

(Notes as Taken by Juan Carlos Zuniga):

IETF Homenet Interim Meeting

Philadelphia, PA

6th and 7th of October, 2011

Chairs: Mark Townsley <townsley@cisco.com>  
Ray Bellis <ray.bellis@nominet.org.uk>

Area Director: Ralph Droms <rdroms.ietf@gmail.com>

Tech Advisor: Acee Lindem <acee.lindem@ericsson.com>

Slides at <http://trac.tools.ietf.org/wg/homenet/trac/wiki>

Note Well

Agenda Bashing

09:30 - 12:00 - Architecture Working Session (Timm Chown,  
Jari Arkko)

Jari - Homenet Architecture draft-chown-homenet-arch

Tim Chown, Jari Arkko, Jason Weil, Ole Troan

Approaches to standardizing homenets: operational (works well for me), implementation commonality (xyz have it), experience (we have enough experience to recommend it), functionality (we need this feature), specification (IETF has to make it)

Authors are in operational experience

List has pushed the envelope a little further

Acee Lindem - Question on requirements: given the discussion on the list, how do we reach consensus?

JA - we have to present the case, understand and then make a decision

Ralph Droms - How do you see the deliverables coming out? IETF Recommendations? In service discovery there are many options. Are we going to pick one and bless it, or are we going to look at all and recommend how to make them work

Michael R - Are we going to look at all flowers?

RD - Any thoughts about the ones that are being implemented and sold today?

JA- We are not here to kill solutions. IETF normally produces specifications and people choose. Here however we should provide a basic recommendation (naming base). We might push the envelope to sensor routing or service discovery

Mark T - 6204 could be seen as homenet v0.0. We are now more focused on how we do it with multiple routers (at home).

RD - most ISPs give a /64. I don't want to make assumptions that this will be everywhere, as we could have more constrained scenarios.

JA - good example. If we want to support multiple networks behind a /64 we might get in trouble.

Jim Witt - this is a deployment case today, but should not be like this.

JA - They provide this today but they are working hard to provide more in the future.

MT - I have not seen a recommendation for /64 at home. Probably there is a 3gpp for handset.

Jim - This is a place where running code should speak.

JA - I meant this by conservative

Lorenzo - Running code today does not do this.

MT - there is, but you cannot get them off-the-shelf. In my company we decided to put 6204

JA - it is reasonable to specify what should be done

Lorenzo - There are routing problems that you cannot solve with RIP, DHCP PD. So we need it

Jim - Development goes in upstream OAM projects and often code is 5 years old, e.g. linux kernels

MT - Some companies do it differently. We do it and Apple does it too

Lorenzo - are you suggesting write code and feed it in the upstream

Jim - Yes

Michael - you can have "orange and blue" versions for specific problems

AL - a lot can be accomplished if the box at the edge it treated as such

Michael - 1% of users with problems cause 50% of support calls

Jim - badly home boxes create these problems. Cheap boxes use open source codes. If we can create this code then "hopefully" the small guys will pick it up

Michael - As ISP, we often replace the box at the users' premises by a less costly box that actually works.

Jari's view

Probably we will create recommendations on  
Things that exist  
Things that are there but not used  
Things that need to be written

Possible recommendations (2 slides):

What I have: Run local DNS servers...

What I should do: Prefix Distribution from ISP, Simple  
security 6092

Michael: How to select when I connect my phone to the PC,  
they both connect and they do tethering.  
You can only solve this if your support VM

Basic Network Architecture

6204, v6ops-ipv6-cpe-router-bis, draft-baker-\*

Network Topology

Multihoming

Mainly a problem for the host

The home should support multiple exist routers, and should  
let the host know which one to use

JA - 484 ingress rules we can rely on

Michel - tow home network

Alex R - So are we considewring the Femto case out of  
scope?

Lorenzo - the netorks cannot deal with end to end. This is  
a role of the host

Fred Baker- I don't agree, you can use 6296 (nat66) and  
maintain connectivity.

Lorenzo - how about applications that carry IP address in  
payload?

FB - they are broken

G - we have to make sure we replicate walled-garden  
services "bring the information all the way to the host"  
and let it decide (e.g. DNS server selection for walled  
garden services from two providers)

Jason Livingston 4191

Principles - Assumptions (Tim Presenting now)

Dual stack (of course IPv6-nly too)

IPv6, but not damaging to IPv4 operation  
Transition tools are covered elsewhere

RD – We don't want to rule out the case when you have a NAT  
(even if connected by mistake)

AL – Do you want to do NAT64 in home or at provider?

JA – We will not address transistion and will not overlap  
with other groups. We can however have IPv6 only  
recommendations for home networks that mention what other  
should be used or not.

Lorenzo – Are we taking a versioned approach?

RD –

JA – Not saying that we cannot solve everything. WE ahave  
to carefully think what we can do o

Lorenzo – What users are going to find useful

JCZ – 802 service discovery?

MT – for a Linksys discovering if you are connected to  
another Linksys or to the internet

It would be interesting (e..g do I turn on Firewall, NAT or  
not?)

Alex – FF and WBA We are going to assume how to connect the  
two networks

MT – CEA, BBF, DLNA, 802 (service discovery)

RD – ZigBee alliance, looking at service and naming  
discovery heavily

(presentation) We should allow for different practices on  
the ISP side

/48, /56, but what about /64?

Static/dynamic prefix delegation?

Intelligent policy

Do not hardcode addresses or security policies – problem  
with security and prefixes

Transparent end-to-end communications

May expect to use PCP or uPnP

What about ULAs with the use of NATs?

MT – moving from link-local to multiple links in home, we can use ULA  
Ole – Are we considering a homenet without an ISP?  
RD – You can setup ULA without ISP and then as soon as you get ISP you disable ULA  
MT – I don't agree, ULA could still work  
Alex – how about DHCP  
Michael – Similar to ULA  
Lorenzo- your printing job will be stopped when your DSL comes back... not good

We have identified advantages but there are drawbacks  
6204

MT – if v6 can detect if this is a WAN or LAN port  
RD – Are we going to require the v4 topology in v6? Eg v bridging v4 routing? If they have to be congruent, no issue. If they don't have to be congruent then we have to look at that  
??? (black shirt) – this may not be possible due to chipsets in the market  
Michael – in cheap boxes this is mainly in SW  
Jim – today chipsets do VLAN and many other things.

Existing protocols are subnet scoped (mDNS, LLNR, DNS-SD)  
JA – How about gaming?

MT – Missing elements?

(Black) – Energy aware elements in the protocol?

Erik Nordmark (remote) – ??

MT – ??

EN – Can you construct loops that break v4?

JA – This is a problem on itself and we cannot fix it

MT – We need to understand how complex the solution is if we want to handle loops and then decide if we want to address this problem or not

Lunch

13:15 – 15:15 – Prefix Configuration Working Session (Ole Troan)

## Requirements list

Are names specific?

MANET and RPL

OT – you might want to know if you want to assign a prefix a link to out again

JA – Do you need this?

MT – Defection of boundary affects PA distribution and many other things (blue and red cables)

Arbitrary Topology – Must (for real cases) Nice (for pathological cases)

Multiple sources – Must (support for more than one ISP, Femto, etc)

Stable prefix – Must

JA – resilient to crashes, reboots and lease expiry

Jim – Stable addresses are not meaningful. Name stability is a critical

## Multilink subnet routing

Are we ok supporting /64 only?

MT – Erik and Ole should analyse if this is valid without host changes

Jason – we are safe assuming multiple subnets and more than /64

Michael – if one /64 addresses 90%, can we assume two /64 address 99.9, or ISPs prefer /63?

Lorenzo – current HW would override old PA and would assign the new one

## Hierarchical DHCP

Flat DHCP prefix delegation draft-baker-homenet + RFC 3633

RD – You can assume that routers will not assign orefixes on links that already have

FB – multiple routers on the same link

JA – there is a fundamental question whether resources are taken from a pool, or a distributed algorithm chooses resources not being used

## Zeroconfig OSPF

Lorenzo, AL – Allows border discovery, could help passing key information

JA - only OSPF?

MT - you could make it work in other protocols

Jim - writing information every few minutes is fine.

Writing every few seconds is bad

Lorenzo - By widely delayed we mean open implementation that has been accepted.

MT - Timecheck: go back to table? Continue?

Jim - OLSR is of use, as there are mesh wireless deployments in Europe using these WiFi.net. The battlenets run one vs the other

Ole - Question is do we overlay routing protocol and prefix assignment?

LC - link state is very useful.

LC - how difficult is zeroconfig with OSPF?

AL - not much

JA - can we have one single thing that does all for you?

MT - this would make it easier to implement

JA - still you have to define extensions and options

MT - but still one spec

Break

15:45 - 17:30 - Routing Working Session (Fred Baker)

Routing presentation - Fred Baker

Some link layers have special requirements (e,g, 802.15.4)

Analysis documents (IPv4 examples)

Use cases

Single router (simplest)

Multiple routers (needs to be addressed in case user adds them even without realizing)

Multipath networks

Not great for RIP

Good point for OSPF or ISIS - these can help prefix allocation

Jim/Fred: we all have built this

Multihoming

Multiple UL

If you assume BCP 38: RFC 3704 :fix problem of packets on wrong router

ISP will filter packets by source address (BCP 38)

FB: source-based routing would wolve this. You can change the routing at the host, or you can re-send the packet

Lorenzo: do you assume IGP?

FB: You need to solve the problem when you have OSPF lifetime of ~45 min. You can detect loss in 40 sec, which is much better than 45 min. In case of router failure, the best would be for the remaining router to advertise RA with zero lifetime

JA: there are other requirements for multihoming: utility devices require a specific ISP, IPTv from one ISP and Internet from another.

Lorenzo: We might need to do source address routing

Michael: we need to support the case when one ISP goes down ?? (squared shirt): E911 needs to be handled

Lee: the second ISP will get the call and you don't want it

MT: we are back to 3484 and Japan problem

FB: the argument tells me that we want to stay away from RIP and get closer to OSPF/IS-IS

## Protocol Styles

AL/Jim/Lee: the issue with the footprint of the code depends on the vendor and business case. However, the code would eventually would need to be made available

## Proactive vs reactive

RIP, OSPF, IS-IS vs AODV and RPL (useful when routes change and appropriate for 802.15.4 networks)

Michael: There could be a use case for both

FB: Today, I would recommend RIPng, OSPF and RPL

Jim: why are the wireless use OLSR

Lorenzo: if we wna tto be compatible with PA we need to support a protocol with TLV

FB: Then IS-IS?



AL: You could do it with OSPF if the information is different to the existing one

Lee: so do we need a dynamic routing protocol?

FB: RIPng would solve use cases 1 2 and 3. For 4 we have the timing issue.

Lee: 30 mint acceptable, but 3 probably. I don't know if we need to go to the seconds

RD: If you don't want the full dynamic protocol we have to compare to a DHCPv6 and small protocol, in which case probably the size is not much different

Jim: the implementation details matter here. We need to see the size of the compiled code and the amount of RAM it takes

JA: I'm ok with supporting OSPF and PA distribution. I'm not sure I'm ok with all the multihoming requirements

MT: If you don't care about these, then what would wyou care

Lorenzo: In v4 no issue, but with v6 this is and will be an issue

MT: Remarks – We are trying to tell how to do v6 in home before the product managers realize they need to do it and decide to do the same that was done in v4

09:00 doors open

09:30 – 12:00 – Architecture Working Session (Timm Chown, Jari Arkko)

12:00 – 13:15 – Lunch

Friday October 7

09:00 doors open  
(ISP issues)

Agenda Bashing

9:40 – 10:15 Security (Mark)

Mark Townsley

Automatic Border Detection is essential  
For service discovery

For prefix assignment and routing  
For security  
Default filters (ULAs?)  
Firewall

Mike: Smart grid borders?

FB: Utility might want you to use RPL

RD: Smart grid is an application

Lee: At least service discovery and security would be nice to have in the same place

AR:

MT: A link between two routers need to be identified.

Home-home

Home-ISP

Home-Utility network

FB: We can have the home become a customer of a number of networks (upstream). This means ISP and Grid are outside. Outside can be to 1) Internet and/or 2) services  
Japanese TV still a different animal

MT: We might want to give some naming to these borders

Detect incoming packets.

Allow incoming connections from your home

Allow incoming connections from internet

Providing borders with ULAs we can define this

"Local" security boundary defined by:

ULAs

Link-local

Prefix pushed down by R

Magic?

Jim: state has to be distributed in case there are multiple routers

Mike: LAN party? With smartphones today people can roam to your home LAN and get access to your content (e.g. photos)?

Lee: we need an authorization mechanism for L2 connectivity

Lorenzo: Security is L7. For L3, visitor in your house is the same network

Ole:

Mike: RAs will have to carry the ULAs

JA: If service differentiation exist. Do we have equipment

James W: Some applications make sure that they are talking to a link local address.

Lorenzo: Today with v4 we cannot limit access to link local  
Square???: xbox would do application level discovery

MT: we are talking about single-homed, multiple routers, multiple networks, that's why we are comparing ULA instead of link-local

## Advanced Security

Lee: statefull firewall

MT: It is a smart Firewall. IPS: Intruder Protection Security

FB: You don't need to standardize this

JA: you might need to standardize the format to configure

James Woodward: Is security needed if it doesn't work

Lee: existing mechanisms are insufficient, however, they are useful

JA: we are defining how to transport magic, not the magic itself

MT: even if we define the content, it still has a place in the architecture definition

Mike: I see UPnP and PCP.

Lee; are we discussing whether we need Advanced security or whether it is in scope defining it?

MT: We agreed that we need to detect borders in general. We might need to define how the borders communicate their existence to neighbours with the protocols.

(\*\*\* MT to check this wording and perhaps expand)

(Inserted by MT):

I believe we agreed that the homenet arch document would outline 3 possibilities:

1. "Transparent mode" or "end to end security" mode which may have basic filters for ULAs, but does not restrict traffic flow based on global addresses (e.g., "no firewall")

2. Simple Security

3. Advanced Security

There was little hope that the group would come to consensus on which of the 3 options would be recommended when a border (e.g., homenet to ISP) was detected, but that at least we should document the 3 alternatives and associated tradeoffs.

10:15 – 12 Naming/Discovery Working Session (Ray and Stuart via Skype)

Michael: Why the difference to sometimes mention devices, sometimes services?

FB: we might need to discover both

FB: in this homenet, these should be resolvable from anywhere within the homenet

Existing Protocols

DNS-SD / mDNS

LLMNR RFC4795

SSDP (UPnP)

SLP RFC2608

All are typically link local only

Unicast DNS

Anything else?

RD: Multicast DNS

Jim: in some cases you might want to do site-local m'cast propagation and in some cases you might not

RD: Extension to site-local addresses for multicast might be needed

FB: then you might need PIM

RD: the question is the group definition

Ole: This is how m'cast works

Tim: well known addresses are defined, but the scope still is not

Ole: what does this have to do with discovery?

Ray: we need discovery for these

JA: DNS could use a trick to discover some addresses only, in a m'cast fashion

Michael: Are there bozes that work as proxied between unicast DNS and bonjour?

Ray: uncertain

mDNS uses ".local"

Do we recommend something like a ".site"?

Mike: you can get site local names?

Ray: Names can be resolved even if they cannot be browsed. Given its service, how to reach it, and how to operate it

More namespace

?: should not IETF register .site with IANA?

Ray: at the RFC yes, but this is about single-label names

RD: devices at the edge might name themselves as  
edge.router.com

JA: is this in scope?

Ray: if we encourage single-labels it would help

JA: Are we then suggesting a change to networks and hosts?

Ray->Olafur: is it worth encouraging their obsolescence

Ray: Recommendation for a single register TLD

Michael: are we going to discuss wall gardens as part of  
this?

MT: they touch everything

Michael: you don't know the name, you don't know the  
address. One ISP is no issue.

RD: MIF is addressing this issue about DNS selection

Michael: we have to make names resolvable

Lorenzo: if two wall gardens register "television" there is  
nothing you can do. If they register television.ntt then  
you can resolve. You cannot have these registered

Michael: Trying to reduce the pain of using wall gardens

AR: connection managers can solve this issue

Lorenzo: Only if you control the OS

Michael: but how to get the walled garden info

Lorenzo: if you pay, you get it

Lee: but these services are outside the home

Michael: but if I have to resolve a name, should I resolve  
the ISP, the NTT television, or my own television. Even if  
they overlap, it should be the site local television that  
wins

Break and off-line conclusions

Lunch

Jim Witt — Running Code presentation

Why CeroWrt?

Test platform for bufferfloat work

Development running in OpenWRT

D-link: All APs do m'cast over unicast for 802.11, since you have to associate anyway.

DNSsec issue for time, as you need it to run time, and you need time to run it

<http://Cero2.bufferbloat.net/cerowrt>

Lorenzo Colitti – Architecture Strawman

Connectivity goals:

No host changes

No NAT anywhere

Allows communication between hosts even if ISP links are down

Routing goals:

Automatic configuration

Addresses

Firewalls

Support arbitrary topologies including loops

Survive loss of any router for any time period

Works regardless of router boot order

No tree topologies like ULA

External connectivity goals

Support multiple uplinks and ingress filetering

Support uplinks that provide partial routing

Walled garden networks

Reasonable startup times

Architecture (possible solution)

One connected network, one naming scheme

Routing protocols hold together

ULA is "inside" the network

Links not in routing protocol (e.g.. ISP with DHCPv6 PD)  
constitute security boundary

(figure)

Two accesses (CPE), two prefixes /56

ULA /48

Prefix assignment

Mesh routers from /64 from aggregates

Use routing protocol collision detection

Assign /64 to routed interfaces

When aggregate goes away, deprecate /64

Border routers inject aggregates into mesh

DHCPv6 PD

Walled garden

ULA aggregate for local /48

Anchored to one of the routers

RD: Why can't we do bridging anywhere but except in special cases?

L: because there are special cases

MT: I mentioned bridge where you can and route where you can, but in Quebec people oppose heavily

JA: But I don't have ULA nor bridging and it works

L: You are not a typical example

Michael: Whatever routing protocol we propose will be well received by enterprises

Routing

Border routers inject destinations into mesh

Ech route has "Acceptable" prefix

Do src+dst routing (s+d, or s, or d)

Security

If you are a BR, you turn on the firewall

Block ULA across border, etc (simple security)

If you hear routing protocol on interface, join the mesh

Secure protocol MD5

Hash WPA

(one password or two buttons is acceptable to a user)

Guest networks

Add attribute to TLV

Indicates string as "guest"  
Rely on routers to firewalls between realms?  
Or just rely on src+dst routing

MT: this guest network can be the utility, or your neighbor. This is a third case in which between two networks you share some things

Naming

I know nothing about naming

Two possibilities

Local naming anchor similar to ULA anchor, DDNS

Multicast DNS, anyone who has a name answers

For external, walled garden names

BRs get DNS domain from DHCPv6

Inject into mesh, like routing aggregates

RD: this is similar to what MIF is doing

Michael: is the host getting connection to several DNSs?

L: Yes, you have to tell the truth to the host

M: once we can provide the prefix, then we can provide the address that will provide the service

L: might need to multicast

L: no way to mandate source address

L: do we require a authDNS? How to get the info to the host about what DNS to use?

Routing protocol

Zero-config

Obviously

Link-state

State convergence

Gives clear idea of who is in a who is out

Src-dst routing

Necessary for multihoming

AL: not easy to have one password because you have to have the same one in all devices

L: I think that if you have a password then you can do it.

Lee: you can say one thing is authorized without telling "the thing" that it is authorized

Michael: your model is that users need to type one password on every box. My model is that they plug things and type only one password. If they need more functionality, then they have to type passwords and do more

AL: random id generation



Japanese activity about home router guidelines  
Akira

Guidelines version 1 and 2 (will send URL)

13:30 – 15:30 Architecture Working Session II

Architecture Discussion-Conclusions  
Jari Tim

Similar to Lorenzo, but no ULA

Key conclusions and non-conclusions

Conclusions

Focus on running code plus some improvements

We could do baseline version and then add improvements  
later

Route where you had an IPv4 NAT seems acceptable

Running IPv6 only requires documenting additional  
considerations

We understand the requirements for prefix assignment in a  
home network

M: ISPs with two gwys in v4, v6 can survive. We can do  
things that do not work for v4

JA: it is fine, but we should not cause v4 to break

M: we have to document it

MT: it is in the charter

R: is suggestion to only run v6 if v4 doesn't work

JA: it would be bad to say what to do in v4, but ok to show  
the problem

Link state routing protocols (OSPF) seems potentially  
doable to solve prefix assignment and other

LLN, virtual machines, etc can participate or map their  
internal mechanisms

AL: if you have lossy interfaces they will not participate

Ole: borders across links or borders across nodes

M: It will probably be OSPF

M: I agree LLN, but virtual machine is a router that should  
be part of the homenet

JA: Agree

M: how to get loop avoidance if devices have another non-  
homenet below?

L/JA/RD: not an issue

M: machines running multiple instances shown no t use multiple prefixes  
If multihoming support, primarily about using right source address and avoid ingress filtering, the rest is up to hosts and applications

AR: Should we not include the DNS?

JA: part of the walled garden etc, yes, we should add that this is the MIF

Not happy with Simple Security

Need to Discover borders

JA: Not clear is labeling is sufficient

MT: you cannot avoid people making errors, so we need an automatic mechanism. Eth to ISP and Eth to PC look the same (LAN/WAN)

JA: are you happy with Lorenzo scheme?

MT: it should be able to advertise something in the protocol

JA: there is nothing today. Is Lorenzo's solution enough?

M: can't you try all optional on all ports and then find out?

?: like try a routing protocol and then decide depending on whether it works/not?

JA/MT: Lorenzo's proposal is good, but need to analyse if there are border cases

Jim: how about if the upstream is a shared link?

Apparently this might be a problem. Need to analyse

Need to do discovery and naming across subnets

Non-conclusions

(To Be Continued by another minute taker...)

End Notes Taken by Juan Carlos Zuniga

Begin Notes Taken by Michael Richardson <mcr@sandelman.ca>

security: we need to come up with a nomenclature for the different

kinds of borders.

instead of ISP/SmartGrid.

Lets talk about home being a customer of some number of networks.

In each case we have a border.

Network:Upstream

What about apartment buildings where the north interface might be shared

across many tenants?

Also what about university residents.

Jari Key Non-Conclusions

NOTE: we need End to End Security (MUST)

Simple Security (MAY)

Advanced Security (MAY)

Recap of name discovery vs boundary system.

What do we know about the problems of using ULA?

We either make ULAs work or rip them out of 6204.

What is the relationship between the Multicast and Unicast, do the

information passed between, and is there a cache for dozy/sleepy devices.

How do local names that want to be in the global name space, get

pushed up. Could this be zeroconf in a vendor's ".net" space (e.g. ".mac")

Yari's take-2 slide, discussion softened some of recommendation,

particularly Simple and Local DNS.

DTCP - television transmission.