Using 464XLAT in Residential Networks

GTER 43, Foz do Iguaçu
Maio de 2017

Jordi Palet (jordi.palet@consulintel.es)

Do you know …

- We already run out of IPv4?
- How you keep deploying Internet access to your residential customers?
- Are you using IPv4 to deploy IPv6?
  - such as tunnel broker, 6RD and so?

Once upon a time …

- IETF was considering to solve this problem by more tunneling …
- So we build up softwires, which decided to use L2TP, so we could do
  - IPv6 in IPv4, IPv4 in IPv6
  - (as well IPv4 in IPv4 and IPv6 in IPv6 for multicast in unicast)
- As a result we have, among others:
  - DS-Lite
  - Carrier Grade NAT (AFTR)
  - lw4o6

NAT444

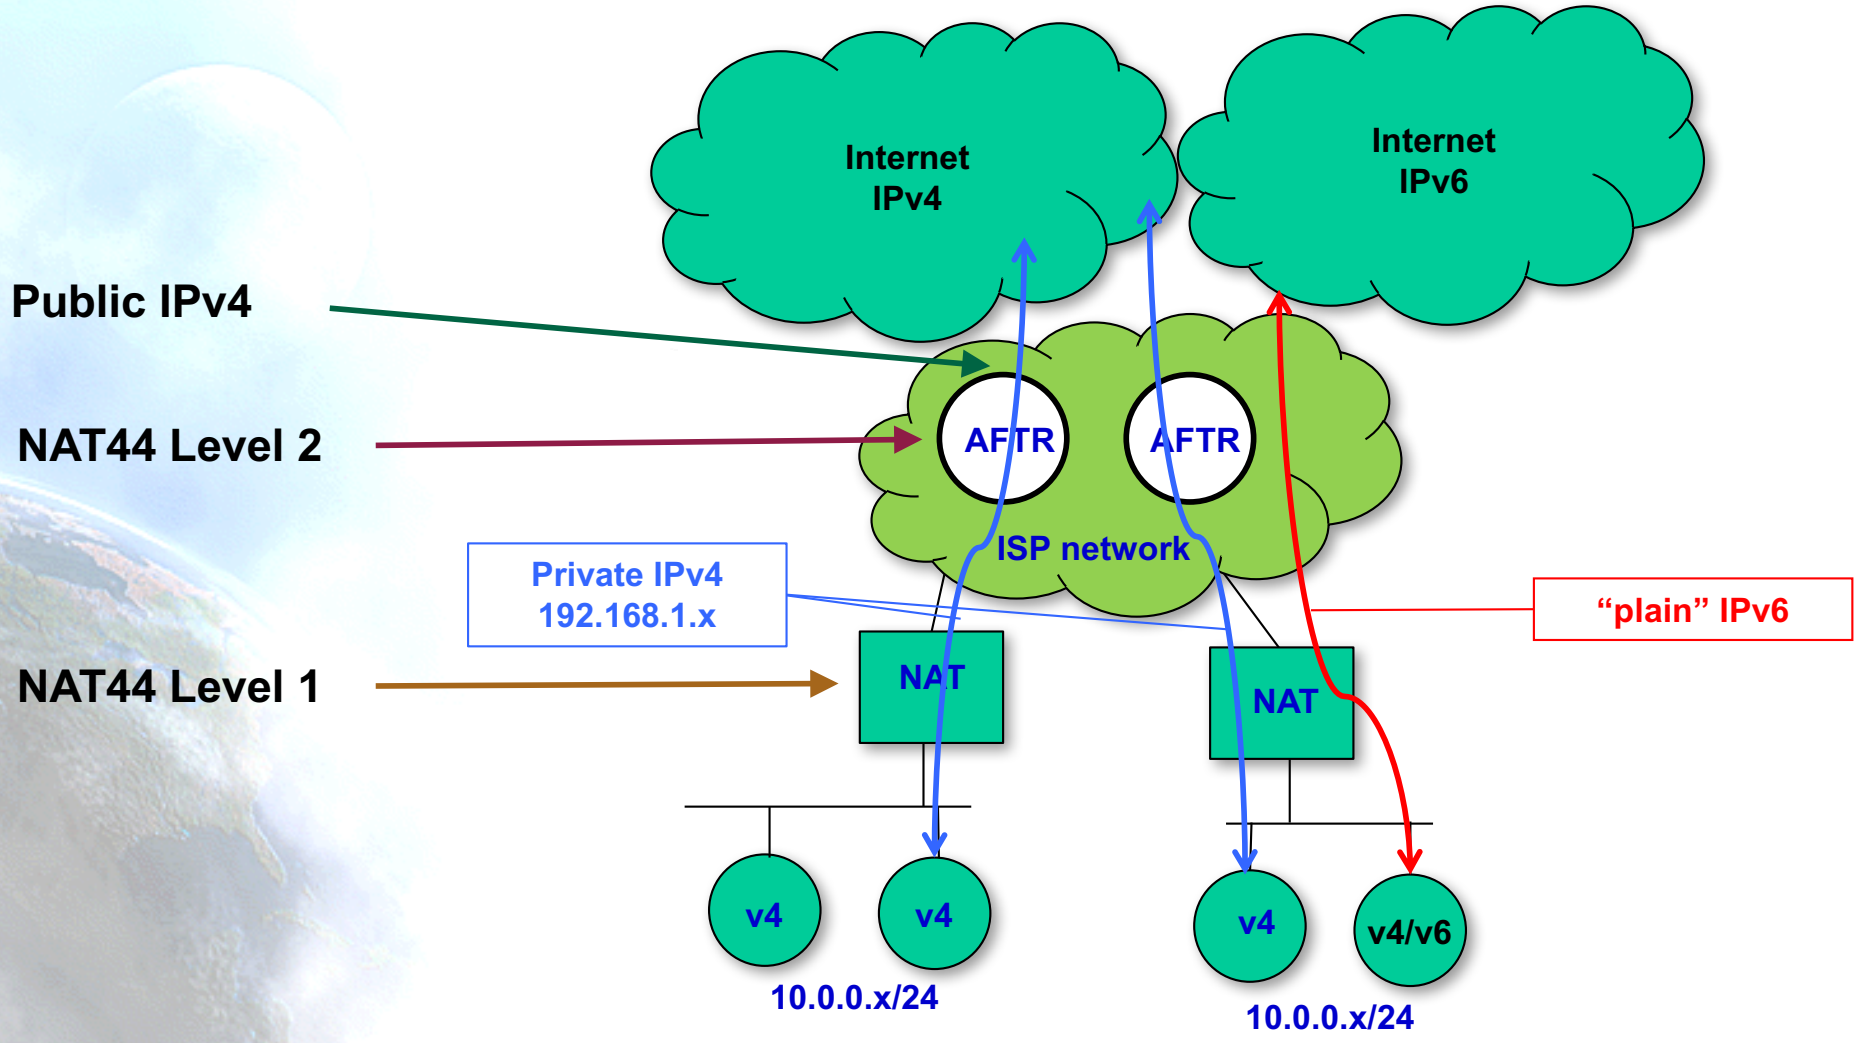
Internet
IPv4
Internet
IPv6
Public IPv4
NAT44 Level 2
AFTR
AFTR
ISP network
Private IPv4
192.168.1.x
"plain" IPv6
NAT44 Level 1
NAT
NAT
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24

DS-Lite

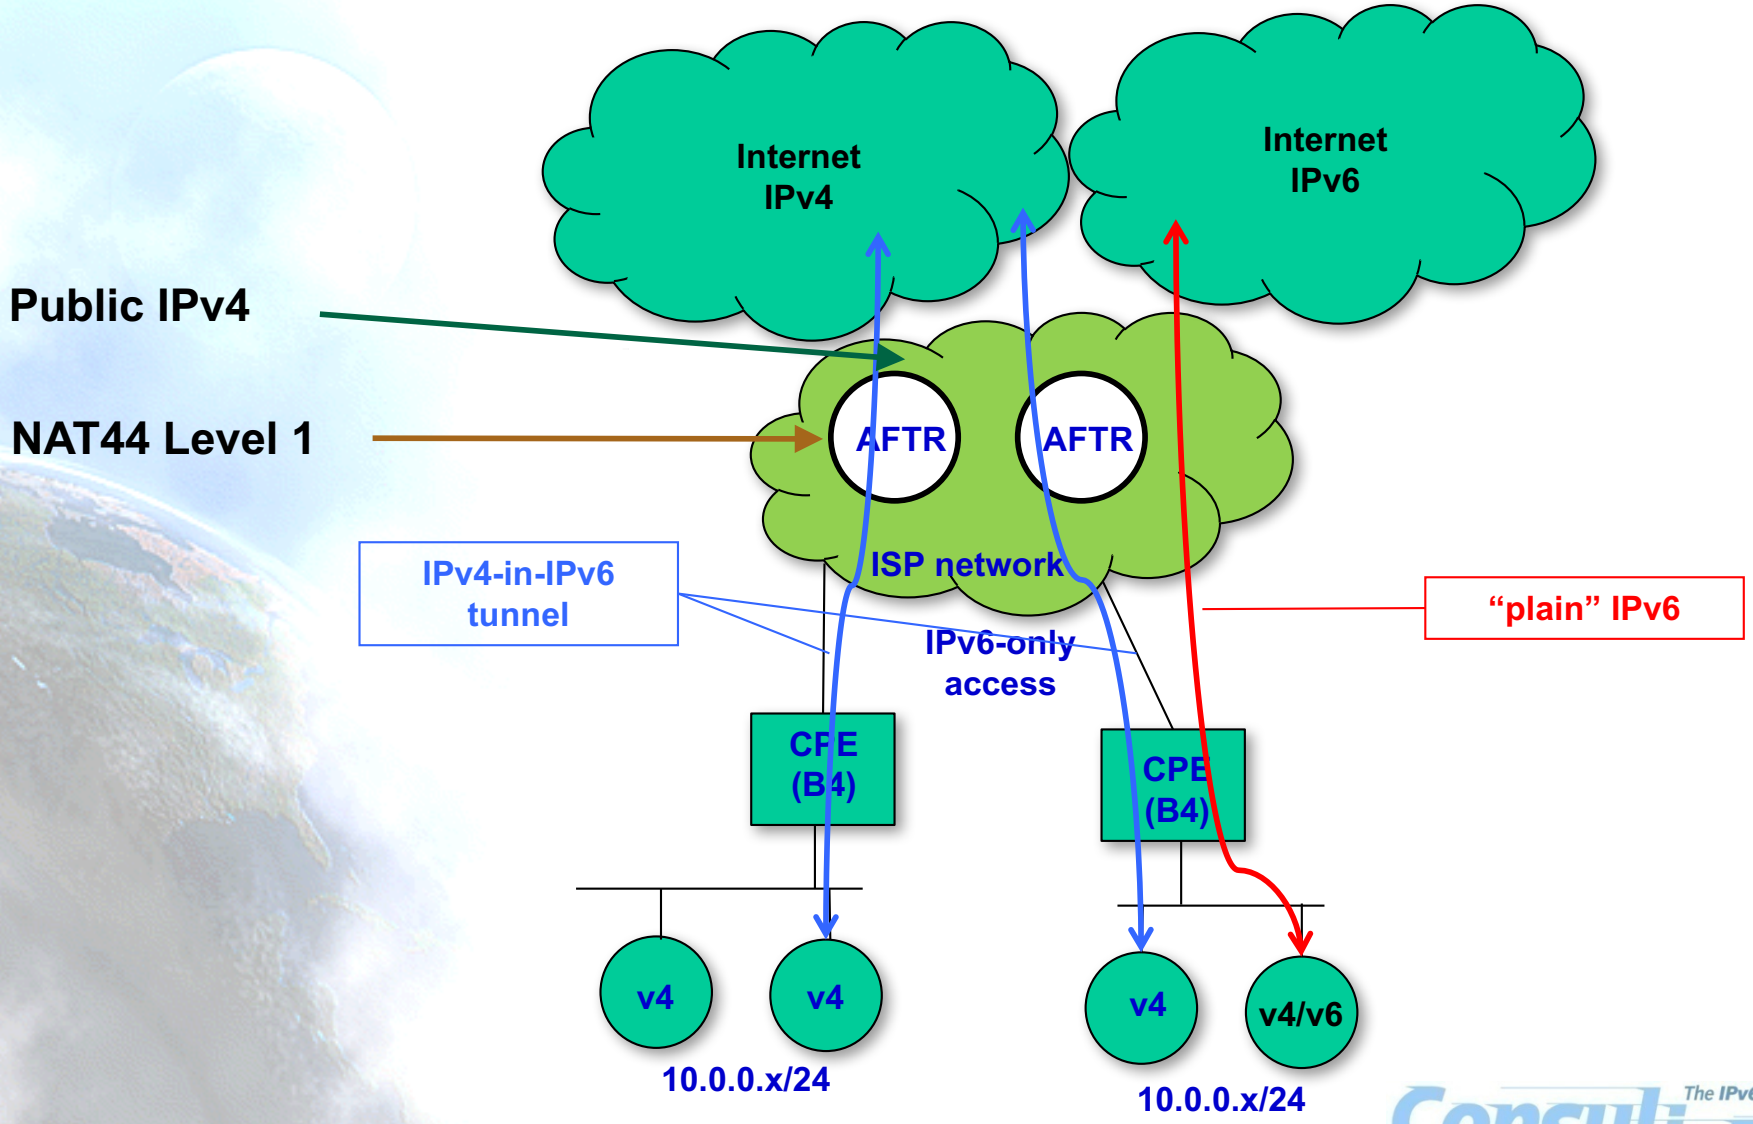
Internet
IPv4
Internet
IPv6
Public IPv4
NAT44 Level 1
AFTR
AFTR
ISP network
IPv4-in-IPv6
tunnel
IPv6-only
access
"plain" IPv6
CPE
(B4)
CPE
(B4)
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24

lw4o6

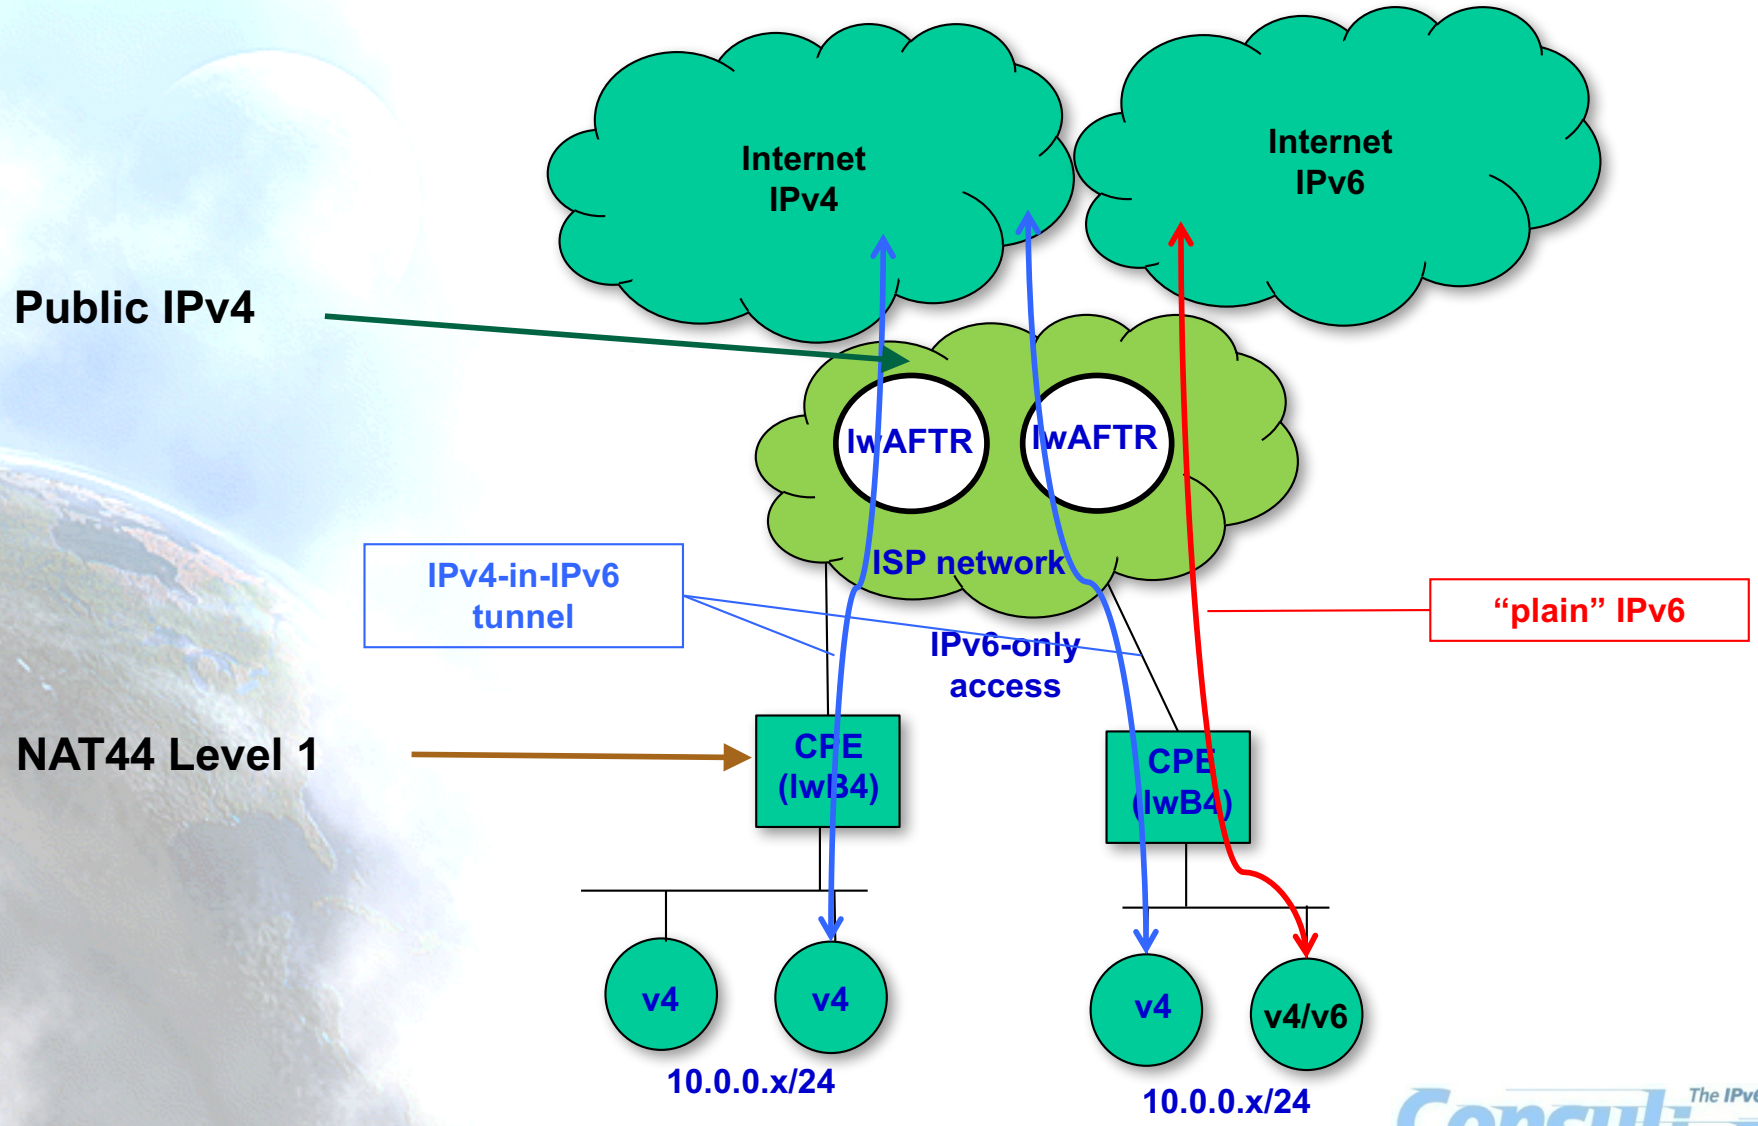
Internet
IPv4
Internet
IPv6
Public IPv4
lwAFTR
lwAFTR
ISP network
IPv4-in-IPv6
tunnel
IPv6-only
access
"plain" IPv6
NAT44 Level 1
CPE
(lwB4)
CPE
(lwB4)
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24

Tunnels per subscribers

- DS-Lite/lw4o6

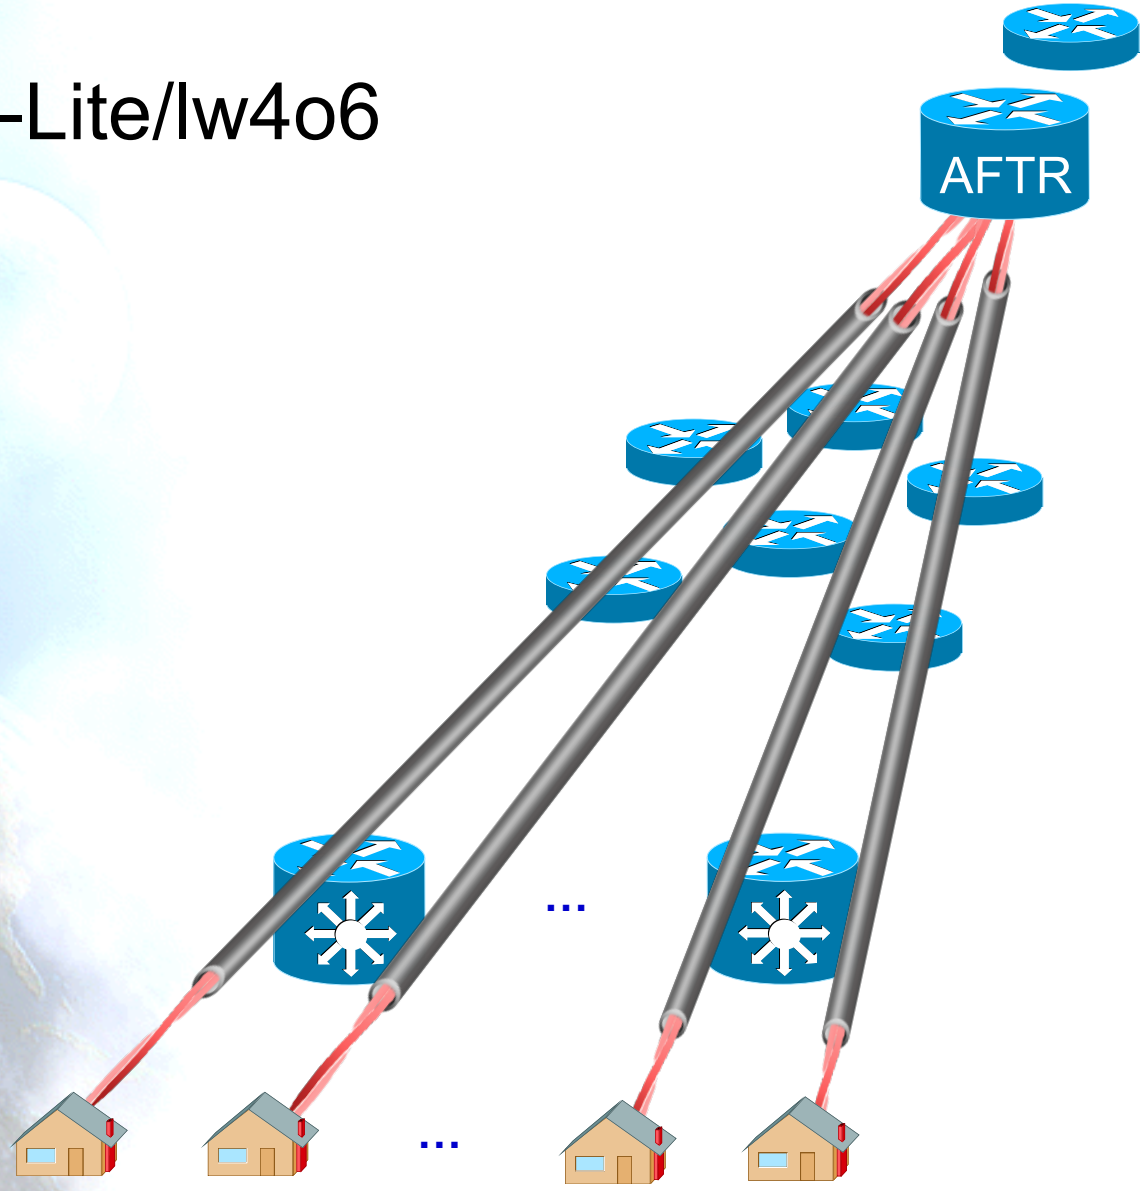

BGP prefixes: Tens

Tunnels: Millions

IGP prefixes: Hundreds

BNG routes: Thousands

Subscribers: Millions

CGN breaks …

- UPnP-IGD (Universal Plug & Play - Internet Gateway Device protocol)
- NAT-PMP (NAT Port Mapping Protocol)
- Other NAT Traversal mechs
- Security
- AJAX (Asyncronous Javascript And XML)
- FTP (big files)
- BitTorrent/Limewire (seeding – uploading)
- On-line gaming
- Video streaming (Netflix, Hulu, …)
- IP cameras
- Tunnels, VPN, IPsec, ...
- VoIP
- Port forwarding
- ...

NAT64

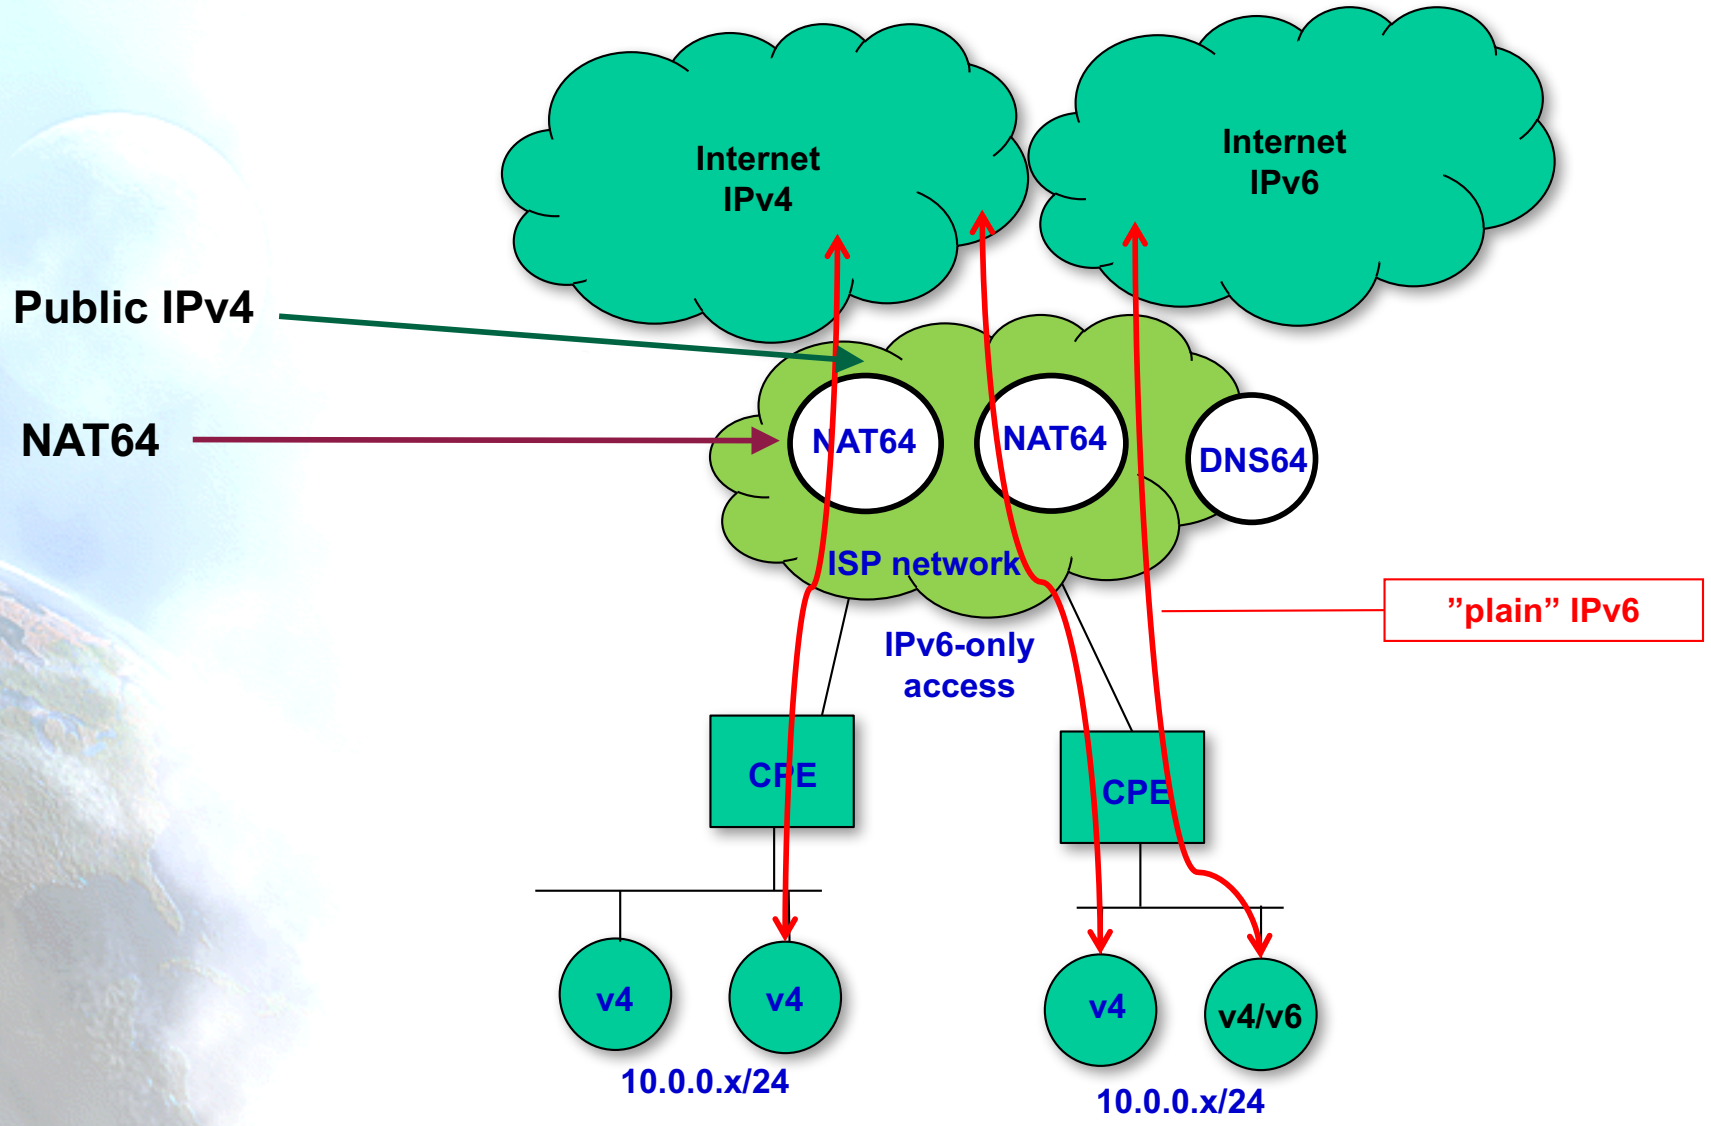
Internet
IPv4
Internet
IPv6
Public IPv4
NAT64
NAT64
NAT64
DNS64
ISP network
"plain" IPv6
IPv6-only
access
CPE
CPE
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24

NAT64 breaks …

App Name	Functionality	Version	464XLAT Fixed
connection tracker	Broken	NA	NA
DoubleTwist	Broken	1.6.3	YES
Go SMS Pro	Broken	NA	YES
Google Talk	Broken	4.1.2	YES
Google+	Broken	3.3.1	YES
IP Track	Broken	NA	NA
Last.fm	Broken	NA	YES
Netflix	Broken	NA	YES
ooVoo	Broken	NA	YES
Pirates of the Caribean	Broken	NA	YES
Scrabble Free	Broken	1.12.57	YES
Skype	Broken	3.2.0.6673	YES
Spotify	Broken	NA	YES
Tango	Broken	NA	YES
Texas Poker	Broken	NA	YES
TiKL	Broken	2.7	YES
Tiny Towers	Broken	NA	YES
Trillian	Broken	NA	YES
TurboxTax Taxcaster	Broken	NA	
Voxer Walkie Talkie	Broken	NA	YES
Watch ESPN	Broken	1.3.1	
Zynga Poker	Broken	NA	YES
Xabber XMPP	Broken	NA	

*T-Mobile

464XLAT

- 464XLAT (RFC6877): RFC6145 + RFC6146
- Very efficient use of scarce IPv4 resources
  - N*64.000 flows per each IPv4 address
  - Network growth not tied to IPv4 availability
- IPv4 basic service to customers over an-IPv6 only infrastructure
  - **WORKS** with applications that use socket APIs and literal IPv4 addresses (Skype, etc.)
- Allows traffic engineering
  - Without deep packet inspection
- Easy to deploy and available
  - Commercial solutions and open source

464XLAT

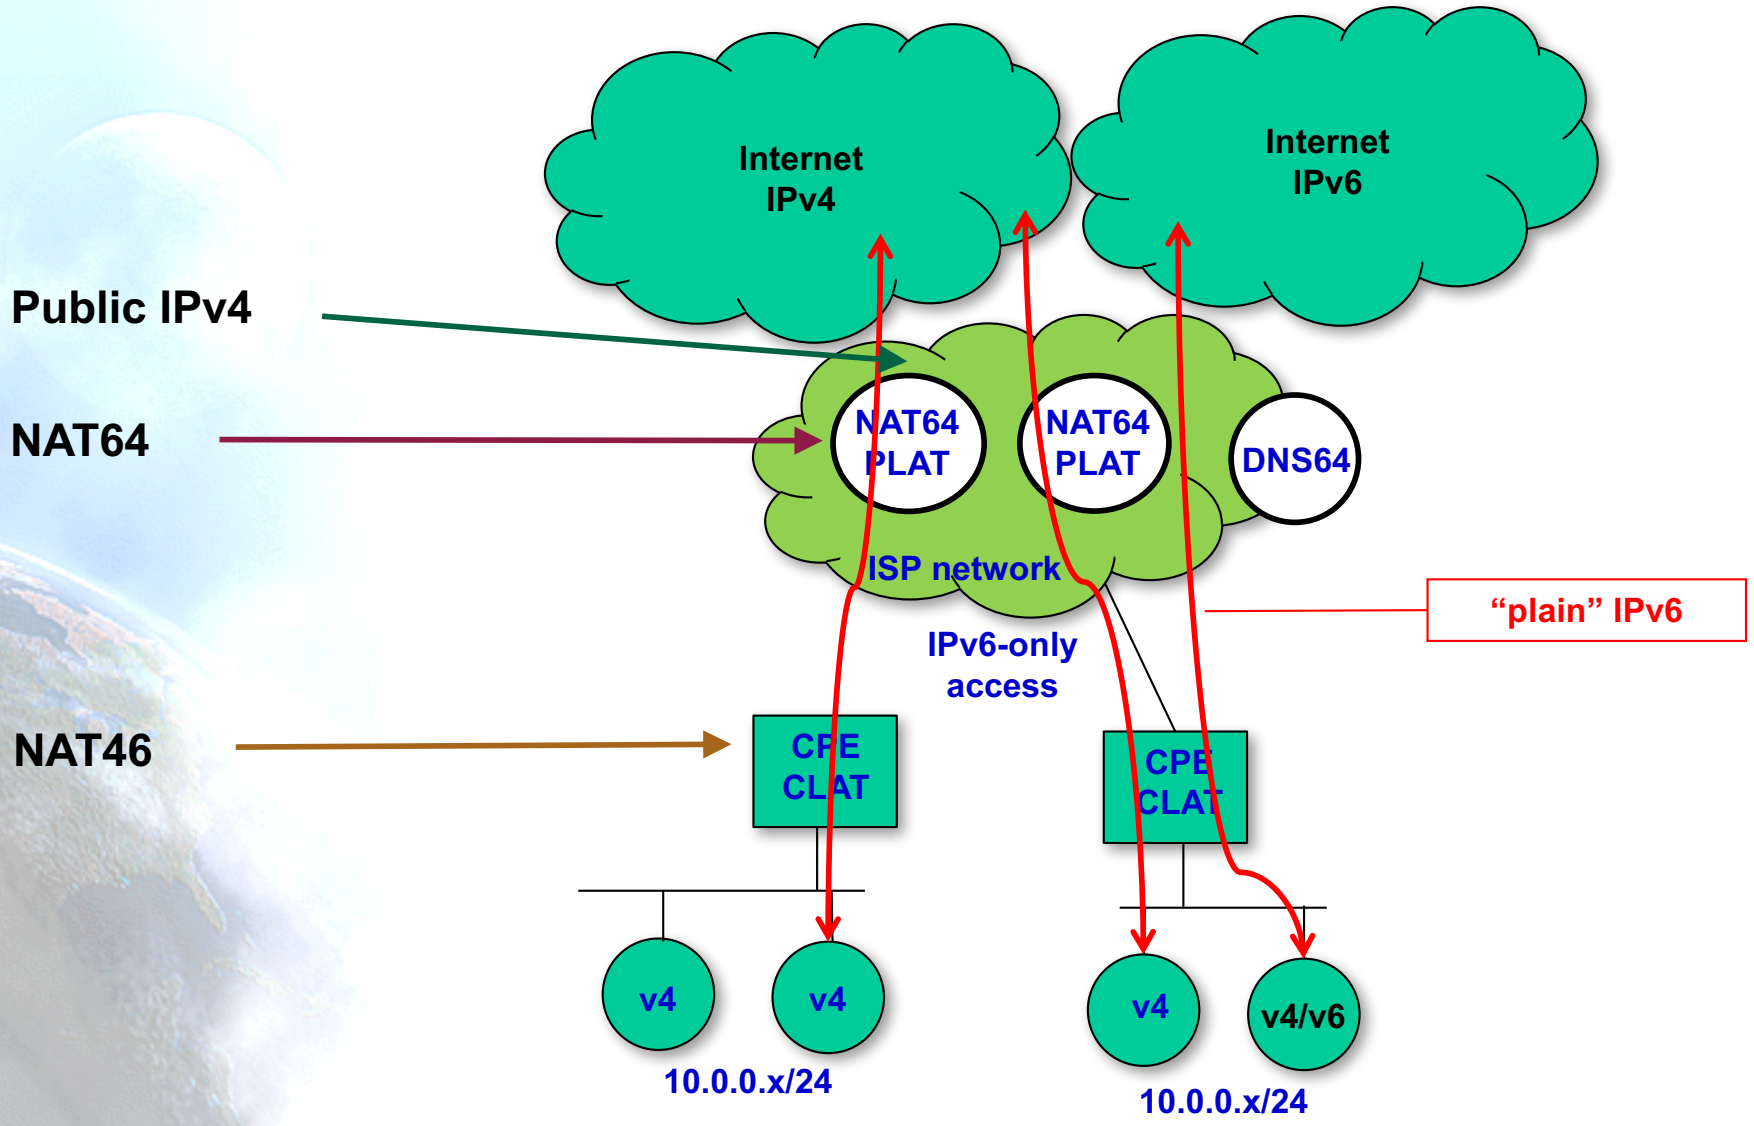
Internet
IPv4
Internet
IPv6
Public IPv4
NAT64
NAT64
PLAT
NAT64
PLAT
DNS64
ISP network
"plain" IPv6
IPv6-only
access
NAT46
CPE
CLAT
CPE
CLAT
v4
v4
v4
v4/v6
10.0.0.x/24
10.0.0.x/24

How it works 464XLAT?

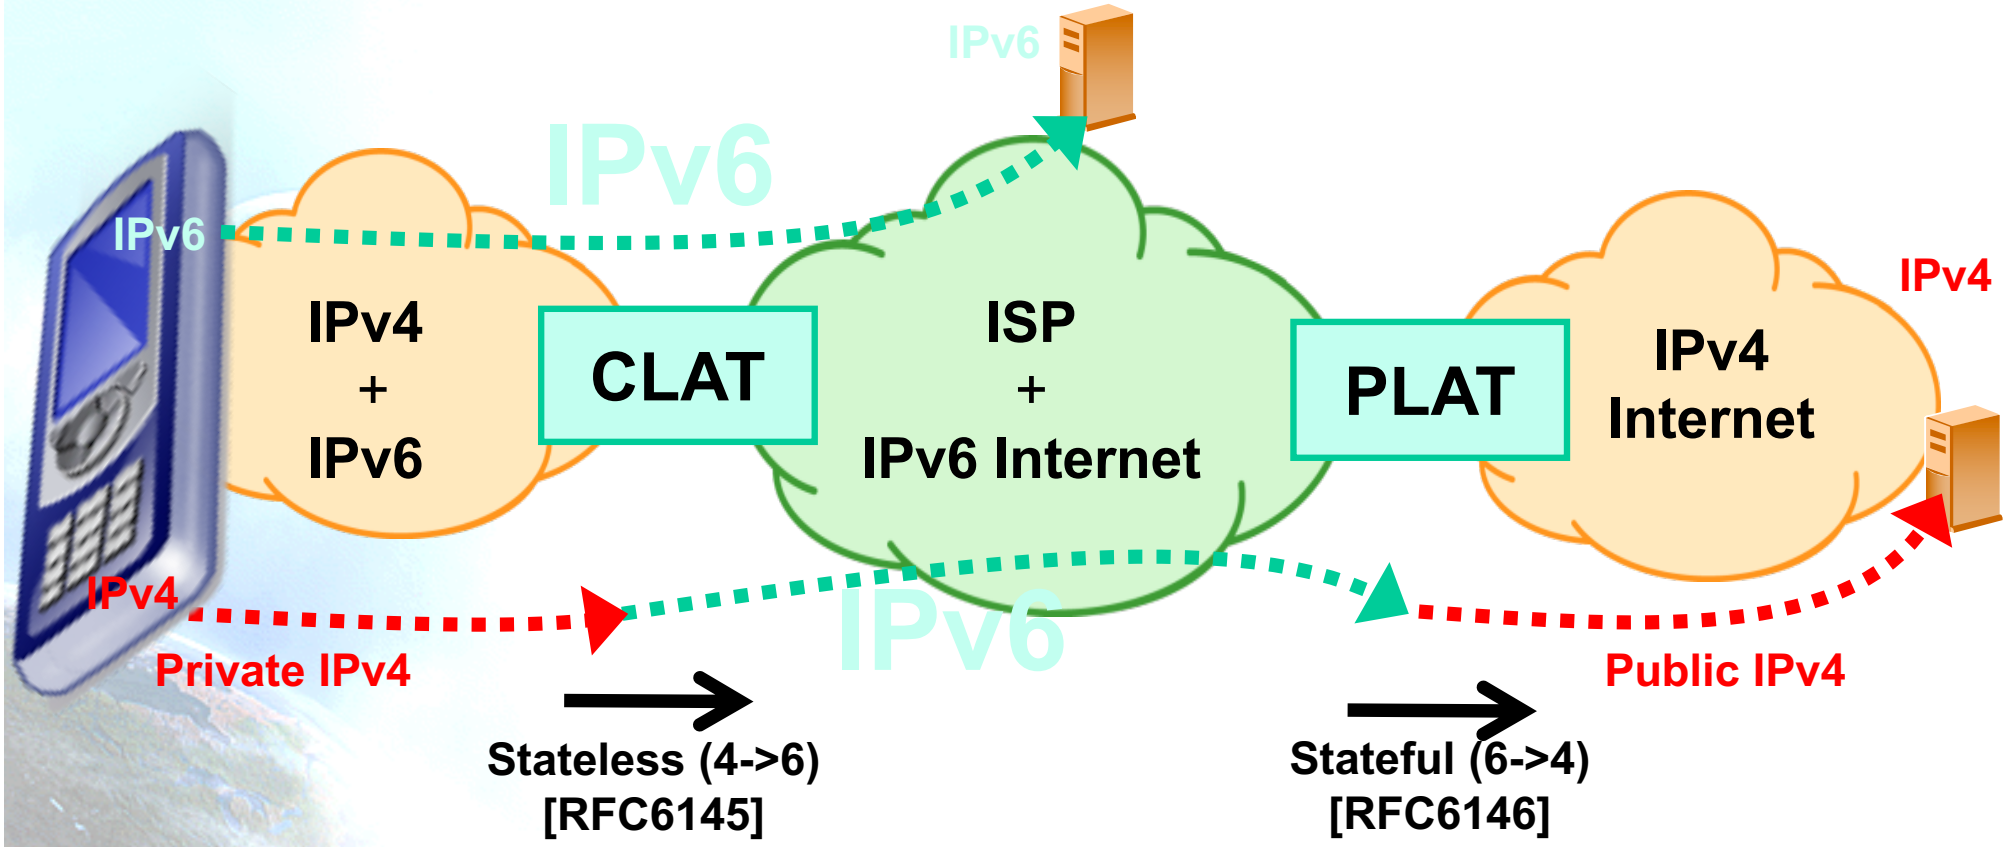

CLAT: Customer side translator (XLAT)

PLAT: Provider side translator (XLAT)

Possible “app” cases

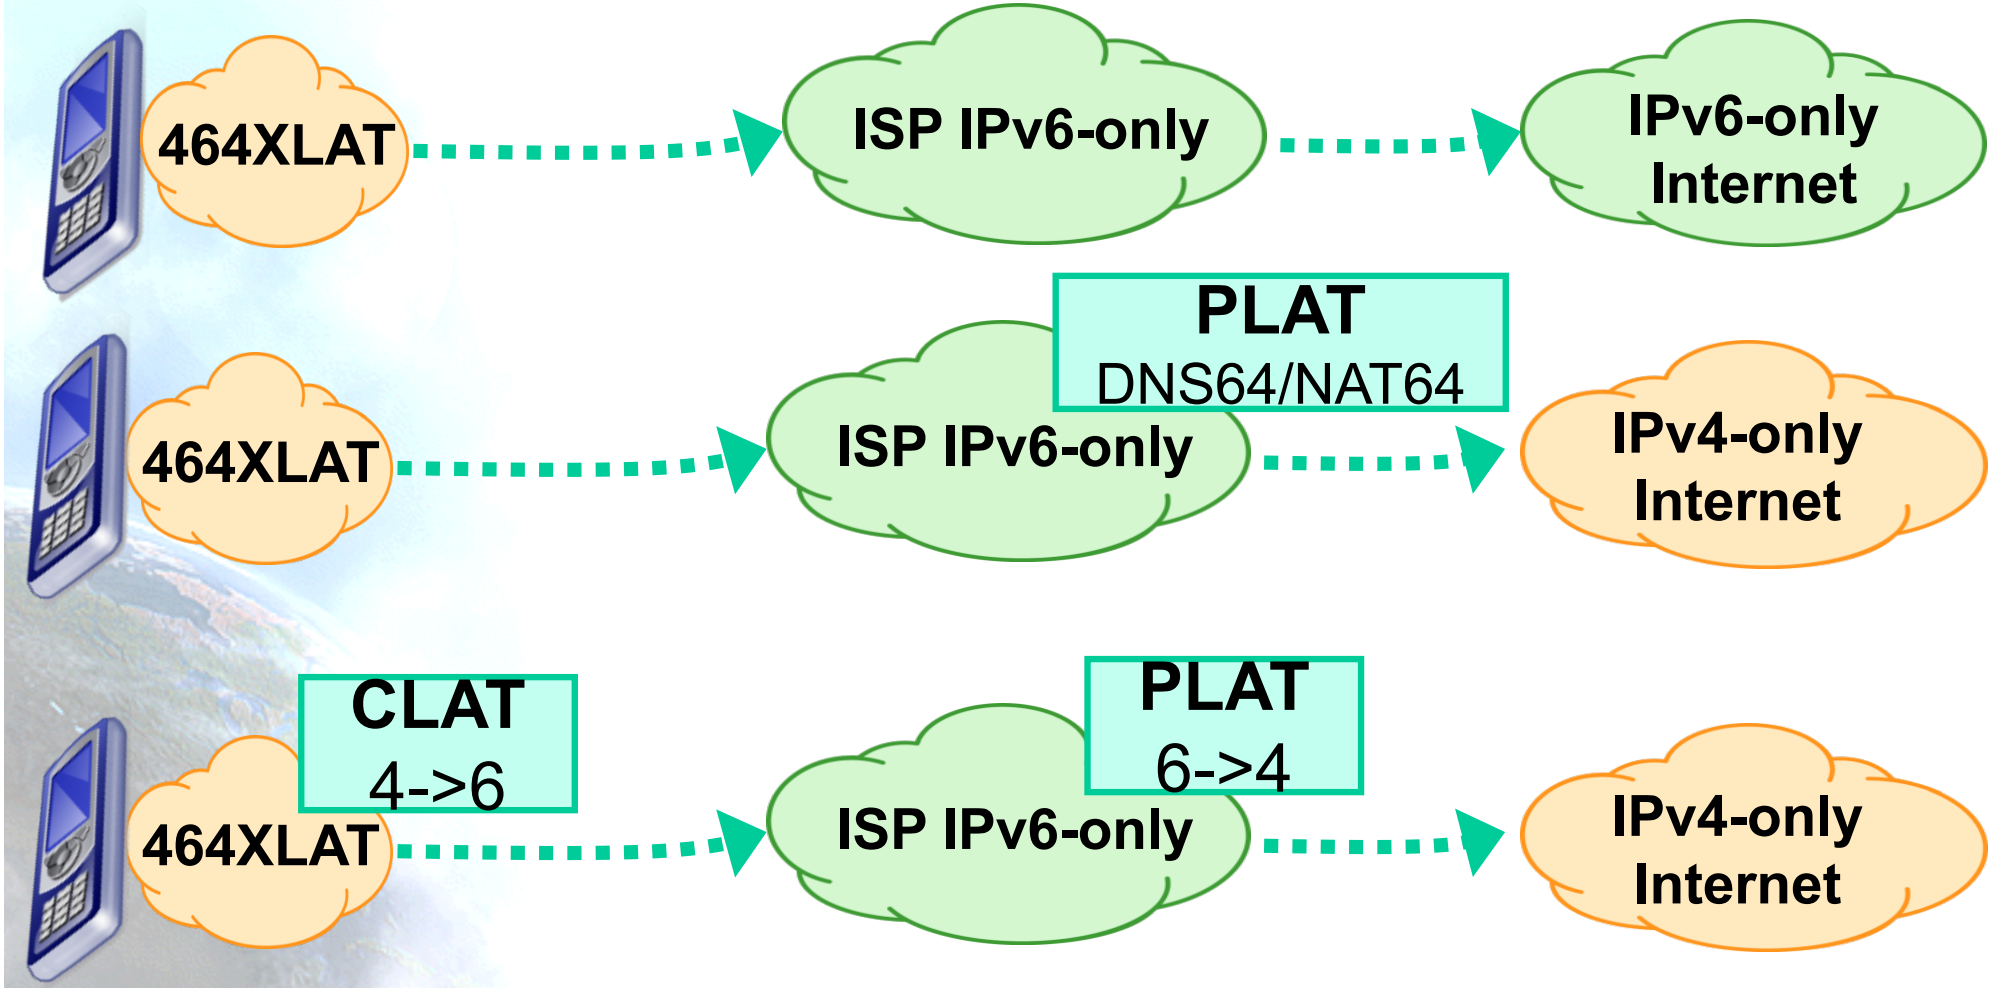

Multiservice Network

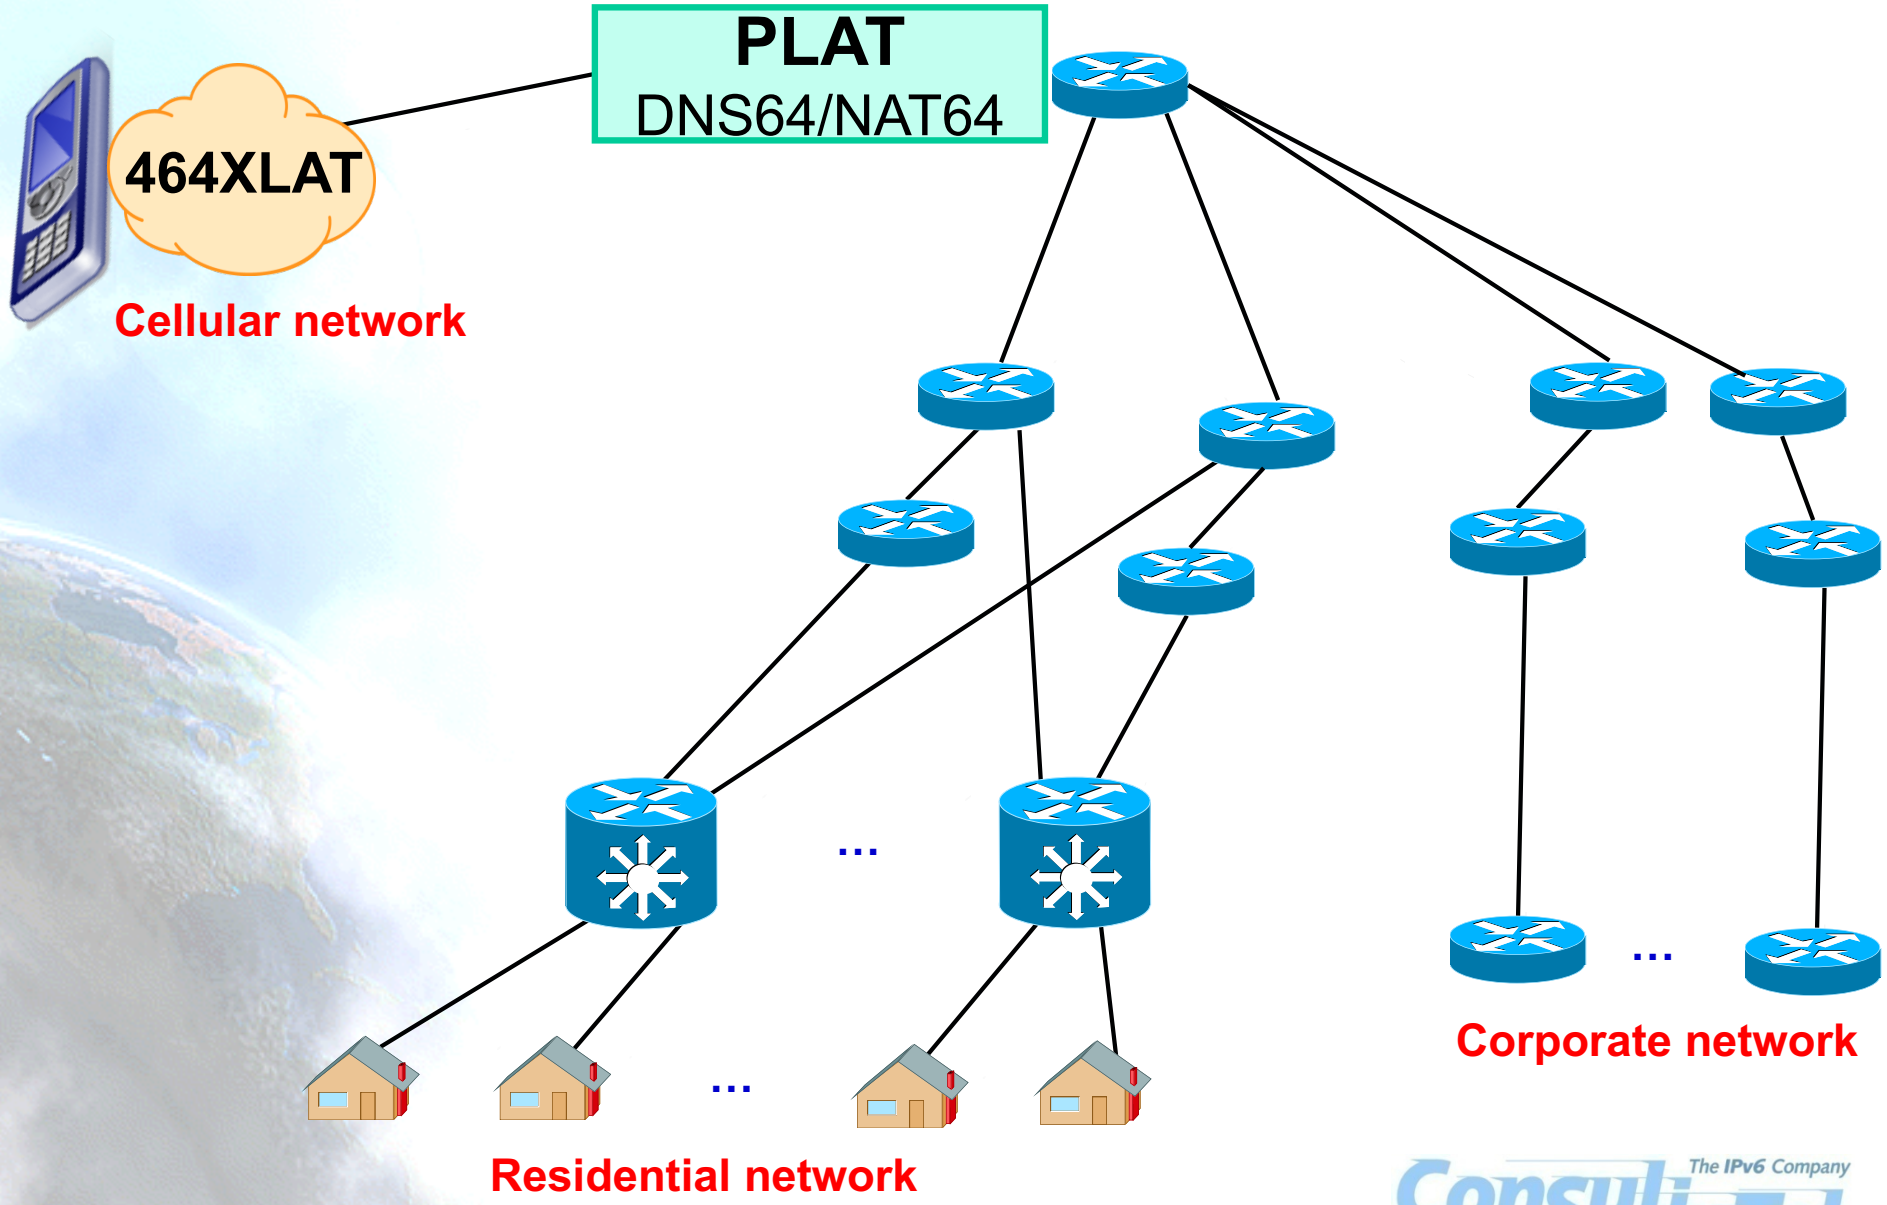

Example Residential Customer

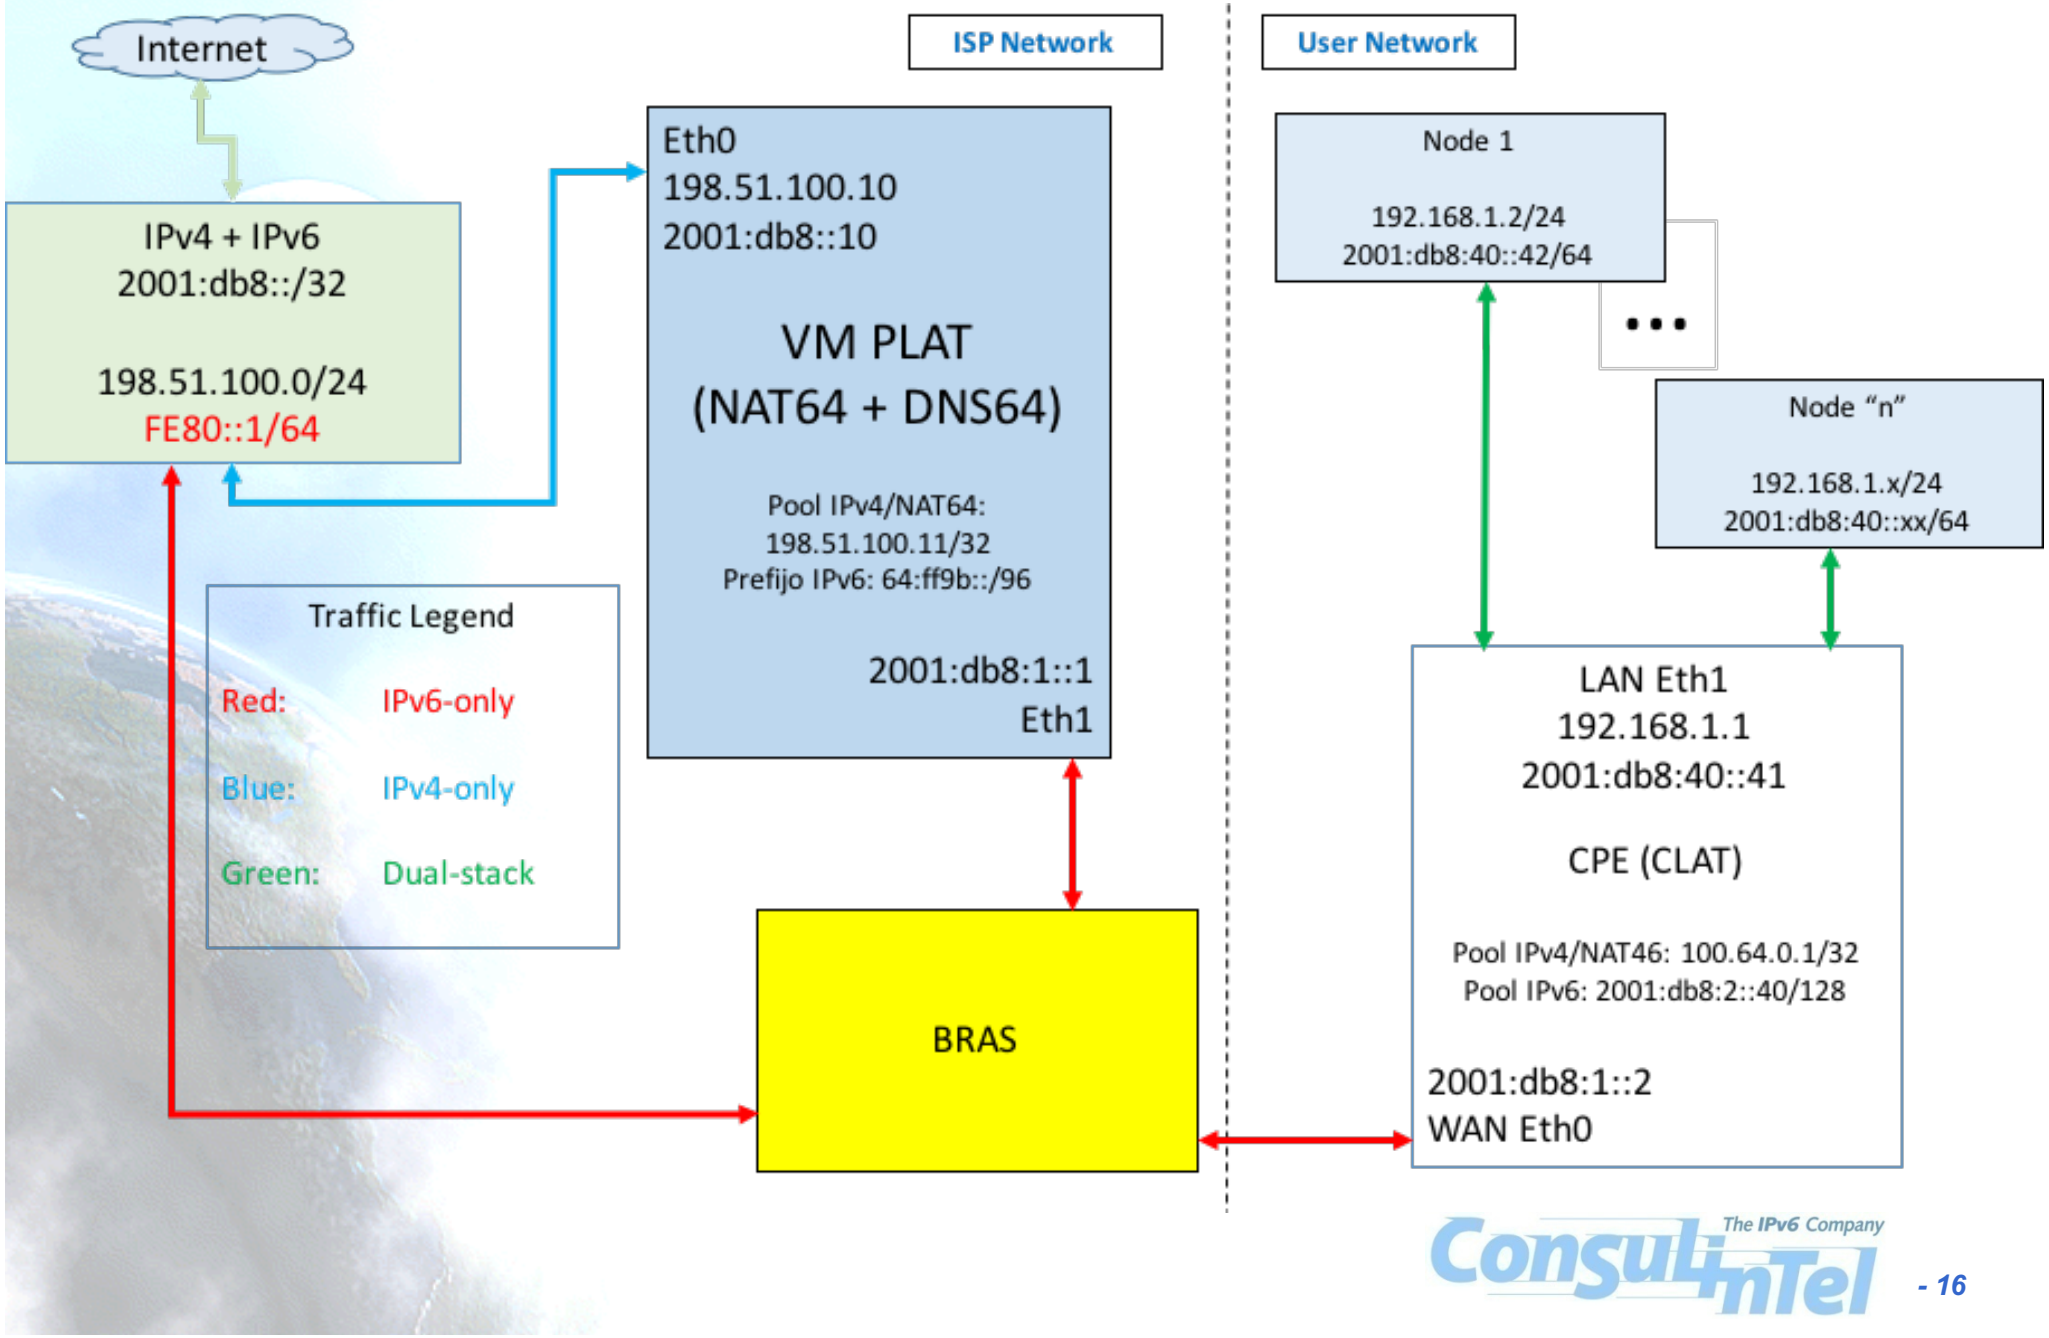

IPv6 in Cellular/US

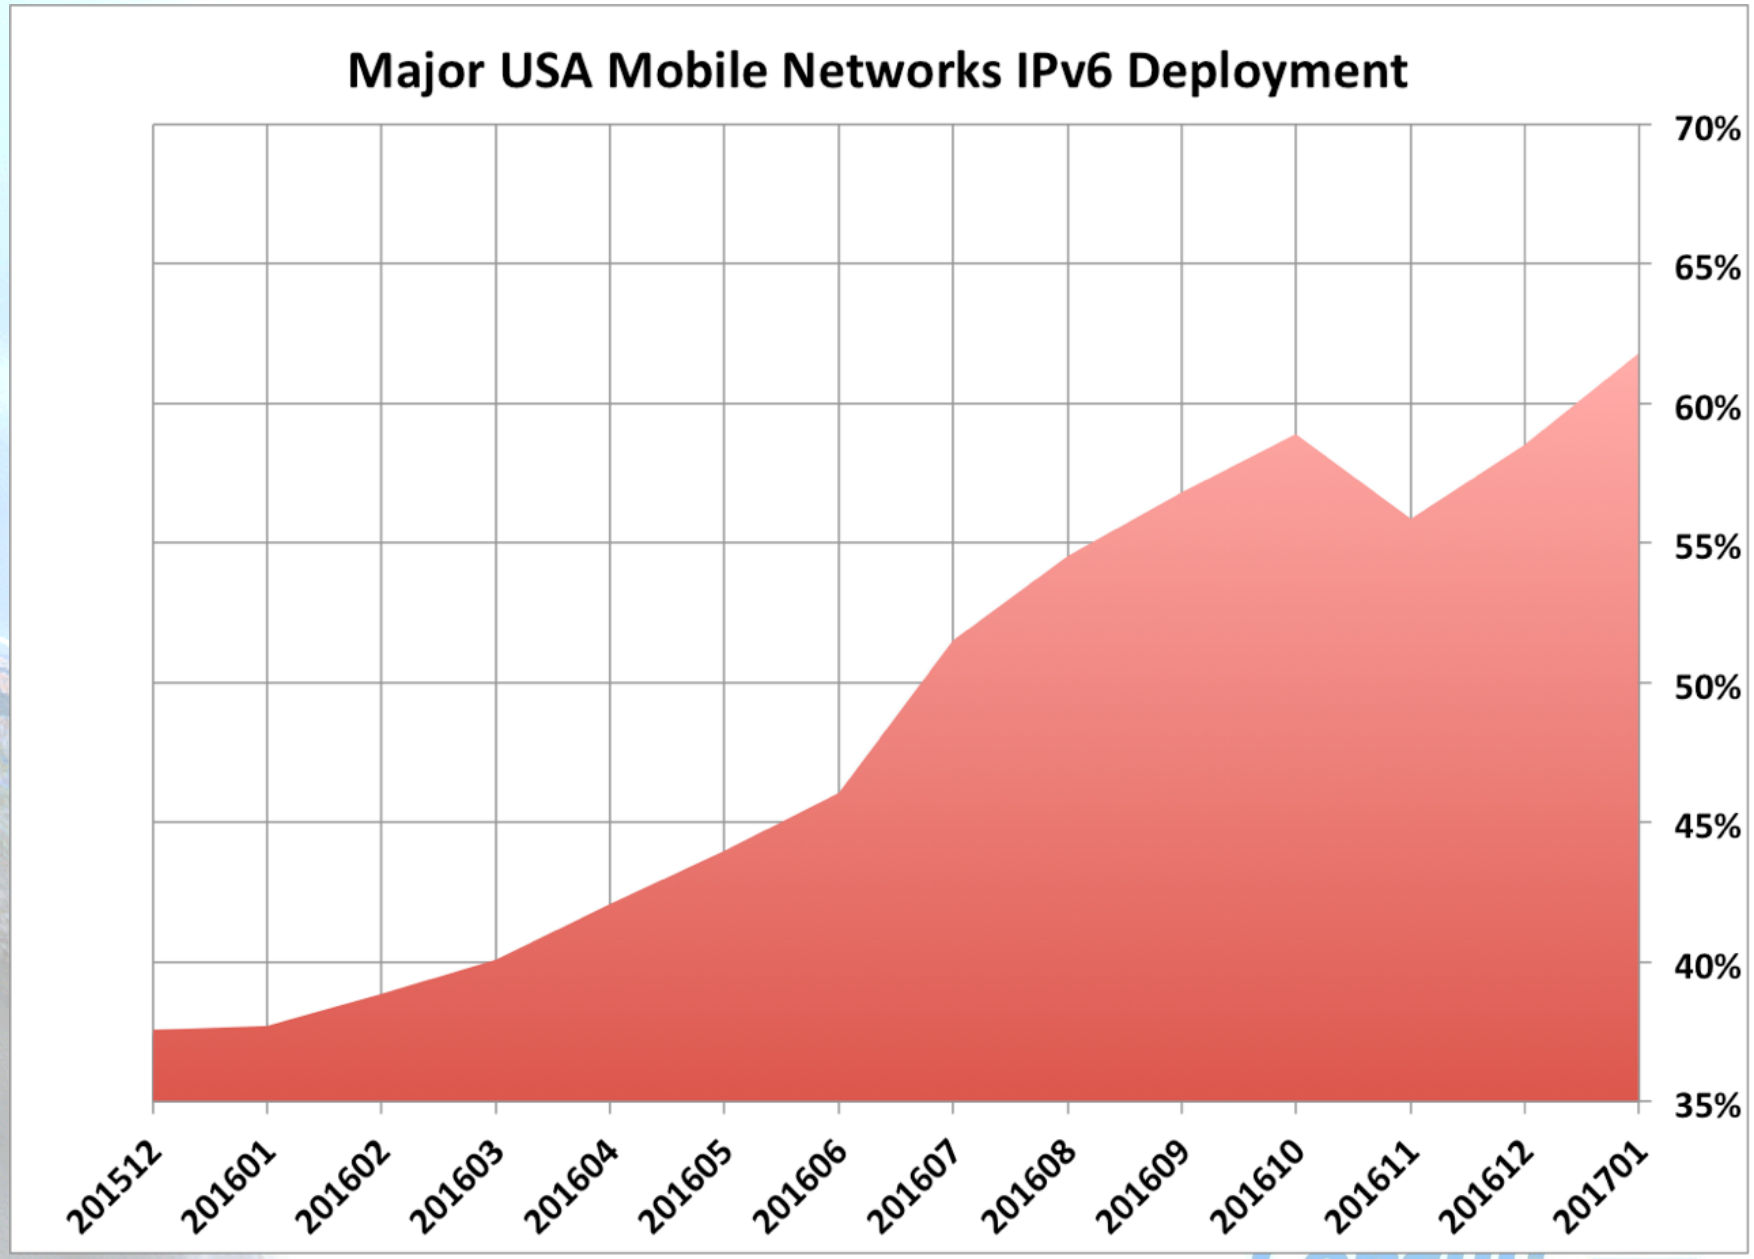

*ISOC/World IPv6 Launch data

464XLAT deployment

- NAT64:
  - A10
  - Cisco
  - F5
  - Juniper
  - NEC
  - Huawei
  - Jool, Tayga, Ecdsys, Linux, OpenBSD, …
- CLAT
  - Android
  - Nokia
  - Windows phone
  - NEC
  - OpenWRT
- Commercial deployments:
  - T-Mobile US: +68 Millions of users
  - Orange
  - Telstra
  - SK Telecom
  - …
  - Big trials in several ISPs (thousands of users)

Performance

US Mobile Performance – Dual Stack Provider iOS

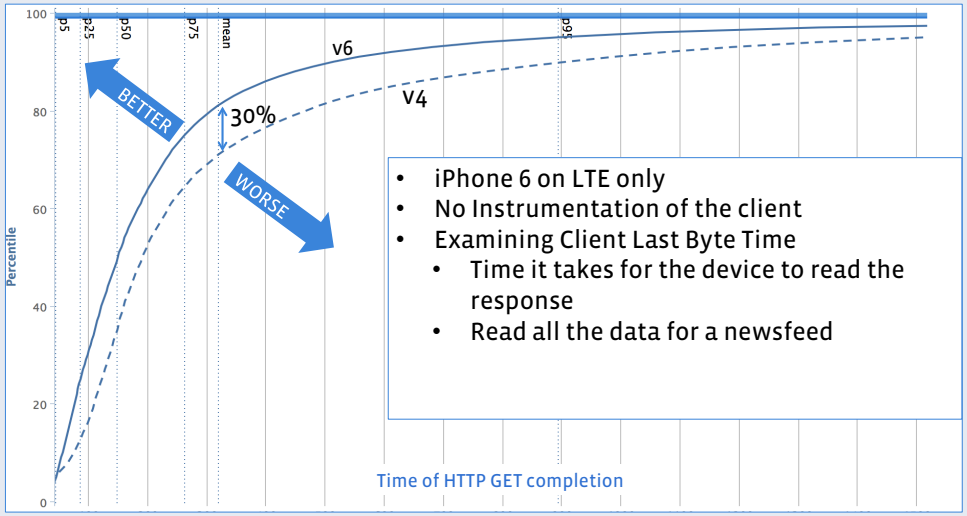

- iPhone 6 on LTE only
- No Instrumentation of the client
- Examining Client Last Byte Time
  - Time it takes for the device to read the response
  - Read all the data for a newsfeed

US Mobile Performance – Dual Stack Provider Android

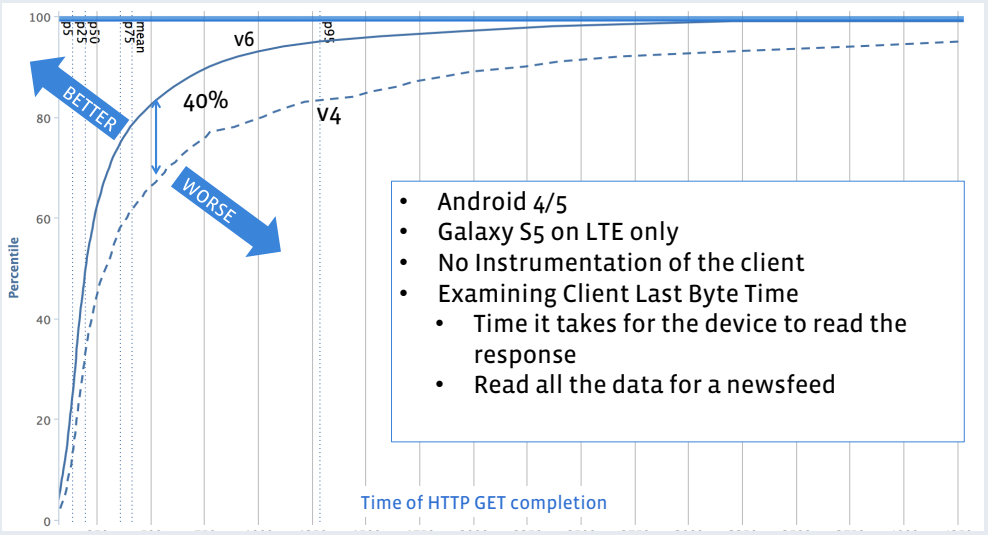

- Android 4/5
- Galaxy S5 on LTE only
- No Instrumentation of the client
- Examining Client Last Byte Time
  - Time it takes for the device to read the response
  - Read all the data for a newsfeed

US Mobile Performance – Dual Stack Provider iOS

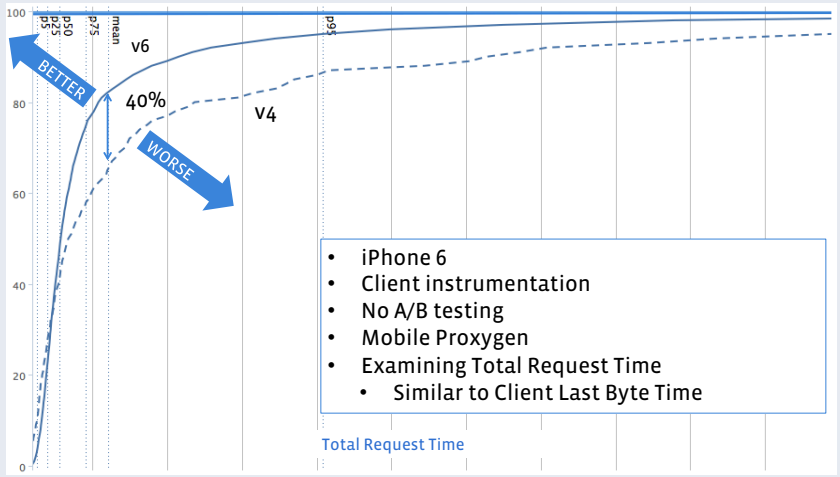

- iPhone 6
- Client instrumentation
- No A/B testing
- Mobile Proxygen
- Examining Total Request Time
  - Similar to Client Last Byte Time

***FaceBook data**

(17/3/2015)

Update of RFC7084

- Basic Requirements for IPv6 Customer Edge Routers
  - Originally include support only for 6RD and DS-LITE
  - Being updated to include support for 464XLAT, MAP T/E, lw4o6, …
- https://tools.ietf.org/html/draft-ietf-v6ops-rfc7084-bis

IPv6 Deployment Survey (Residential/Household Services)

How IPv6 is being deployed? (October 2016)

Jordi Palet

(jordi.palet@consulintel.es)

Consulintel, CEO/CTO

The IPv6 Company

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”

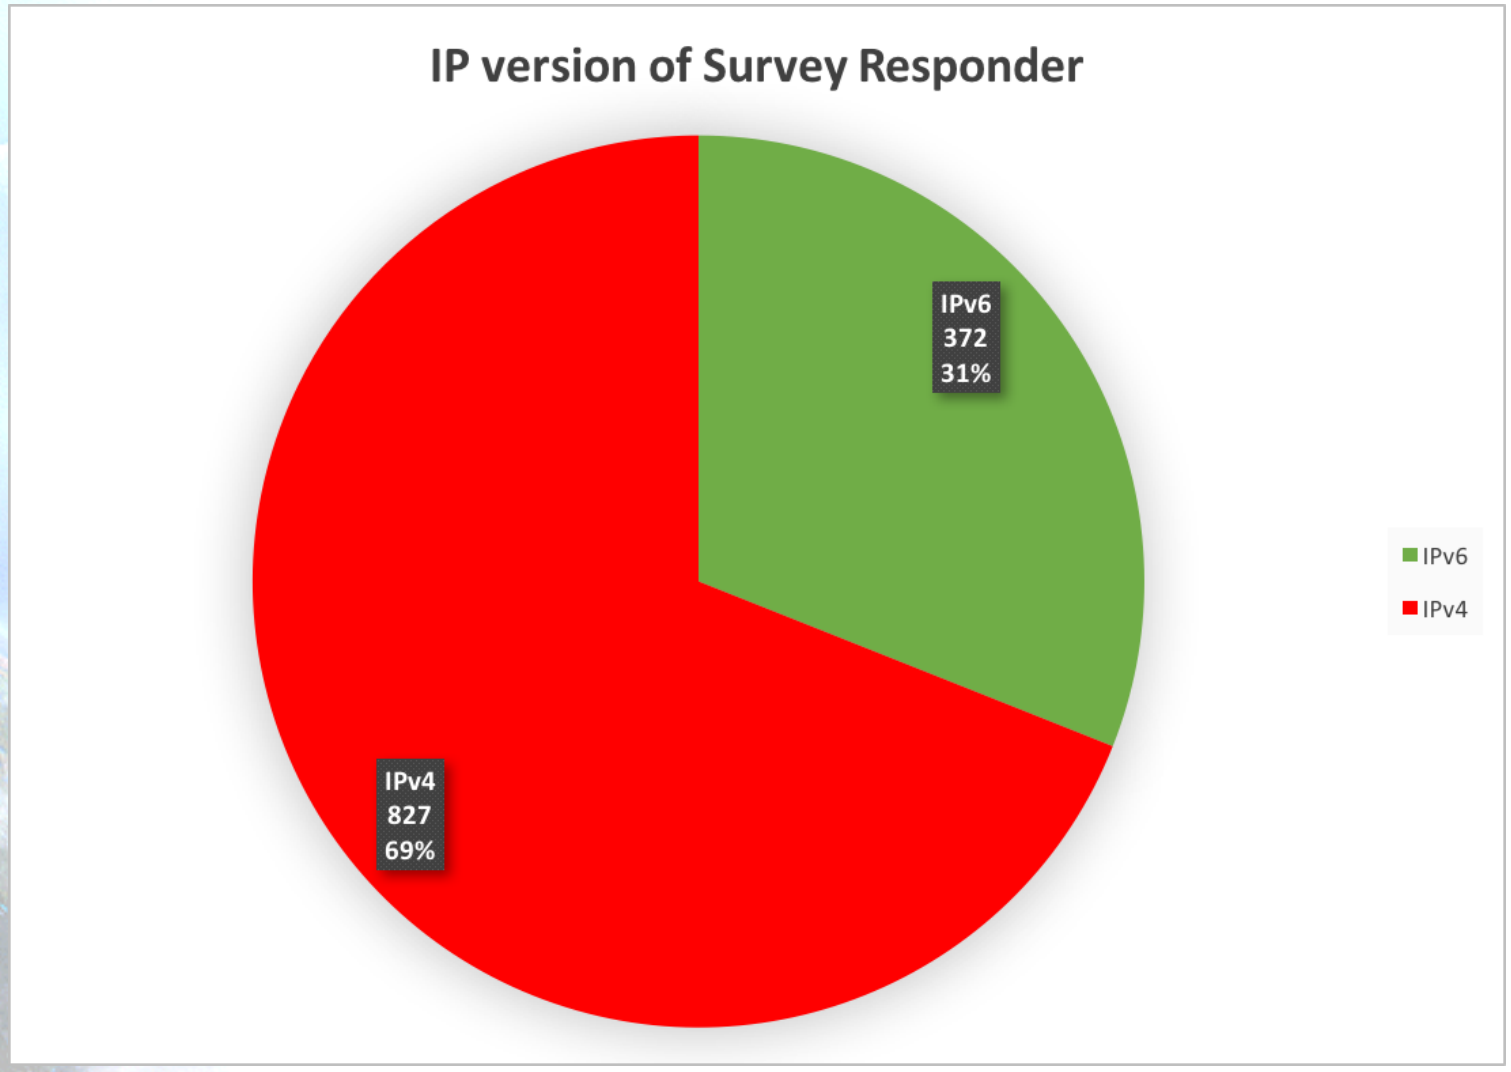
IP version of Survey Responder
IPv6
372
31%
IPv4
827
69%
IPv6
IPv4

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?

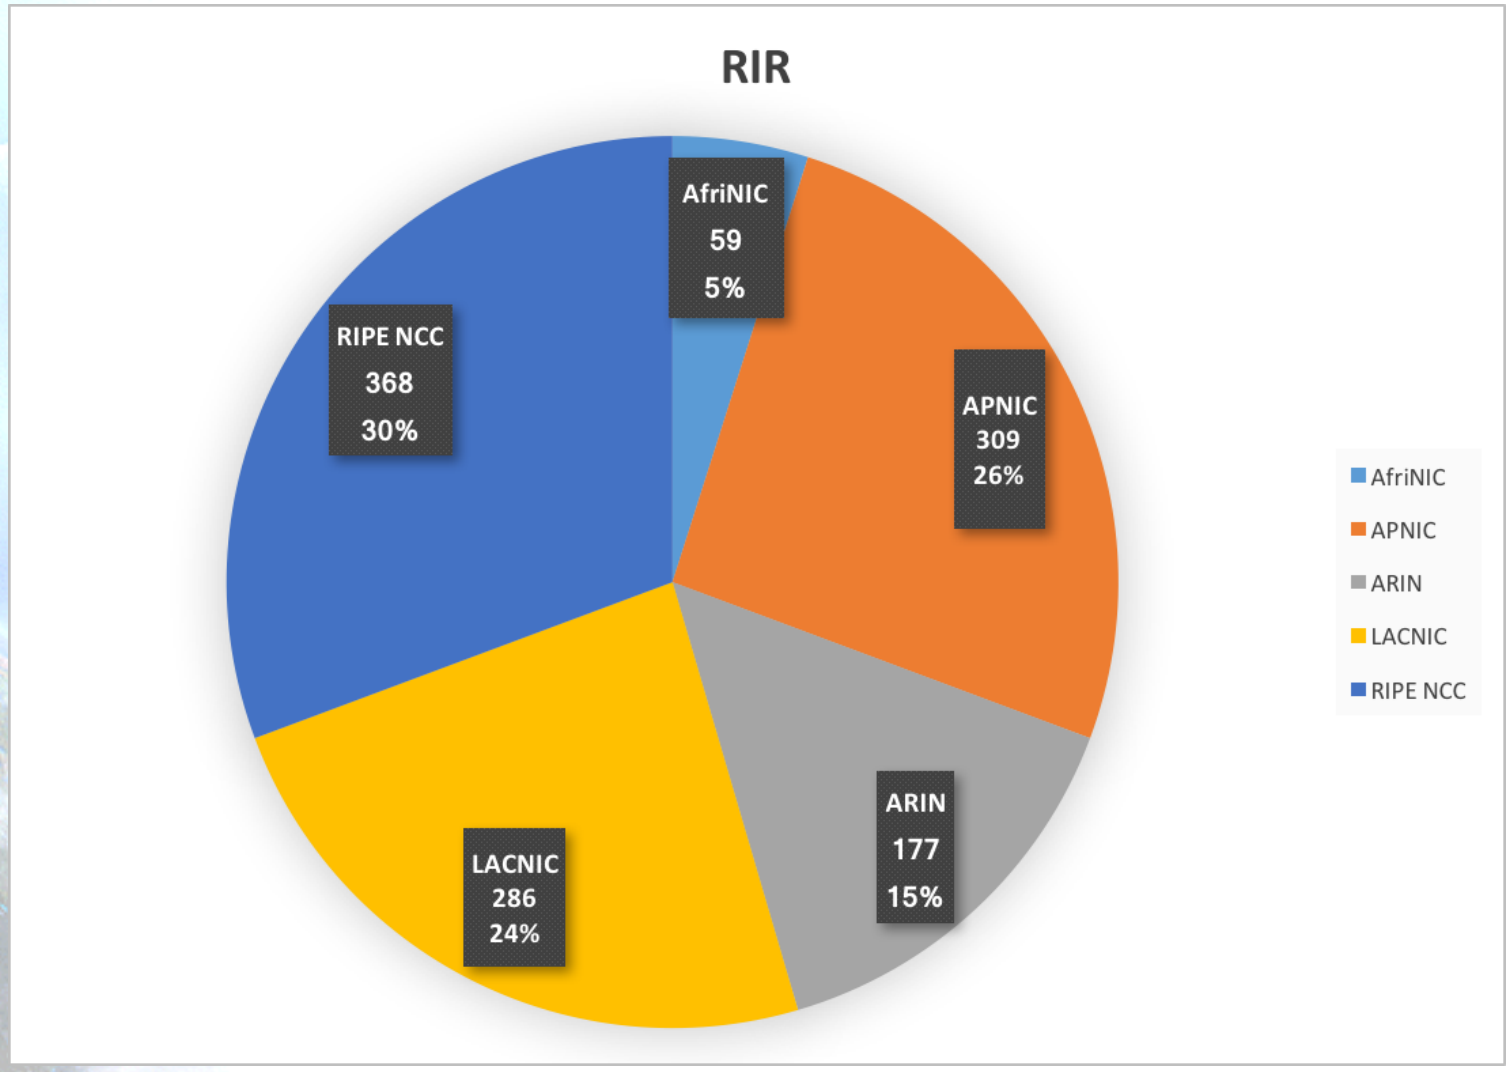
RIR
AfriNIC
59
5%
APNIC
309
26%
ARIN
177
15%
LACNIC
286
24%
RIPE NCC
368
30%
AfriNIC
APNIC
ARIN
LACNIC
RIPE NCC

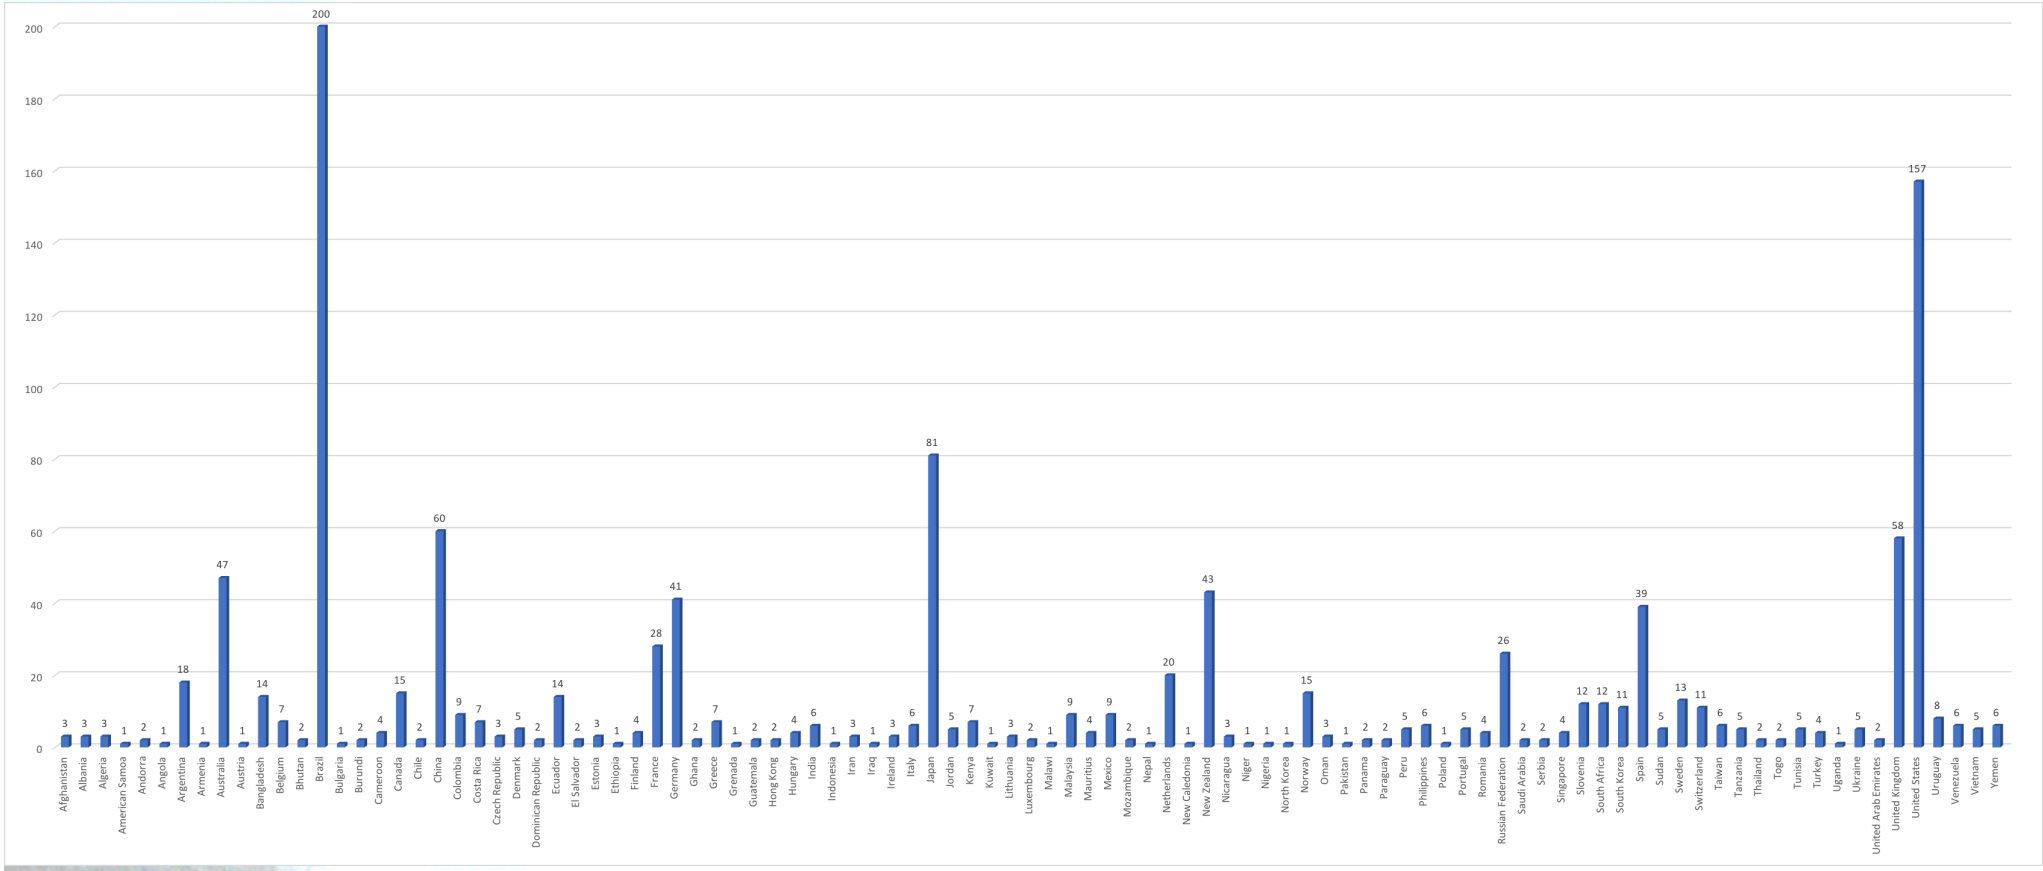

- Responses from 100 countries

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

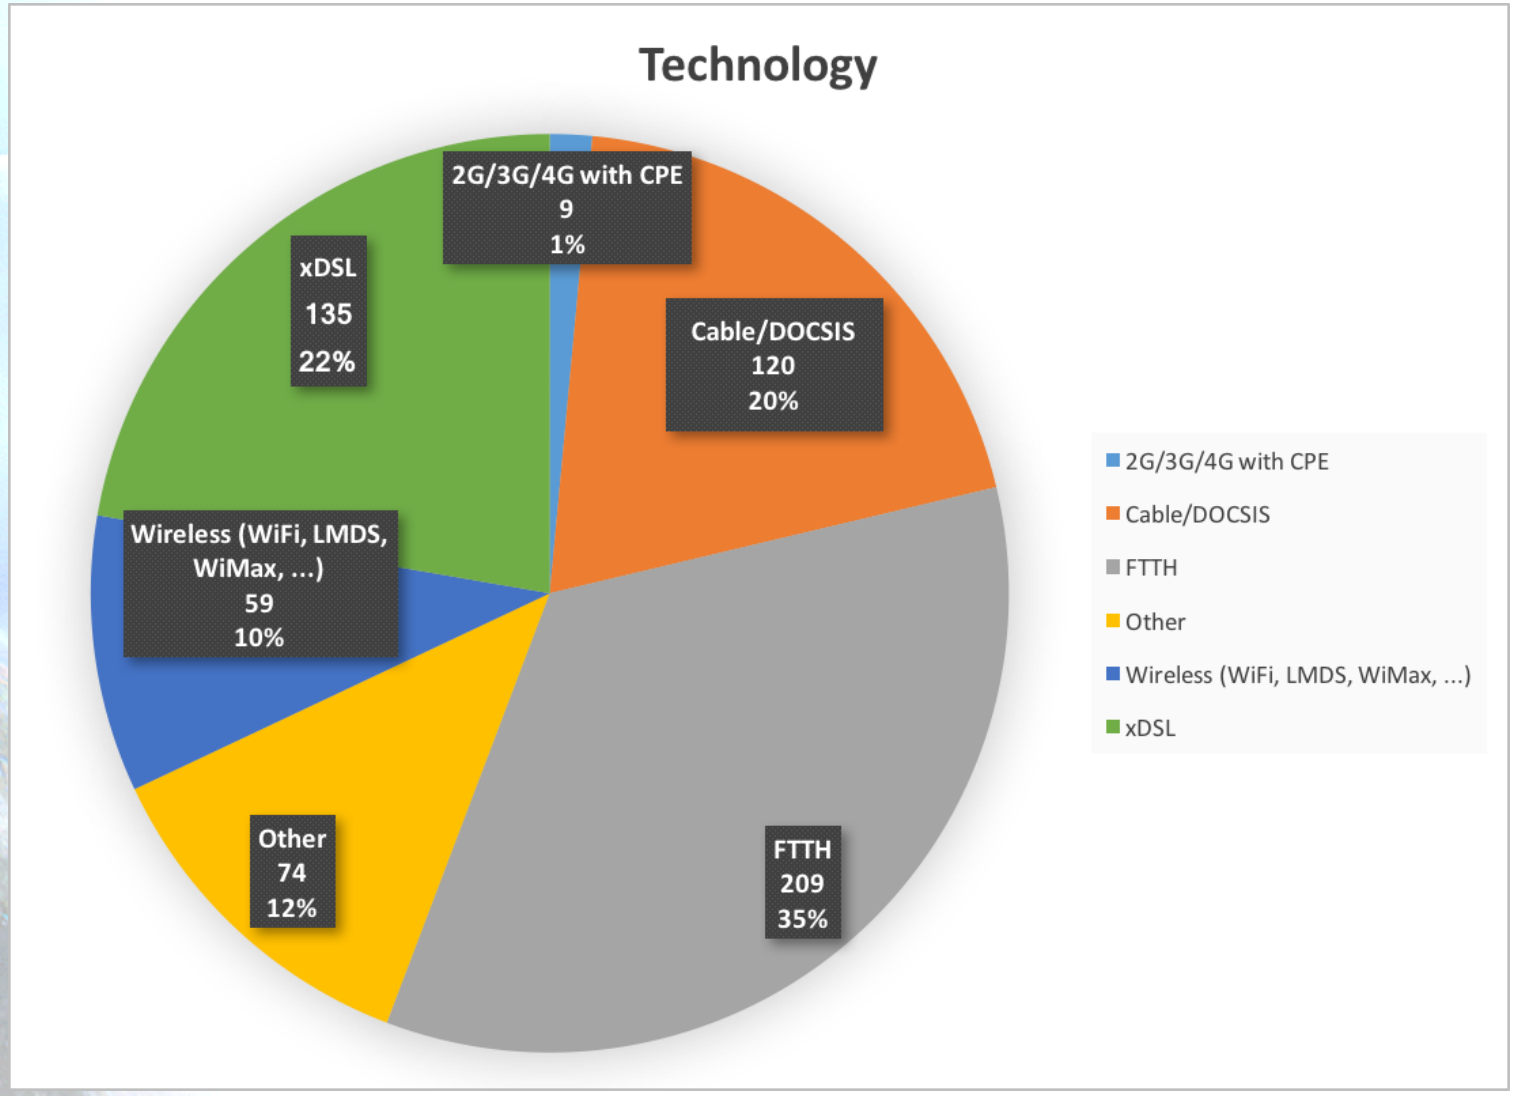
Technology
2G/3G/4G with CPE
9
1%
Cable/DOCSIS
120
20%
FTTH
209
35%
Other
74
12%
Wireless (WiFi, LMDS, WiMax, ...)
59
10%
xDSL
135
22%
2G/3G/4G with CPE
Cable/DOCSIS
FTTH
Other
Wireless (WiFi, LMDS, WiMax, ...)
xDSL

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

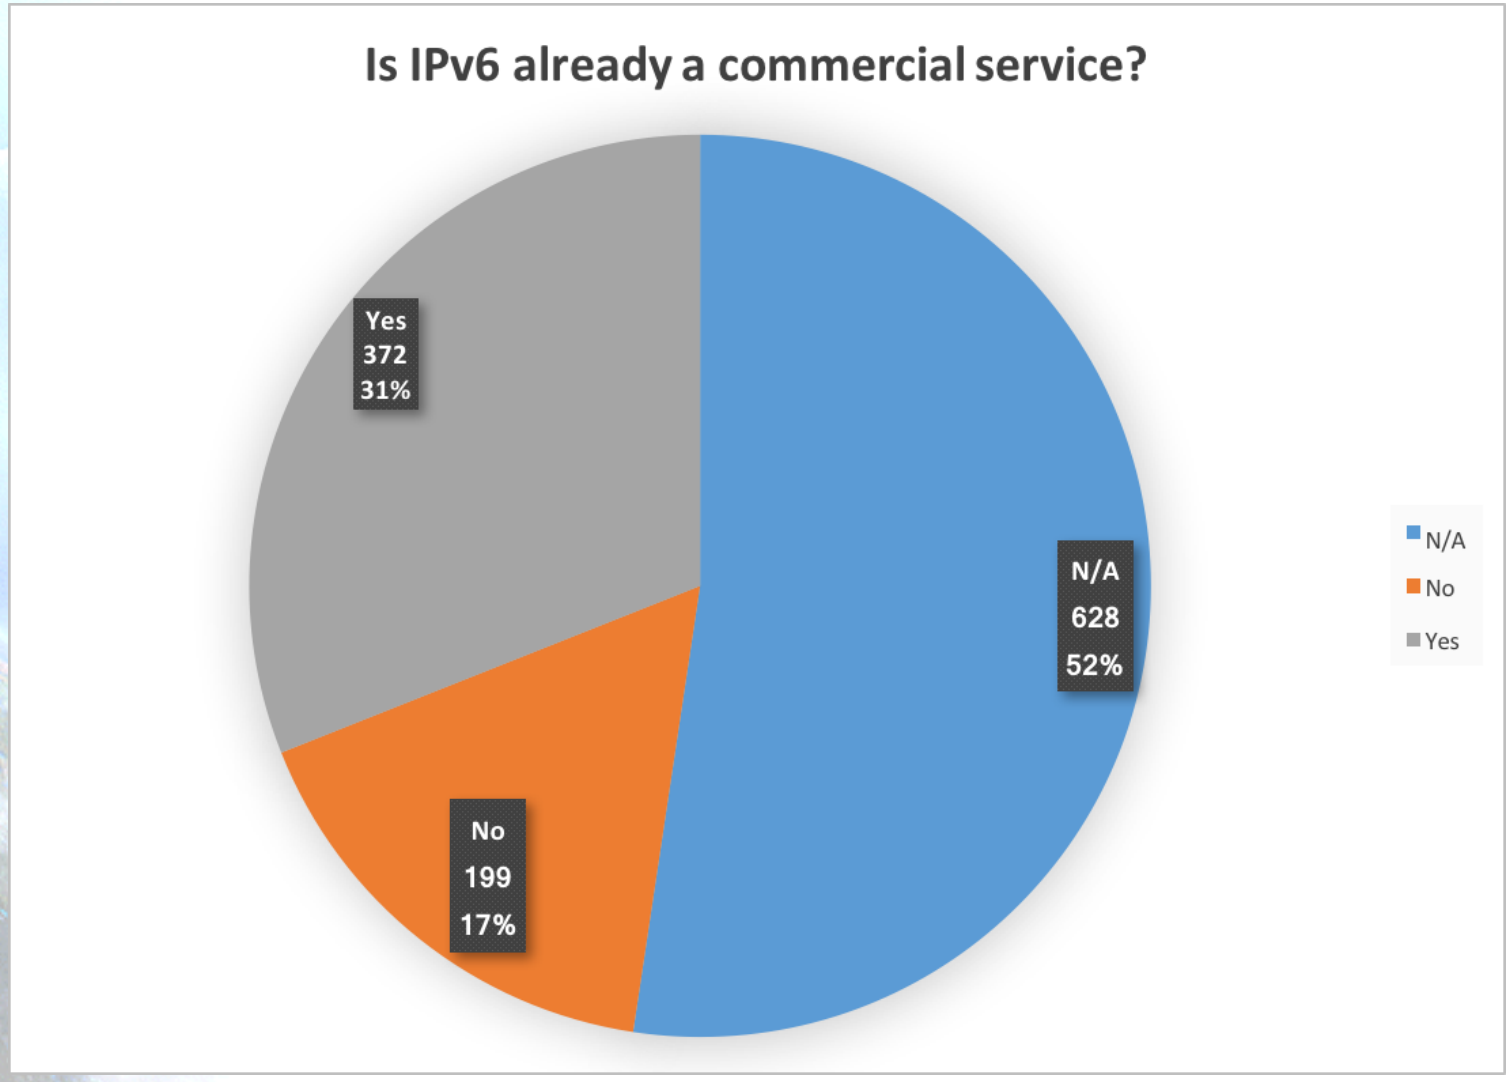
Is IPv6 already a commercial service?
Yes
372
31%
N/A
628
52%
No
199
17%
N/A
No
Yes

Why still not commercial?

- 52% responses –> No Answer, mainly customers or even employees of ISPs which really don't know
- 31% Yes, already commercial
- 17% 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

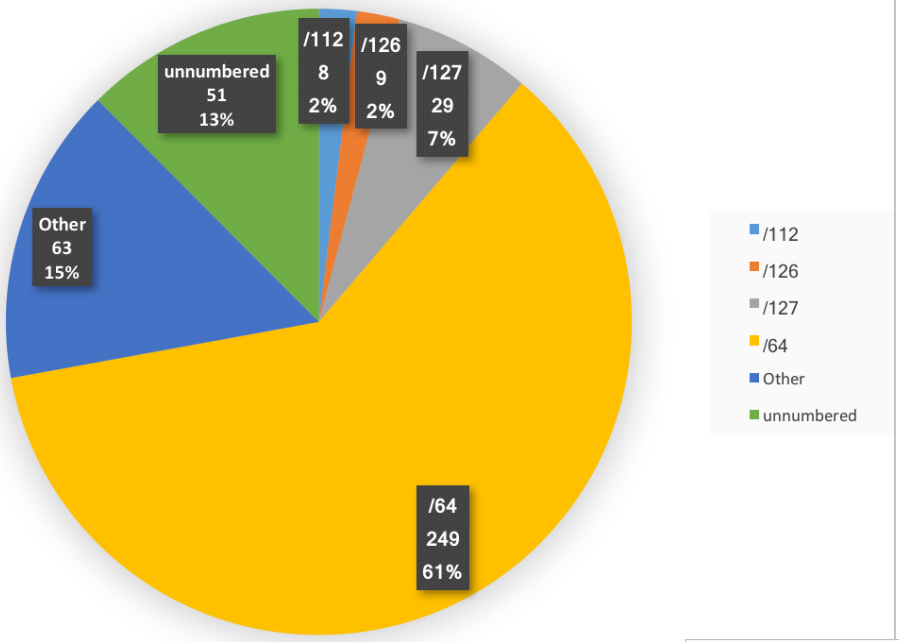

WAN Prefix Stable

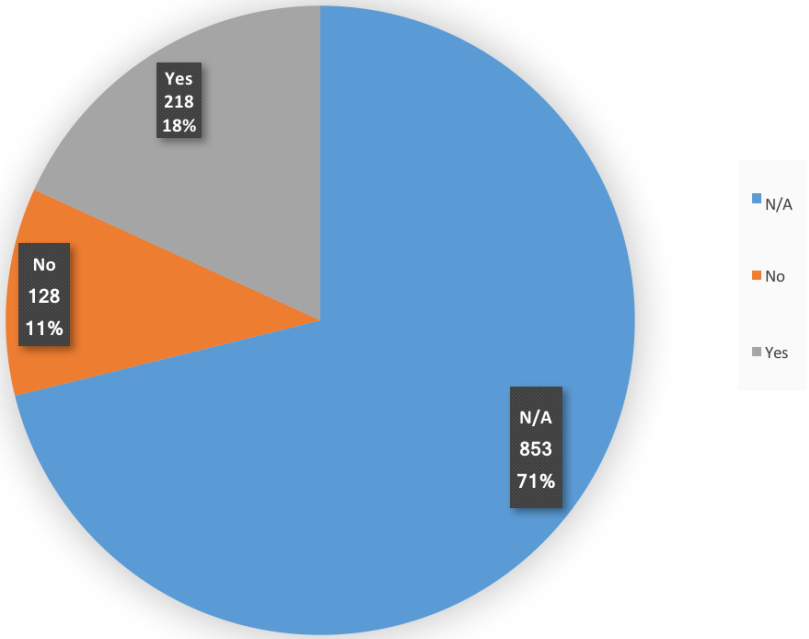

WAN Addressing Type

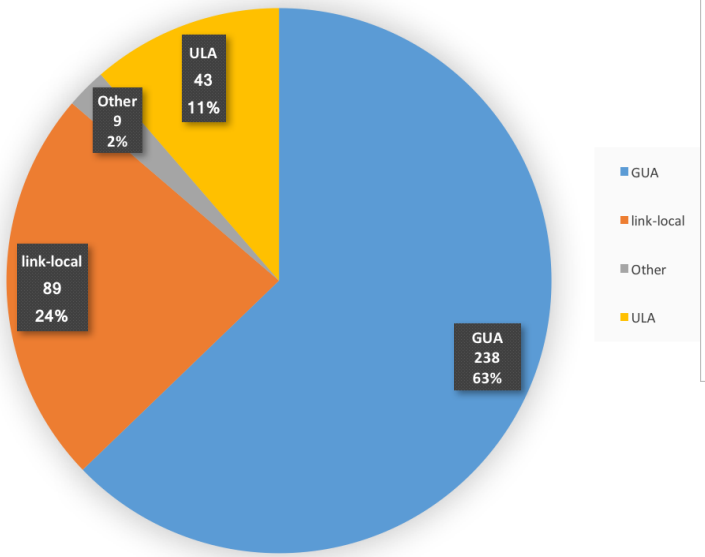

WAN from same pool as customer prefixes

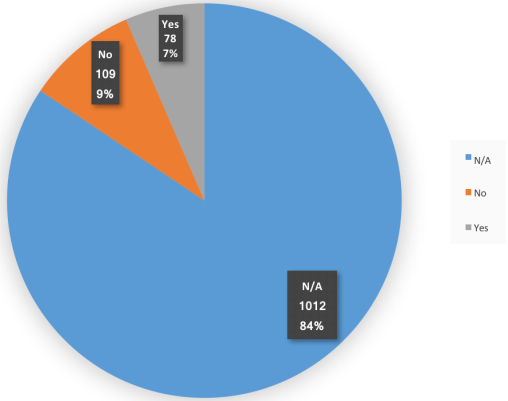

WAN /64 from customer prefix

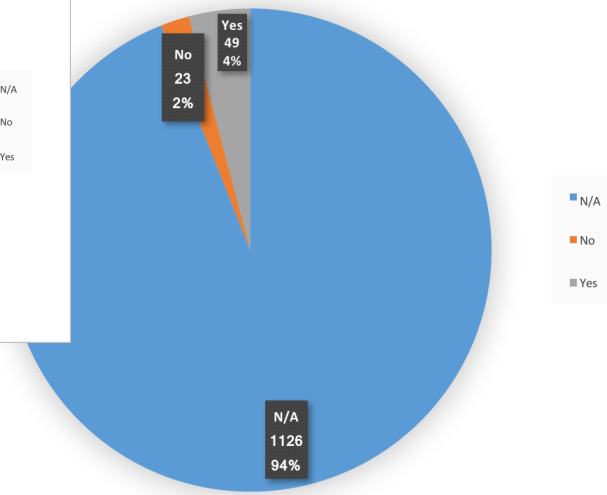

WAN prefix issues

- Remarkable -> /64 61%
- What means other?
  - /128, /62, /60, /56, /48, /32 ... No comments
- Why not stable (11%)? -> Note 71% no answer
  - Provisioning systems?
- 63% using GUA
- Interesting figures about using the /64 from the customer allocated prefix
- Distribution of those technical aspects not related to any specific country/region

LAN Prefix Size

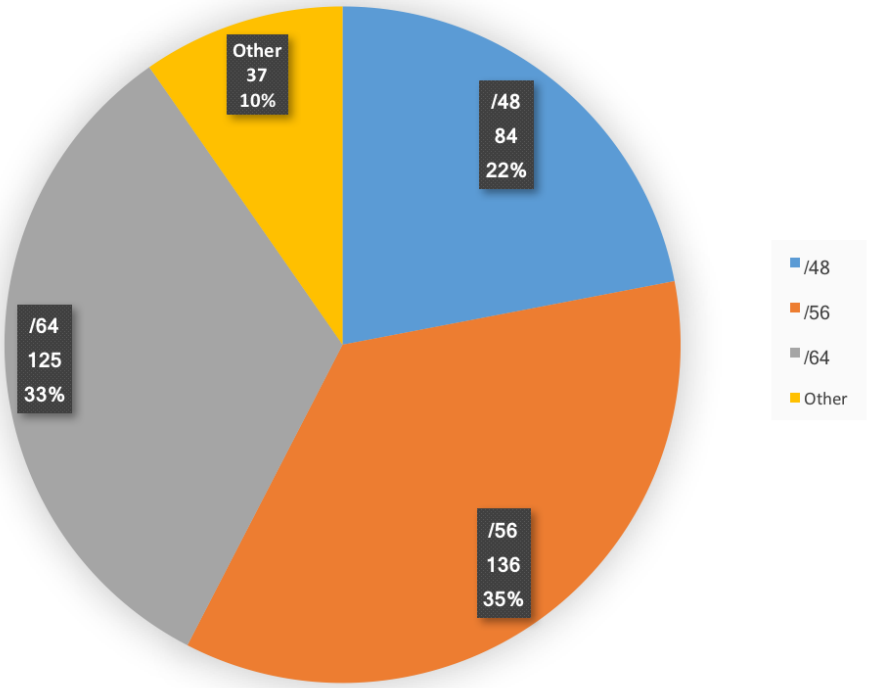

LAN Prefix Stable

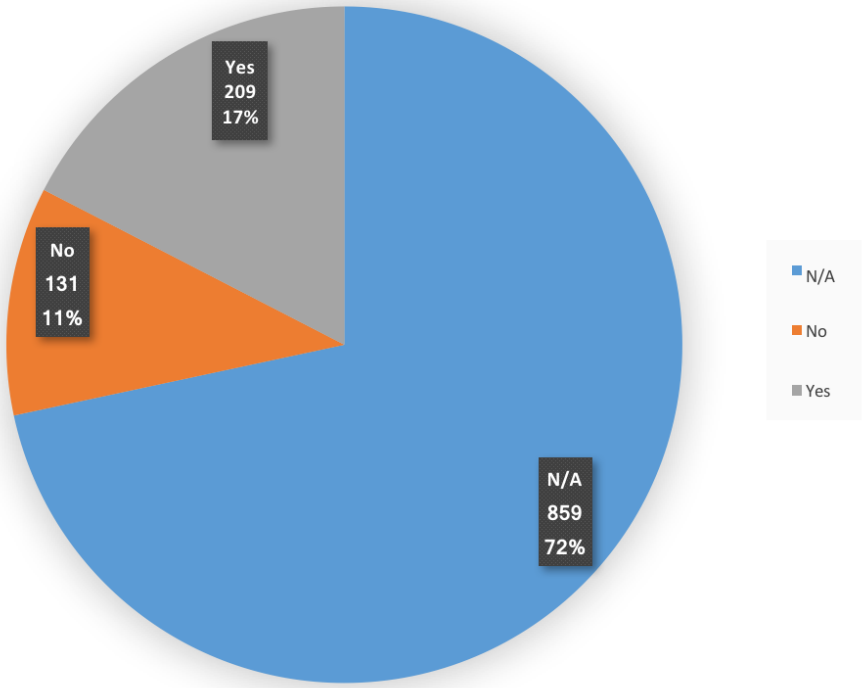

Can the customer opt to have it "stable"?

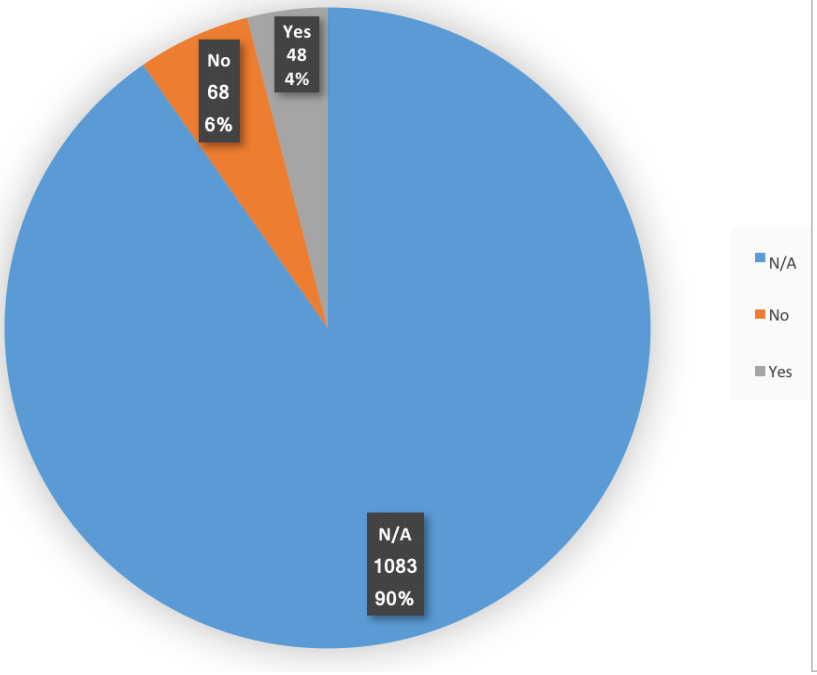

Extra cost (on top of stable IPv4)?

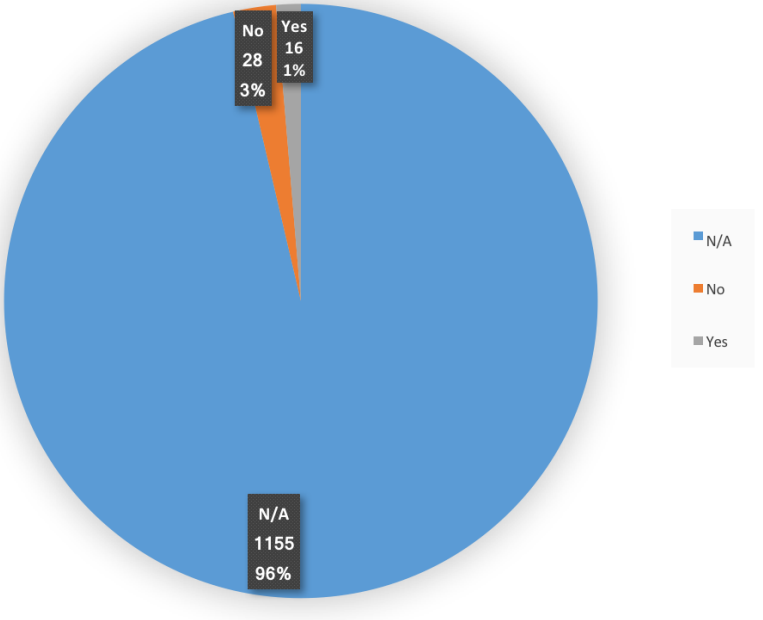

LAN prefix issues

- 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
  - 22% /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?

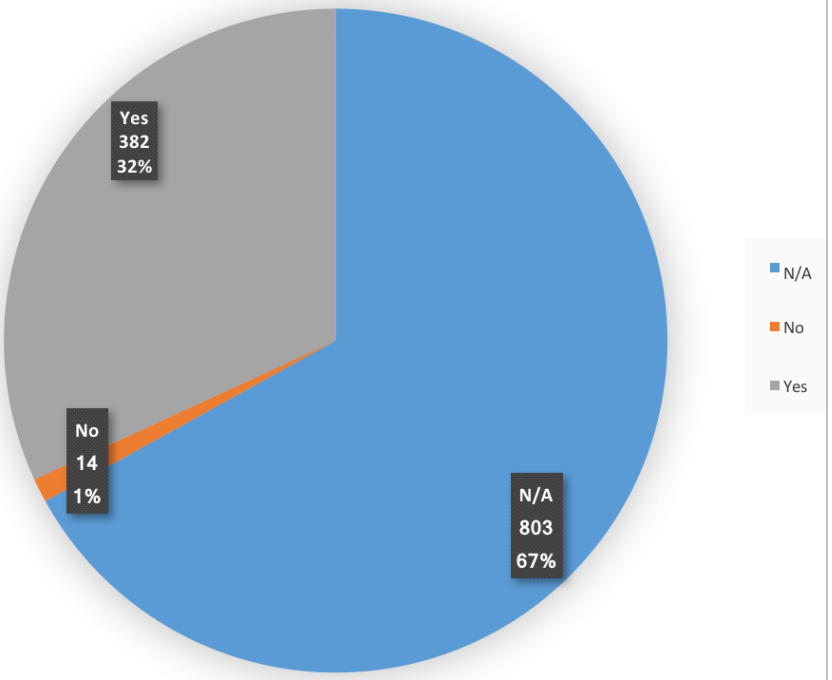

Public IPv4 address at CPE WAN?

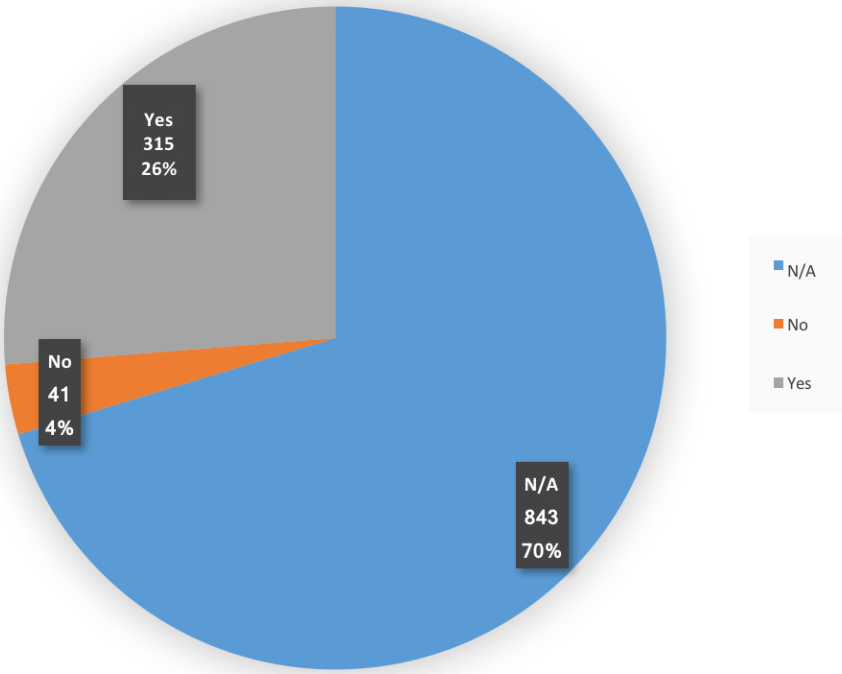

IPv4 address is "stable"?

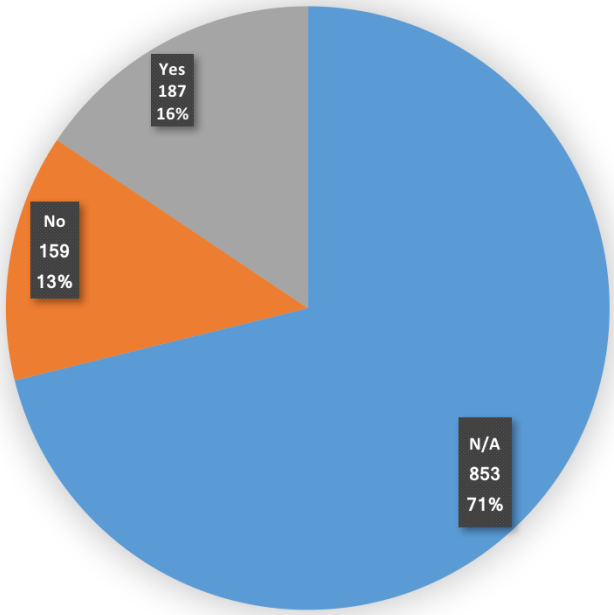

Can the customer opt to have IPv4 "stable"?

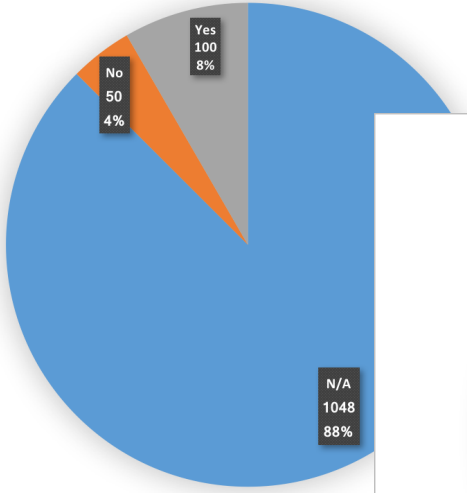

Extra cost for stable IPv4?

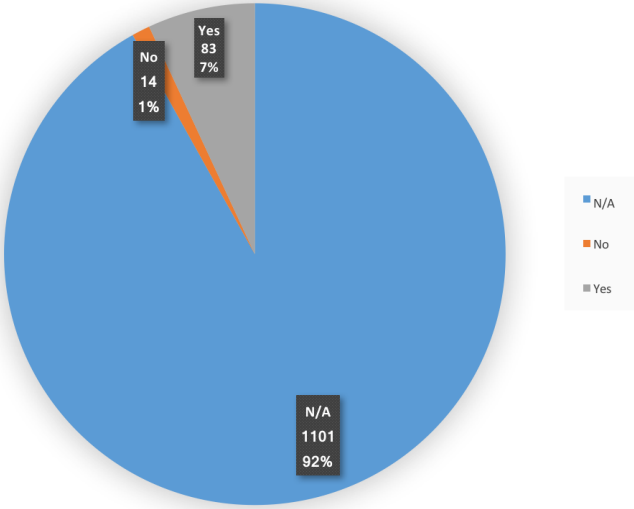

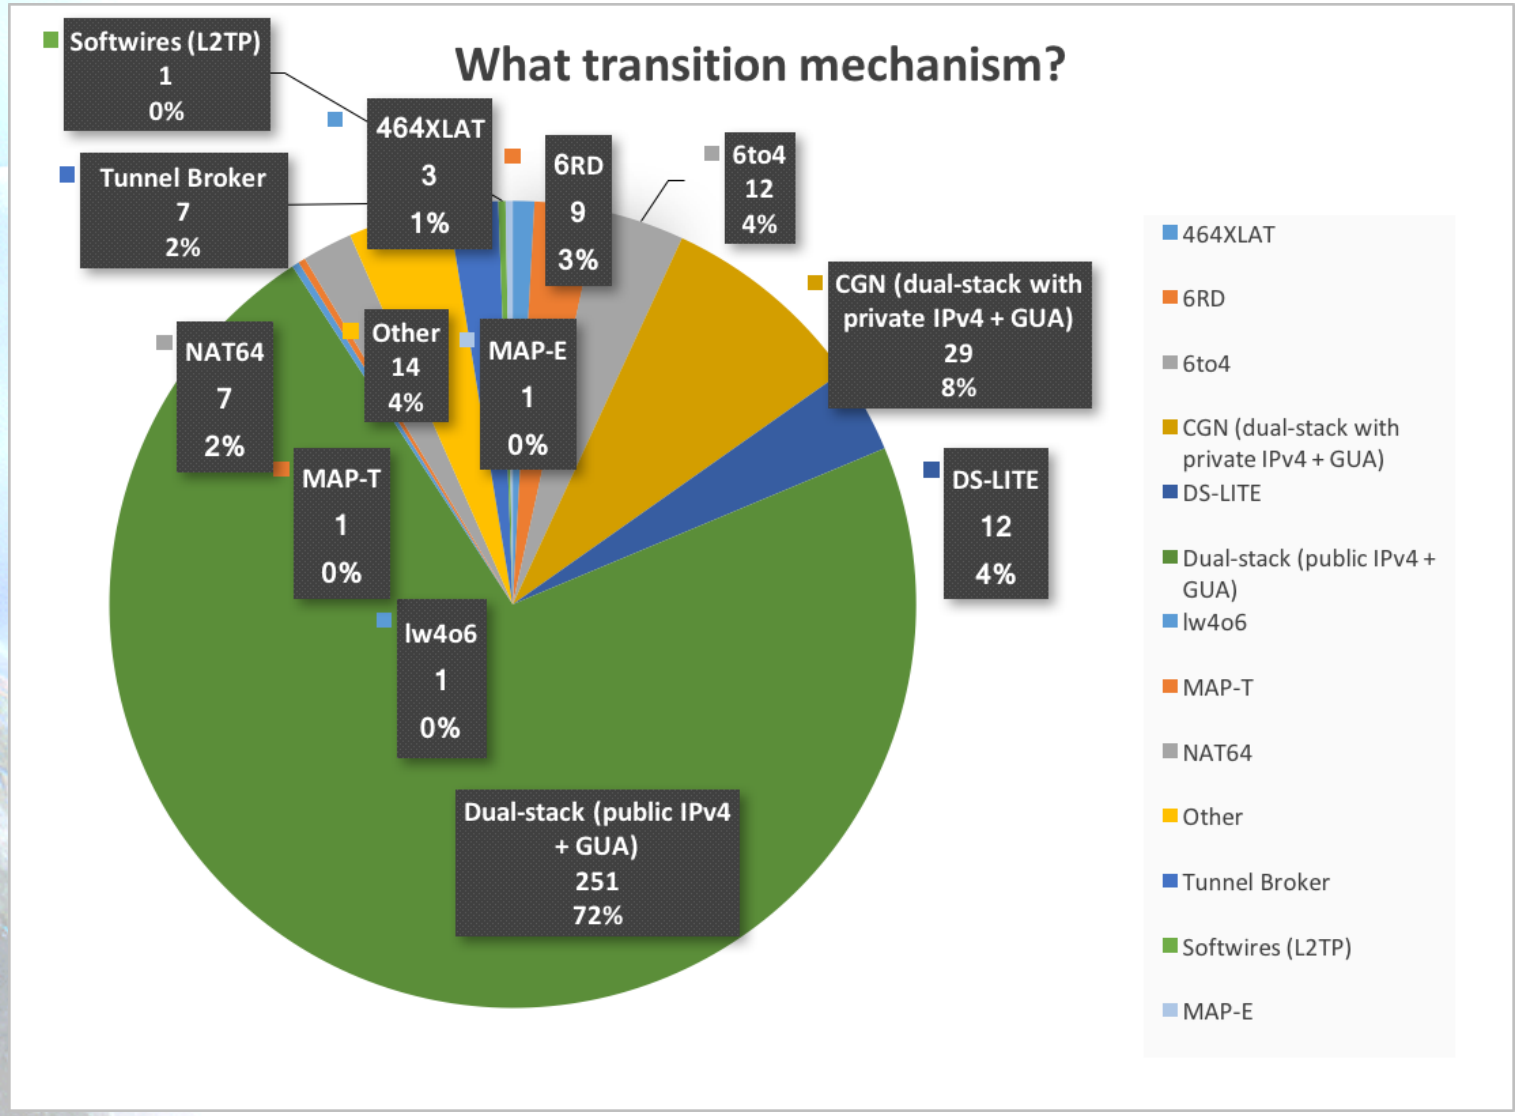
What transition mechanism?
Softwires (L2TP)
1
0%
464XLAT
3
1%
6RD
9
3%
6to4
12
4%
Tunnel Broker
7
2%
CGN (dual-stack with private IPv4 + GUA)
29
8%
NAT64
7
2%
Other
14
4%
MAP-E
1
0%
DS-LITE
12
4%
MAP-T
1
0%
lw4o6
1
0%
Dual-stack (public IPv4 + GUA)
251
72%
464XLAT
6RD
6to4
CGN (dual-stack with private IPv4 + GUA)
DS-LITE
Dual-stack (public IPv4 + GUA)
lw4o6
MAP-T
NAT64
Other
Tunnel Broker
Softwires (L2TP)
MAP-E

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?

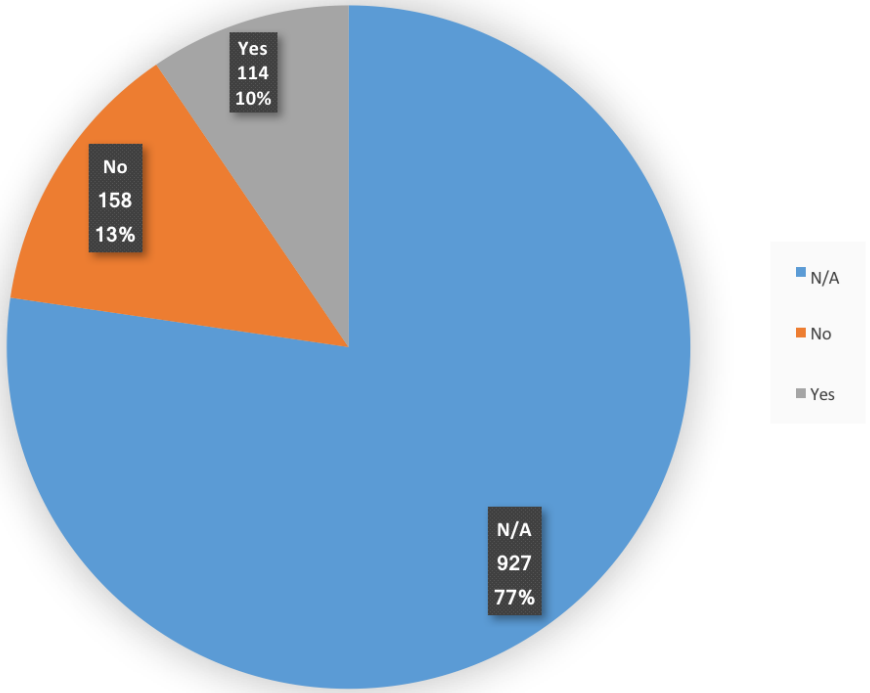

NS Delegation for stable IPv6 prefix?

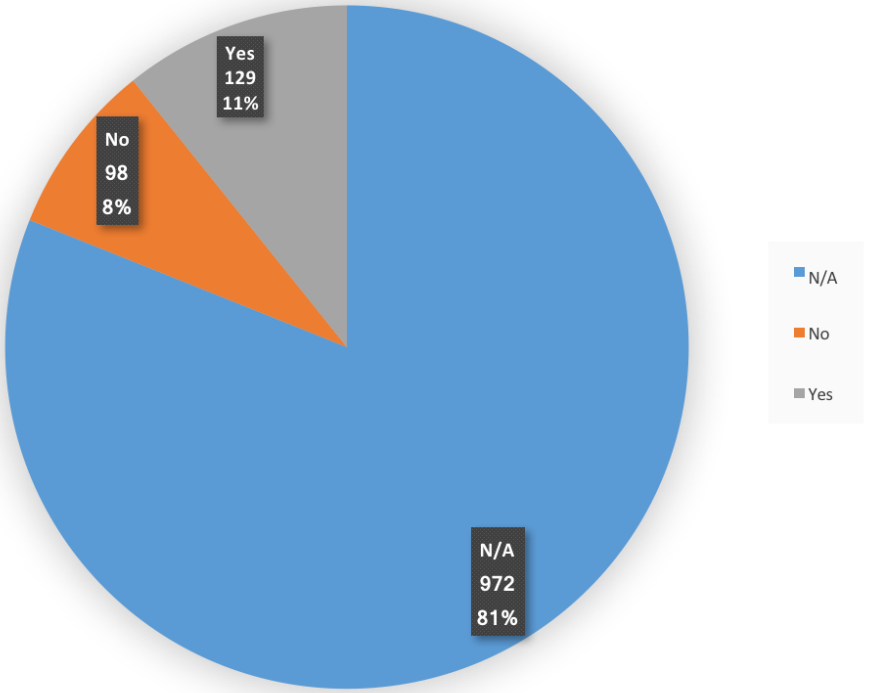

DNAME for non-stable IPv6 prefix for PTRs?

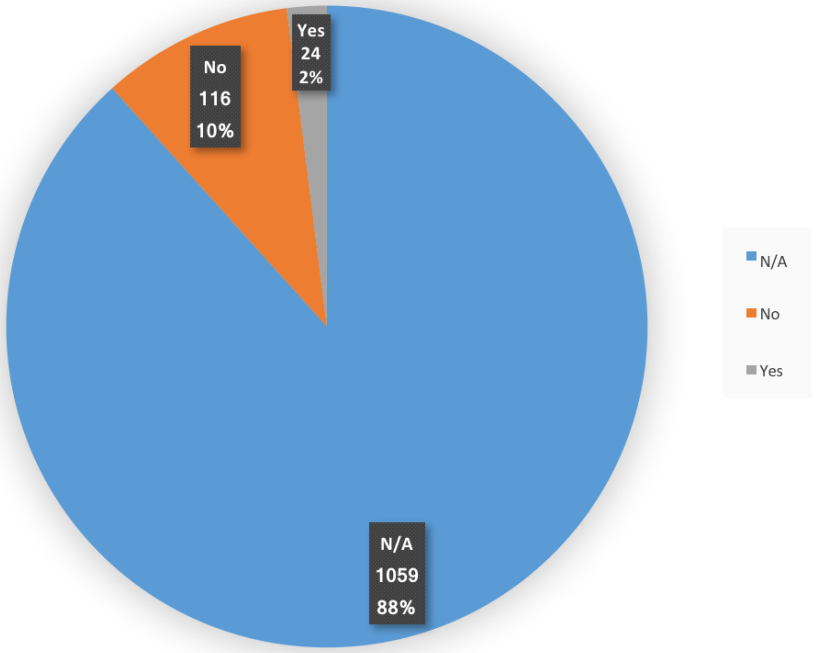

DNS

- 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 (Consulintel): jordi.palet@consulintel.es**