Campus Networking In The Post IPv4 Exhaustion World

Alan Whinery - U. Hawaii System

I2 Tech Exchange San Francisco

October 17, 2017

(Join SSID: IPv6 Solutions, find something it won’t do)

Join The “IPv6 Solutions” SSID

- This is an “IPv6 Only” network
  - Meaning that clients only have IPv6 connectivity to the Internet
  - Even though they have local-wire IPv4 addresses, any IPv4 that appears on the wire is translated into IPv6 and sent to a NAT64
- Running DNS64/464XLAT
  - A means of providing IPv4 As A Service
- It’s meant to demonstrate IPv6 actually solving a real-world problem
  - Bet you didn’t see *that* coming.

A Brave New World (and not a little Huxley-esque)

- IPv4 Address Supply From RIRs Has Run Out
  - You can be on a long waiting list
  - You can seek to buy from someone who has too many
  - Even if you can buy addresses, doesn't mean you can use them
- There are industries for which revenue depends on the perpetuation of IPv4
  - IPv4 brokers
  - Carrier Grade NAT vendors
- BYOD and IoT (BYOT) are still driving growth of your address corpus
  - Corpus = the addresses you need to have alive at a given moment
- In 2017, the North American IPv6 Summit consensus seemed to be
  - That the idea of IPv6-*only* networks is coming of age
  - That going forward, working solutions are an important goal

Your Options

- Buy more IPv4 space
  - Will undoubtedly appreciate for a few years, then depreciate
  - Exhaust/buy more
  - Remember when a /16 was an inexhaustible pool?
- Build bigger and bigger v4 NAT networks
  - support growth in NAT going forward
  - On the pro side, NAT makes more efficient use of “real” addresses than per-device assignment
    - On the order of 64,000 active flows per NAT-pool address
    - Versus per device which initiates about 0.2 to 10 flows per second, keeps active flows on the order of dozens at a time, at most
- Build IPv6-Only Networks, provide IPv4 As A Service

IPv4 Only, NAT and CGN

- When you do “regular” NAT or carrier-grade NAT, you assign clients non-global (usually RFC 1918, and/or RFC6598) addresses
- You place “real” global addresses in the NAT Pool, where they can handle about 64,000 each
- As your number of devices grows, your NAT pool gets more use, and edges toward needing more “real” addresses
- There’s also no exit strategy for the Big NAT scenario.
- With CGN specifically
  - End-user functionality, user experience are negatively affected
  - Identity of users gets confusing

IPv6 Only, IPv4 As A Service

- You assign real IPv6 addresses to clients
- You assign RFC1918 addresses to clients *
  - Any local v4 traffic is translated on wire to IPv6 and sent to NAT64
  - (Or alternatively, simply NAT them)
  - It’s looking like an increasing number of devices can exclude this step.
- You place “real” IPv4 addresses in a NAT64 IPv4 pool
  - The load on these addresses increases with number of devices
  - BUT it also declines as more content becomes available in IPv6
- The way to mitigate address requirement growth is to use IPv6 more

IPv6 Solutions SSID

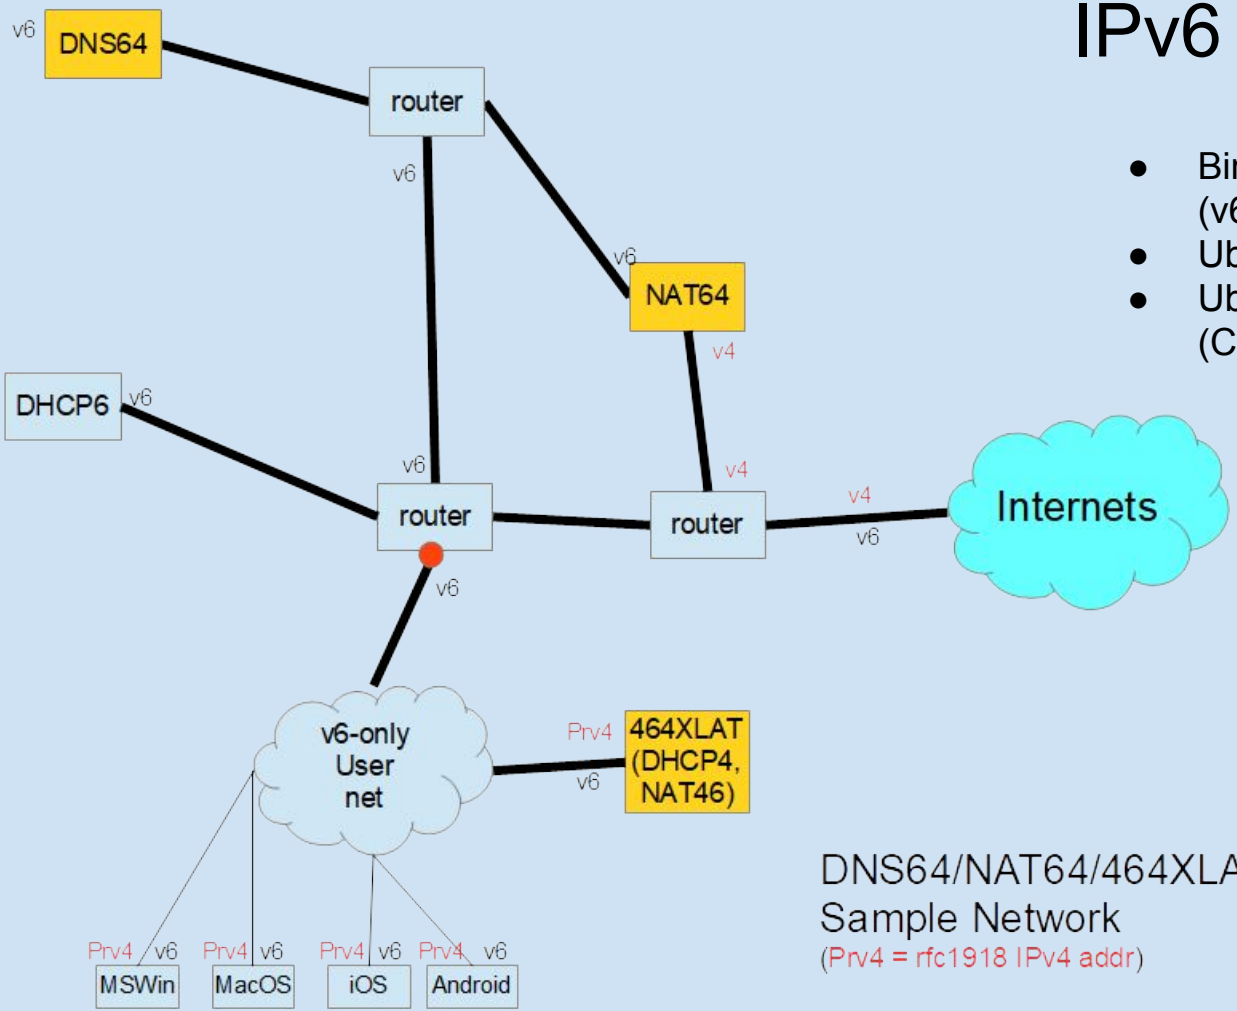

- Bind9 DNS64 forwarder (v6-only)
- Ubuntu16/Jool 3.5.4 NAT64
- Ubuntu16/Jool3.5.4 SIIT (CLAT)

DNS64/NAT64/464XLAT
Sample Network
(Prv4 = rfc1918 IPv4 addr)

IPv6 Solutions

- So-named because it was configured for the “IPv6 Solutions Tutorials” last Sunday
- The idea is that -- what if -- after all this time -- IPv6 actually started solving problems?
  - Too crazy to contemplate
- Seriously, this DNS64/NAT64/464XLAT conglomeration is complicated
  - Yet not at all inter-twined, very modular.
  - DNS64 drives traffic to NAT64
  - You could easily add more redundant DNS64 and/or NAT64 without a rebuild
  - You could turn it all off and outsource to a regional PoP or a IPv4AAS vendor
    - Without changing clients or client addresses
  -

NAT64: NAT That Shrinks, Rather Than NAT That Grows

- Better than NAT44 for various reasons:
  - Begins as an enhancement to a v6 network
  - The more IPv6 grows, the less NAT you have to do.
  - Promoting Native IPv6 reduces use of NAT v4 address/port space
- Better than dual-stack because the more you shift to IPv6 only networking, the less you spend precious IPv4 addresses on building an IPv4 network
- As IPv4 addresses become more expensive and more scarce, providing native IPv4 addresses on per-interface assignments becomes unthinkable.

Stateful NAT64 Manufacturer Support

- Cisco In IOS-XE on certain router platforms
- Cisco ASA (incl DNS64)
- Juniper routers (MX)
- Juniper Security Appliance (SRX)
- Microsoft Forefront Unified Access Gateway (incl DNS64)
- Palo Alto IPS/Firewalls
- A10 Networks (incl DNS64)
- F5
- This list is incomplete

Linux Router NAT64

- Jool (jool.mx)
  - Project of Nic.mx and Tecnológico de Monterrey
  - 2 linux kernel modules (1 for stateful, 1 for stateless)
    - Which kidnap packets fromt the netfilter chain and put them back later.
  - Userspace programs
    - Management for the modules
    - Joold for inter-NAT64 communications
      - Supports redundancy and clustering
  - Logs translations for tracking/accounting
- Is the NAT64 and SIIT engine for the “IPv6 Solutions” SSID in the room today
- Both boxes running at nominally 0% CPU load.

Other Interesting Linux Routers With NAT64

- Linux Routers:
  - VPP (https://wiki.fd.io/view/VPP/NAT#Stateful_NAT64)
  - Any Linux box with routing software

DNS64

- Enables v6-only connected clients to access v4-only resources through a NAT64
- By returning v6 addresses in AAAA answers, with the NAT64’s translation prefix prepended on the target’s IPv4 address.
- Tends to push more flows toward IPv6

DNS64

- Comes as a feature in BIND, Infoblox, PowerDNS, Microsoft, others
  - Also some DNS64/NAT64 Combos shown on Manufacturer slide
- Google Public DNS64
  - https://developers.google.com/speed/public-dns/docs/dns64

464XLAT CLAT via SIIT (AKA Stateless NAT64)

- Picks up stray literals and/or v4 resolver preferences
- Many common OSes do this without the one on the wire
- Another use-case to use SIIT is to make IPv4-only services appear as if they’re IPv6 connected
    - Probably better in most cases to just install IPv6 on the server, but there are those edge cases

http://jool.mx/en/intro-xlat.html

What's Available Over IPv6

CNN.com*, Netflix.com*, YouTube.com*, FoxNews.com, Aljazeera.com
Google.com (perhaps not with external auth), Yahoo.com (incl email)
Wikipedia.org. Xkcd.com, www.acm.org
arin.net, ripe.net, apnic.net, lacnic.net, internet2.edu, aarnet.edu.au, nsf.gov, isoc.org, csiro.au, es.net, ietf.org, nanog.org,
apple.com, microsoft.com, cleveland.com
applebees.com (but not locator map)
Juniper.com (with search), Cisco.com (no search), Brocade.com (no search), www.a10networks.com, paloaltonetworks.com, jool.mx
(includes streaming)*
Also: The Woodwork: Akamai, Limelight, Cloudflare, Limelight AWS, GCP

Who’s Accessing Content With IPv6

- Various Internet2 Members, research organizations
- Comcast, Spectrum (TW,BH), Cox, Google Fiber
- T-Mobile, Verizon Wireless, AT&T Wireless
- Various large corporate networks

U. Hawaii Manoa Wireless
(dual stack)

Flows: v6 ~ 30%

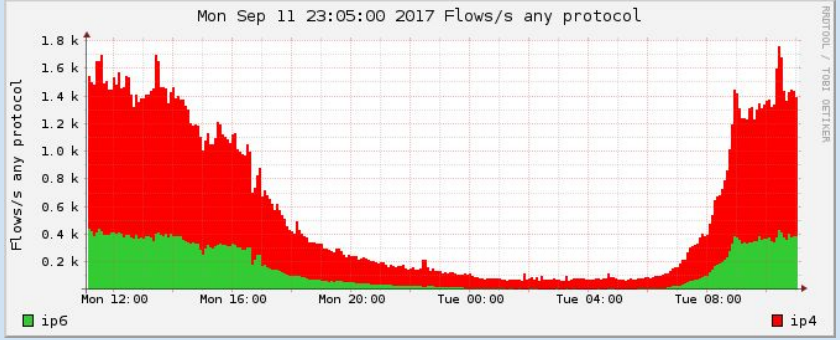

Packets: v6 ~ 60%

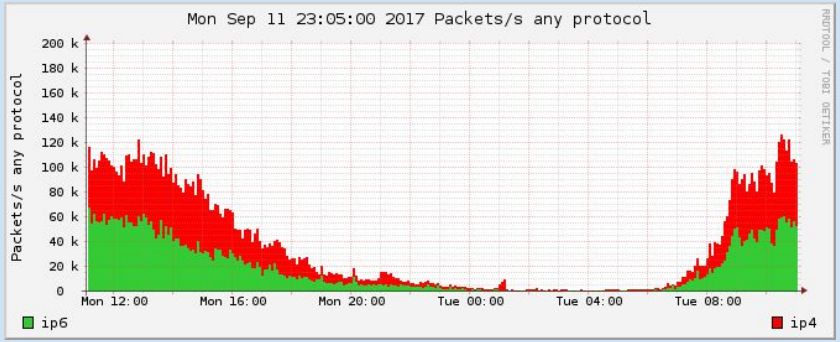

Bits: v6 ~ 50%

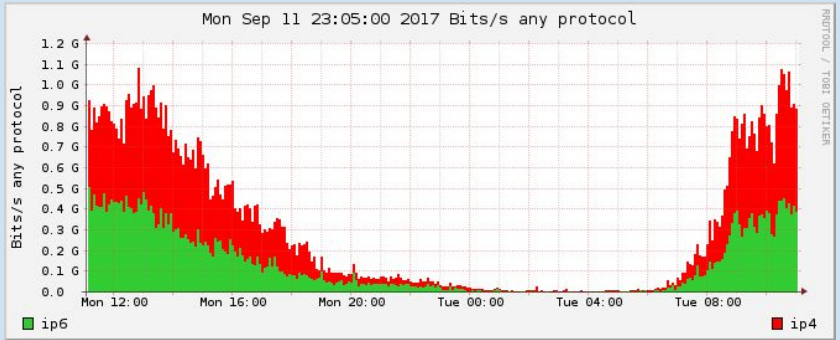

But Wait, There's More

- ARIN constituents who intend to employ IPv6 transition technologies can (pretty easily) receive an IPv4 block, between /28 and /24 inclusive, to use for that purpose, if they can demonstrate need based on prior assigments use %
- ARIN set aside an IPv4 /10 for this, described in the Number Resource Policy Manual (NRPM) coincidentally-numbered section 4.10
- https://www.arin.net/policy/nrpm.html#four10
- In NAT64, 256 addresses would support (on the order of) 16 million NAT mappings

Accountability and CGN

Secure | https://www.ispreview.co.uk/index.php/2017/10/europol-calls-internet-providers-end-cgnat-ip-address-sharing.html

Home | Articles | ISP List | Reviews | Top 10 | Forum | Speedtest | Broadband | Complaints | Contact

Home » ISP News »

Europol Calls on Internet Providers to End CGNAT IP Address Sharing

Tuesday, October 17th, 2017 (1:55 pm) - Score 212

Email | Link News 7 Comments

Europol, which helps the 28 member states of the EU (inc. UK) to fight serious international crime and terrorism, has called on broadband and mobile providers to end the use of Carrier Grade NAT (CGN) in order to "*increase accountability online*" and stop people "*sharing the same IP address as a criminal.*"

Generally everybody needs an Internet Protocol (IP) address to go online and your ISP is responsible for assigning one to your connection (it's the internet equivalent of a phone number). Most fixed line ISPs tend to use **Dynamic IP** addresses for domestic connectivity, which changes each time your broadband link is disconnected and isn't

Latest UK ISP News

- Genesis and NGA Sign Deal to Deploy UK Copper mBond Broadband Tech
- Europol Calls on Internet Providers to End CGNAT IP Address Sharing
- Ofcom CEO's "Key Test" for Openreach is a Fibre Co-Investment Deal with Rival
- Fixed Wireless Broadband ISP Connexin Boosted by £10m Investment
- Superfast Essex Project Opens Community Wi-Fi Scheme to All Suppliers

RSS | Twitter | News Archive

Promotion

*link

Single layer NAT, IPv6 is accountable

- If you run single-stage NAT - not CGN, you can still fully account for user identity and track DMCA, subpoena issues, if you run CGN, less so.
- Jeff Harrington’s IPv6 Security Tutorial talk spoke to doing accounting in an IPv6 network
  - NDP tables
  - SNMP
- This is a commonly misunderstood, undersolved problem for many people, but it appears to be do-able.

Outsourcing IPv4 As A Service

- Can’t endorse, but worth mentioning --
- Retevia.com
  - **IPv4 as a Service**
  - “Our flagship offering is IPv4-as-a-Service (v4aaS). You can provide IPv6 to your users, and send all of their legacy traffic to us over IPv6, using NAT64, MAP-T, MAP-E, Dual-Stack Lite, or 464xlat. Your network can streamline toward single stack IPv6 without losing access to legacy IPv4 content.”
- Essentially, you build a v6 network and let the vendor deal with IPv4.
- This has several attractive components, among which is that when someone blacklists your NAT pool address, it’s the vendor’s problem

IPv6 Adoption: The Landscape Changes, While The Song Remains The Same

- Big Content has really stepped up. (Google, YouTube, Wikipedia, Netflix, Facebook, Akamai, Cloudflare, Limelight, etc.)
- Numerous large residential providers are providing dual-stack to homes
- Mobile carriers:
  - T-mobile: IPv6-only especially since iOS 10.3 (http://www.rmv6tf.org/wp-content/uploads/2017/04/04-IPv6-NAv6TF-Langerholm-1.pdf)
  - Verizon: dual-stack
  - AT&T: Weird proxy matrix, addresses look like: 2600:387:4:804::9
- Most journalism about IPv6 transition is parroting something they were told 5 or 10 (or 15) years ago, or they're asking your boss, who's parroting something he/she heard 10 years ago
- Much IPv6 writing is out-of-date, or referring to things that are out-of-date

IPv6 Adoption | **Per-Country IPv6 adoption**

https://www.google.com/intl/en/ipv6/statistics.html

Per-Country IPv6 adoption

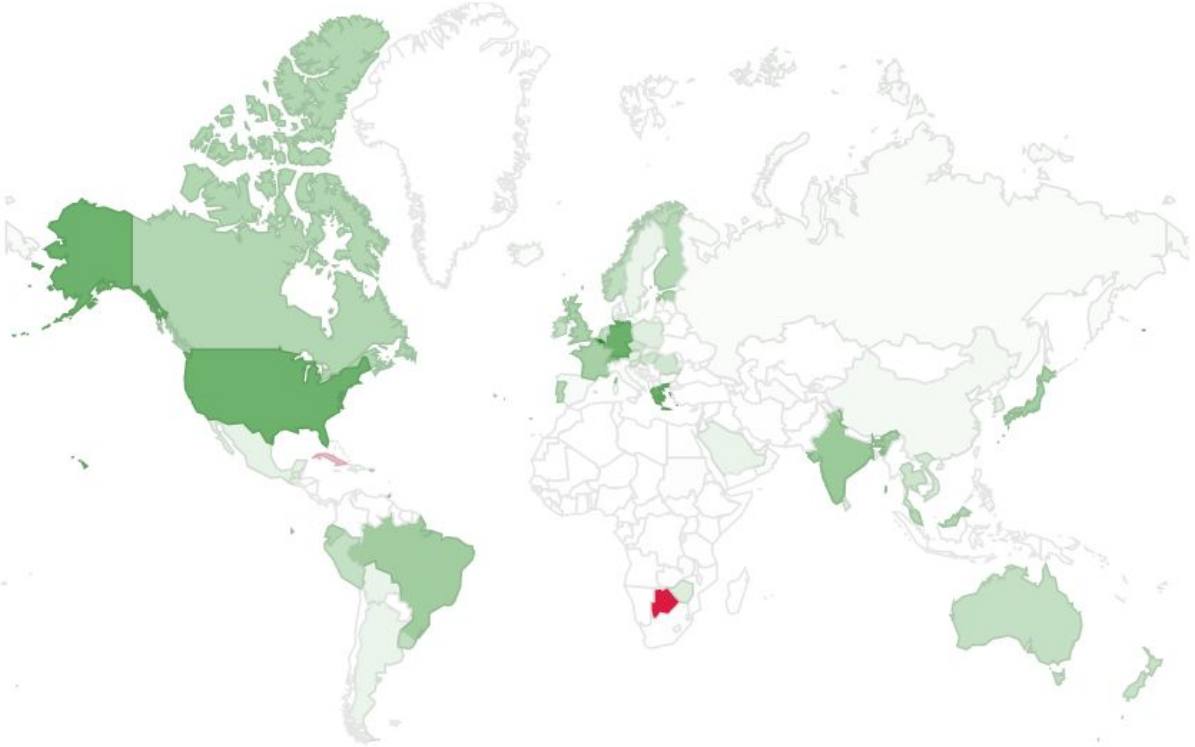

World | Africa | Asia | Europe | Oceania | North America | Central America | Caribbean | South America

The chart above shows the availability of IPv6 connectivity around the world.

IPv6 Connectivity Around The World

- Looking at a 3 months of IPv6 flows at U. Hawaii
  - The IPv6 addresses UH interacted with during that time
- The most frequent countries with which we interacted were predictable

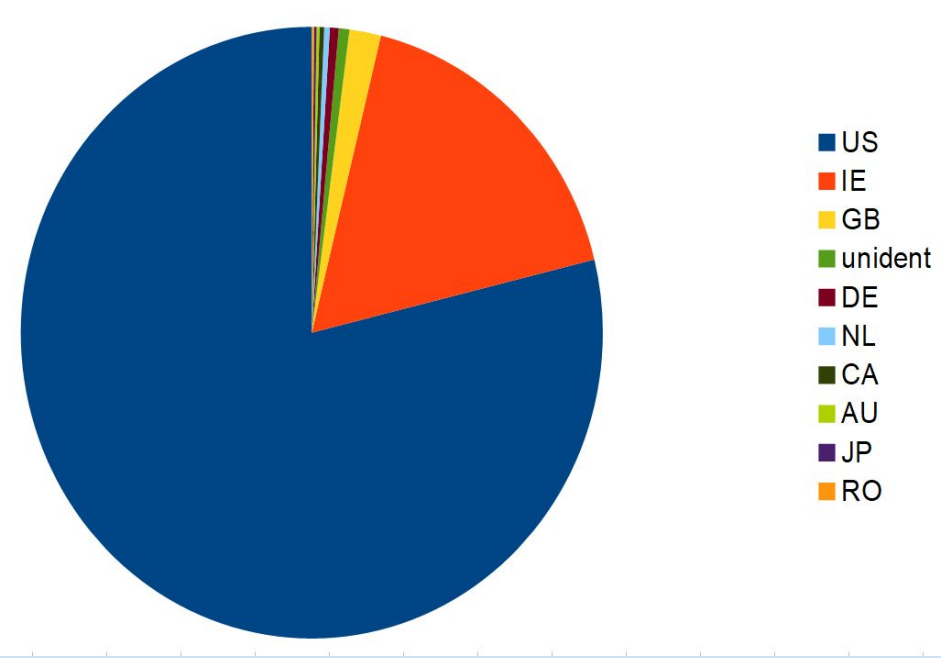

LATENT IPv6 Connectivity Around The World

- But in the long tail, there are unexpected revelations --
  - Example1: Country code BB (Barbados)
  - “0% adoption”, according to Google & Akamai’s IPv6 adoption pages
  - 1 hit out of 146 M recorded
- One of our hosts interacted with (apparently) a Facebook cache in Barbados on IPv6
  - ARIN whois, traceroute output, MaxMind GeoIP agree that this address is in Barbados
  - Below from Akamai: https://www.akamai.com/us/en/about/our-thinking/state-of-the-internet-report/
  - Eric Vyncke’s site shows a more complete picture
    - https://www.vyncke.org/ipv6status/detailed.php?country=bb

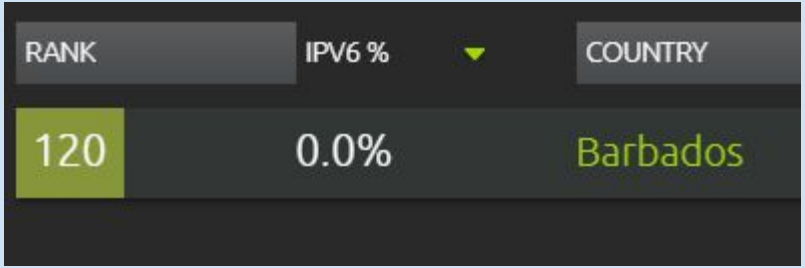

LATENT IPv6 Connectivity Around The World

(other places with an IPv6 address)

- Likewise (infinitesimal “adoption %”, but connected, nonetheless):
  - Guernsey
  - Brunei
  - Bonaire
  - Marshall Islands
  - Guam
  - Aruba
  - Gabon
  - Seychelles
  - Equatorial Guinea (Malabo, although perhaps not continental)
  - Jamaica
  - Senegal
  - Cuba
  - Botswana
  - Many others I haven’t checked

Point is

- IPv6 connectivity exists to many places that have “0%” adoption
- If we start a conversation about how to use IPv6 to solve real problems, it incentivizes further deployment.
- If IPv6 is a solution to problems, a lot of map can get greener in a short time

IPv6 Measurements

http://www.potaroo.net/ipv6/

http://www.worldipv6launch.org/measurements/

https://usgv6-deploymon.antd.nist.gov/govmon.html

https://www.akamai.com/us/en/about/our-thinking/state-of-the-internet-report/state-of-the-internet-ipv6-adoption-visualization.jsp

https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption&tab=per-country-ipv6-adoption

https://www.mrp.net/ipv6_survey/

End