Whatever happened to IPv6?

Geoff Huston AM
Chief Scientist, APNIC

Projecting IPv6 Adoption

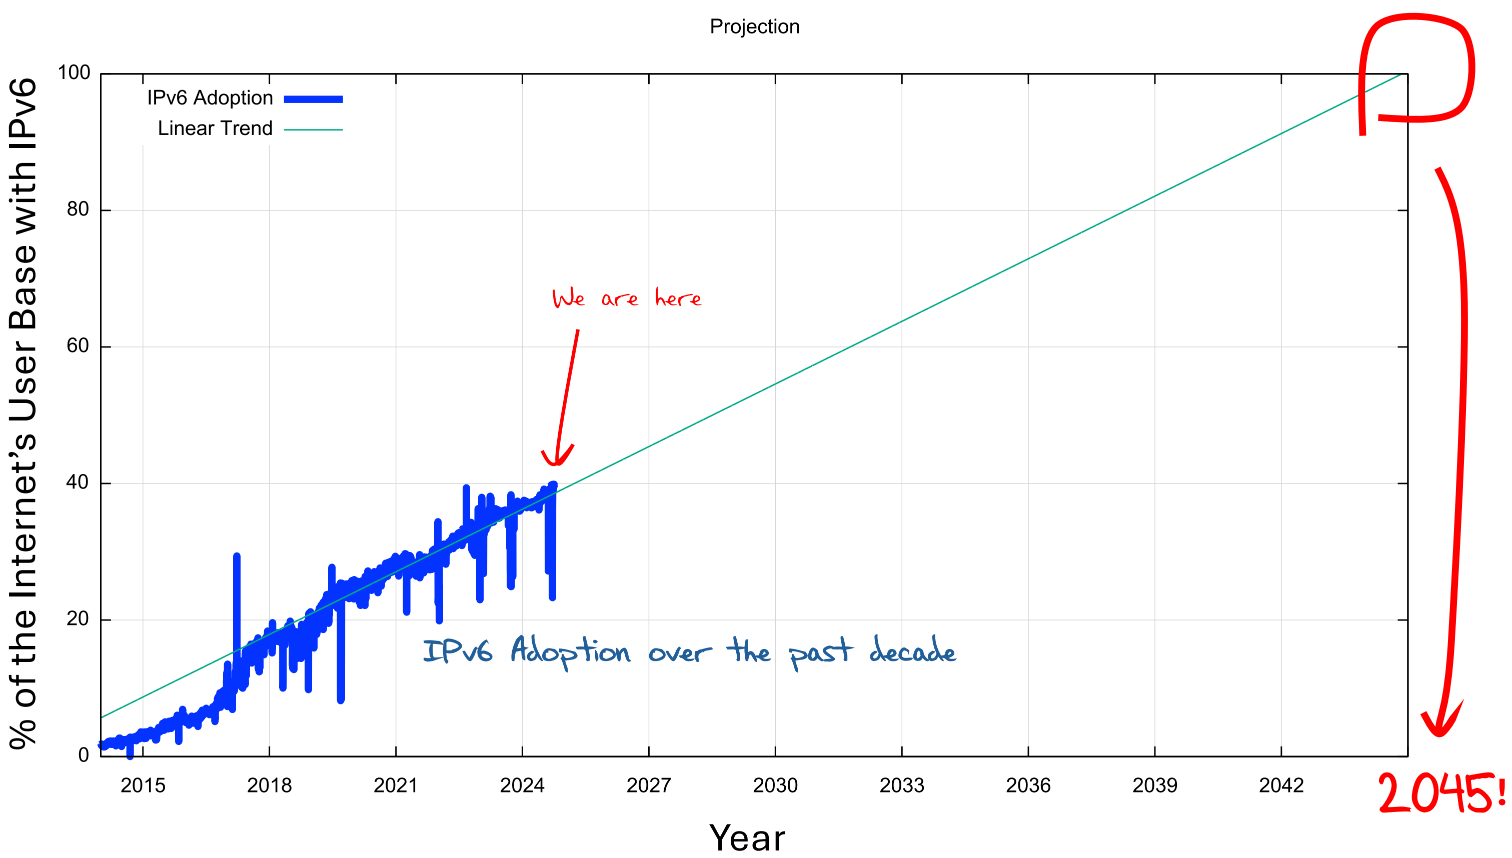

This is unexpected

- Back in the early nineties when the Internet was just picking up momentum NOBODY could conceive that a transition to IPv6 would take longer than five years - tops!
- A total timeframe to complete this transition from start to finish of fifty years was unthinkable!
- But that is where we are
- What went wrong?

We started early

- The Internet was only just gathering momentum in 1990 we were told that the address plan had just a few years to run before the Class B address pool would be fully depleted

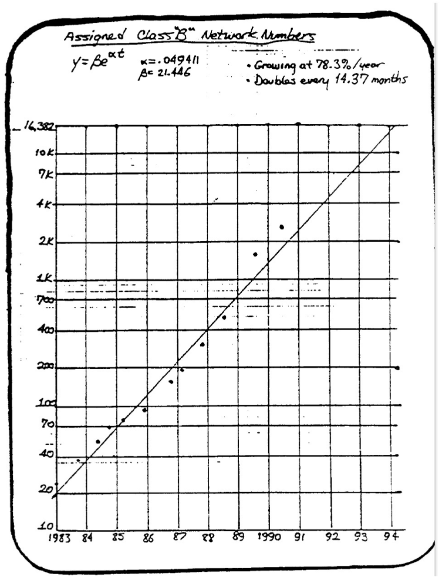

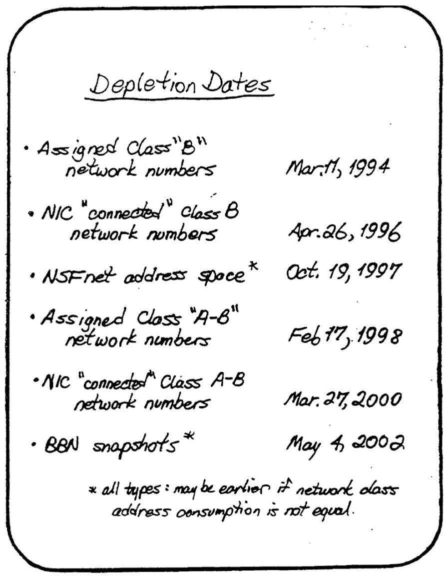

"Internet Growth" Frank Solensky, Proc. IETF, Aug 1990

We worked quickly

- The ROAD effort in the IETF had produced candidate *IPng* protocols by 1992
- And by 1994 we had managed reach a rough consensus on what to do…
  - Reduce the speed of address runout by dropping Class A, B and C structured addresses in address deployment and routing and replacing it by variable sized network/host identifiers
  - Adopt a successor IP protocol with a larger address field…

IPv6!

RFC 1883, December 1995

IPv6 was incremental

Minimal changes to IP:

- Expand the address fields four-fold to 128 bits
  - 64-bit network prefix, 64-bit interface identifier
- Remove packet fragmentation-on-the-fly
- Replace ARP with Multicast

- BUT
  - It was NOT backward compatible with IPv4!

Transition using Dual Stack

- The plan was that we needed to run some form of a “dual stack” transition process
    - Network-level proxies/translators were deemed to be too insecure
- Which meant that we needed to equip EVERY host and EVERY network with two protocol stacks
- But the network was too big to ”just do it” so we devised a dual stack transition plan that allowed for piecemeal adoption

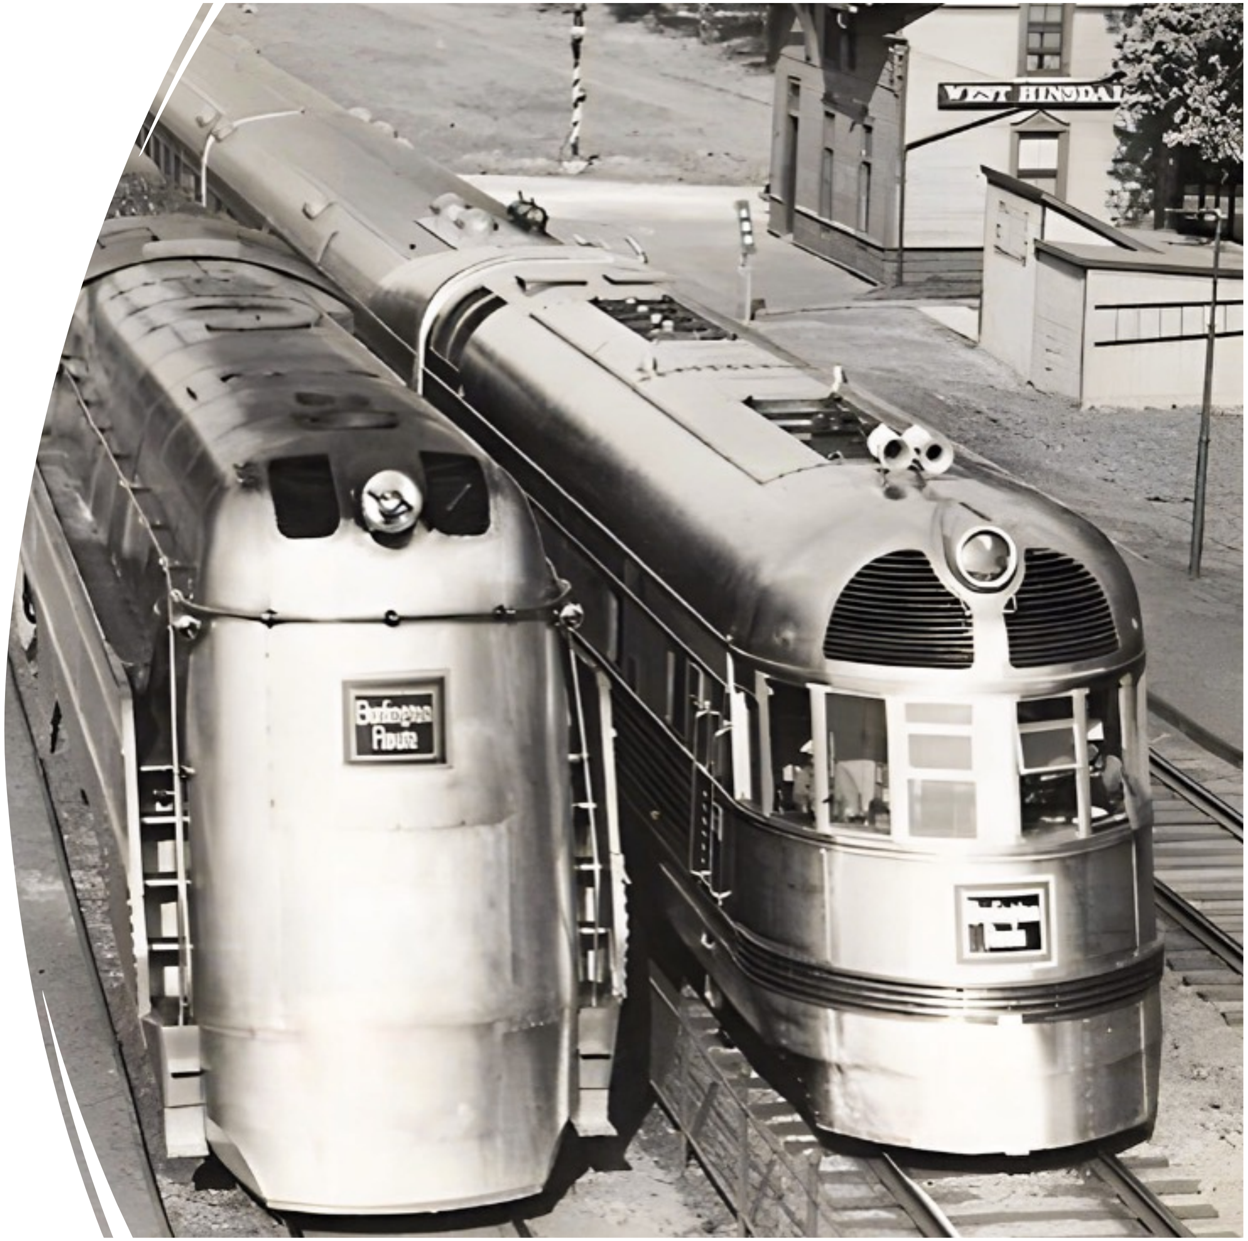

Problem Solved!

- We had a technology solution to address depletion
- Hosts preferred to use IPv6 when there was IPv6 available
- The transition would operate automatically as networks enabled IPv6

- So, we then shifted our collective attention elsewhere!
- For the next decade or so
- Until...

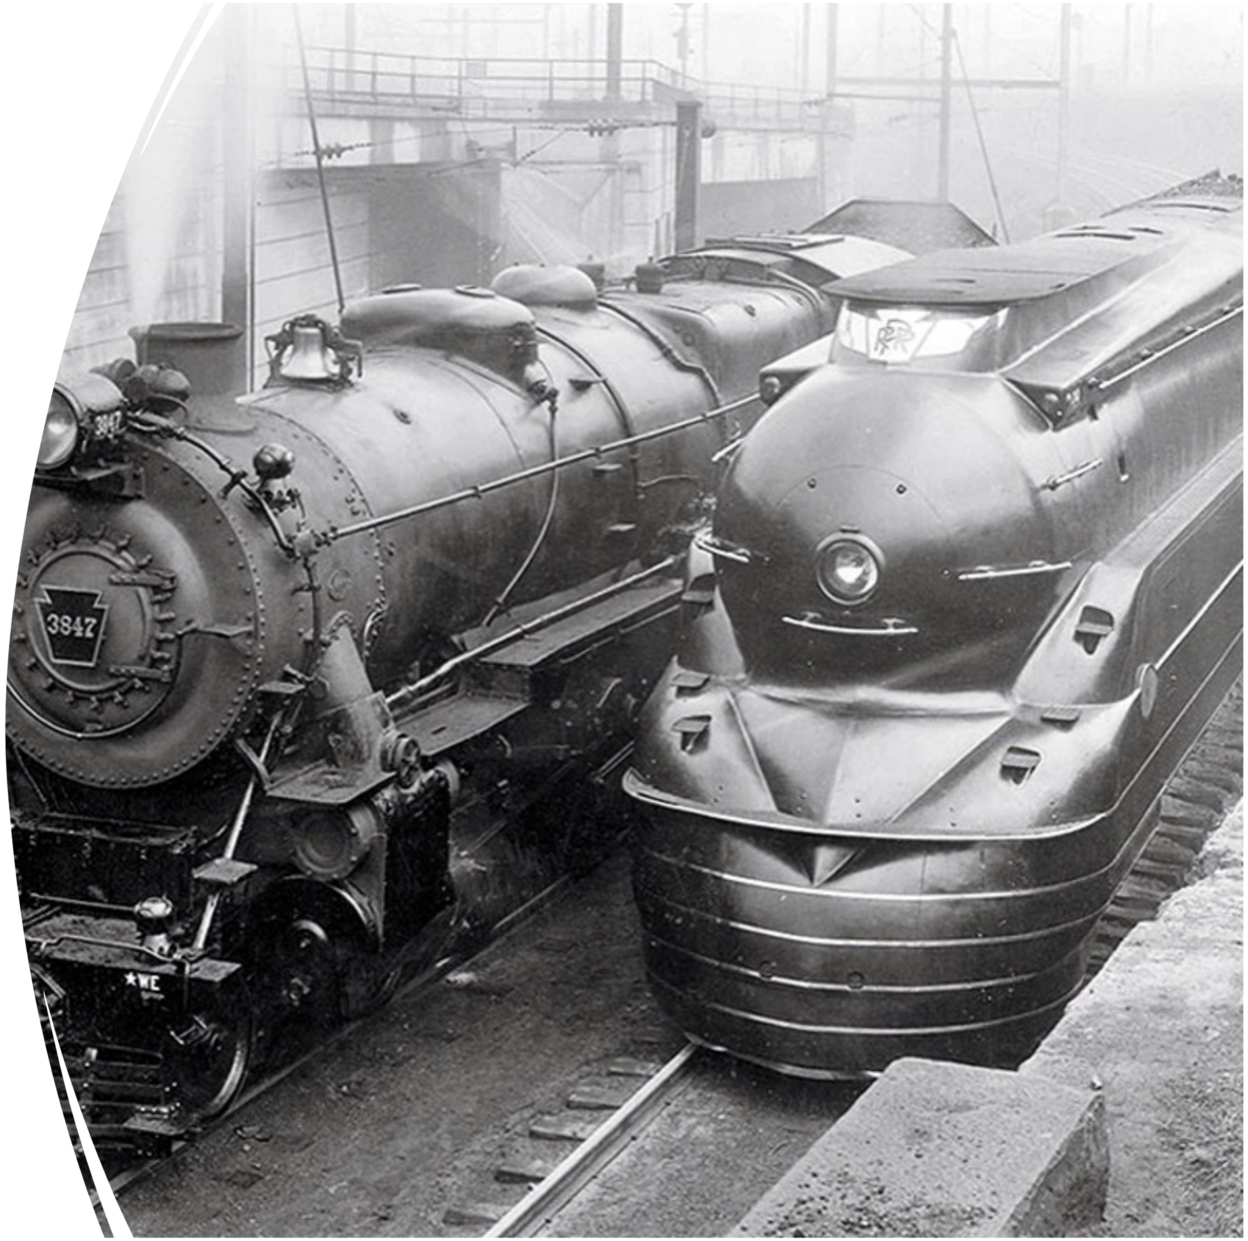

Accelerating Growth

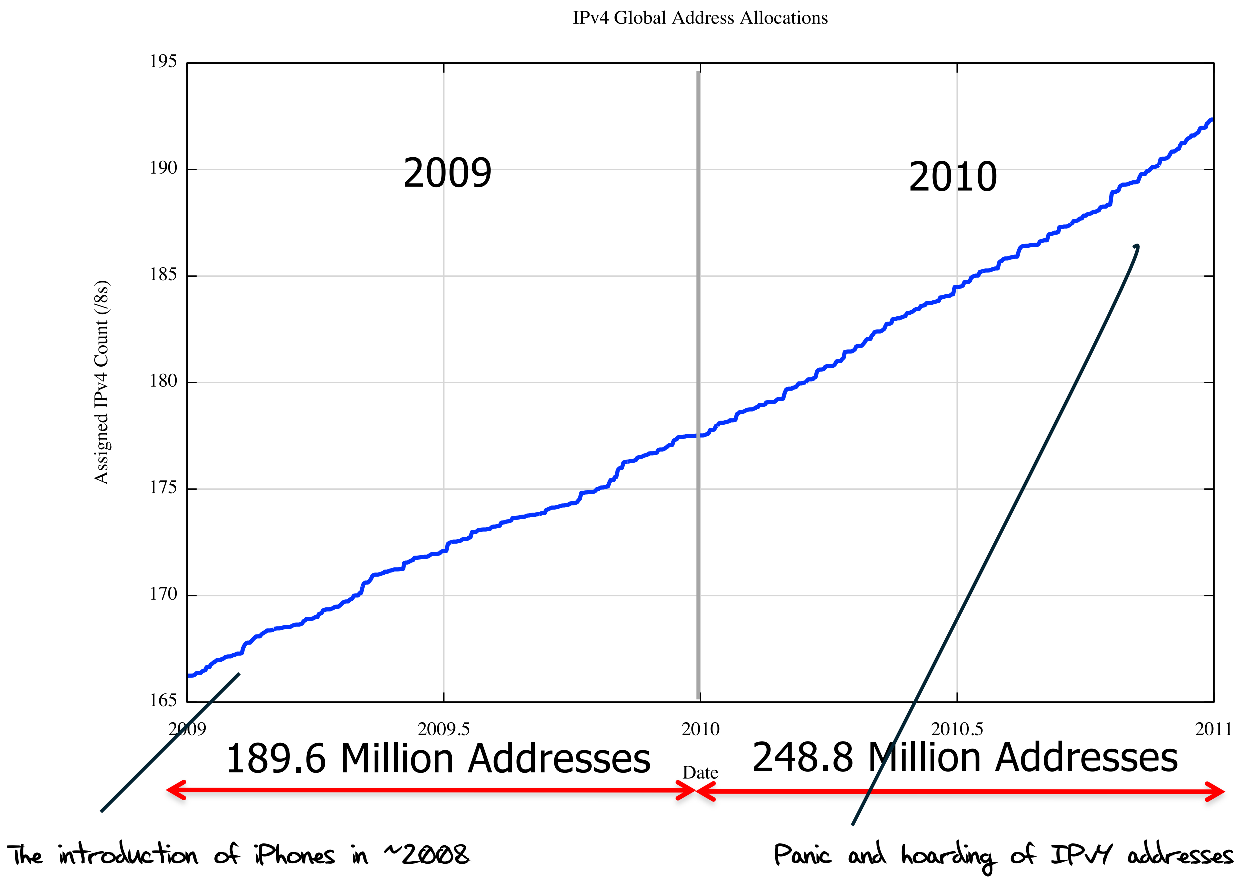

We had this plan …

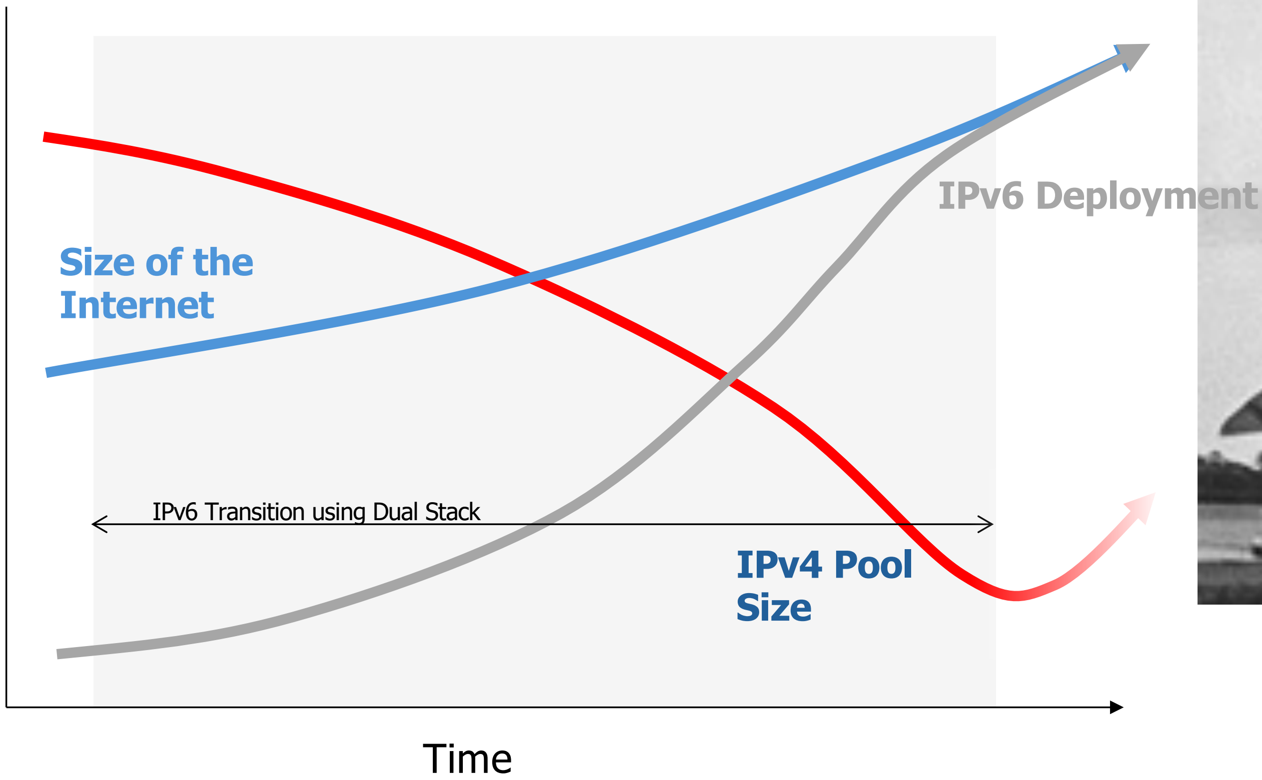

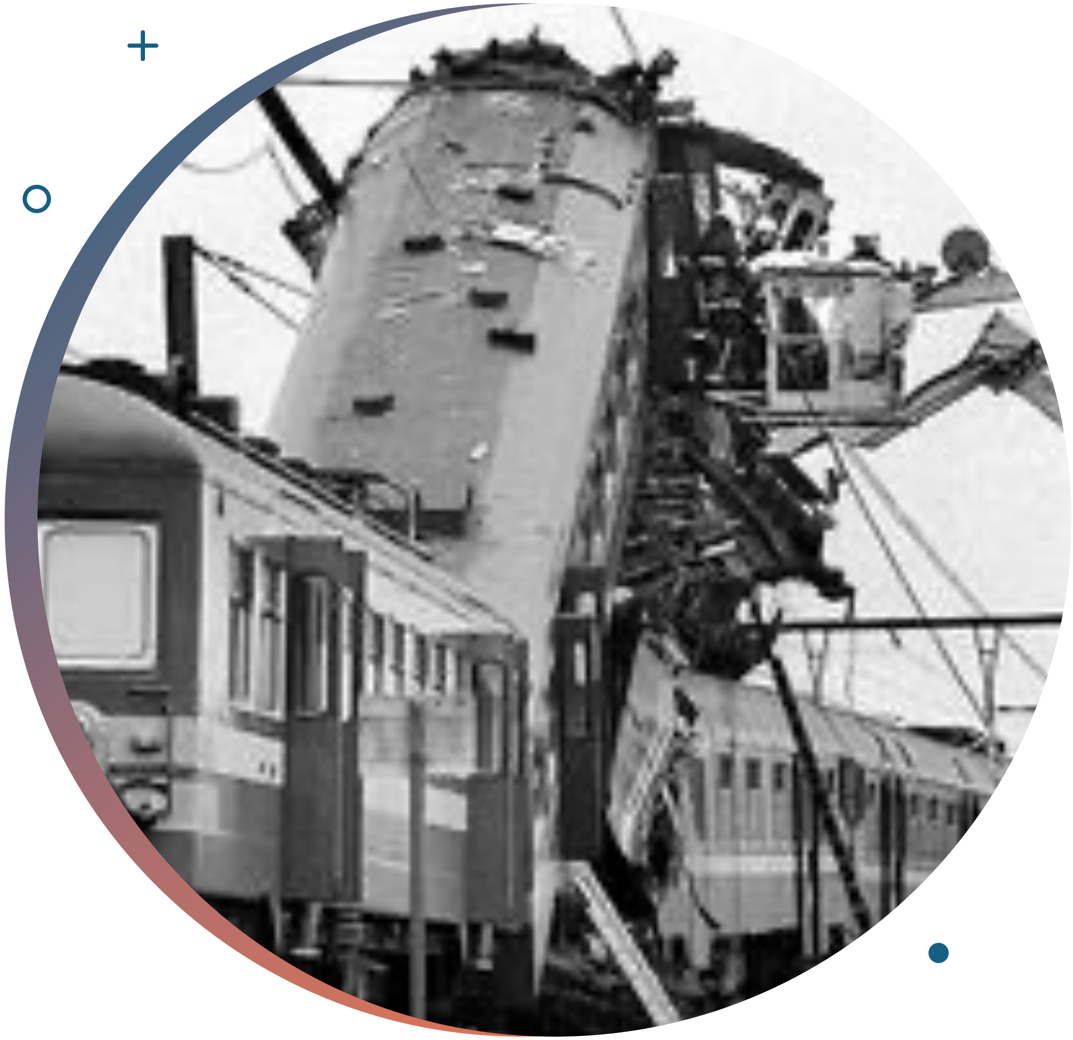

Something wasn't right!

We were meant to have completed the transition to IPv6 BEFORE we completely exhausted the supply channels of IPv4 addresses

The 2012 IPv6 Transition Plan

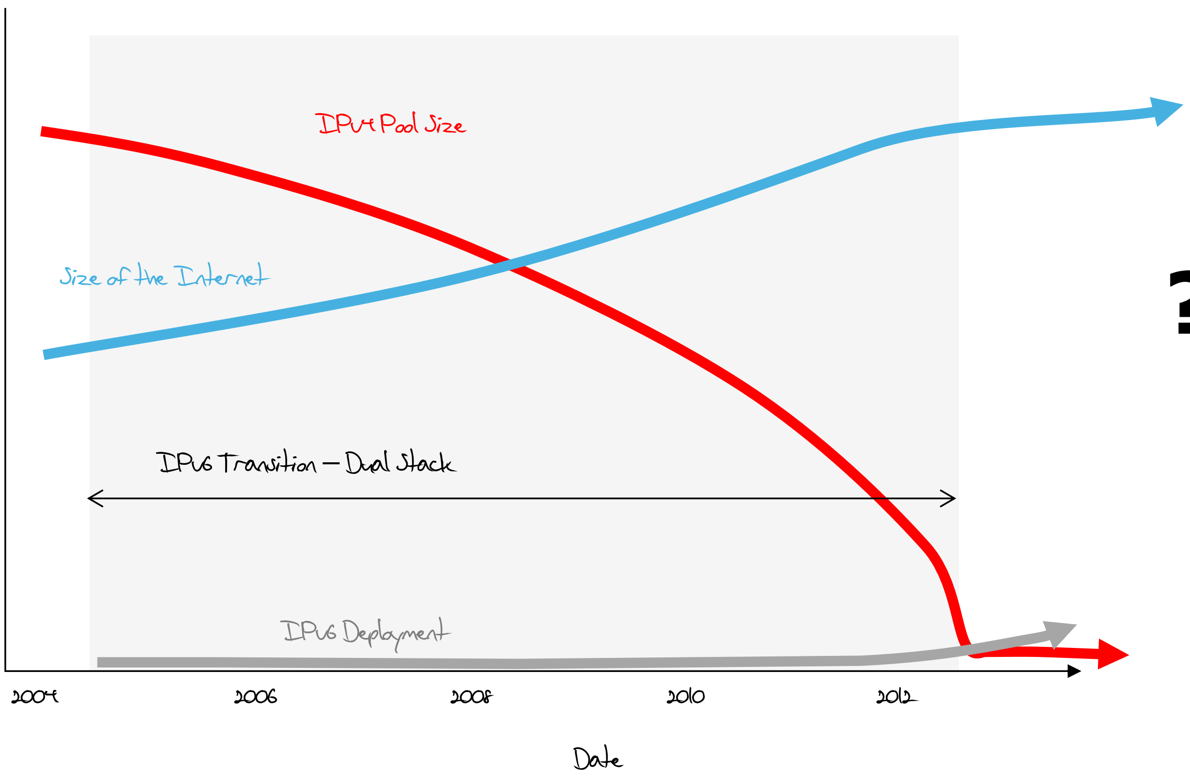

?

What now?

Despite the whinging from IETF purists over the compromise of a pristine end-to-end model there really was no other option:

The answer was NATs!

NATs

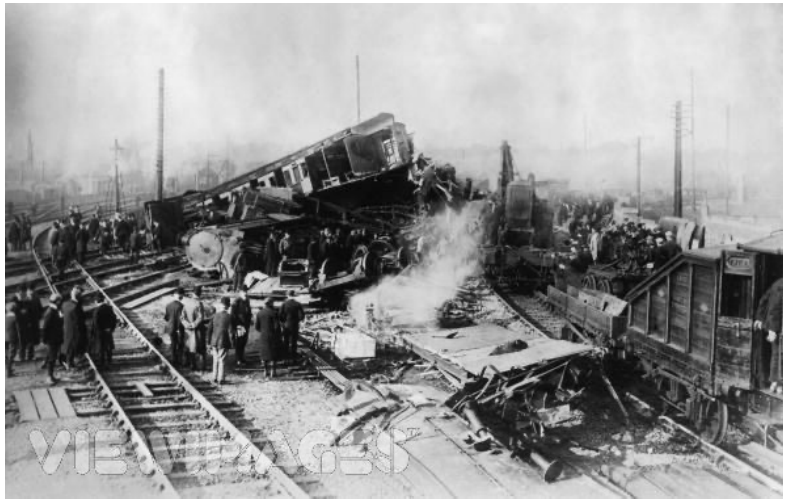

- This low friction response to IPv4 address depletion had been used for more than a decade in client/server network architectures
- **Clients** initiate a service transaction and only need an external address/port binding for the duration of the transaction
- **Servers** sit in central data centres and share platform IP addresses using name-based distinguishers

Implications

- IPv4 addresses continued to be in demand far beyond the exhaustion of the RIR's free space pools
  - In the transition environment, all new and expanding network deployments needed IPv4 service access and IPv4 addresses for as long as we were in this dual track transition

- But the process was no longer directly controlled through RIR' s address allocation policies
  - Address access for IPv4 addresses is mediated by market pricing
  - And the large CDN actors appear to be dominating this space

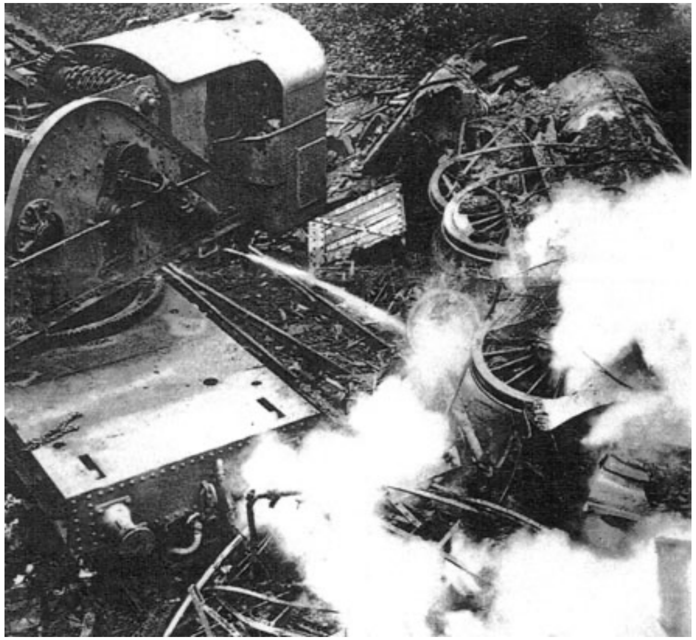

Not everyone is feeling the pressure to move to Dual Stack

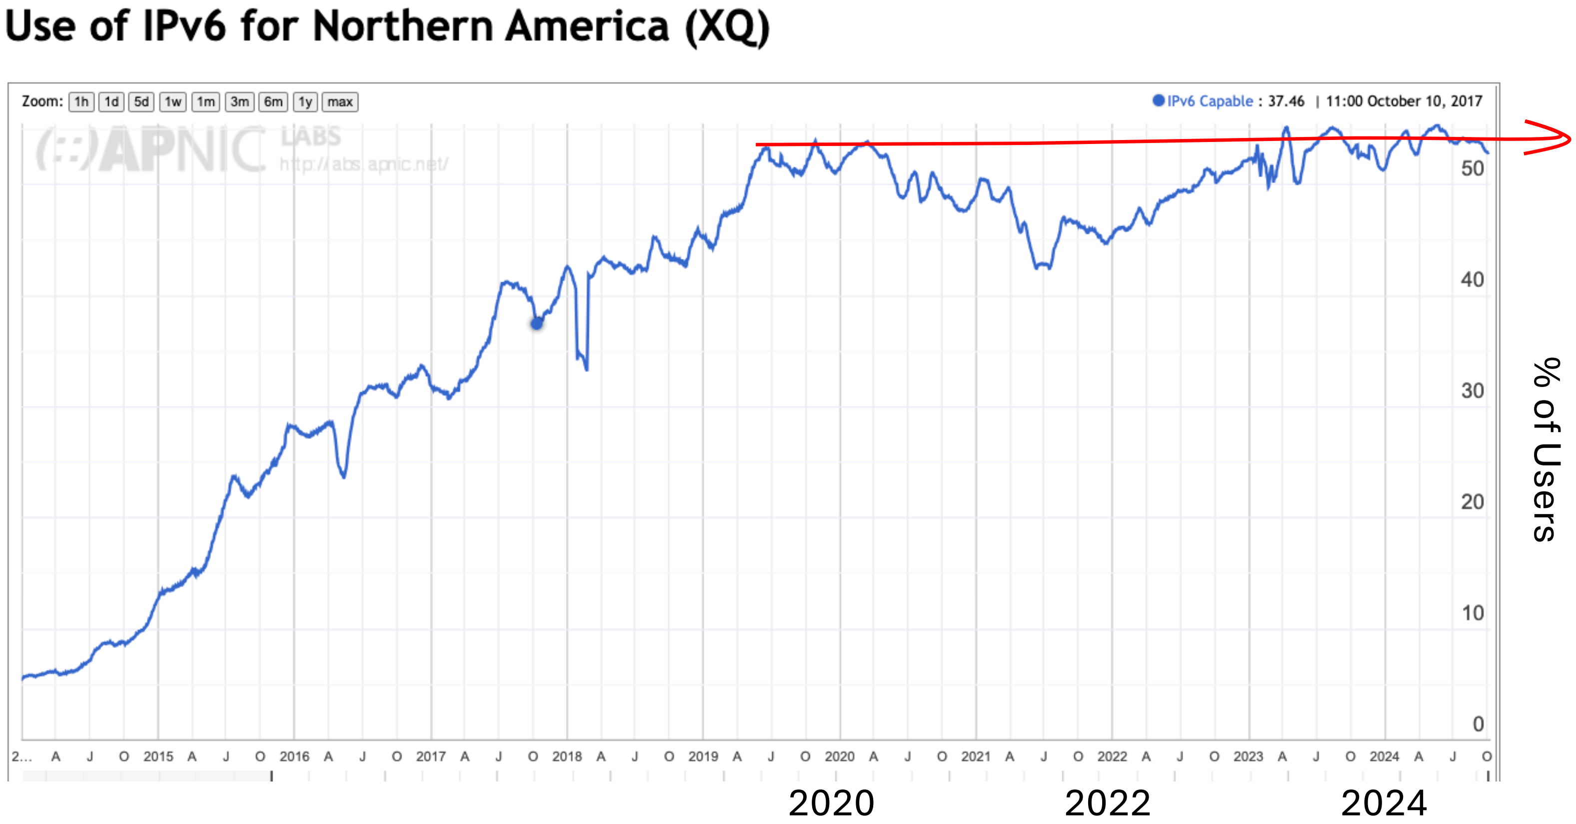

https://stats.labs.apnic.net/ipv6/XQ

Why not?

- Because we no longer operate within a strict address-based network architecture
  - Clients no longer use a permanent unique public IP address to communicate with servers
  - Servers no longer use a permanent unique public IP address to communicate with clients
- Address scarcity takes on a different dimension when you don’t need public addresses to uniquely number every host and service

What's driving change today?

- **From scarcity to abundance**!
- For many years, the demand for communications services outstripped available capacity
- We used price as distribution function to moderate demand to match available capacity
- But this is no longer the case – available capacity in the communications domain far outpaces demand

Abundant Capacity

Fibre cables continue to deliver massive capacity increases within relatively constant unit cost of deployment

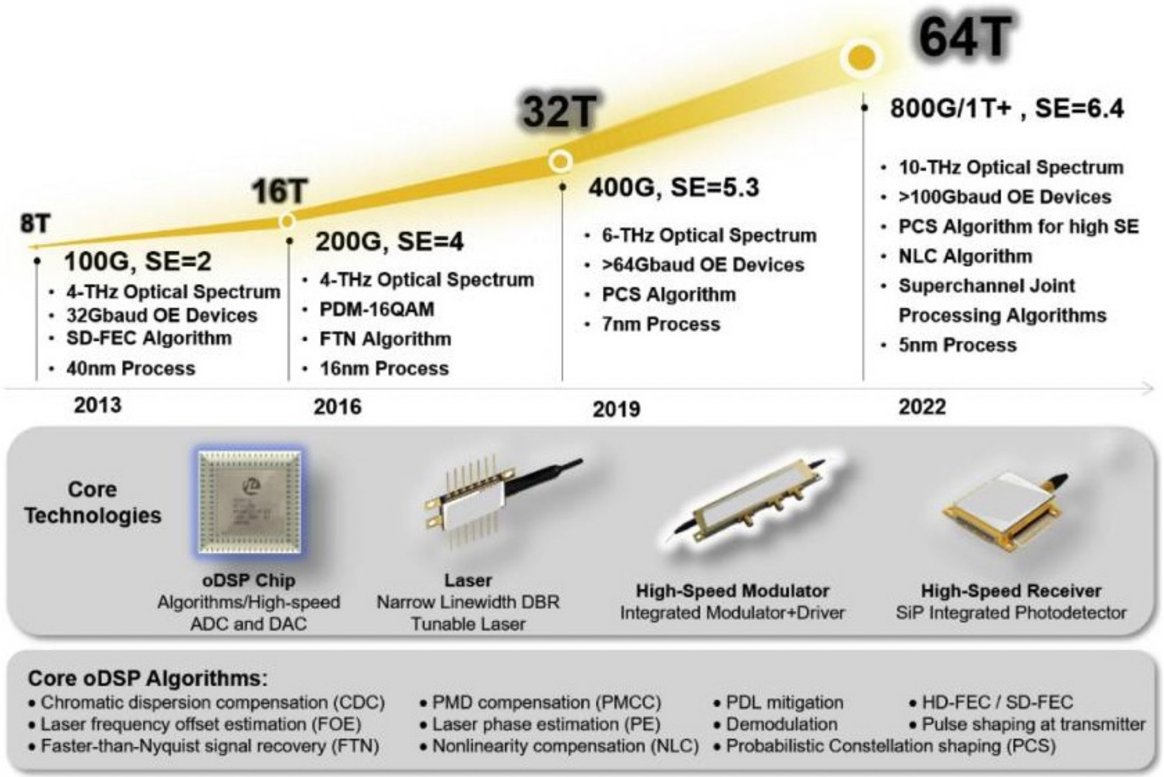

https://www.ncbi.nlm.nih.gov/

(That 2022 number is probably low – at the end of 2022 we can pull 2.2T per lambda with a 190Gbd signal rate, giving a fibre capacity of 105T)

Abundant Compute Power

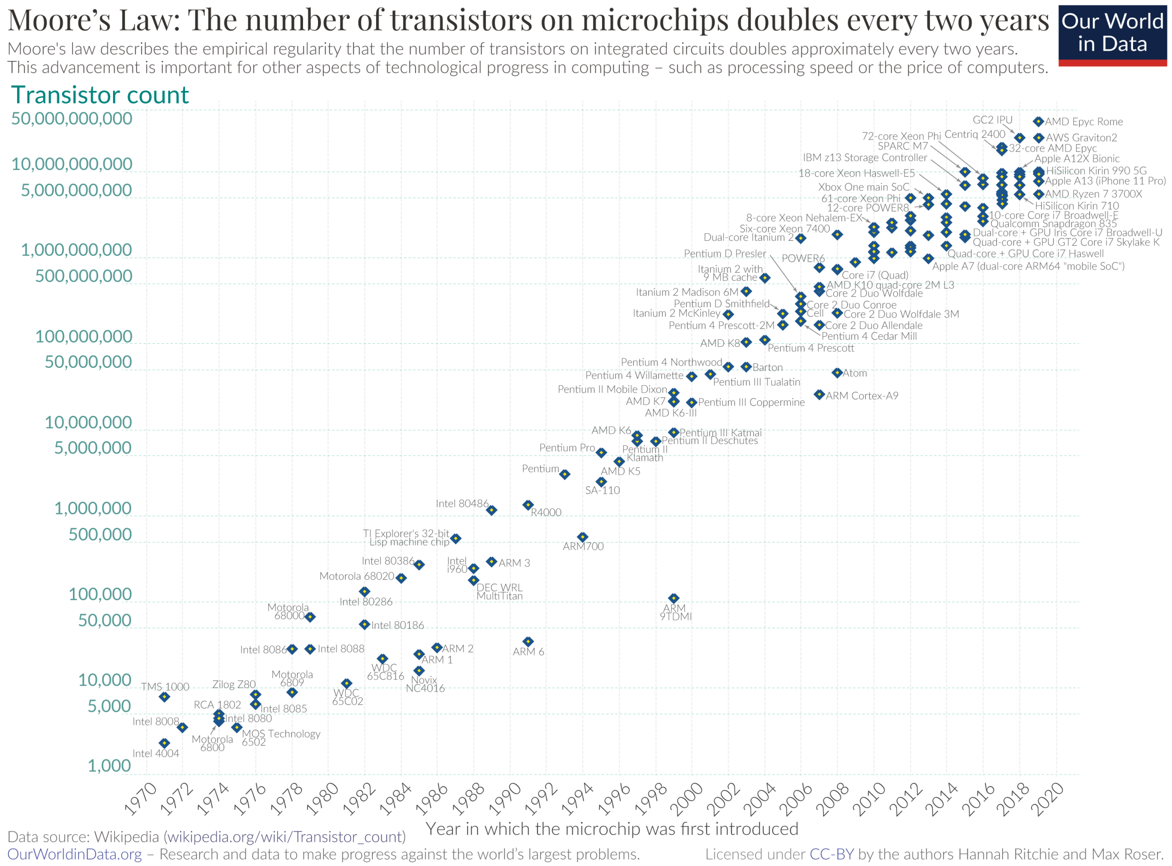

By Max Roser, Hannah Ritchie - https://ourworldindata.org/uploads/2020/11/Transistor-Count-over-time.png, CC BY 4.0, https://commons.wikimedia.org/w/index.php?curid=98219918

Abundant Storage

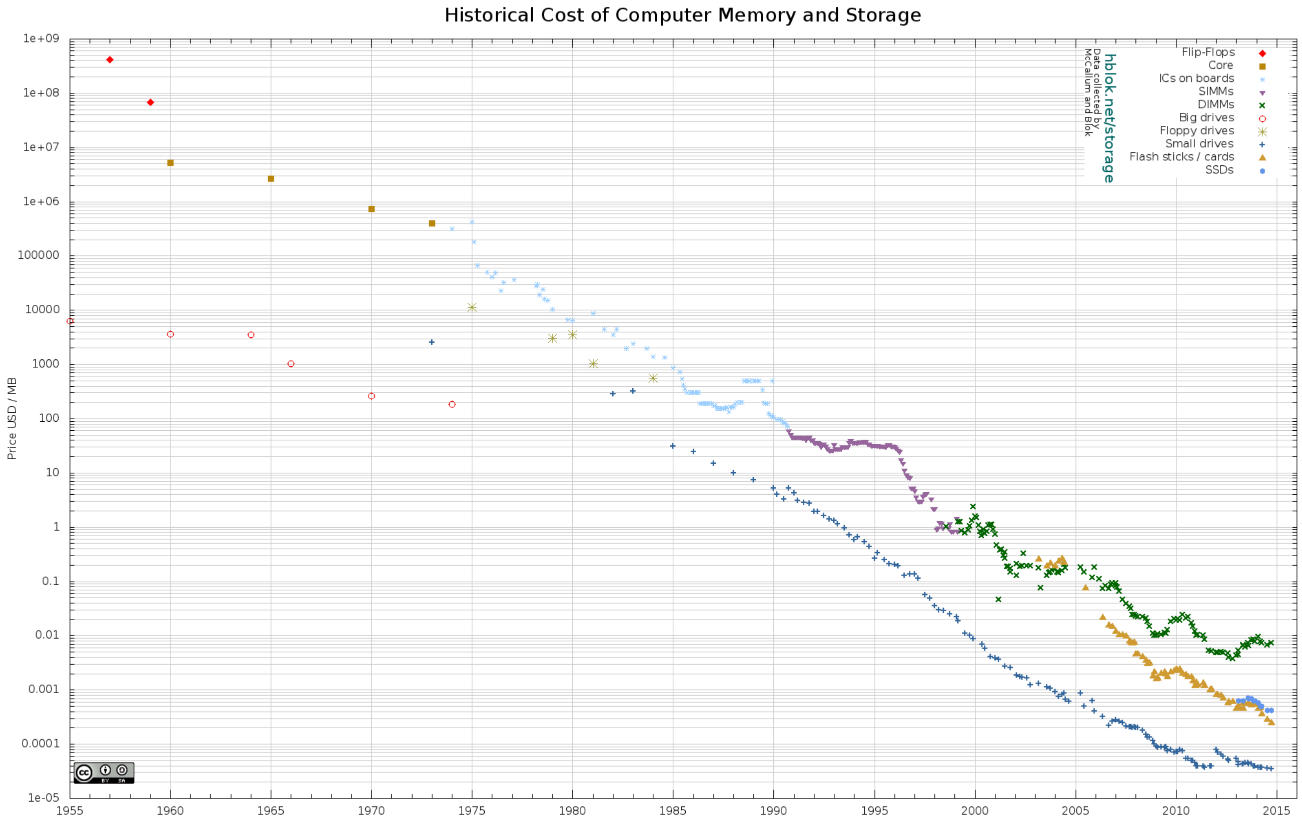

http://aiimpacts.org/wp-content/uploads/2015/07/storage_memory_prices_large-_hblok.net_.png

How can we use this abundance?

- By changing the communications provisioning model from ***on demand*** to ***just in case***
- Instead of using the network to respond to users by delivering services *on demand* we've changed the service model to provision services close to the edge just in case the user requests the service
- With this change we've been able to eliminate the factors of *distance* from the network and most network transactions occur over **shorter network** spans
- What does a ***shorter*** network enable?

Bigger

- Increasing **transmission capacity** by using photonic amplifiers, wavelength multiplexing and phase/amplitude/polarisation modulation for fibre cables
- Serving content and service transactions by distributing the load across many individual platforms through **server and content aggregation**
- The rise of high-capacity mobile edge networks and mobile platforms add massive volumes to content delivery

- To manage this massive load shift we've stopped pushing content and transactions across the network and instead **we serve from the edge**

Faster

- Reduce latency - stop pushing content and transactions across the network and instead **serve from the edge**
- The rise of CDNs serve (almost) all Internet content and services from massively scaled distributed delivery systems.

- The “Packet Miles” to deliver content to users has shrunk - that’s faster!
- The development of high frequency cellular data systems (4G/5G) has resulted in a highly capable last mile access network with Gigabit capacity
- Applications are being re-engineered to meet faster response criteria

- Compressed interactions across shorter distances using higher capacity circuitry results in a much faster Internet

Better

- If "better" means "more trustworthy" and "more privacy" then we are making progress at last!
  - Encryption is close to ubiquitous in the world of web services
  - TLS 1.3 is moving to seal up the last open TLS porthole, the SNI field
  - QUIC is sealing up the transport controls from the networks
  - Oblivious DNS and Oblivious HTTP is moving to isolate knowledge of the querier from the name being queried
  - The content, application, and platform sectors have all taken the privacy agenda up with enthusiasm, to the extent that whether networks are trustable or not doesn't matter any more – **all network infrastructure is uniformly treated as untrustable!**

Cheaper

- We are living in a world of abundant comms and computing capacity
- And working in an industry when there are significant economies of scale
- And it's being largely funded by capitalising a collective asset that is infeasible to capitalise individually – the advertisement market
- The result is that a former luxury service accessible to just a few has been transformed into an affordable mass-market commodity service available to all

And in all this, the money moved up the stack

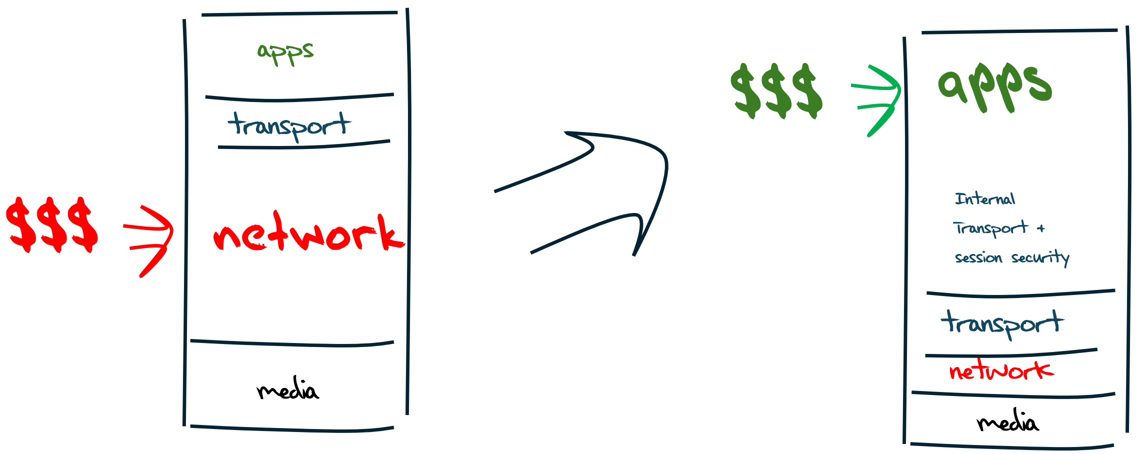

So, who pays?

- Networks need to make the investment to switch to a dual stack mode that includes IPv6
- But neither the user base nor the content world really care
  - And they are certainly not going to pay a premium to the network operator for IPv6
- And in the application service world, IP addresses are **not** the critical resource
- We’ve changed the “currency” of networks!

A Network of Names

- Today's public Internet is largely a service delivery network using CDNs to push content and service as close to the user as possible
- The multiplexing of multiple services onto underlying service platforms is an application-level function tied largely to TLS and service selection using SNI
- The DNS is now used to perform "closest match" service platform selection, supplanting the role of routing
  - Most large CDNs run a BGP routing table with an average AS Path Length that is intended to converge to 1!

Is it Routing? Or Switching?

Let me repeat that, because it's important:

- Most large CDNs run a BGP routing table with an average AS Path Length that is converging to a value of 1!

- The DNS is now used to perform "closest match" service platform selection, supplanting the role of routing
- By volume, most of today's Internet traffic is switched, not routed across the inter-AS space

A new Internet Architecture

- We’ve moved from end-to-end peer networks to client/server asymmetric networks
- We’ve replaced single platform servers-plus-network to replicated servers-minus-network with CDNs
- Clients aren’t identified with a unique public IP address – clients are inside NATs are uniquely identified only in a local context
- Individual services aren’t identified with a unique public IP address – services are identified in the DNS

A new Internet Architecture

- We've moved from end-to-end peer networ… …nt/server asymmetric networks
- We've replac… …to replicated ser…
- Clie… …a unique public IP address – clients are i… …are uniquely identified only in a local context
- Individual services aren't identified with a unique public IP address – services are identified in the DNS

We've moved from address-based networks to name-based services

What am I saying?

- The slow uptake of IPv6 is not because this industry is chronically stupid or short sighted
- There is something else going on here…

What am I saying?

- IPv6 alone is not critical to a large set of end user service delivery environments
- We’ve been able to take a 1980’s address-based architecture and scale it more than a billion-fold by altering the core reliance on distinguisher tokens within the network from addresses to names
    - There was no real lasting benefit in trying to leap across to just another 1980’s address-based architecture (with only a few annoyingly stupid differences, apart from longer addresses!)

Today's Internet:

- Names Matter
- The DNS Matters

- Addresses - not so much
- Address-based Routing - not so much

Thank You!