Self-organizing

Avoid manual configuration where possible

May be “secrets” to set for shared area, WLAN, etc.

Allow for differing ISP practices

May get varying prefix lengths by PD

Allow internal operation independent of ISP

Keep using the prefixes you have, even if connectivity goes down

Can take this even beyond leases... inappropriate but known to work...

ULAs could bring connectivity before the first ISP connection

Or a cleaner way to deal with beyond-lease connectivity problems

Prefix Assignments

Protocols TBD, only interested in the concepts and requirements here

Usable Prefix – Global IPv6 prefix delegated to a home

Assigned Prefix – A /64 automatically assigned to a given part of the home network

Main requirements

Assignments should be stable across reboots, power cycles, software updates, and preferably, simple modifications

Stability across major network reorgs is not a requirement

Reasonable efficiency may be necessary – One assigned prefix per usable prefix per physical network

Discover Borders

Different types of borders

Homenet:ISP

Private:Guest

Route:Bridge hop

Affects prefix assignment, firewalling

Some ways to discover the ISP border

Manual configuration

“Connect this port to the ADSL modem”

Probing, e.g., ISP interface has DHCPv6 PD but no routing

ISP-managed CPE router knows this a priori

Virtual CPE routers residing in the ISP premises know this a priori

Others?

Other principles

Prefer to re-use existing protocols

Conservative approach

Small enhancements towards auto-configuration

Routing protocol

May compromise availability with functionality

Multihoming

Just focus on source address selection problem

May imply routing based on src+dst

All other aspects of multihoming are out of scope

Avoid making future renumbering harder

The architecture

So what can we say so far about the properties of our homenet architecture?

Need to agree these to move forward

Will steer the deep dives into the five areas

Architecture (1)

Support multiple subnets and routers

- Route IPv6 where use IPv4 NAT today

- Maximize subnet size

Use link-state routing protocol (e.g., OSPF)

- May be able to leverage for prefix assignment (e.g. a la zOSPF)

LLNs, VMs, etc. can attach to home networks

- Either participate in the same manner or map to their internal mechanisms

Architecture (2)

Internal stable and efficient prefix assignment
/64 for internal subnets (and possible NAT64 use)

“Simple security” (RFC6092) +PCP +extensions
User-friendly security associations desirable

Local DNS servers and cross-subnet mDNS
Cross-subnet service discovery

Today constrained to local subnet

Open issues?

Completely arbitrary topologies?

Or make the least assumptions possible?

Discovering (security) borders?

Is multihoming part of Homenet v1?

Not happy with “Simple security”

How should we include “Advanced” security?

ULAs needed?

Discovery and naming across subnets

Relationship between unicast and multicast DNS