Best Current Operational Practice for operators:

IPv6 Prefix Assignment for end-customers – persistent vs non-persistent and what size to choose

Jordi Palet

jordi.palet@theipv6company.com

Authors:

- Andrew Alston
- Gert Doering
- Jan Žorž
- Jen Linkova
- Jordi Palet
- Kevin Meynell
- Lee Howard
- Luis Balbinot
- Mark Townsley
- Primož Dražumerič
- Sander Steffann

Draft v2 meeting:

RIPE BCOP TF

https://www.ripe.net/participate/ripe/tf/bcop

https://www.ripe.net/publications/docs/ripe-690

Table of Content

Executive Summary

- Making wrong choices when designing your IPv6 network will sooner or later have negative implications …
  - **IPv6 is not the same as IPv4**. In IPv6 you assign a short prefix to each end-customer site, so they are able to have as many subnets (/64s) as they need.
  - It is **strongly discouraged** to assign prefixes longer than /56. If you want a simple addressing plan, /48 for each end-customer.
  - In order to facilitate troubleshooting and have a future proof network, you should consider **numbering the WAN links using GUAs**.
  - Non-persistent prefixes are considered harmful in IPv6 as you can't avoid issues that may be caused by simple end-customer power outages, so assigning **persistent prefixes is a safer and simpler approach**.

BCOP and Why?

- Describe best **actual** practices
- Target: ISPs deploying IPv6
- Lack of experience or following IPv4 practices bring unexpected or unwanted results
  - IPv6 “brokenness” = Content providers rejection of your AS
  - **Lack of compliance with new standards such as Homenet**
    - Complete **production** network renumbering, etc.

Size of end-customer prefix

- /48, /56 or something else?
- Change your mind, this is not IPv4!
- IPv6 has been designed to assign prefixes not addresses
- Tony Hain “maths”:
  - IPv6 lifetime over 480 years, and keep doing that several times
  - Scarcity of addresses is not going to be our next problem

/64 ?

- **DO NOT DO THAT!**

/64 ?

- **DO NOT DO THAT!**
  - **–NEVER!**

/64 ?

- DO NOT DO THAT!
  - NEVER!
    - NO WAY!

/64 ?

- DO NOT DO THAT!
  - NEVER!
    - NO WAY!
      - BROKEN!

/64 ?

- DO NOT DO THAT!
  - –NEVER!
    - NO WAY!
      - –BROKEN!
        - »VERY BAD FOR YOU

/64 ?

- DO NOT DO THAT!
  - –NEVER!
    - NO WAY!
      - –BROKEN!
        - »VERY BAD FOR YOU
        - »BAD FOR YOUR CUSTOMER

Numbering the WAN link

1. /64 out of the end-customer prefix
2. /64 out of a dedicated pool
3. Unnumbered
4. ULA

/64 from customer prefix

- Use the 1st /64 from the customer prefix
  - draft-palet-v6ops-p2p-from-customer-prefix-01
  - Simplifies routing and provisioning
- Some CPEs may not support RFC6603
  - Prefix exclude option for DHCPv6-PD
- Even being required by RFC7084
  - Basic Requirements for IPv6 CPEs

/64 from dedicated pool

- Most common scenario
  - Dedicated pool for WAN links
- CPE performs router discovery
  - If it is a host (PPPoE), setup is completed
  - If it is a router, will request a prefix (DHCPv6-PD)
- /126, /127, /112 or /64?
  - RFC6164 suggest /127
    - Not all hardware supports it
    - /64 is future proof
    - Hardware limitations for longer than /64 prefixes
    - Allocate /64, use /127 to prevent ND attacks
- If there is *always* a CPE, you can apply security policies w/o harming customers

Unnumbered

- Don’t use GUAs
  - Instead use Link-Local
- Doesn’t work for all the devices, which can’t request DHCPv6-PD
  - No GUAs means no traffic …
- Complicate troubleshooting
  - Not able to traceroute the point of failure
- Not suitable for unknown CPEs or non-CPEs attached to the WAN link
- End-host will stay unnumbered
- Some hardware may consume additional resources for numbered links

ULA

- Strongly discouraged
- ICMPv6 from the CPE to outside ISP
  - ULA source address will not traverse filters
  - PMTUD will break
  - IPv6 connection will break if Path MTU is not the same

WAN link summary

- /64 GUA is the recommended choice
  - From the customer prefix if RFC6603 is supported
- It may be even required when more that 2 endpoints
  - Managed bridges
  - Repeaters
  - Redundancy (VRRP, multiple routers)
  - Monitoring/troubleshooting devices

Prefix assignment options

- Align the size of the delegated prefix with a nibble boundary (multiples of 4 bits), so it match DNS reverse zone delegations
- A single customer network is /64
  - A single /64 is plain wrong
  - IETF work allows a single /64 for an interface
- Multiple /64 must be the rule
  - RIR policies allow /48

/48 for business, /56 residential

- Some operators do this
  - Rationale -> Marketing/Sales differentiation
- Advanced home users may have problems with this
  - You're not able to use all the 4 digits (/48-/56)
- Some may have already an addressing plan with /48 (ULA, TB, transition, etc.)
  - /56 forces to redo it + renumbering
  - /48 just means changing the prefix
- Alternatively, reserve /48, assign /56
- Are you considering SMEs?

/48 for everybody

- Most practical and pragmatic
- Less call-centre time to sort out problems
- Single “flat” provisioning system
- Same prefix size as ULAs, transition, etc.
  - Direct mapping of existing addressing plans

Less than /56

- Not recommended
  - Technically no reason for that, enough addresses, this is not IPv4!
    - Over 134 million /56 in a /29
    - Over 16 million /56 in a /32
- Ask for more space to your RIR if required
- Never assign a single /64
  - Except for cellular phones (1 /64 for each PDP)
- LTE modems still require /56 or /48

Persistent or non-persistent

- Persistent typically by means of AAA or custom provisioning system
  - At customer connection they always get the same prefix
- Non-persistent by means of a big pool in each termination point
  - At customer connection they get a random prefix
  - If persistent, the lease time may provide days, weeks or even months

Non-persistent is easier?

- Less effort to deploy
  - Issues come later
  - It comes from IPv4 practices, DHCP
    - But we have NAT!
  - Looks easier for aggregation
  - Not looking for “customer” portability
    - May be an extra service
- Commonly using DHCPv6-PD
  - Each end-customer device has a GUA

However … non-persistent is harmful

- In case of power failure, CPE hang-up, …
  - Common even in highly-developed countries
- CPE doesn't send prefix valid lifetime = 0
  - End-customer devices keep the old prefix
  - Will try to use it, will fail
    - Customers claims to the call-centre
- Content providers measure IPv6 brokenness
  - Will ignore your IPv6 traffic
- Power outage often happen several consecutive times …
- Non-persistent prefixes force a logging system

Best choice: Persistent or non-persistent

- Allow broadband services provided by the customer and the ISP
  - Allow stable DNS names
    - camera1.username.ispname.com
  - New business/apps/services, new incomes
- Key for non-residential customers
- Avoid having a logging system
- The WAN link still can be non-persistent

Questions?

Thanks!

IPv6 Deployment Survey (Residential/Household Services)

How IPv6 is being deployed? (January 2018)

Jordi Palet

(jordi.palet@theipv6company.com)

Consulintel, CEO/CTO

Survey Contents

- Basic ISP data (name, country, RIR)
- Technology of the customer link
- Is it a commercial service or a “pilot”
- IPv6 WAN link
- IPv6 customer addressing
- IPv4 service
- Transitioning and provisioning
- IPv6 DNS services
- Other data (optional contact details)

Note: Survey not intended for service to mobile phones, however, 2G/3G/4G response can be provided for service via a “CPE/modem”

IP version of Survey Responder

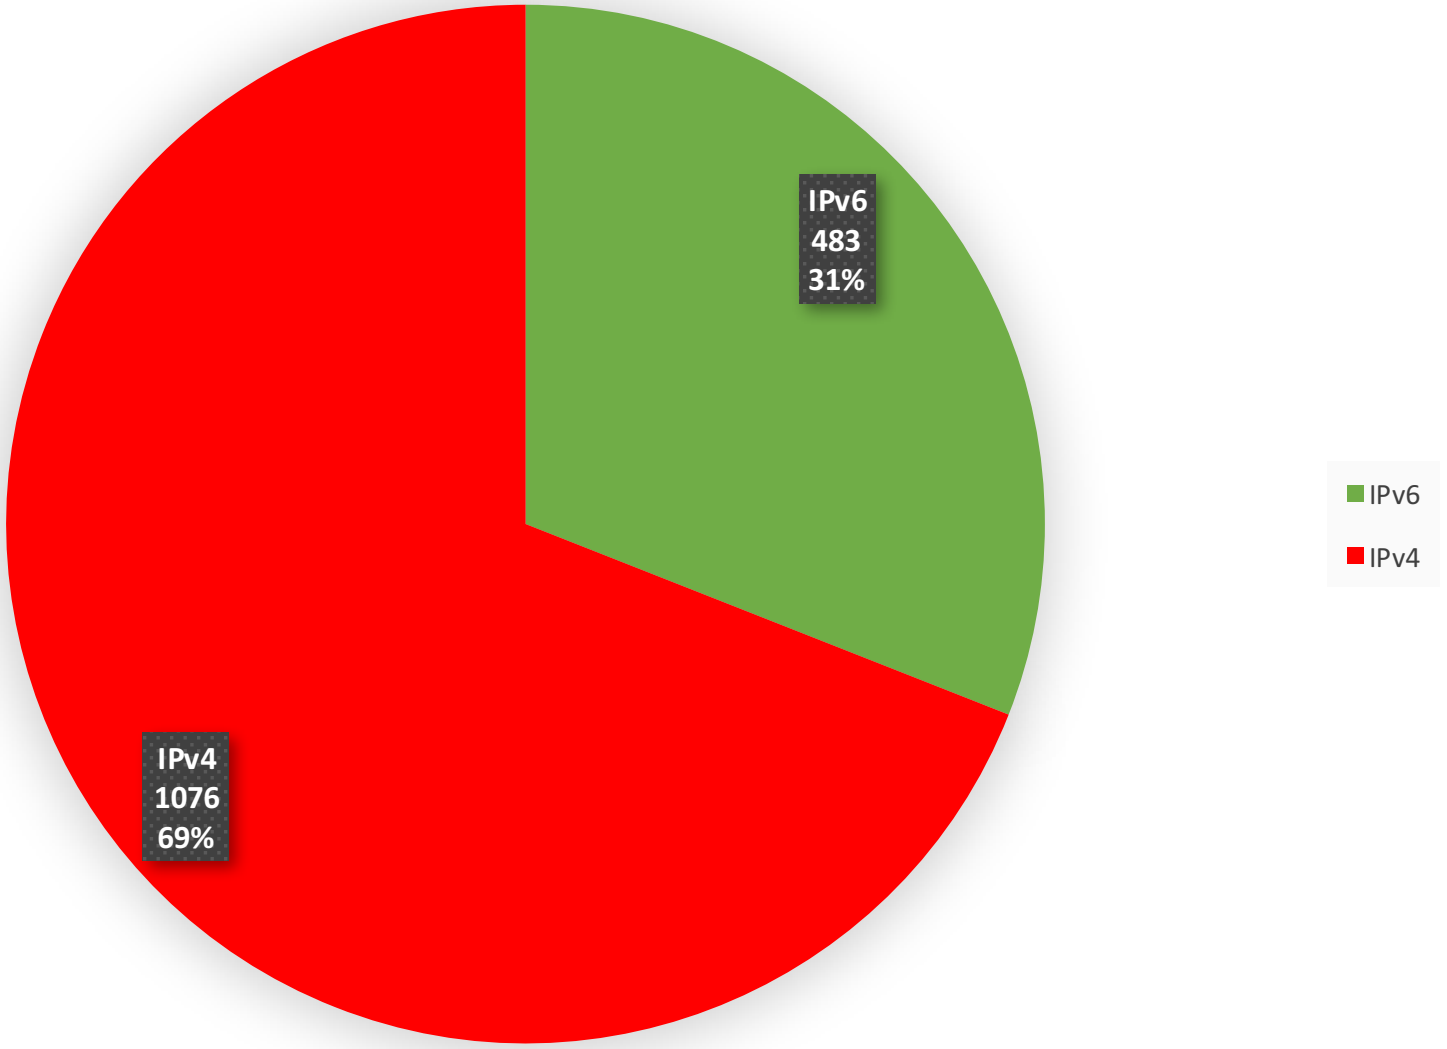

Who is responding?

- Looking at whois …
- ISP employees
  - From their own network most of the time
- Customers
  - Most of the time from their own residential networks
- Most of the responder “networks” have both IPv4 and IPv6 allocations
  - Responding with IPv4 from ISP network probably means, even if they have deployed IPv6 to residential customers, may be not in (all) the corporate LANs.
- Other observations, looking at bind and apache logs:
  - Happy-eye-balls timeout …
  - Is that anymore needed? Time to retire it?
  - Hiding IPv6 network problems?

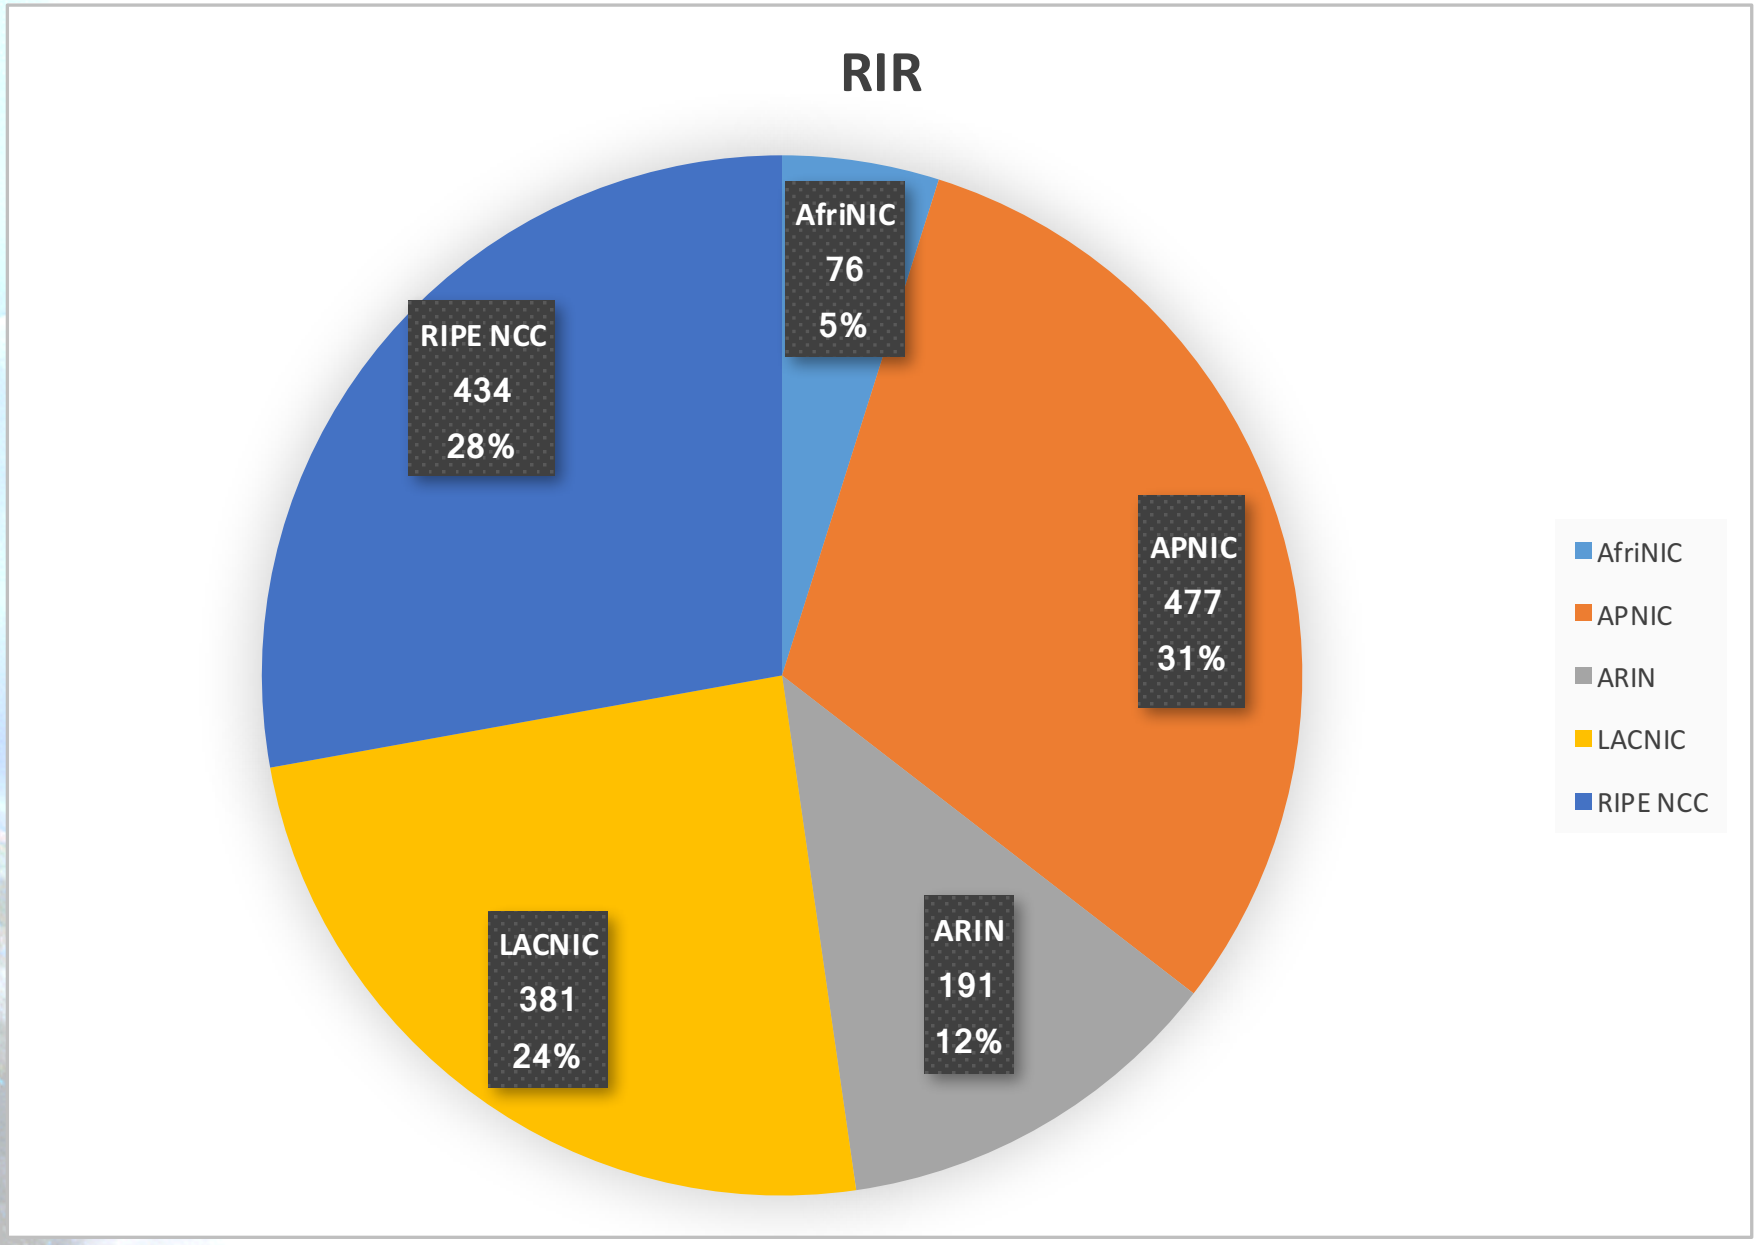
RIR
AfriNIC
76
5%
APNIC
477
31%
ARIN
191
12%
LACNIC
381
24%
RIPE NCC
434
28%
AfriNIC
APNIC
ARIN
LACNIC
RIPE NCC

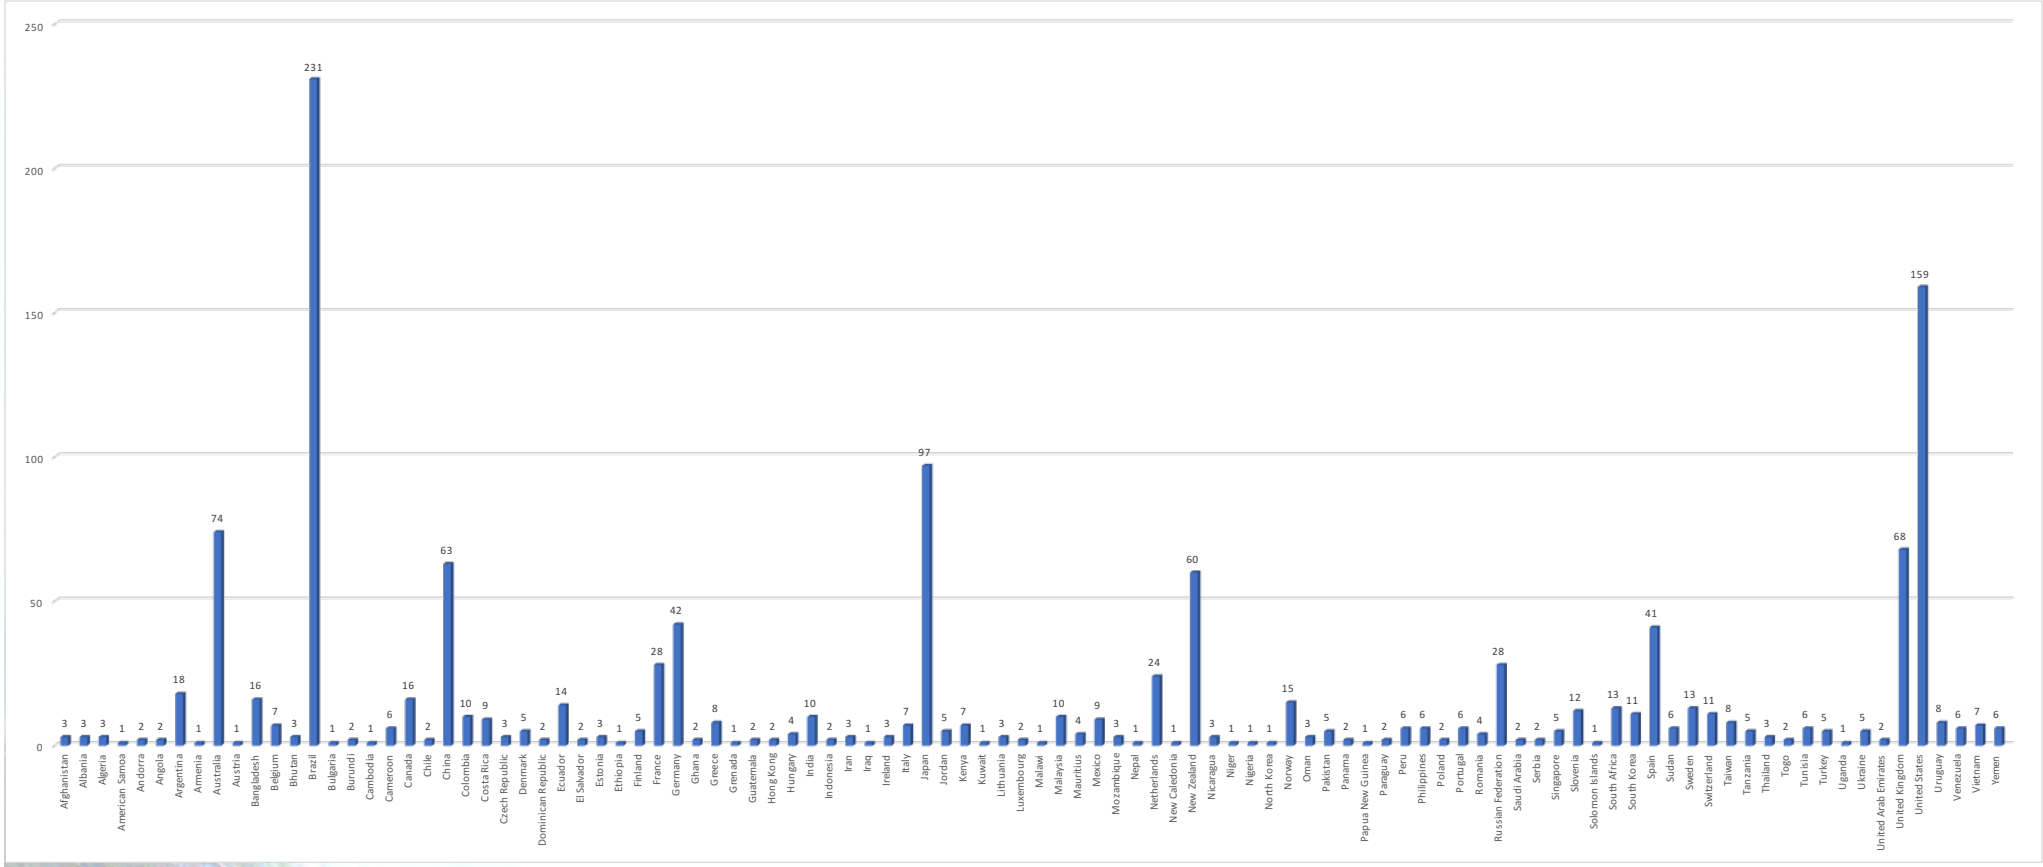

- Responses from 105 countries

Regional/Country analysis

- Is this meaning there are some regions/countries with a higher degree of residential deployment?
  - APNIC (Australia, China, Japan, Malaysia, New Zealand). Missing responses from South Korea, India.
  - ARIN (US, Canada)
  - LACNIC (Argentina, Brazil, Colombia, Guatemala, Paraguay, Peru, Venezuela). Missing responses from Mexico.
  - RIPE NCC (Belgium, Denmark, Finland, France, Germany, Greece, Luxembourg, Netherlands, Norway, Portugal, Romania, Russia, Slovenia, Spain, Sweden, Switzerland, UK)

- Or instead regions/countries not doing it?
  - AfriNIC
  - LACNIC

Technology

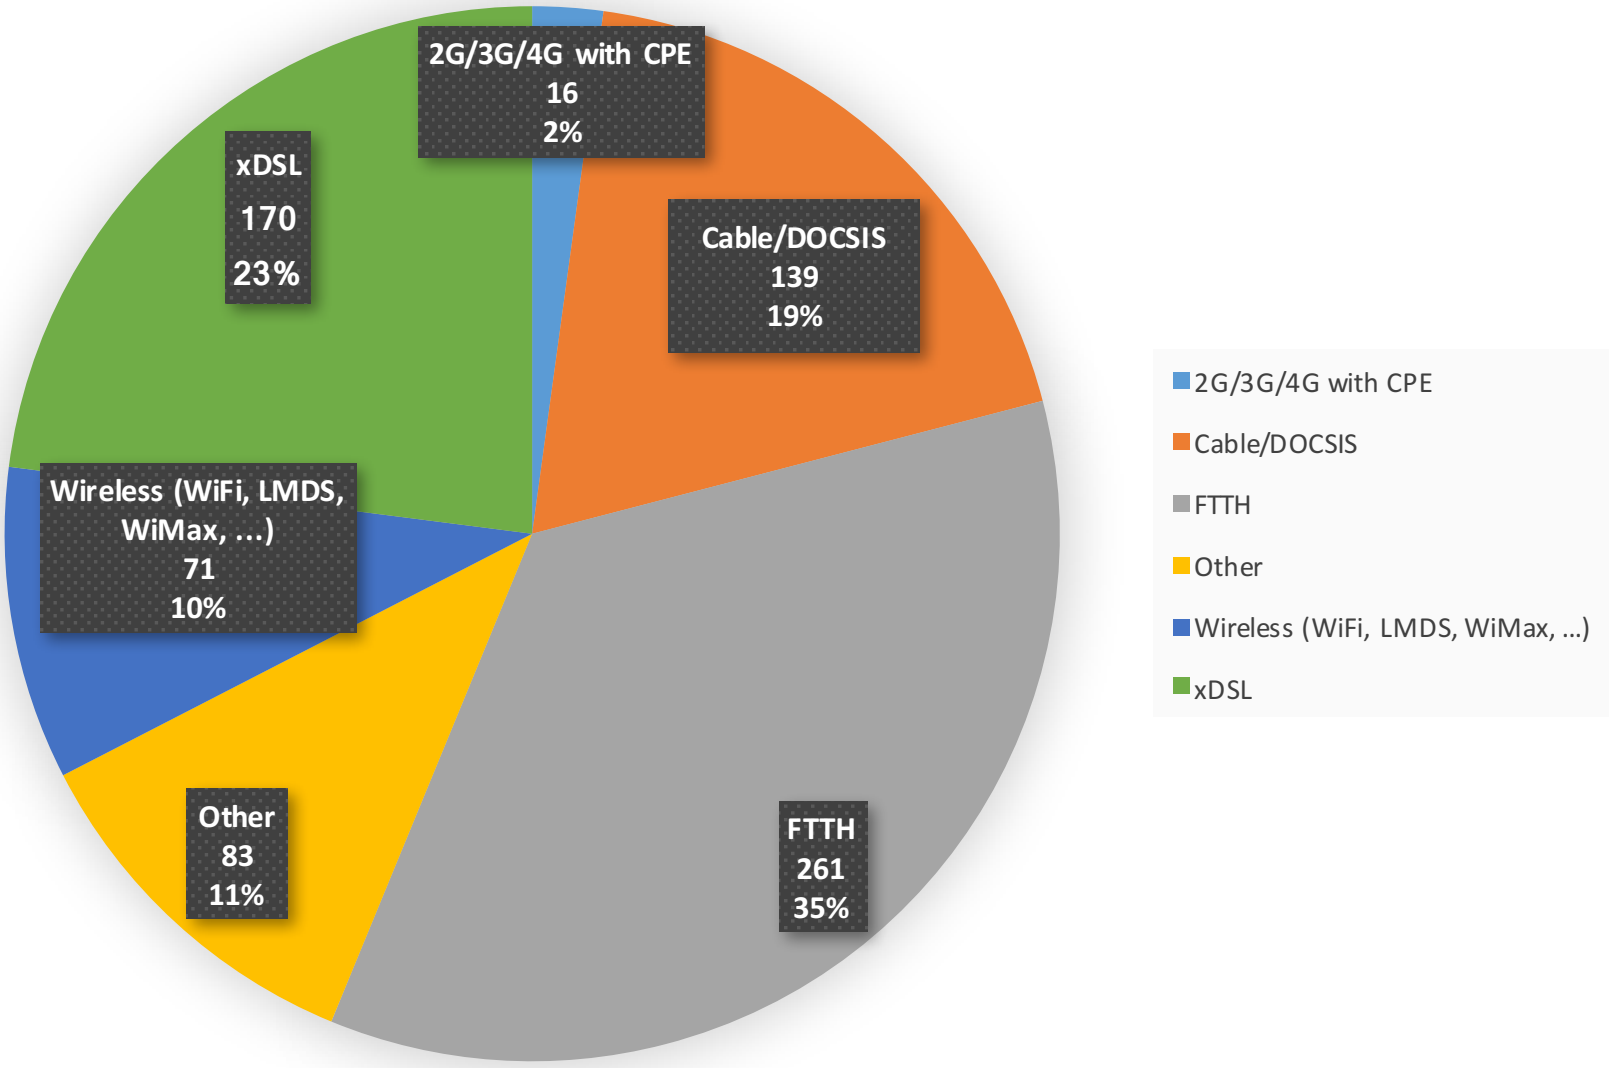
2G/3G/4G with CPE
16
2%
Cable/DOCSIS
139
19%
FTTH
261
35%
Other
83
11%
Wireless (WiFi, LMDS, WiMax, ...)
71
10%
xDSL
170
23%
2G/3G/4G with CPE
Cable/DOCSIS
FTTH
Other
Wireless (WiFi, LMDS, WiMax, ...)
xDSL

Deployment differences by techology

- More deployment by “newer” technologies:
  - FTTH
  - xDSL
  - Cable/DOCSIS
  - Wireless (WiFi, LMDS, WiMax, …)
- → Avoids investing in replacing CPEs

- Are there problems/dificulties with some specific access technologies?
  - According to the responses, I don’t think so …
- Vendor or transition technologies issues with some access technologies?
  - Nothing reported

Is IPv6 already a commercial service?

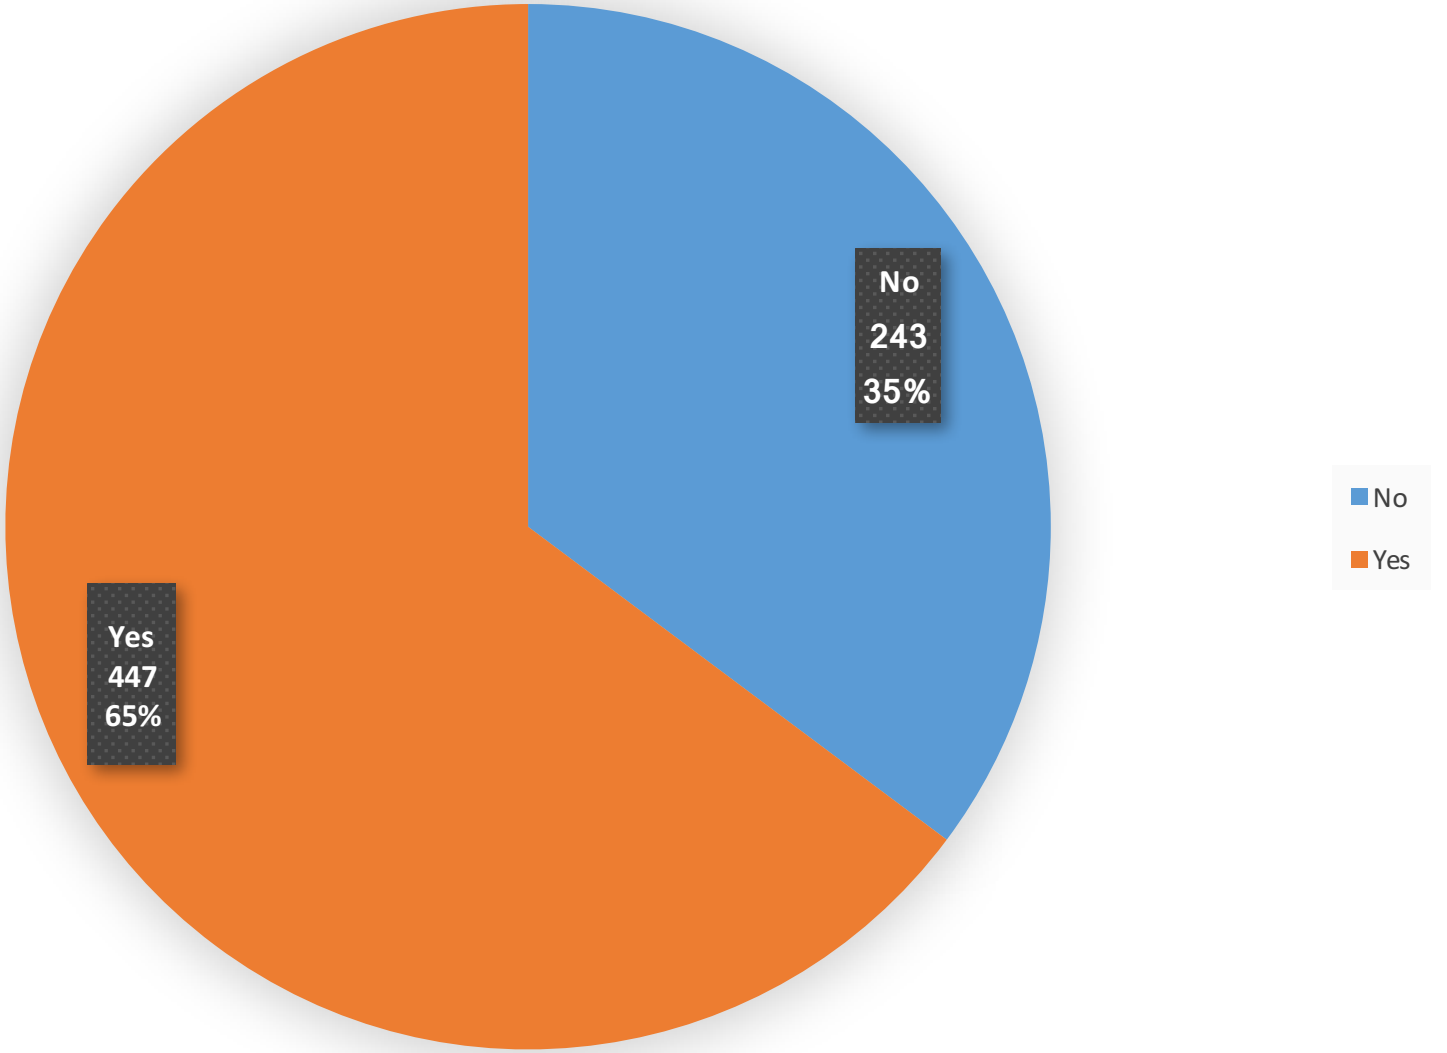

Why still not commercial?

- 65% Yes, already commercial
- 35% No commercial
  - checked with some of the responders, they will go to commercial, typically it is a trial, but they plan to deploy (few months from now)

WAN Prefix Size

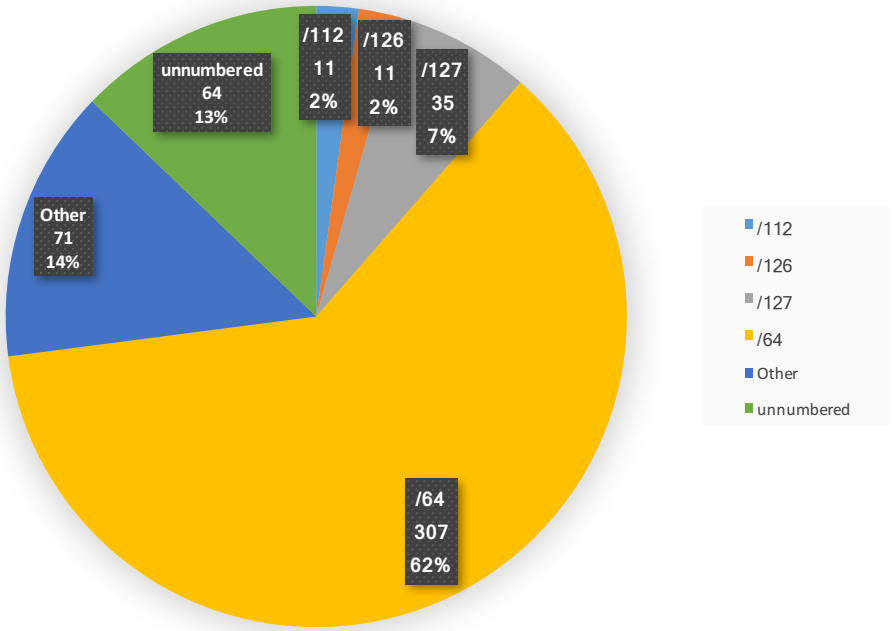

WAN Prefix Stable

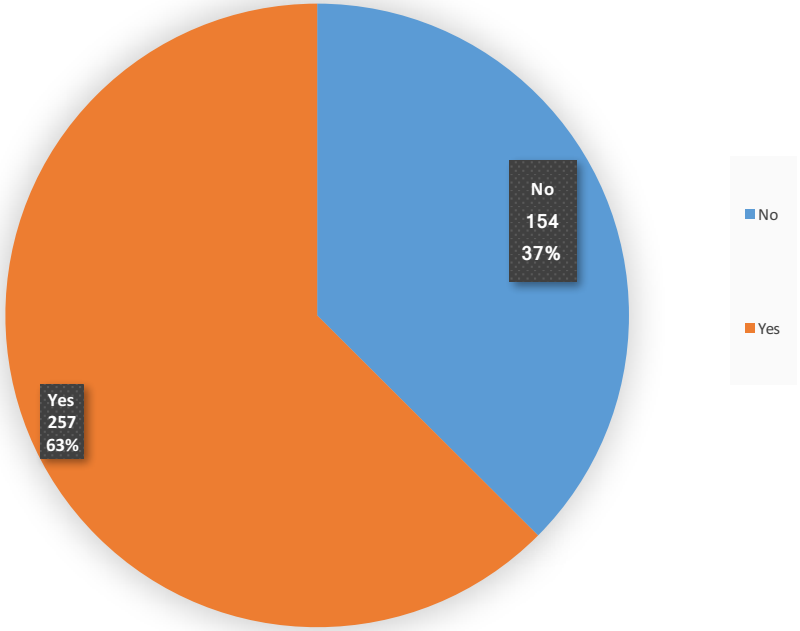

WAN Addressing Type

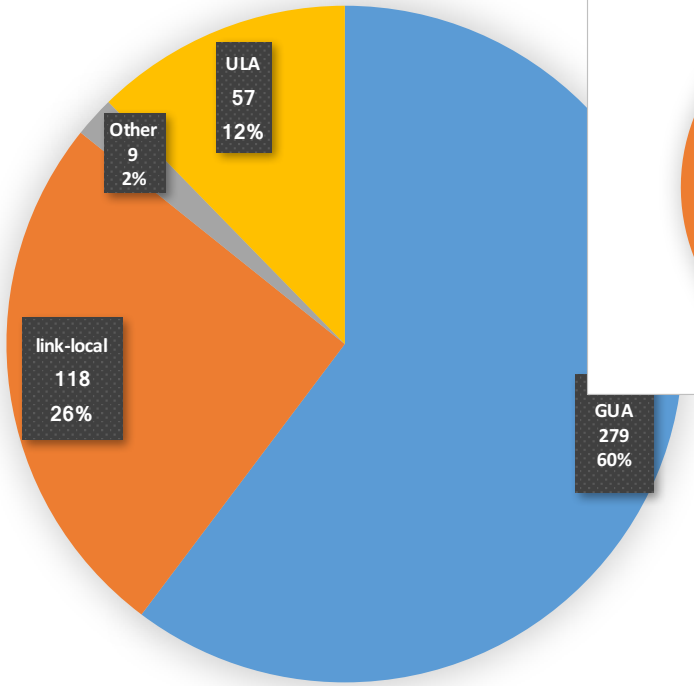

WAN from same pool as customer prefixes

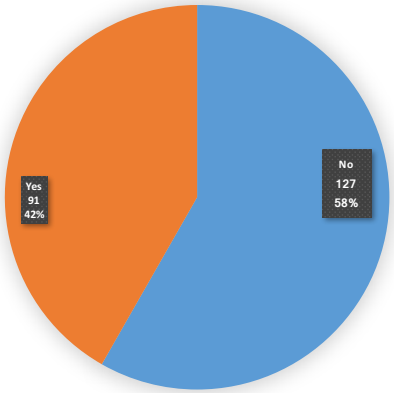

WAN /64 from customer prefix

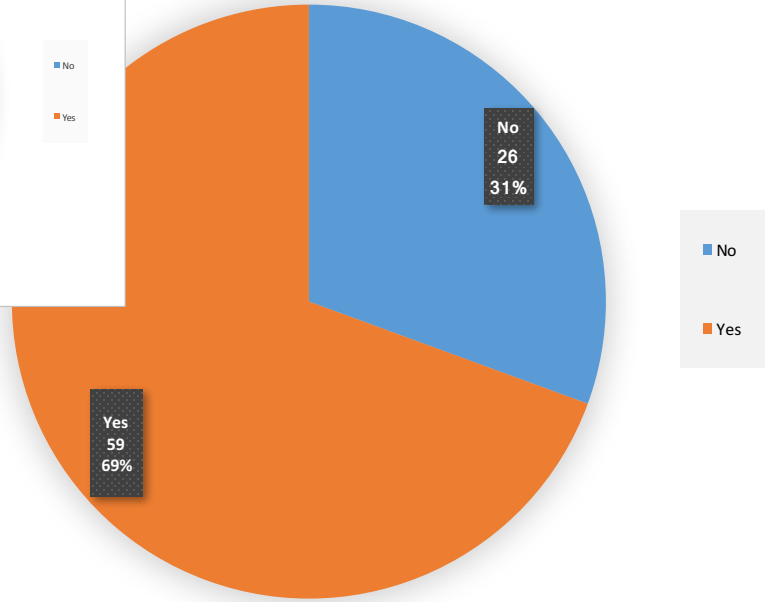

WAN prefix issues

- Remarkable -> /64 62%
- What means other?
  - /128, /62, /60, /56, /48, /32 ... No comments
- Why not stable (37%)?
  - Provisioning systems?
- 60% using GUA
- Interesting figures about using the /64 from the customer allocated prefix (69%)
- Distribution of those technical aspects not related to any specific country/region

LAN Prefix Size

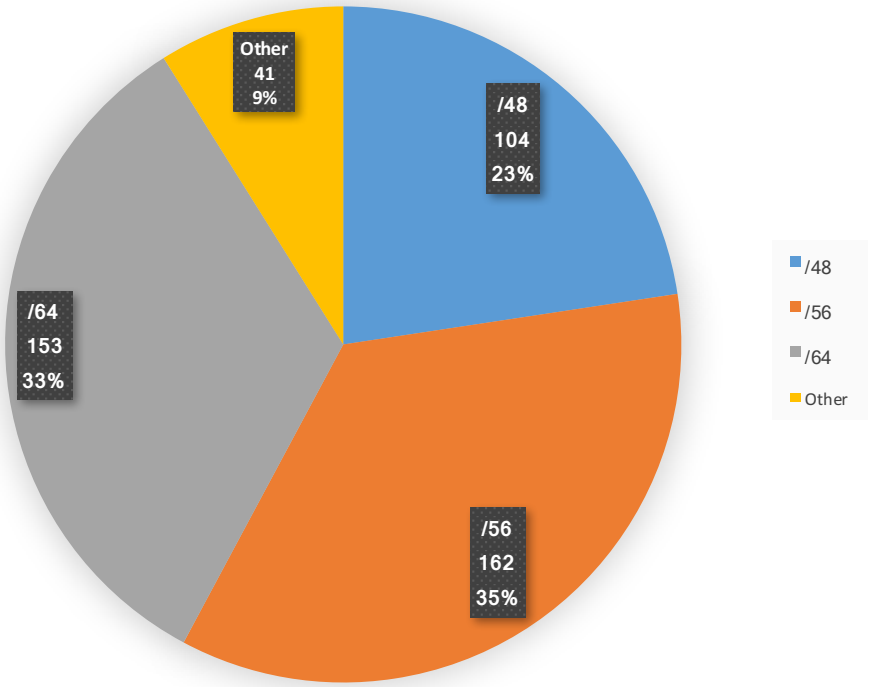

LAN Prefix Stable

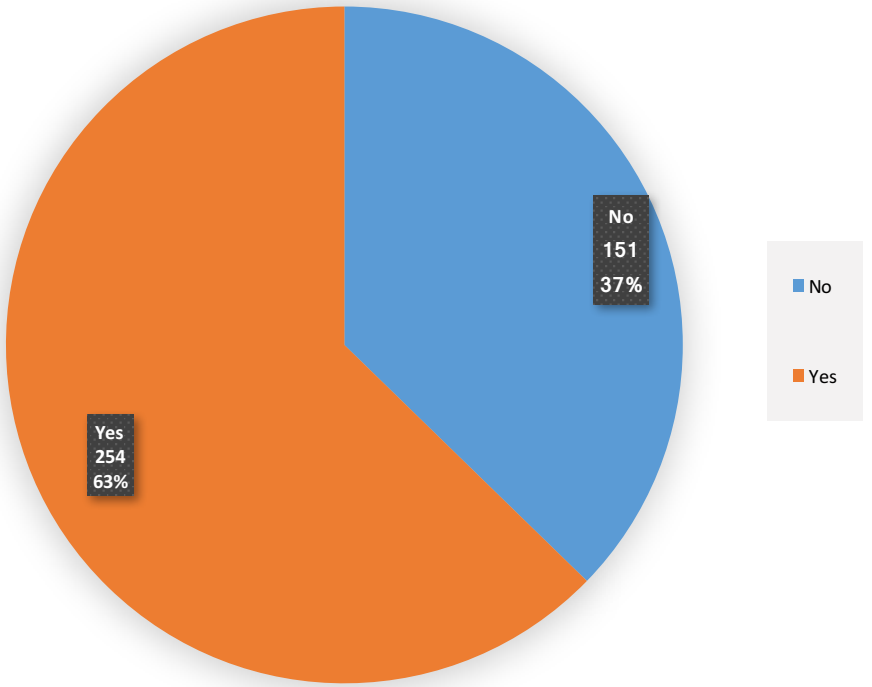

Can the customer opt to have it "stable"?

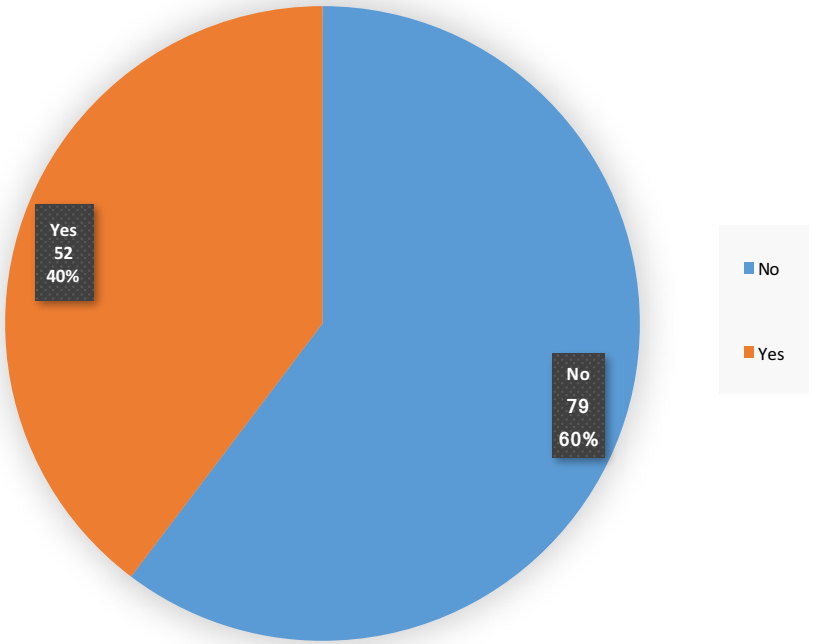

Extra cost (on top of stable IPv4)?

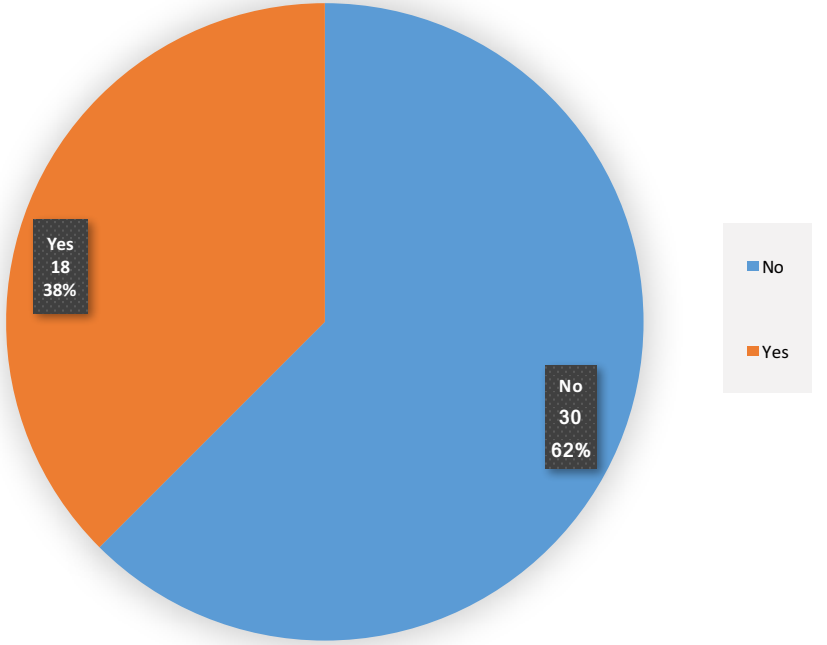

LAN prefix issues

- What are the “other" sizes?
  - A few /60 and /62 (others … /29, /44, /57, /127, /128)
  - Surprising (1) response -> shared /64
- Are we doing right/wrong? It is related to specific regions or countries?
  - 33% /64 mainly in LACNIC, some countries in APNIC
  - 35% /56 ARIN/RIPE NCC
  - 23% /48 mainly “more advanced” countries (Australia, New Zealand, Germany, Finland, Denmark, France, UK, China, Japan)
- Are we realizing that services work better with “stable” addressing?
  - AfriNIC, RIPE NCC and APNIC mainly stable
  - ARIN, mainly not-stable
  - LACNIC, half and half
- Why not allowing stable even as an “extra”?
  - Training issues? IPv4 mind-set?
  - Extra cost, very few

IPv4 service provided?

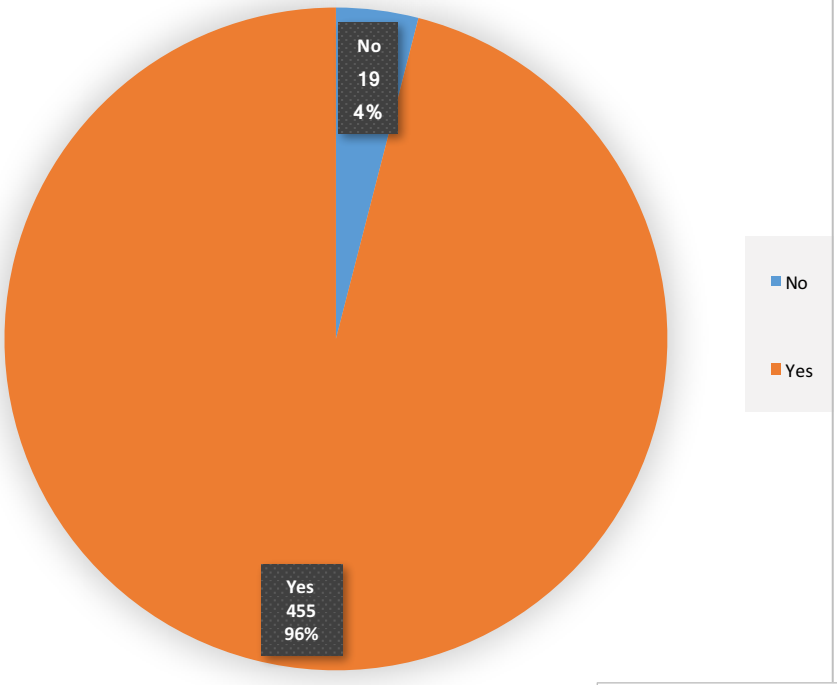

Public IPv4 address at CPE WAN?

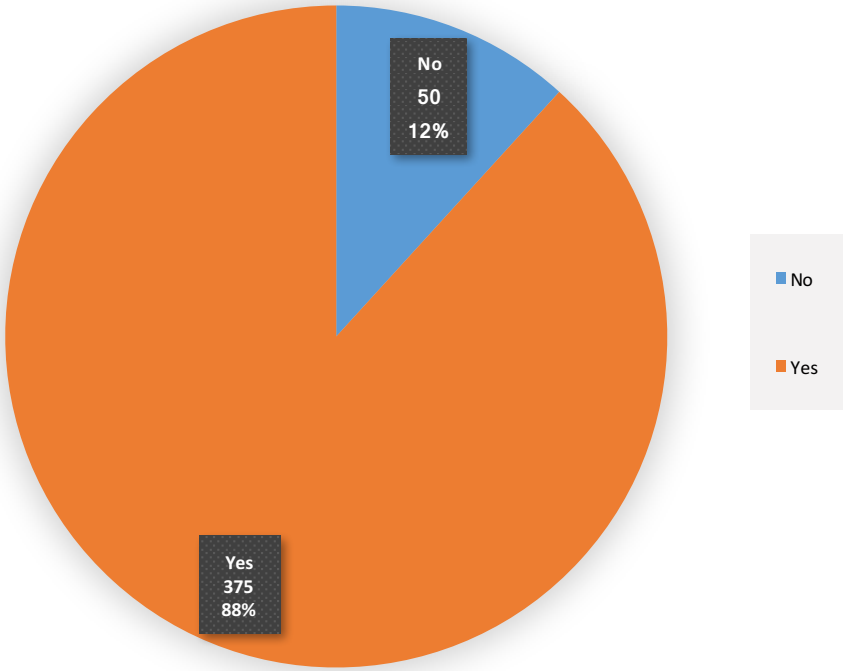

IPv4 address is "stable"?

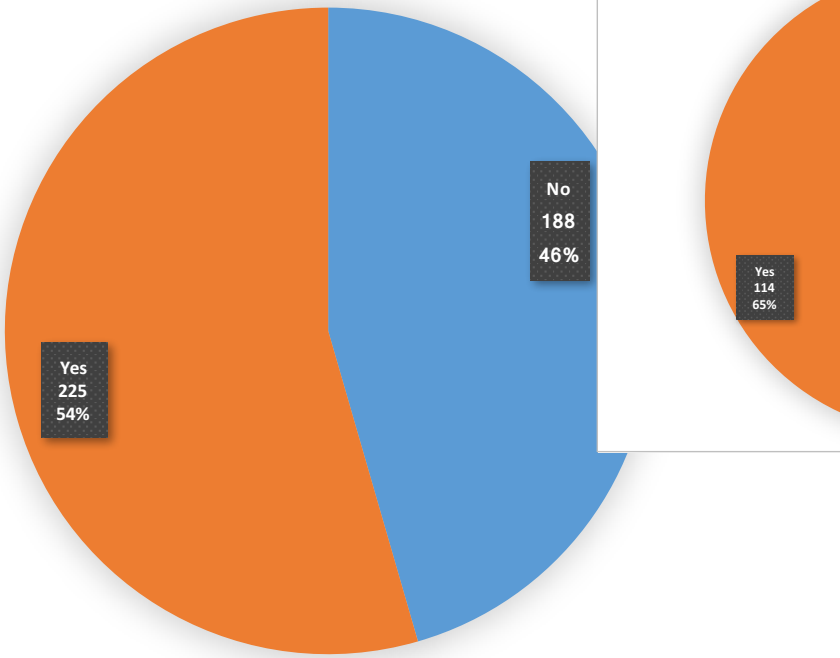

Can the customer opt to have IPv4 "stable"?

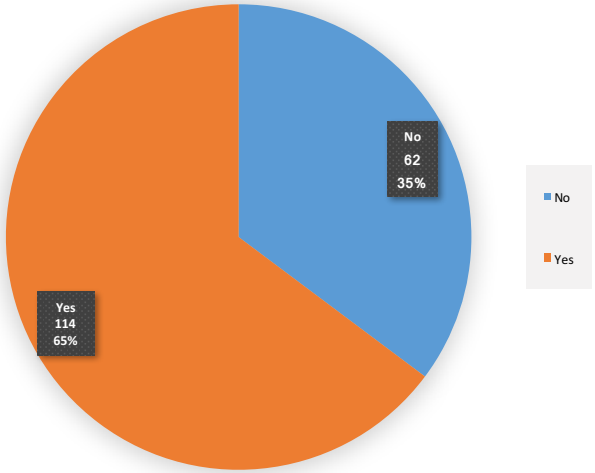

Extra cost for stable IPv4?

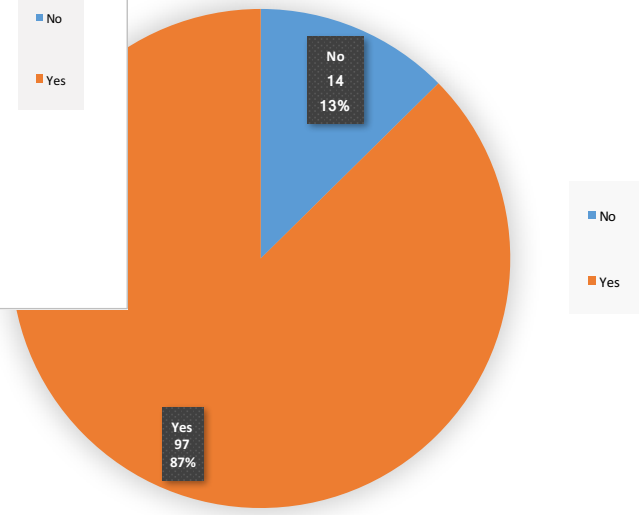

What transition mechanism?

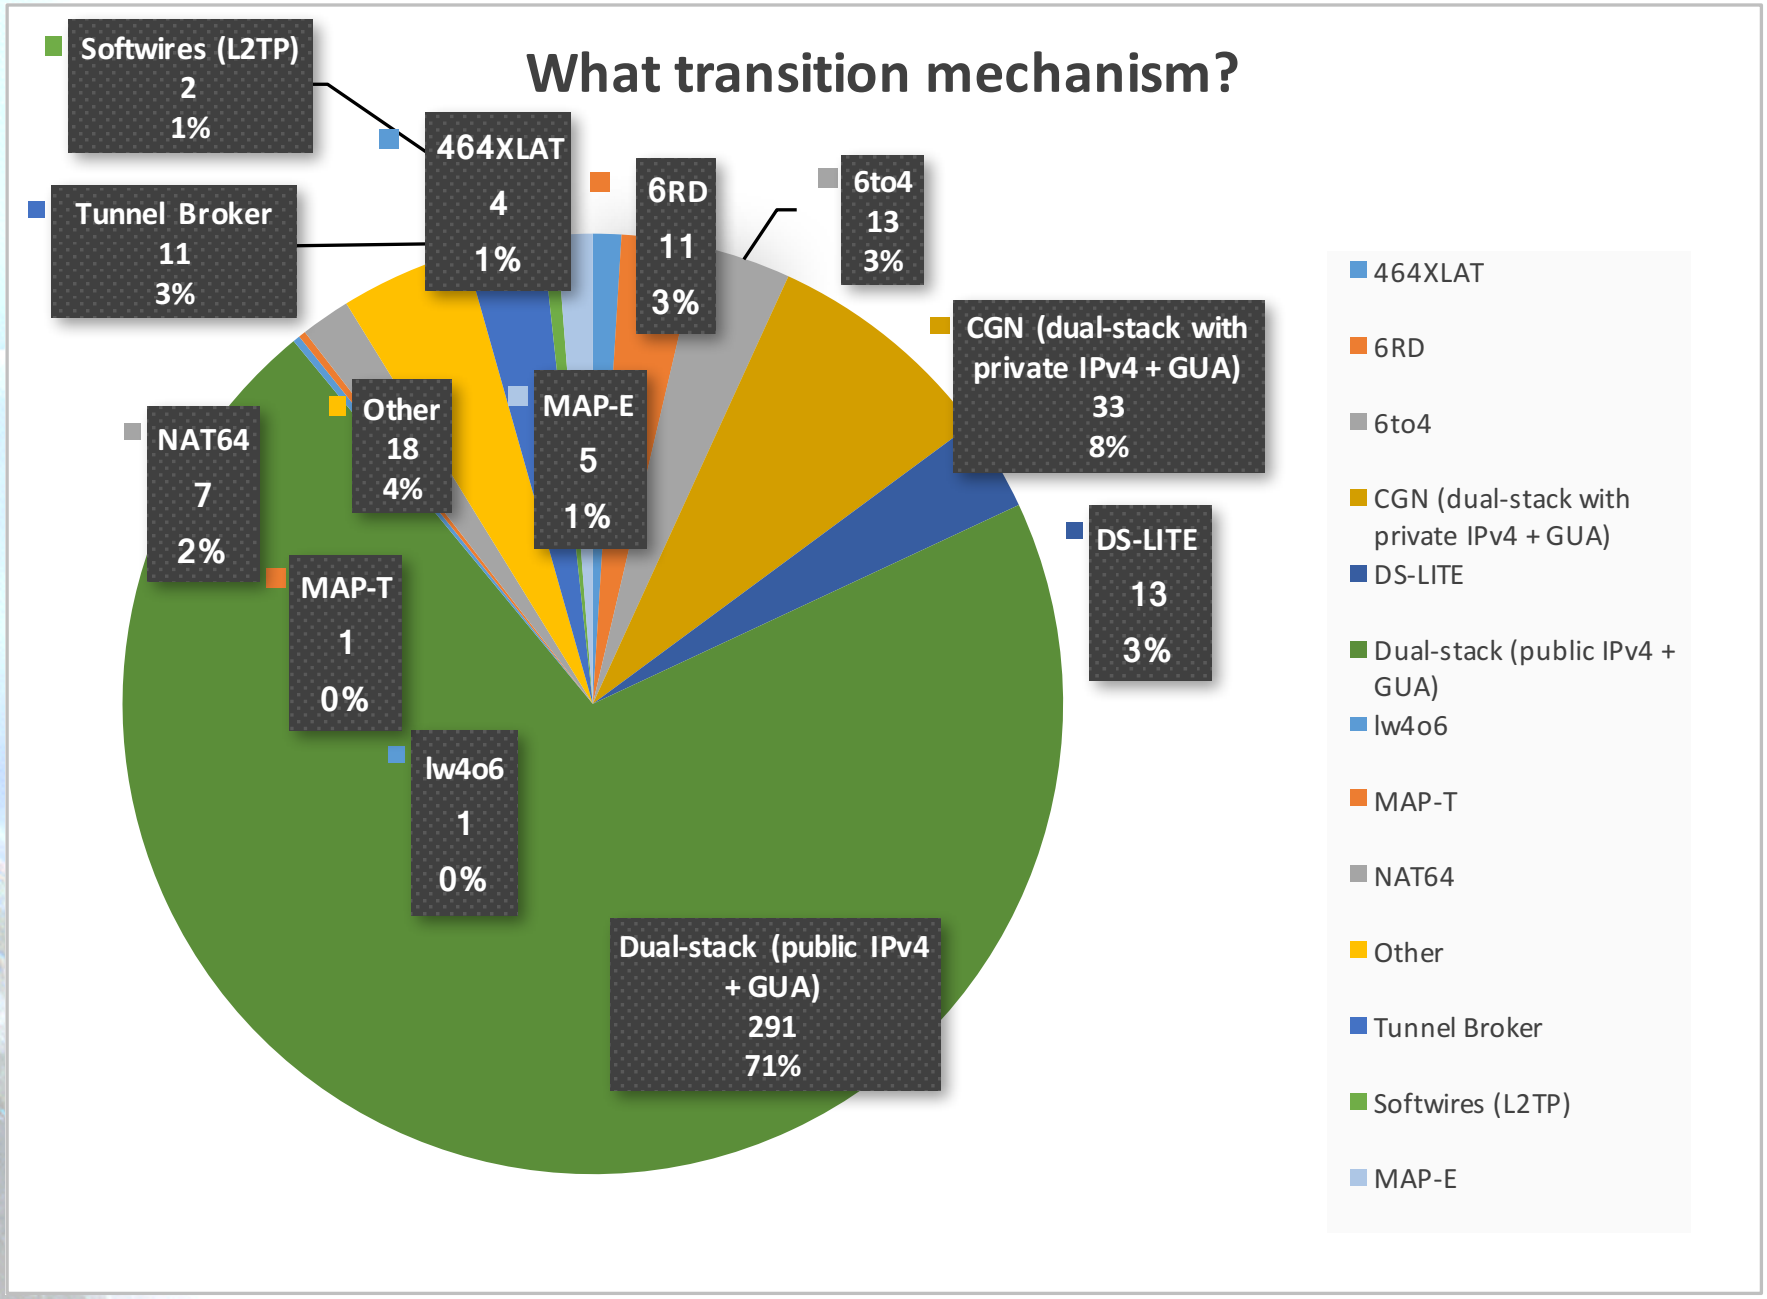

Transition and IPv4 issues

- It is a trend not providing IPv4 in the access?
  - It means some transition technologies being used which don't require IPv4 in the access.
- Not related to specific regions/countries

- What other "transition" technologies?
  - Actually none, just "bad answers"

- CGN deployment increasing clearly increasing ...

IPv6 reverse DNS?

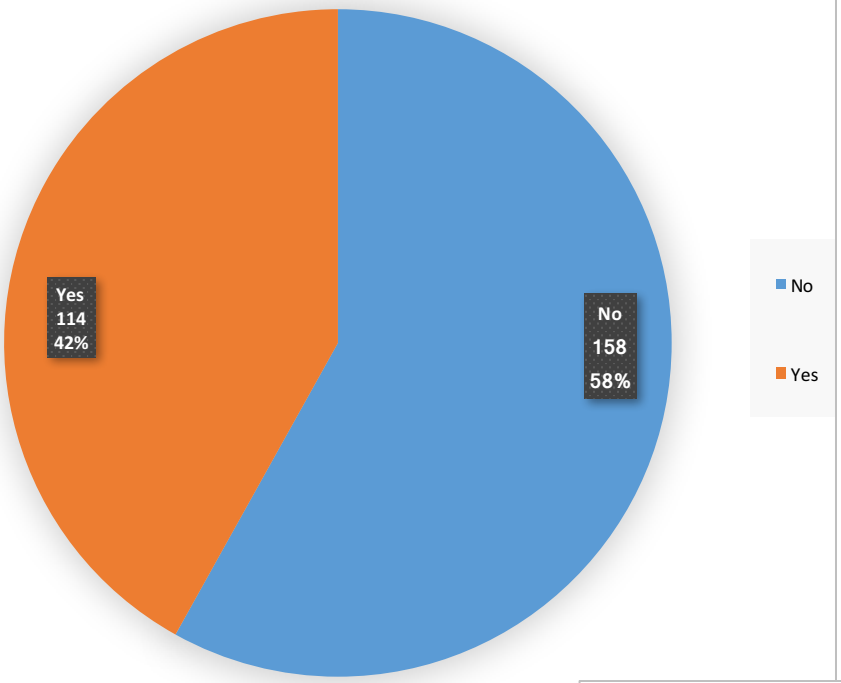

NS Delegation for stable IPv6 prefix?

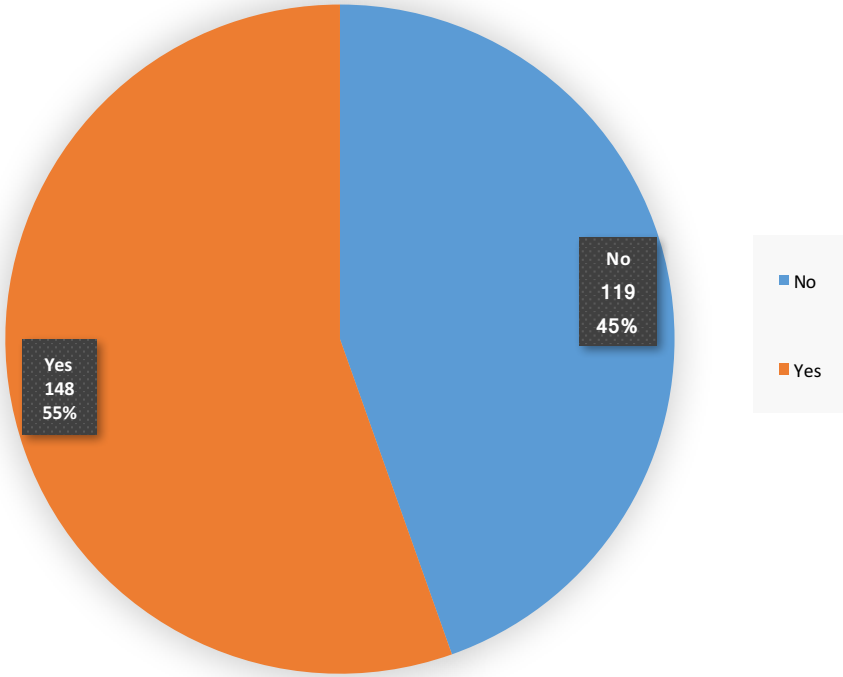

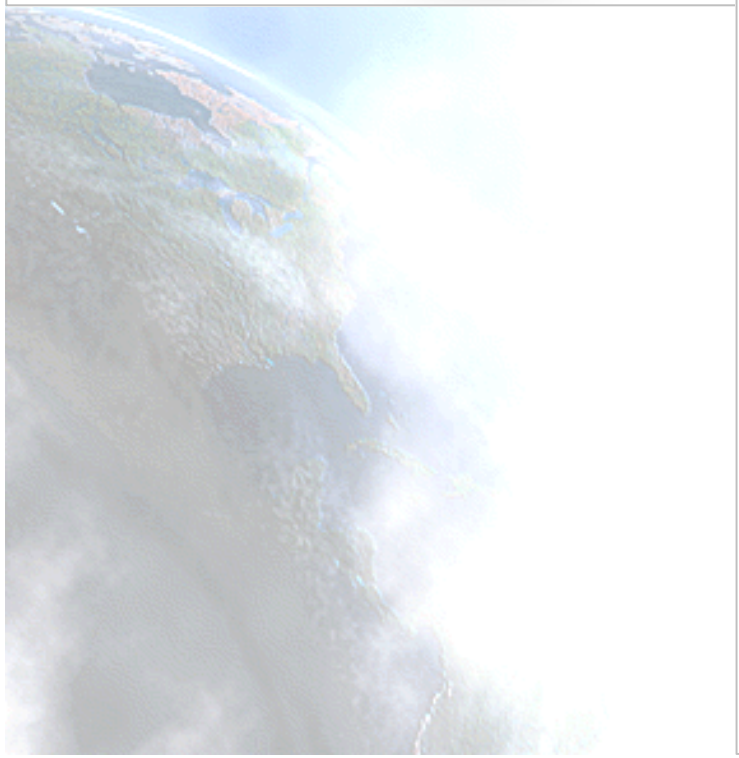

DNAME for non-stable IPv6 prefix for PTRs?

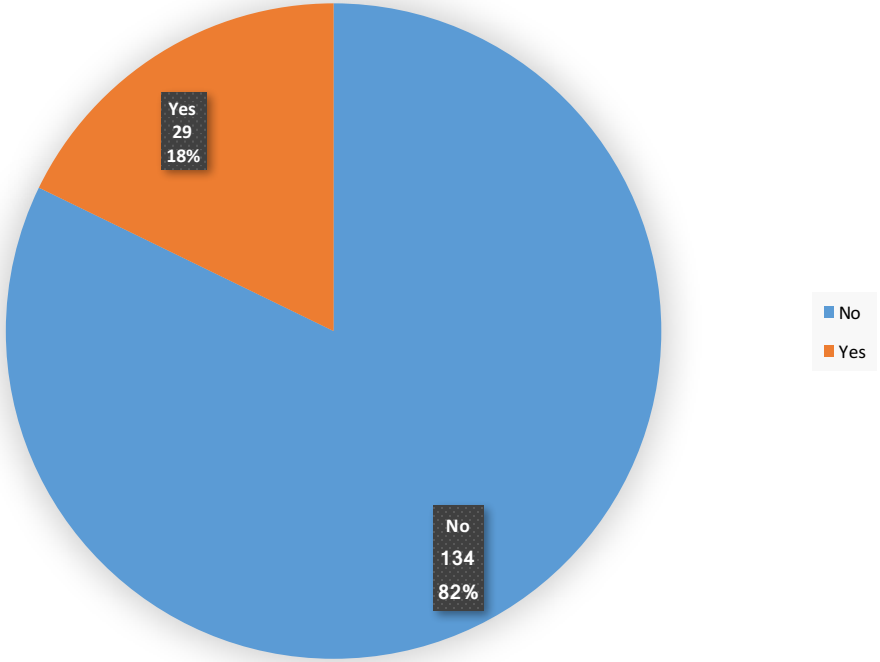

DNS

- Seems to follow “LAN IPv6 stable prefix”
- Reverse DNS as an extra service?

Conclusions

- In general “correct” deployment
  - Some exceptions
  - IPv4 “mind-set” – lack of coherent expert training
- Misunderstandings on IPv6 technology/marketing/other reason:
  - IPv6 prefix size
  - Stability of prefix
- More “advanced” countries seem to do it smartly, less ”misunderstandings”

Thanks !!

Survey link:

http://survey.consulintel.es/index.php/175122

Contact:

- **Jordi Palet (The IPv6 Company):**
  jordi.palet@theipv6company.com