IPv6 Thinking

A quick-and-dirty introduction to IPv6 and its current status

Tom Hill
24th April 2023

Background

I've been working on IPv6 deployments since ~2010

Initially in hosting networks, but I dabbled in access before hitting BT in 2018

A very fresh faced Tom, speaking about IPv6-only data centres at UKNOF29, 2014

In the years since, I've seen and heard every. single. argument., both for and against moving customers and traffic to IPv6

IPv6 101

Some basics to get us all up to speed

Addressing

Hex notation ain’t scary! Easier than counting to 65,536…

128bit address = eight 16bit sections (2 Bytes, “hextets”)

2001:db8::

2001:0db8:0:0:0:0:0:0

2001:0db8:0000:0000:0000:0000:0000:0000

16 : 32 : 48 : 64 : 80 : 96 : 112: 128bits

What’s the smallest address you can think of?

::1 2600::

Addressing

Scopes, and source selection: your devices will have interfaces with many IPv6 addresses, with different scopes!

The source selection algorithm in RFC 6724 determines which one to use for a given connection.

- All IPv6 interfaces use link-local addresses (fe80::) incl. for local resolution
- Some sites use site-local (deprecated) or ULA (‘Unique Local Addressing’)
- ‘Public’ (GUA, or ‘Globally Unique Addressing’) is usually the answer

Important: you do not translate between scopes

There is no IPv6 equivalent of RFC1918!

Addressing

In IPv6, we have two ways for hosts to get addresses and resolvers:

Autoconf/SLAAC

- Used by interfaces for generating link-local addresses
- In response to ‘Router Advertisements’ from local gateways (GUA/ULA)
- Don’t panic about /64 prefixes

DHCPv6

- Everyone’s favourite protocol, reborn
- Can’t really be used without any SLAAC at all
- Often used to provide additional things to hosts (e.g. prefix delegation)

Addressing

On multicast: don’t panic, it’s only layer-2.

NDP is the main use-case (replacing ARP)

- Solicited-node groups keep flooding down & facilitates DAD
- Can lead to some limitations when snooping MLD ‘joins’ – do you need it?
- Some special multicast group use exists in FHRP (e.g. VRRP)
- NDP cache limits/RA guard features are good ideas for all networks
- (Yes, even the ones without IPv6 yet…!)

BT

IPv6? Where??

Dispelling some myths: what's changed?

Deployments

What's changed? We've known about all of this for years!

'You said IPv4 was GONE in 2012!'

'You said it again in 2018!!'

Some outfits might even tell you that there are vast quantities of IPv4 addresses waiting to be repurposed and made available

2^32bit addressing simply will not last for very long.

RIPE tore through an entire /8 in a very short space of time, and that was *with* an intrinsic aim to allocate fairly

Deployments

There isn't any IPv6, right?

Is there anyone using it?

On average, a third of BT's broadband traffic flows, are IPv6.

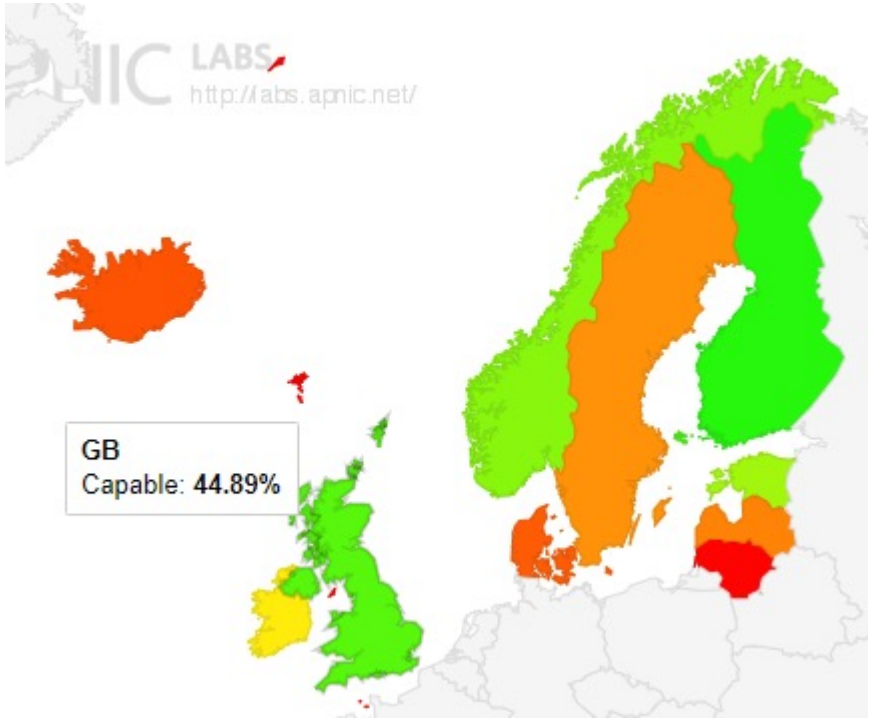

ASN	AS Name	IPv6 Capable	IPv6 Preferred	Samples
AS5607	BSKYB-BROADBAND-AS	93.02%	92.56%	4,421,122
AS2856	BT-UK-AS BTnet UK Regional network	85.94%	85.47%	4,418,896

We spend all day fighting* with Sky over who can get the most IPv6 into consumer homes!

* We do actually collaborate a lot <3

Deployments

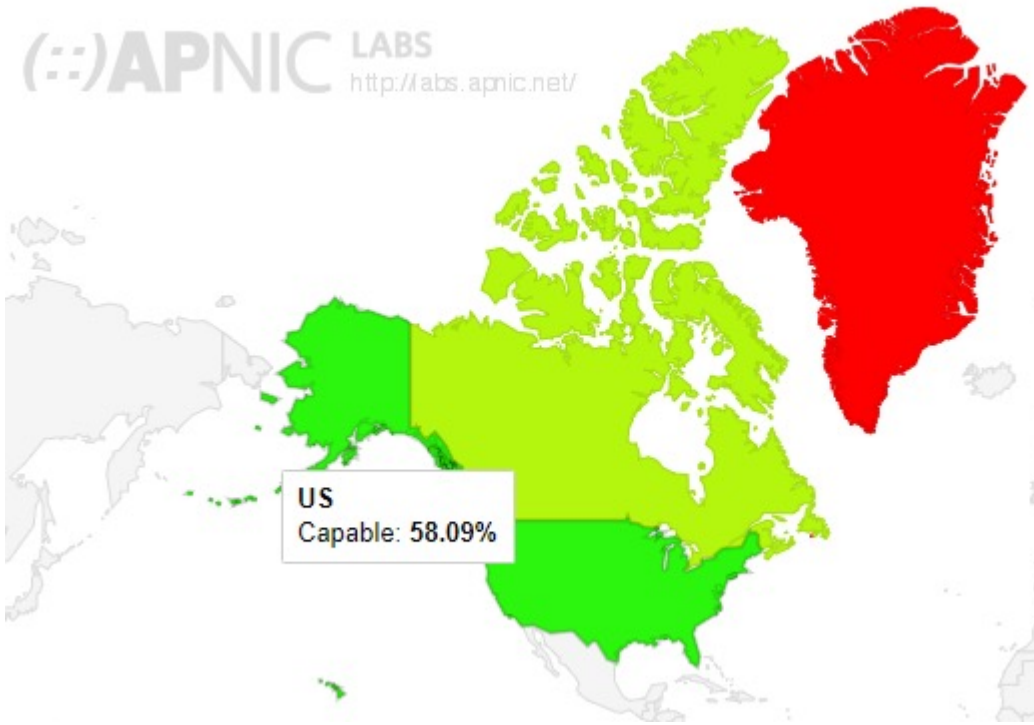

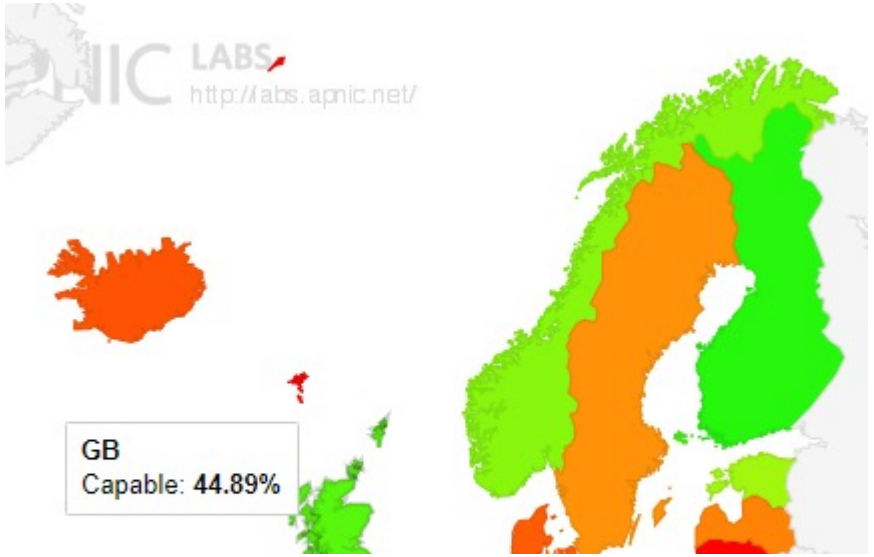

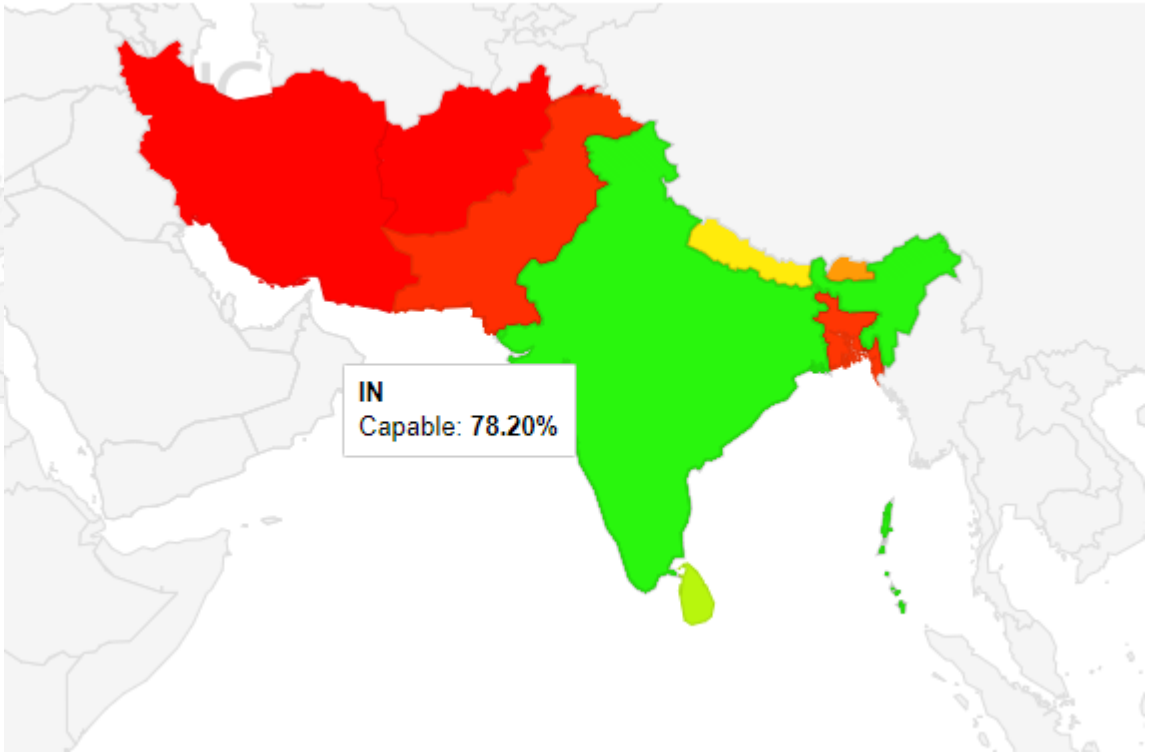

Deployments

The reality in 2023, is that the largest ISPs are *finally* hitting their limits, and they will need to implement address sharing for IPv4

Notably this has been the case for mobile broadband for many years, and the greater goals of 5G necessitate a better connection.

Additionally, many altnets are *starting out* with little-to-no IPv4, and the combination of CGNAT gateways + per-IPv4 prices, means…

… flows using IPv6 are less expensive, and have fewer failure modes.

BT

Deployments

“We have plenty of IPv4, it works fine?”

It probably will work fine for you for some years to come…

… Assuming you only sell services to people with good IPv4 connectivity, of course.

… and that your staff never dial-in from home, the coffee shop, a ship in the Atlantic…

…or from an IETF meeting with IPv6-only WiFi :D

Deployments

“Vendor A doesn’t support IPv6 yet, so we’ll wait”

Fire them. They ought to know better by now.

But seriously: vendors will only do so when asked by customers, so please ask at the very least, and make ‘IPv6-only operation’ a contractual stipulation – the sooner, the better.

The ‘big three’ Cloud vendors have made major strides in IPv6 equivalency in recent years (even months) so this is worth revisiting now, and often.

Some Sundries

A few takeaway thoughts for the rest of the day

~~Food~~ Toast for thought

Is anyone here holding on to large (>22) IPv4 assignments? $$$$$!

No phishing email ever cared about NAT

Conversely, every corporate merger has cared deeply about RFC1918 overlaps

Please do not let anyone tell you that SRv6 is IPv6 ‘plus’ or ‘evolved’

Are you absolutely certain that you’re not already doing some IPv6?

Be sure to inflict some IPv6 on your staff long before your customers

Don’t embark on this journey without AAAA records and a firm discussion with Talkie Toaster

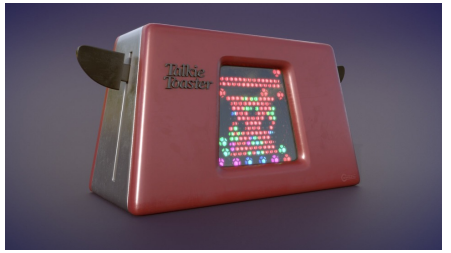

Oblig. XKCD

"I think the IETF hit the right balance with the 128 bits thing. We can fit MAC addresses in a /64 subnet, and the nanobots will only be able to devour half the planet."

https://xkcd.com/865/

Thanks for listening!

Are there any questions?

(Not about bread products of any kind)