Dual Stacking Broadband Subscribers

Lessons Learned

AS5607 Overview

- Primarily residential broadband
- ~6 million fixed-line subscribers
- On-net (LLU):
  - ~80% ADSL2+
  - PPPoE
- Off-net:
  - ~20% VDSL2
  - IPoE
- Alcatel-Lucent^WNokia 7750SR as BRAS/BNG

AS5607 Overview

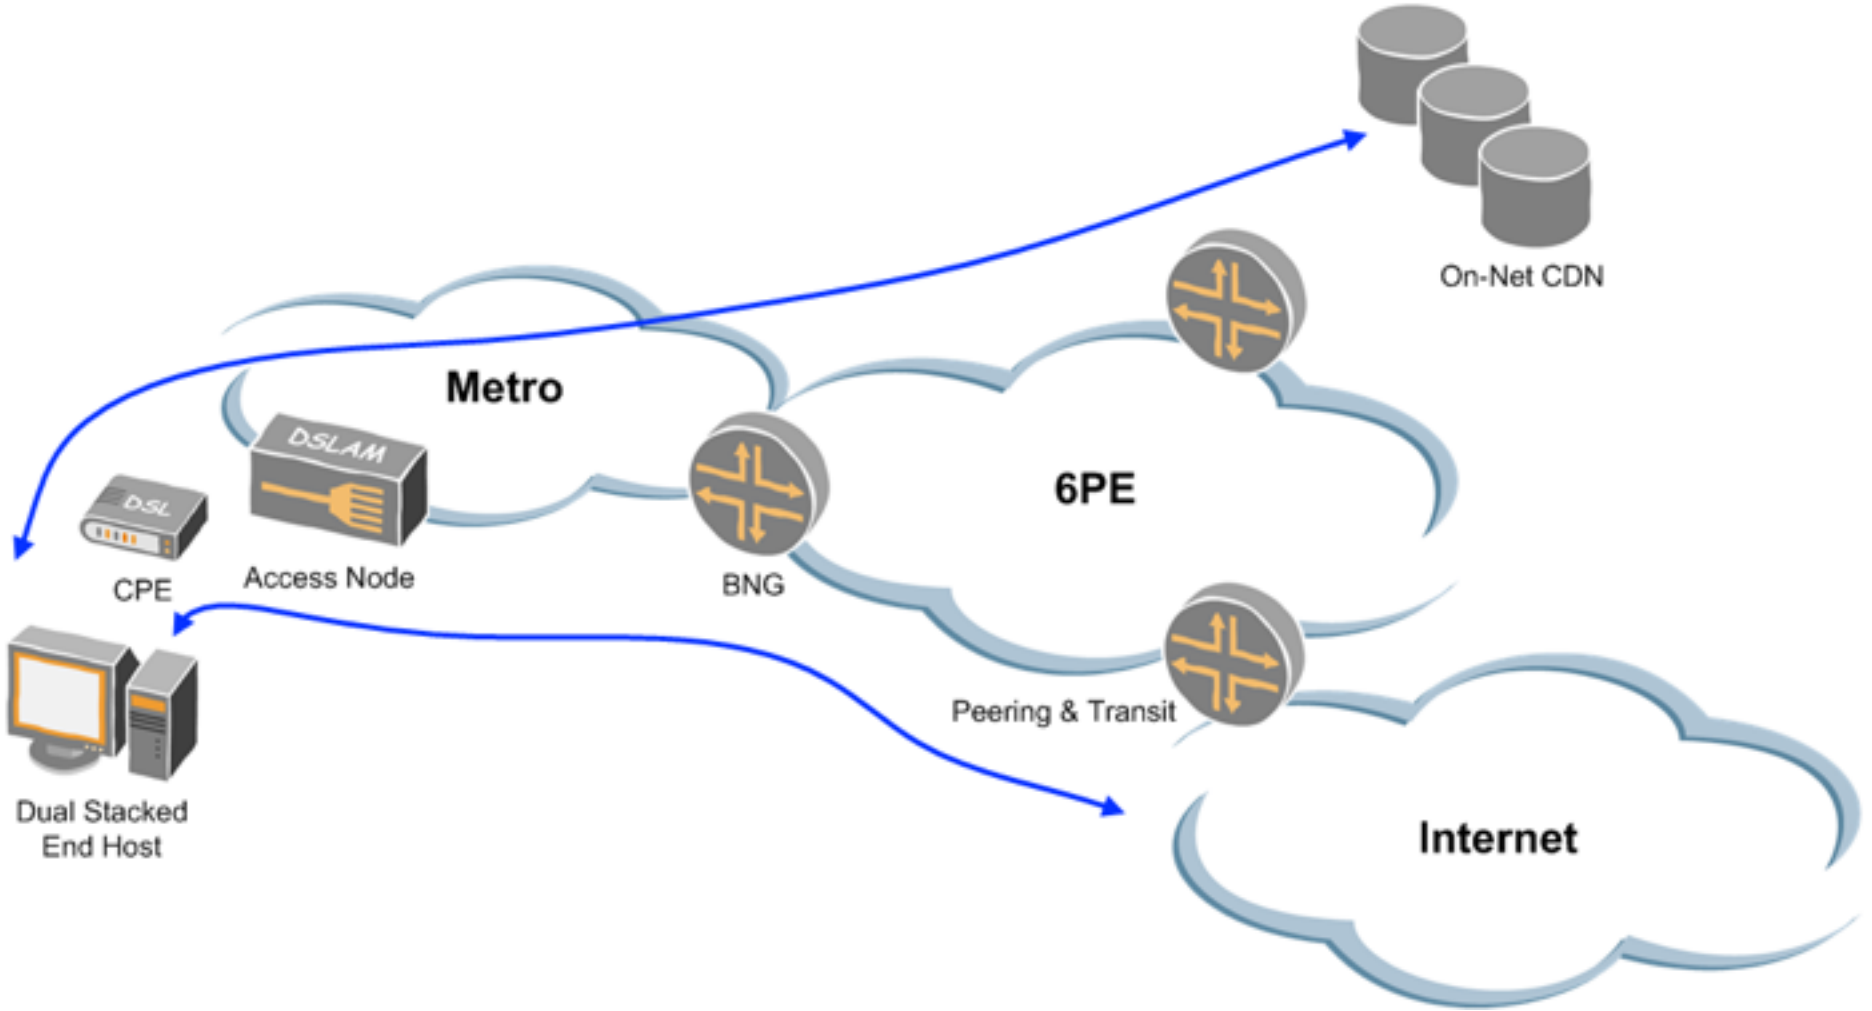

IPv6 Deployment Strategy

- Native dual-stack service
- Link-local WAN addressing only
- /39 per service per BNG
- Single /56 allocated via DHCPv6 PD
  - Dynamic
  - Although “sticky”, unless the client sends a RELEASE or changes DUID
- SLAAC running on CPE for LAN address assignment
  - Using just the first /64 for now
- ULA on LAN side

Enablement

- Pre-requisites:
  - All BNGs capable
  - P&T, CDN dual stacked for comparable v4 routes
  - CPEs upgraded to v6 capable firmware
- Granular, per-user enablement
- Targeted enablement. Not geographic/BNG based.
  - Service Type: LLU (PPP) vs Off-net (IPoE)
  - CPE type/version
  - Ability to exclude misbehaving CPEs/devices.
- RADIUS controlled
  - DHCPv6 pool name returned in RADIUS Access-Accept.

Enablement

- Started soft-launch end of July 2015
- Up to ~1M by September
- Paused for 6 months:

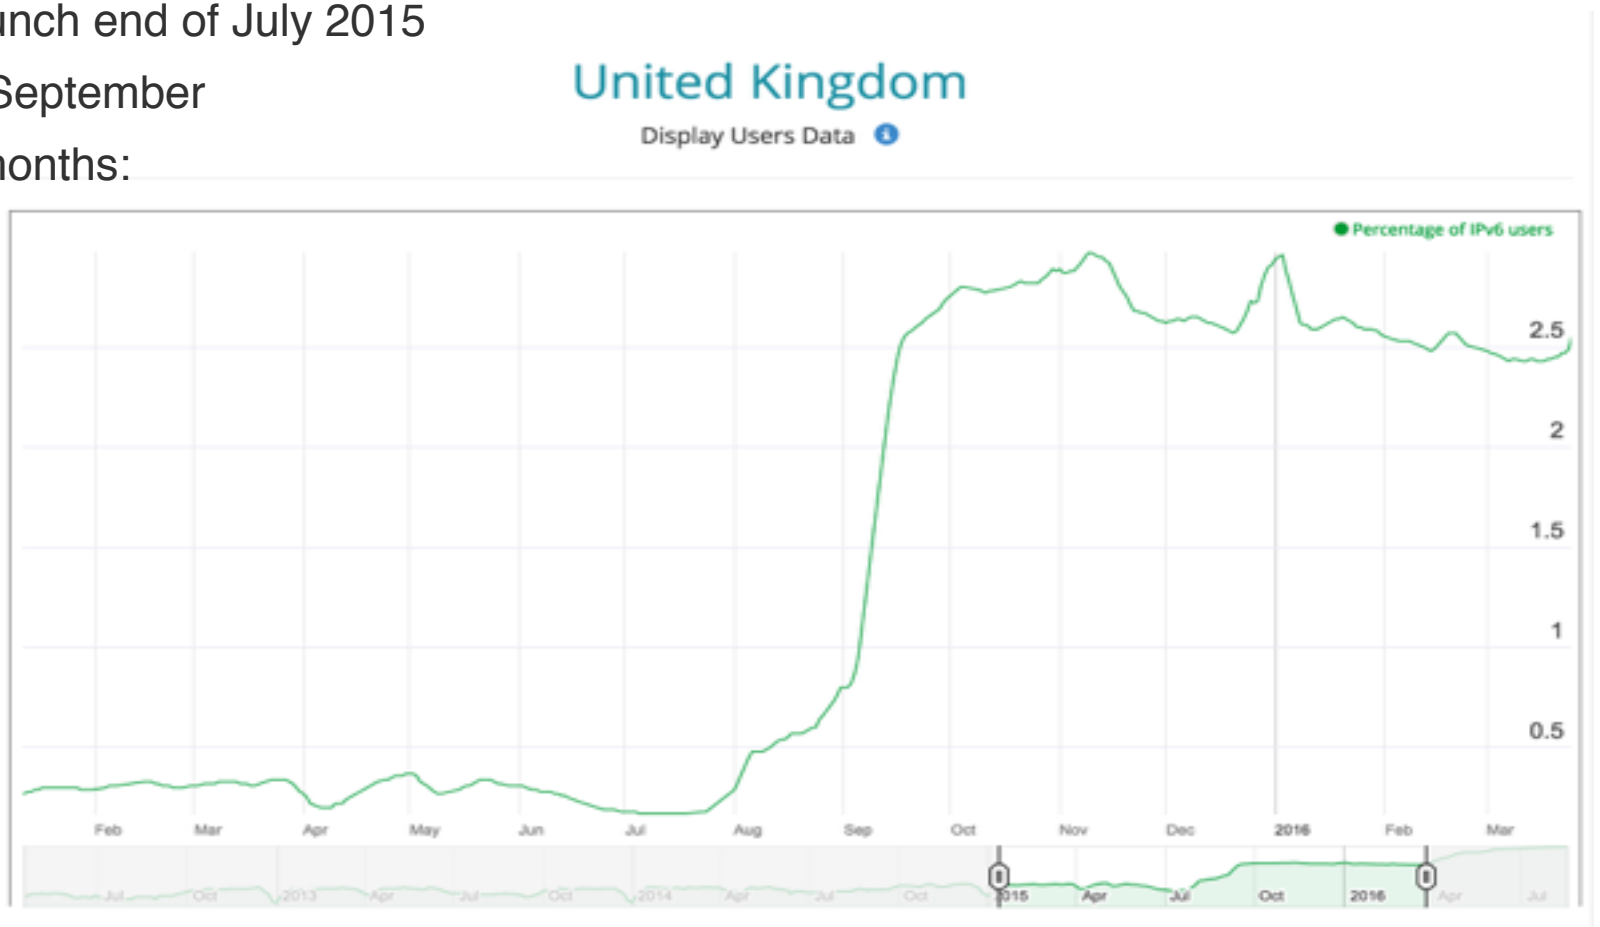

Enablement Paused.
Why?

- Anycast platform load doubled
  - RADIUS Authentication
  - Recursive DNS

- RDNSS
  - Each host lookup now sends A & AAAA queries
  - We expected this, and capacity planned for it
  - Load increase didn't follow our enablement plan linearly
  - AAAA queries were triggered by the CPE firmware upgrade, not IPv6 WAN enablement:

    Windows & OSX will send AAAA queries when assigned a non-link-local address. Eg. ULA

    Android and Chrome will not

Completely Tangential Side-Note

- Don't expect Chrome to work on your private ULA LAN/ Intranet:
- Yes, that's Google's public v6 DNS server hardcoded in as a connectivity check.

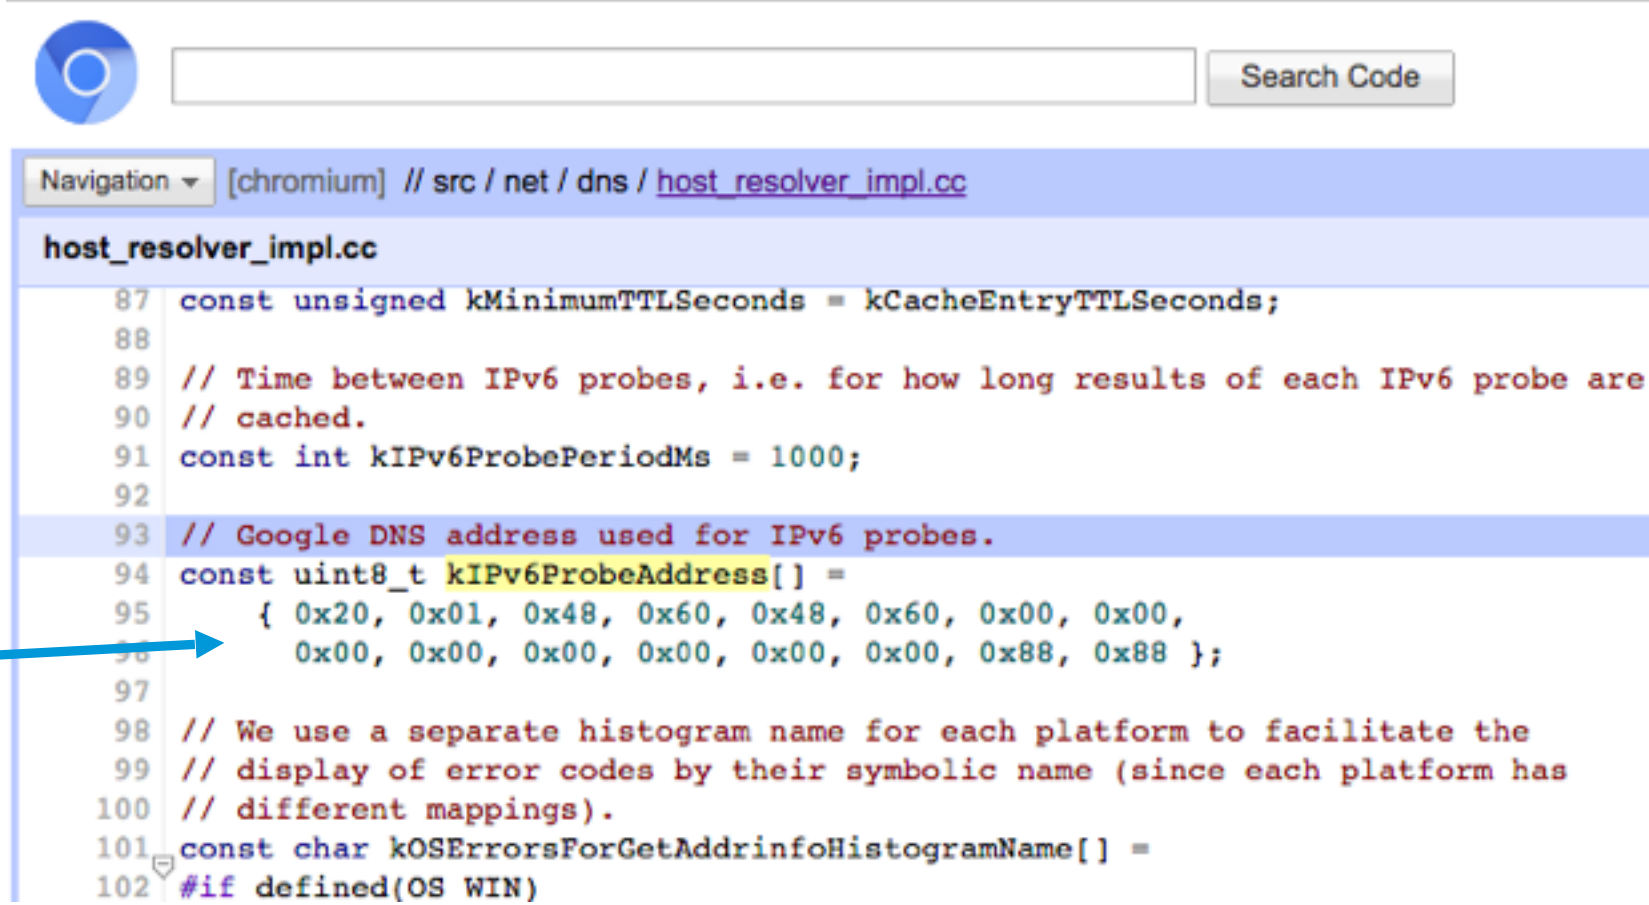

+1 the bug: https://bugs.chromium.org/p/chromium/issues/detail?id=584028

Subscriber Authentication

- PPPoE
  - Authentication based on PADI
  - Remote-ID/Circuit-ID (Broadband Forum TR-101)
  - Ignore CHAP username/password
  - **PPP session comes up before DHCPv6 PD**

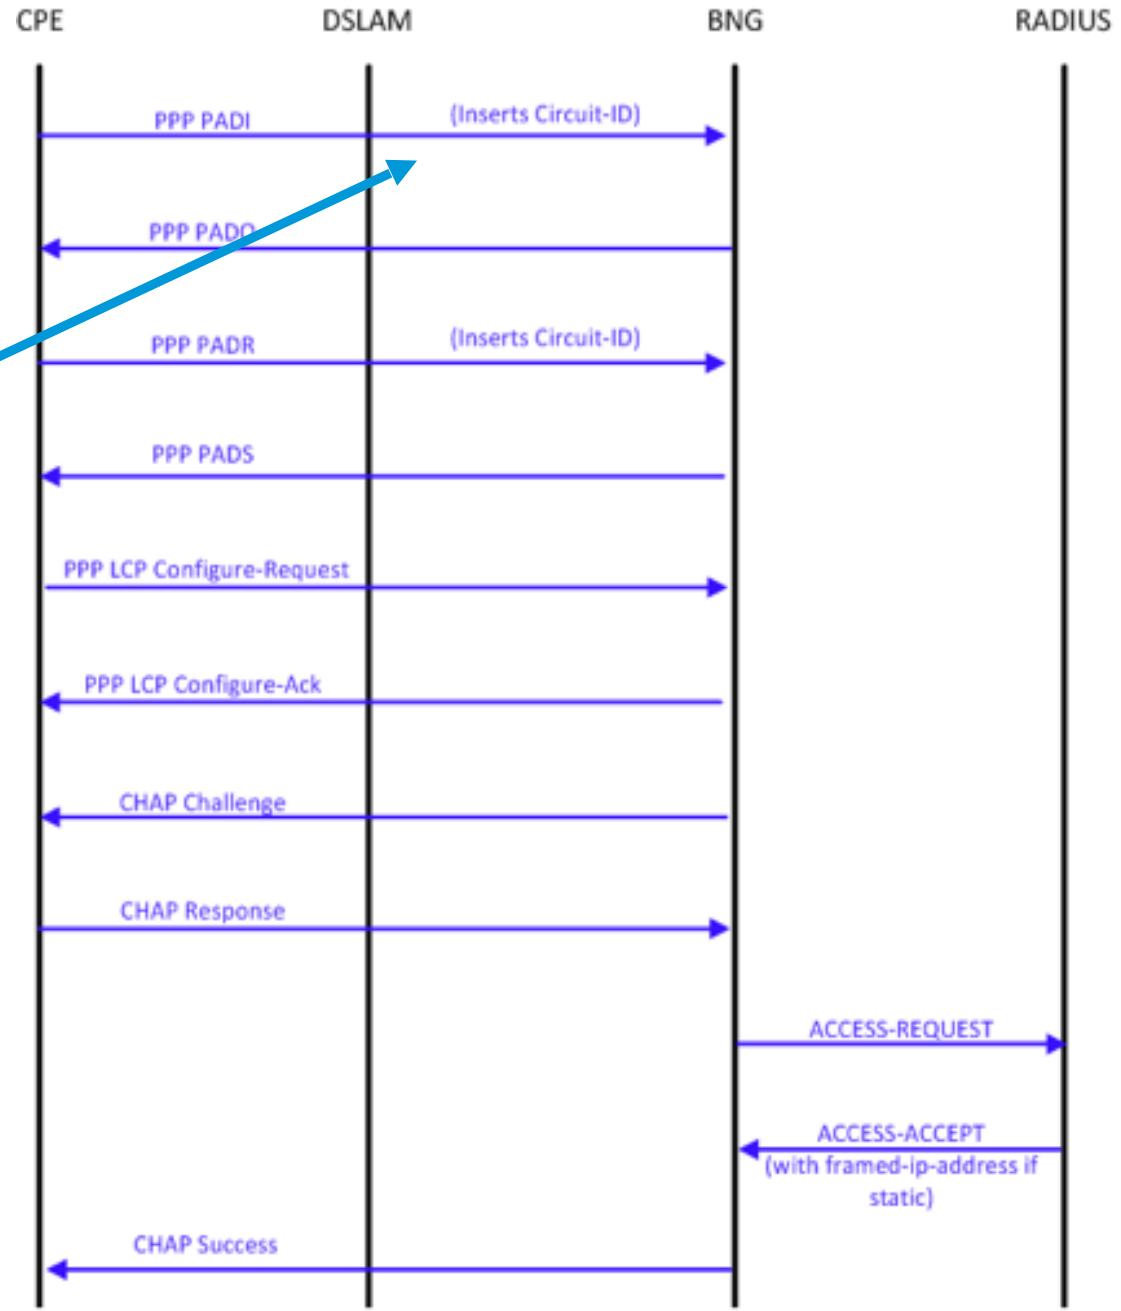

Subscriber Authentication

- IPoE
  - DHCPv4 Option 82.2 (Remote-ID) inserted by the Access Node
  - BT/Openreach do not support the DHCPv6 equivalent on their entire VDSL footprint. (Huawei ISAMs do, but ECI don't)
  - Nokia 7750 SR feature "IPoE Linking"

    Within 10 seconds of v4 session establishment

    Not 100% reliable.
  - Python cache

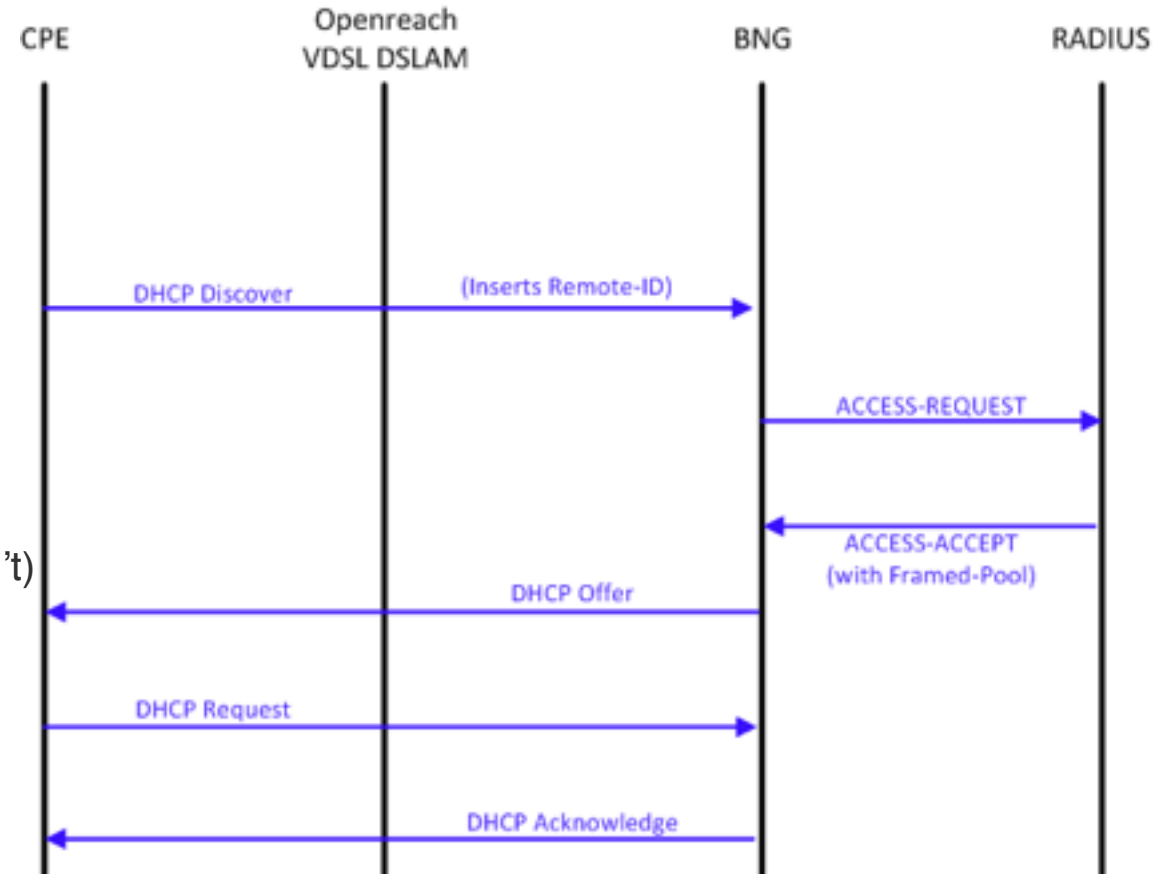

Subscriber Authentication
With DHCPv6 LDRA

- DHCPv4 Option 82.1 Circuit-ID doesn't exist
  - Replaced with Option 18: Interface-ID
- Remote-ID: Option 37
  - Prepended with an extra 4 bytes for Enterprise-ID
  - Gets copied into the same RADIUS Remote-ID attribute
  - Extra logic required in RADIUS to know when an Access-Request has been triggered by DHCPv4 or DHCPv6

3. The Relay Agent Remote-ID Option

This option may be added by DHCPv6 relay agents that terminate switched or permanent circuits and have mechanisms to identify the remote host end of the circuit.

The format of the DHCPv6 Relay Agent Remote-ID option is shown below:

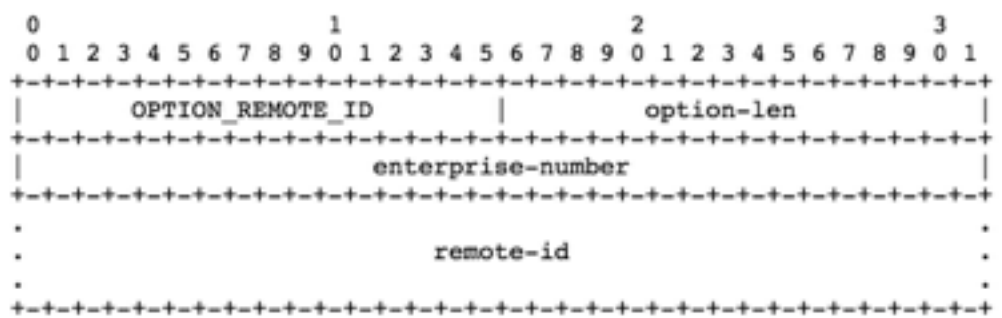

option-code: OPTION_REMOTE_ID (37)

option-len: 4 + the length, in octets, of the remote-id field. The minimum option-len is 5 octets.

CPE

- 8 different CPEs. Majority have been developed in-house.
  - (Both a blessing and a curse)
- DHCPv6 Client Not Backing Off
  - RFC3315 specifies a back-off algorithm when an ADVERTISE **hasn't** been received, but not if an ADVERTISE is received with NoPfxAvailable
  - A lot of extra RADIUS Authentication load
  - Python script to block requests until we were ready to enable v6
- Stateful Firewall or not?
  - Allow unsolicited inbound IPSec, as per RFC6092
  - Lack of IGD:2 support. No UPnP functions for IPv6 firewall manipulation.
- Other devices with lack of IPv6 support
  - Wifi Boosters

CPE etc.

- Firmware upgrade broke 6in4 tunnels. Eg. Hurricane Electric.
- Pfsense wouldn’t send a DHCPv6 SOLICIT until it received an RA.
- Happy Eyeballs can mask potential IPv6 issues.
- Dynamic Prefix Issues
  - Some 3rd party CPEs don’t withdraw stale prefix delegations
  - Static firewall rules

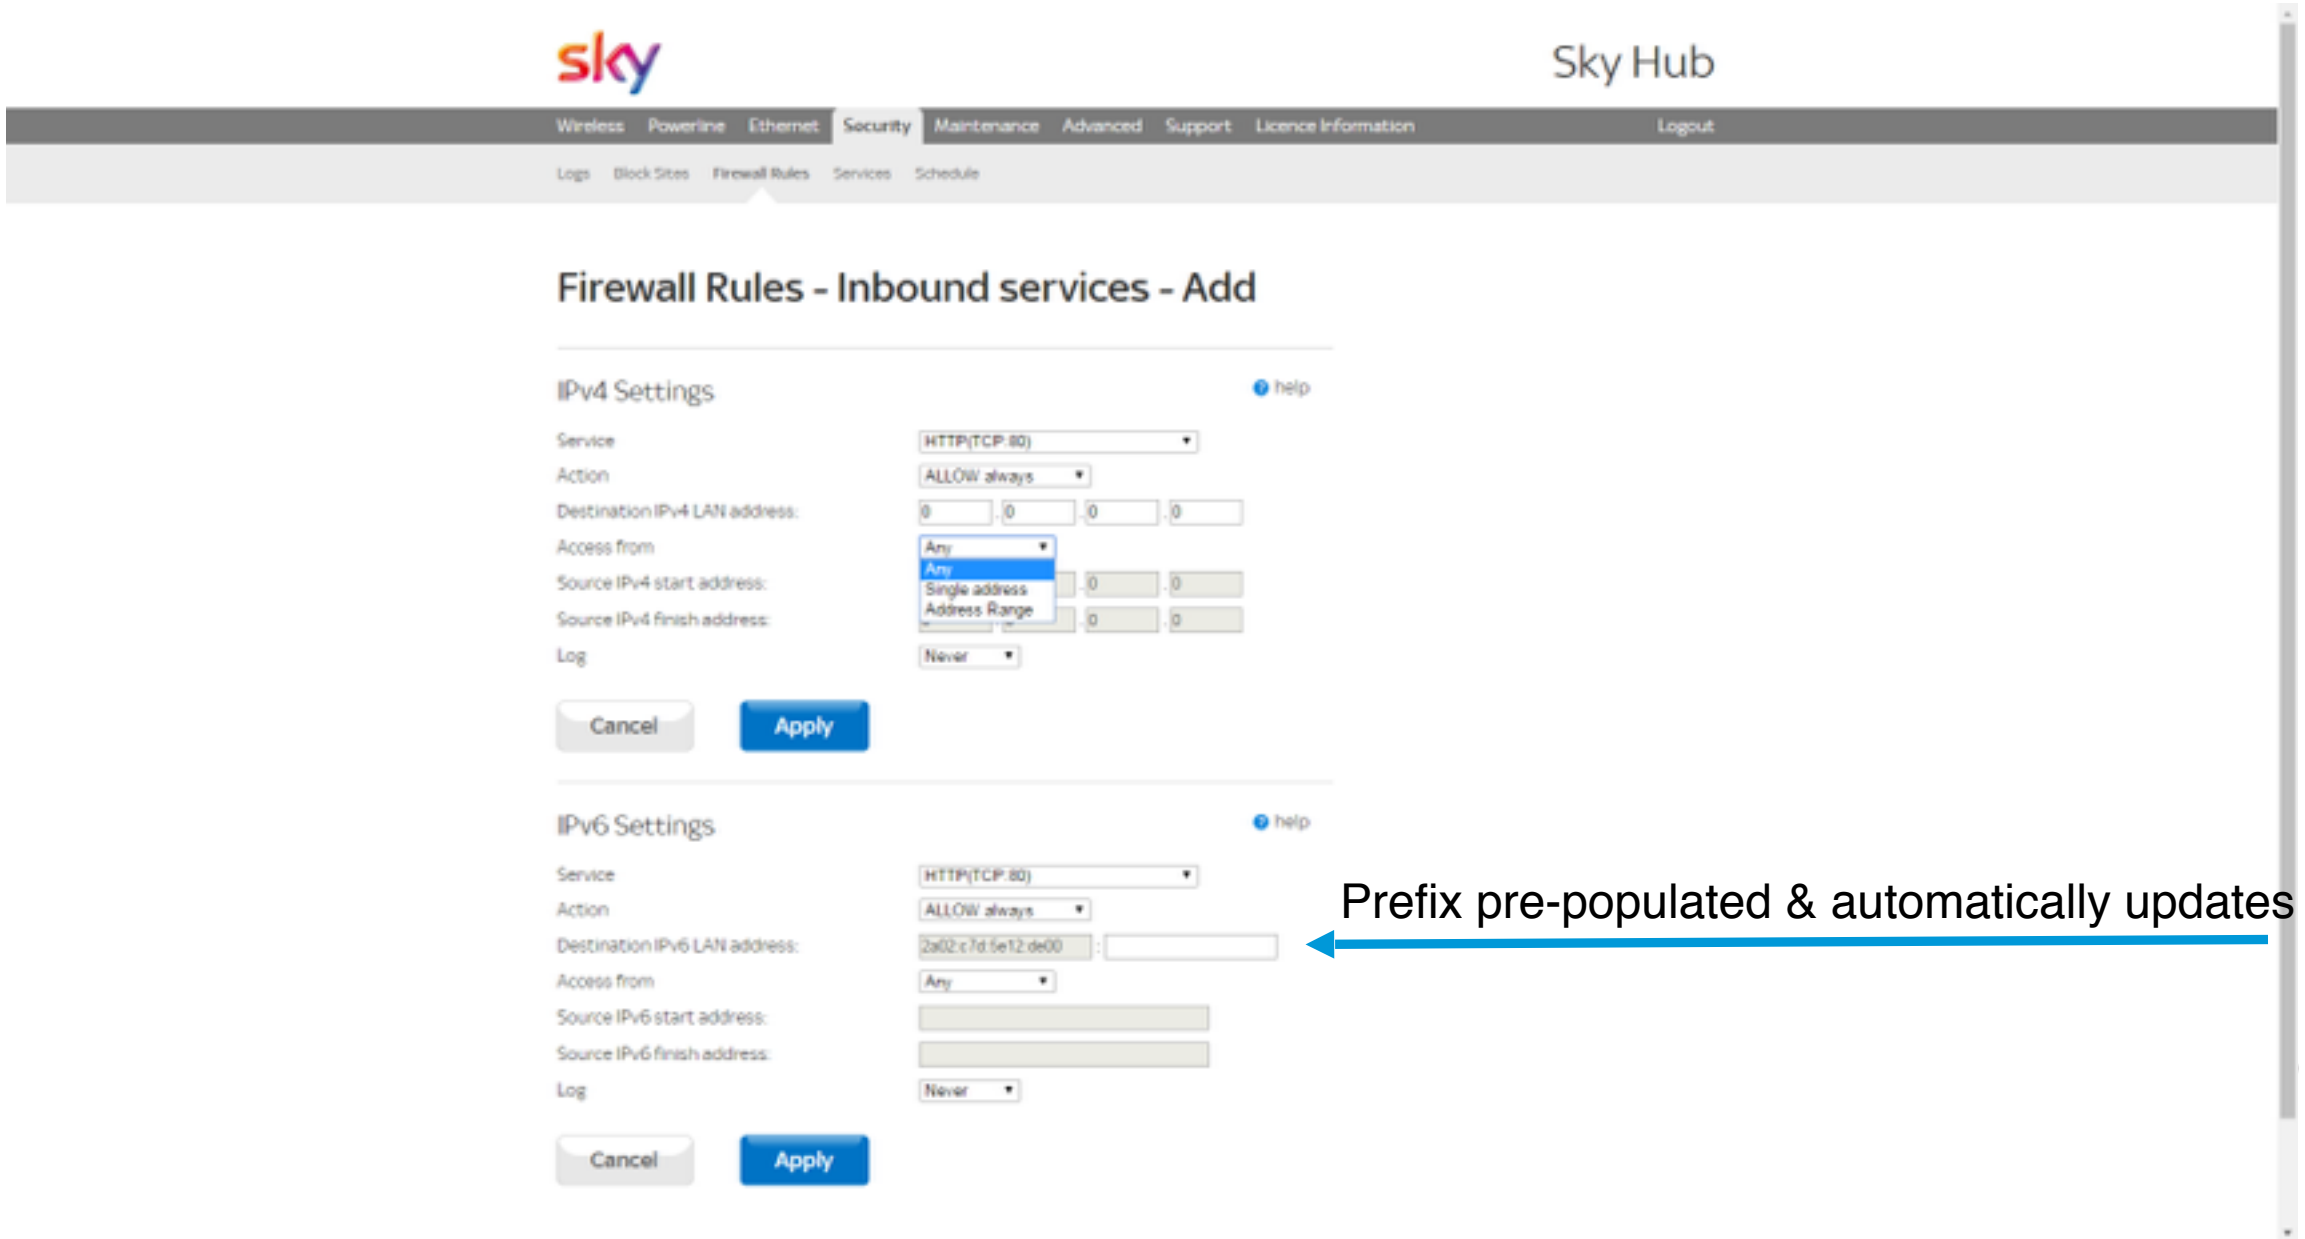
sky
Sky Hub
Wireless
Powerline
Ethernet
Security
Maintenance
Advanced
Support
Licence Information
Logout
Logs
Block Sites
Firewall Rules
Services
Schedule
Firewall Rules - Inbound services - Add
IPv4 Settings
help
Service
HTTP(TCP:80)
Action
ALLOW always
Destination IPv4 LAN address:
Access from
Any
Any
Single address
Address Range
Source IPv4 start address:
Source IPv4 finish address:
Log
Never
Cancel
Apply
IPv6 Settings
help
Service
HTTP(TCP:80)
Action
ALLOW always
Destination IPv6 LAN address:
2a02:c7d:5e12:de00
Access from
Any
Source IPv6 start address:
Source IPv6 finish address:
Log
Never
Cancel
Apply
Prefix pre-populated & automatically updates

CDN etc.

- Google's v6 report showed UK's v6 with +10ms latency
- Google Global Cache switch FDB entries would time out
  - GGC only send unsolicited Neighbor Advertisements
  - Ie. They don't respond to Neighbor Solicits
  - Enabled '**ipv6 nd na glean**'
- Geolocation
  - $CDN reported ::/0 as being in the USA
- Logfiles
  - What happens when a tool expects a v4 dotted decimal address but receives 2001:db8:f00:baa::1?
  - SEGFAULT is not the correct answer

Bonus Anecdotes

- A few customer escalations for broken websites
  - Sites advertising a AAAA, but pointing to a misconfigured server
- $Bank and their F5s
  - F5 & $Bank response: *Sky is unique and the only one affected, it must be your fault.*
  - Silently discard any DNS query for an unknown record, instead of returning NOERROR. Eg. AAAA queries.
  - This is considered by F5 to be ‘best practice’.
  - Our recursors started blacklisting their authoritative servers, for all queries.
  - Quote from our vendor: *“the immediate answer that I got from the engineers contained some swearing and some funny stuff but basically says that what F5 is doing is horrible and is more like a ‘worst practice’ ”*

Bonus Anecdotes

- Finally something good(ish):
  - *6PE was on a different set of route-reflectors from IPv4*

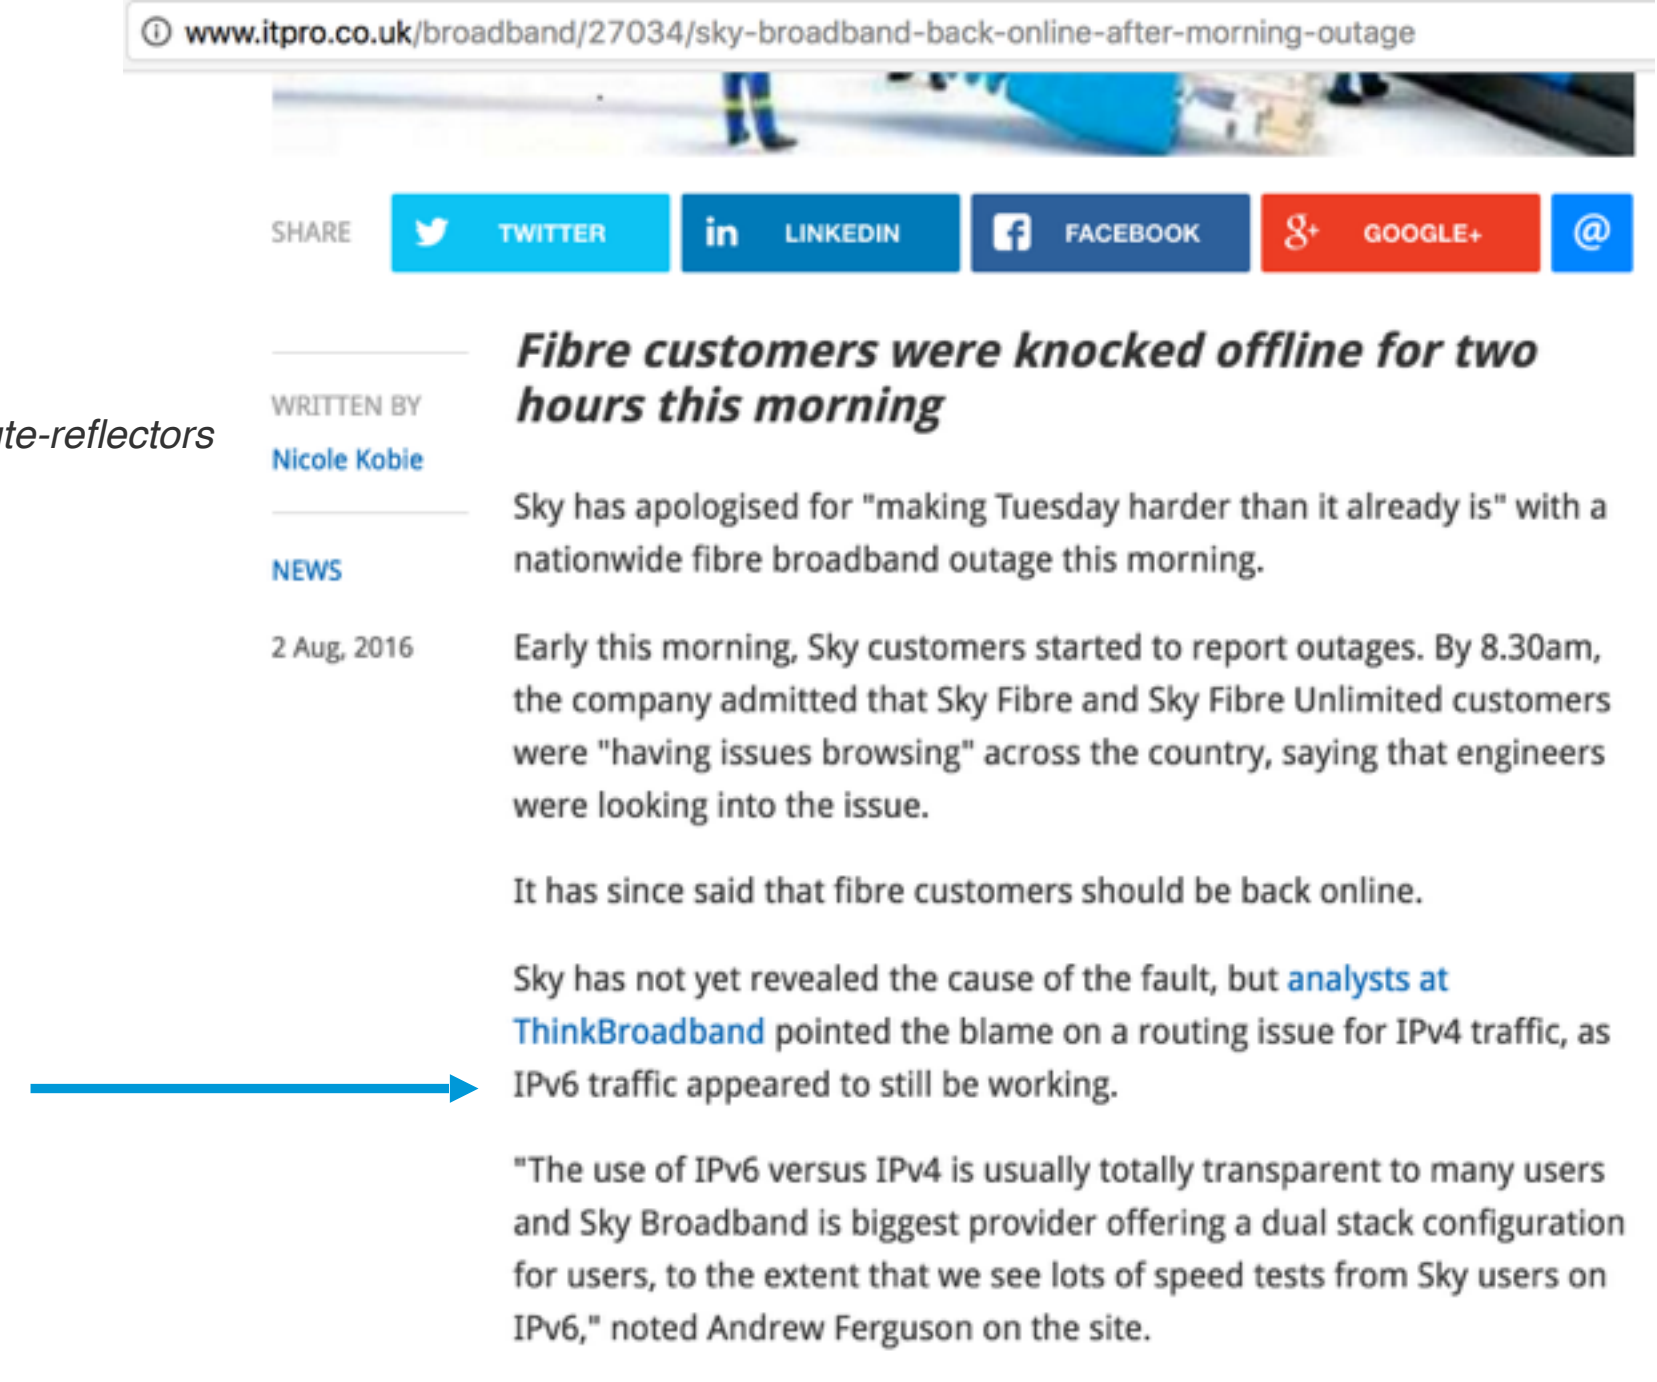

www.itpro.co.uk/broadband/27034/sky-broadband-back-online-after-morning-outage

SHARE TWITTER LINKEDIN FACEBOOK GOOGLE+ @

WRITTEN BY
Nicole Kobie

NEWS

2 Aug, 2016

Fibre customers were knocked offline for two hours this morning

Sky has apologised for "making Tuesday harder than it already is" with a nationwide fibre broadband outage this morning.

Early this morning, Sky customers started to report outages. By 8.30am, the company admitted that Sky Fibre and Sky Fibre Unlimited customers were "having issues browsing" across the country, saying that engineers were looking into the issue.

It has since said that fibre customers should be back online.

Sky has not yet revealed the cause of the fault, but analysts at ThinkBroadband pointed the blame on a routing issue for IPv4 traffic, as IPv6 traffic appeared to still be working.

"The use of IPv6 versus IPv4 is usually totally transparent to many users and Sky Broadband is biggest provider offering a dual stack configuration for users, to the extent that we see lots of speed tests from Sky users on IPv6," noted Andrew Ferguson on the site.

EOF

IPv6 Excuses
@ipv6excuses

Follow

But no one else is deploying #IPv6

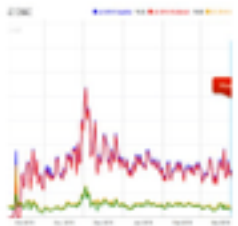

Ars Technica UK @ArsTechnicaUK
Sky now at 80% IPv6 coverage, BT rollout coming early 2017
arstechnica.co.uk/information-te… by @mrseb #ipv6

RETWEETS 4 LIKES 3

4:39 AM - 23 Aug 2016

4 3

Reply to @ipv6excuses