Homenet Architecture

draft-chown-homenet-arch

Tim Chown, Jari Arkko, Jason Weil, Ole Troan

Outline

- Approaches to standardizing home architectures
- Personal viewpoint on what would be useful
- Topology assumptions
- Principles (Section 3.4 from the architecture doc)

Approaches to Standardizing Homenets

- Operational – ”this worked well for me”
- Implementation commonality – ”this seems to be available in most devices today”
- Experience – ”we have enough experience of this to recommend it for everyone”
- Functionality – ”we need this feature”
- Specification – ”the IETF is here to develop new mechanisms”

Thoughts on These Approaches

- The authors are personally mostly in the operational – experience – implementation camp
- On the list we've discussed a number of topics that push the envelope further (multi-homing, sensor routing protocols, ...)
- Our recommendation is to take a conservative approach that starts from something that is known to work and then pushes the envelope a bit from that
  - E.g., add automatic turn-on for RIP/OSPF
  - Or develop PD further for internal prefix allocation
- But additional things could go beyond this
- Homenet versions needed? (WPA 1, WPA 2, ...)

Jari's Personal Viewpoint

- Thinking about the impact on my own network
- Going from my current state to something that would feel an improvement
- Some of the improvements are just a matter of turning on existing functionality, others would require new code
- A homenet recommendation will probably include practices from different categories:
  - Those that are already in use
  - Existing things that should be taken into use
  - New mechanisms that require new code

My Current State of the Art

- Multiple routers – yes
- Multiple subnets – yes (on v4, even more on v6)
- Routing – through RIPv3, but manually turned on
- Prefix assignment from ISP – manual
- Prefix allocation within my network – manual
- Naming – local DNS server(s), manually setup
- Service discovery, zero config naming – not in use / does not work across subnets or NATs
- Firewall – multiple, manually configured

Possible Homenet Recommendations 1/3

What I already do:

- Use an IPv6 router where you have an IPv4 NAT
- Use multiple subnets
- Hierarchical internal allocation of prefixes
- Use RIPv3 (or OSPF or ...)
- Run local DNS servers
- Place a firewall between guest / private, guest / internet, and private / internet

Possible Homenet Recommendations 2/3

What I should also turn on:

- PD from my ISP
- RIP on all routers (currently not on all)
- Simple security (RFC 6092)

Possible Homenet Recommendations 3/3

New things that we could specify:

- Some way to automatically turn on routing
  - Is this just apt-get install rip?
  - Or the interaction between RIP and PD?
  - Or some new on-the-wire negotiation?
- Mechanism to internally allocate /64s out of the shorter prefix we got from the ISP
  - For all the subnets
  - For NAT64 address mapping purposes
- Something to enable zero config DNS across subnets
  - If this works, service discovery should work too (?)

Topology Assumptions

Basic Network Architectures

- See RFC 6204, v6ops-ipv6-cpe-router-bis, draft-baker-*
- One router, one subnet on the home side
- Or multiple subnets – guest and private
- Or multiple routers
- Heterogeneous link layer technology, mixture of old and new devices, routers, servers, and hosts

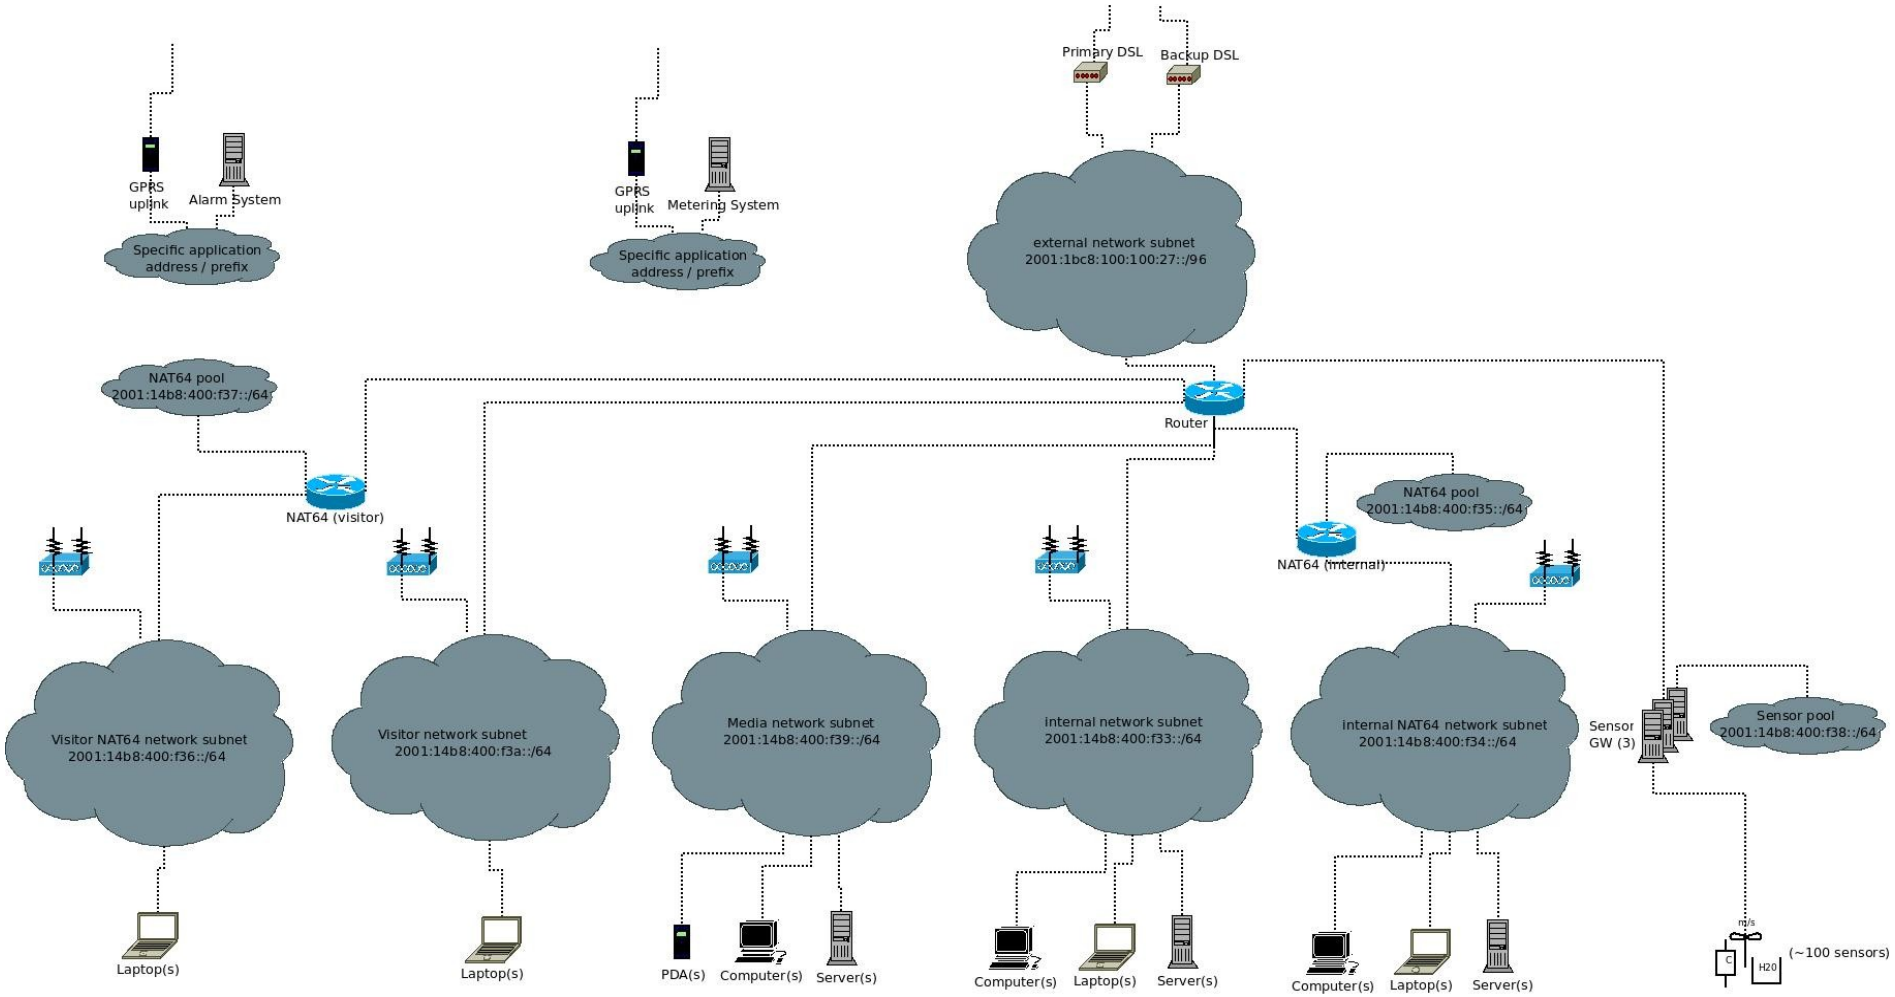
GPRS uplink
Alarm System
Specific application address / prefix
GPRS uplink
Metering System
Specific application address / prefix
Primary DSL
Backup DSL
external network subnet 2001:1bc8:100:100:27::/96
NAT64 pool 2001:14b8:400:f37::/64
Router
NAT64 (visitor)
NAT64 pool 2001:14b8:400:f35::/64
NAT64 (internal)
Visitor NAT64 network subnet 2001:14b8:400:f36::/64
Visitor network subnet 2001:14b8:400:f3a::/64
Media network subnet 2001:14b8:400:f39::/64
internal network subnet 2001:14b8:400:f33::/64
internal NAT64 network subnet 2001:14b8:400:f34::/64
Sensor GW (3)
Sensor pool 2001:14b8:400:f38::/64
Laptop(s)
Laptop(s)
PDA(s)
Computer(s)
Server(s)
Computer(s)
Laptop(s)
Server(s)
Computer(s)
Laptop(s)
Server(s)
m/s
C
H20
(~100 sensors)

Thoughts About the Topology (1/3)

- The actual number of subnets and routers should be left for actual deployments
- Affected by non-technical issues, like buying another device and chaining it
- Affected by technical issues such as possibly needing to separate wired and wireless for efficiency reasons
- Affected by policy reasons: guest, private, smart home, utility, kids

Thoughts About the Topology (2/3)

- The topology is also important in the sense of directing what kind of solutions we can use
- Arbitrary topologies work better routing-protocol-like designs than with delegation designs
  - PD extensions vs. RIP/OSPF extensions
- Looping and arbitrary topologies are IMO not the target of home networks, but do want to deal with the eventual misconfigurations?

Thoughts About the Topology (3/3)

Is multihoming included in scope?

- What kind of multihoming? Site? Host? Truly separate networks?
- Current deployments in home are mostly completely separate networks
- Some non-starters (IMO):
  - Host changes
  - Waiting for ID-locator split to happen
  - Requiring infrastructure in the Internet
  - Launching yet another IETF effort on this
- Possible OK approaches
  - Defer this for now
  - Require source routing (for ingress filtering)

The Principles of the Architecture

Principles

From Section 3.4:

- Dual stack operation
- Reuse existing protocols
- Self-organization & ISP constraints
- Intelligent policy
- Topology assumptions
- Transparent end-to-end communications
- Subnet size
- Routing functionality
- Discovery & proxying vs. extending

Dual Stack Operation (1/2)

- The practical assumption is that we operate in dual stack networks
  - (Of course, IPv6-only can work too)
- Homenet is about IPv6, but it is useful to understand the context
- At least we should not damage IPv4 operation
- The principle of causing the least amount of surprise would lead us to place an IPv6 router hop where today there is an IPv4 NAT or router
- Here we have a potential to do better than IPv4 – routing would allow inbound connections

Dual Stack Operation (2/2)

- Transition tools are probably already sufficiently covered elsewhere
- In any case, homenet is probably not the place to do more transition related work and transition tools are mostly on the border router to the Internet anyway, outside our scope

Reuse Existing Protocols

- Our effort is naturally not about inventing new protocols
- Use existing prefix delegation, routing protocol, etc mechanisms; in some cases extensions may be needed
- Should not preclude new or emerging protocols, however – how far towards the ”new” should we extend ourselves?
- Backwards compatibility with deployed equipment
- Standards & practices from other SDOs

Self-Organization & ISP Constraints

- Home networks need to come up all by themselves (granny connects her new IPv6 router to the wall socket and it should just work)
- We should allow for different practices on the ISP side
  - Prefix lengths: /48, /56, but what about /64?
  - Static / dynamic prefix delegation?
- Internal operation should not depend on ISP

Intelligent Policy

- Do not hardcode addresses or security policies into either the code or configuration of the devices
- This is a particular problem with prefix allocations and firewall functionality
- Do not make future renumbering events harder

Topology Assumptions

- Already discussed

Transparent End-to-End Communications

- IPv6 architecture should allow this
- Note the difference between ”reachable” and ”addressable”
- Security may be applied between nodes
- May expect use of PCP or uPnP
- What about ULAs? Used for all nodes?
  - Lets not associate ULAs with the use NATs

Subnet Size

- Most IPv4 networks are one subnet today
- Initial IPv6 deployments will also be similar
- But there are many good reasons for multiple subnets: policy, cannot be bridged, …
- Principle to use the same topology for IPv6 as for IPv4, but use routing instead of NAT

Routing Functionality

- Much of the list discussion was on ”best” protocol
- Can one protocol be chosen? Or should there be negotiation?
- RIP is the simplest and most widely supported. OSPF/ISIS are better designs, but is more expertise or code needed?
- New MANET/LLNs protocols another alternative
- Need a better understanding of the requirements
- Deeper dive on this topic this afternoon

Discovery & Proxying vs. Extending (1/2)

- Naming and service discovery should work across the home, not affected by NATs or subnet boundaries
- Subject to policy restrictions (e.g., guest)
- Existing protocols are generally constrained to one link and one multicast domain

Discovery & Proxying vs. Extending (2/2)

- Existing protocols are subnet scoped, need some way around this, either by
  - Extending them beyond subnets, or
  - Proxying them beyond subnets
- Does this affect just a couple of protocols (mDNS, LLMR, DNS-SD) or is this a more generic issue with application protocols?
- More analysis needed on requirements and the impacts of solutions