IPv6

Geoff Huston AM
Chief Scientist, APNIC

What's the problem? Up and to the right, yes?

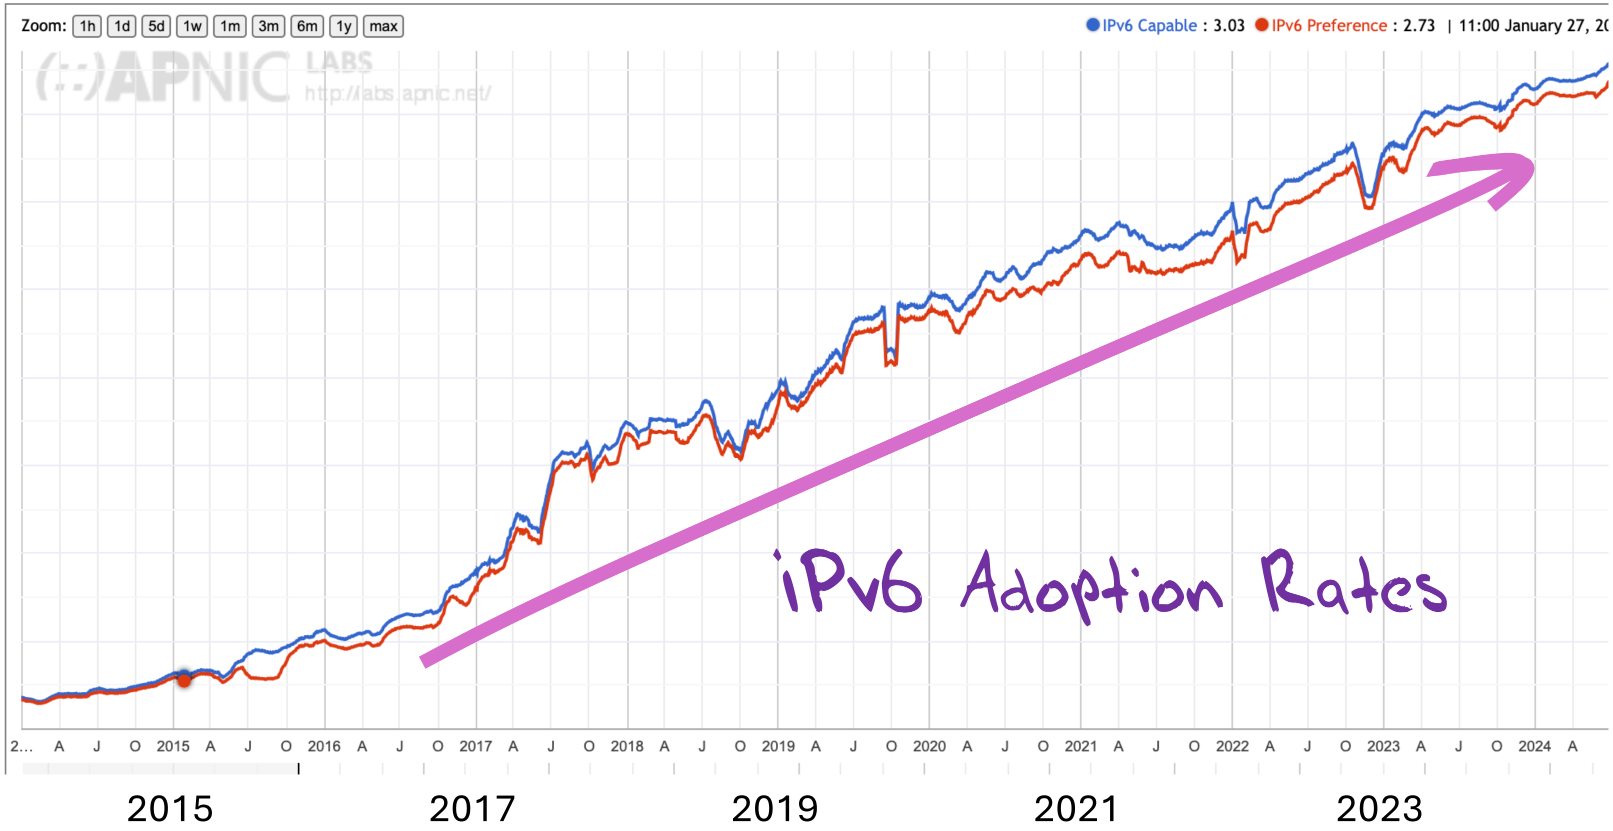

What's the problem? Up and to the right, yes?

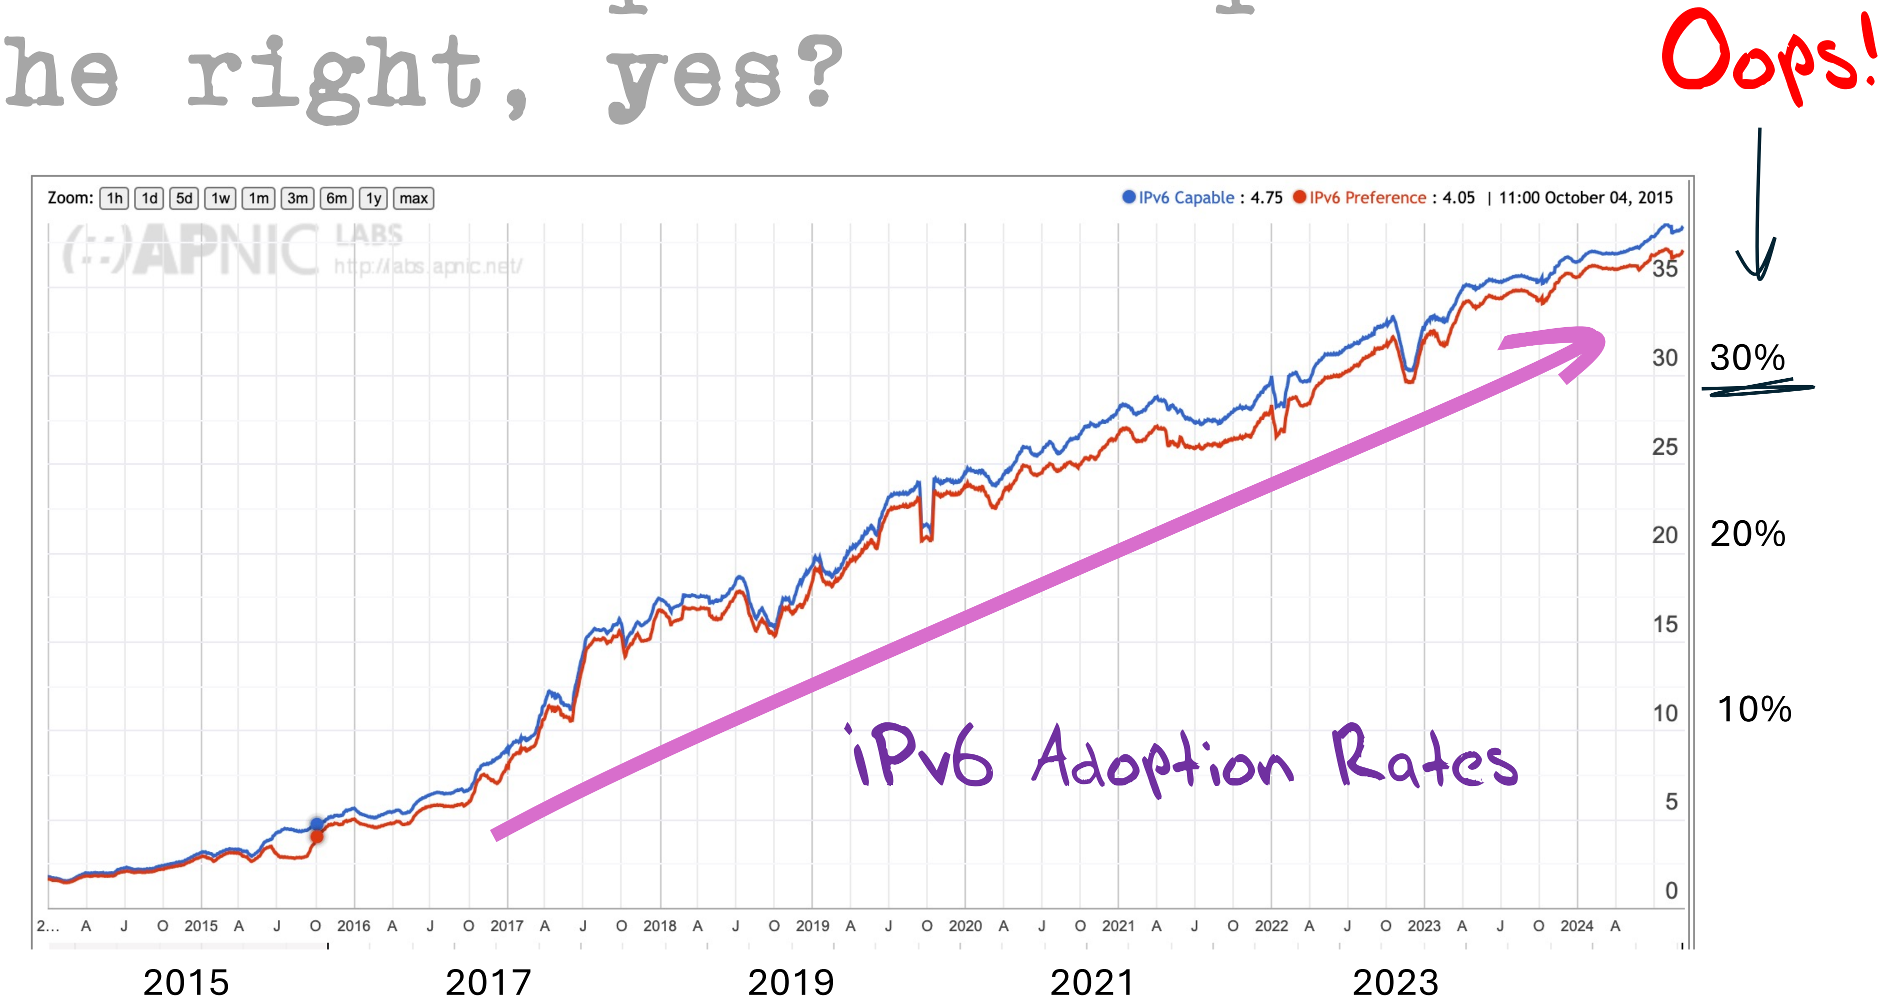

It's been around a decade…

- Since the RIRs handed out their last substantial IPv4 address blocks
- We've been in a state of IPv4 "address exhaustion" for more than a decade
- And yet the global uptake rate of IPv6 is a little over one third of the Internet's user base

- **This is completely unexpected!**

Projecting this Forward

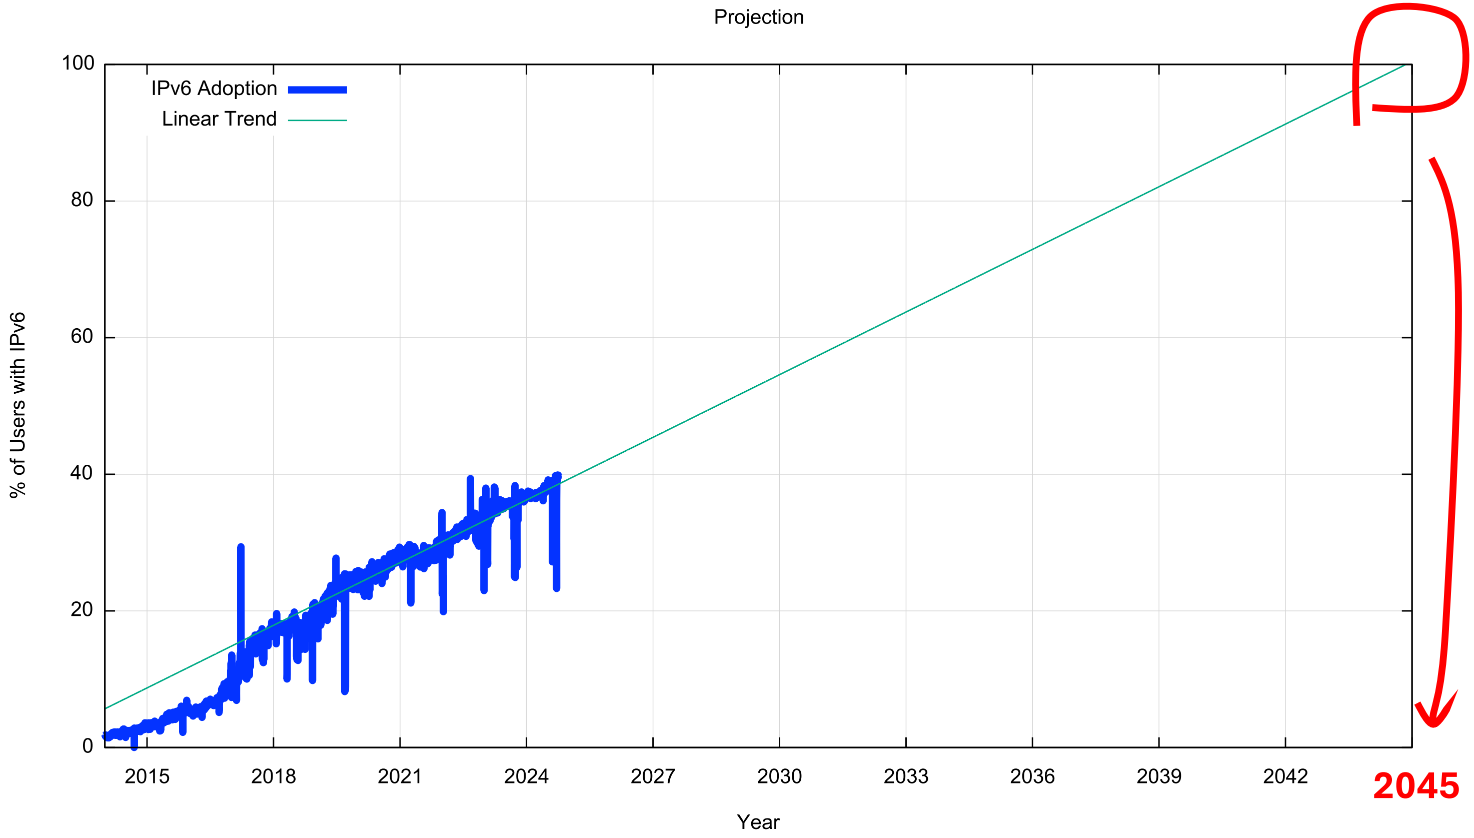

Projecting this Forward

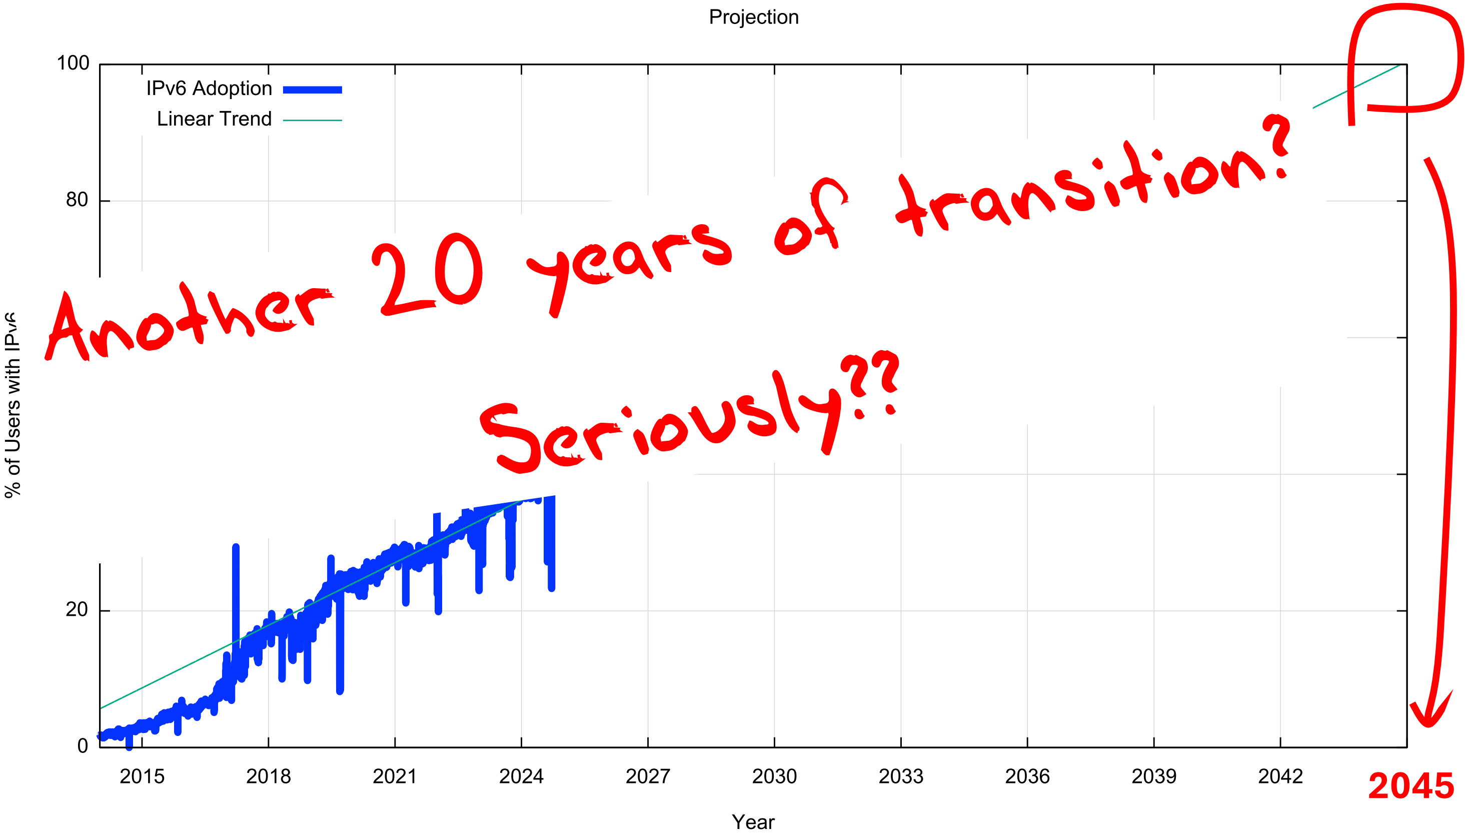

Another 20 years of transition?

Seriously??

What's gone wrong here?

Let's head to the Time Machine

And dial it back by about 35 years!

https://www.flickr.com/photos/tjjohn12/24398810050

1990 - Panic at the IETF!

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

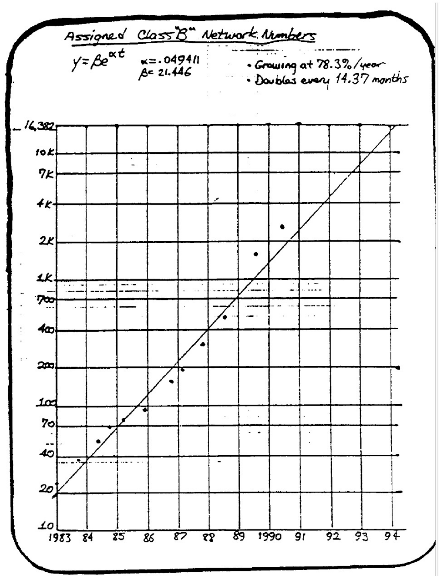

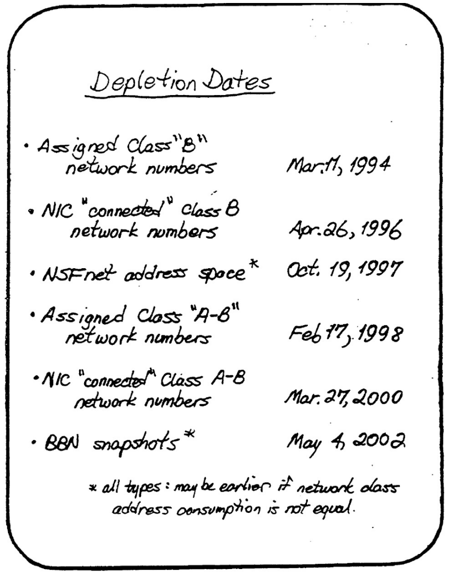

Depletion Dates

• Assigned Class "B" network numbers	Mar. 11, 1994
• NIC "connected" Class B network numbers	Apr. 26, 1996
• NSFnet address space*	Oct. 19, 1997
• Assigned Class "A-B" network numbers	Feb 17, 1998
• NIC "connected" Class A-B network numbers	Mar. 27, 2000
• BBN snapshots*	May 4, 2002

* all types: may be earlier if network class address consumption is not equal.

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

1990 - Panic at the IETF!

- The Internet was only just gathering momentum we were told that the address plan had just a few years to run before the Class B address pool would be fully depleted
- Short term hack – remove the Class A/B/C structure from IP addresses (and routing)
- Longer term solution – develop a new network design that encompassed a far larger scale of use…

IPv6!

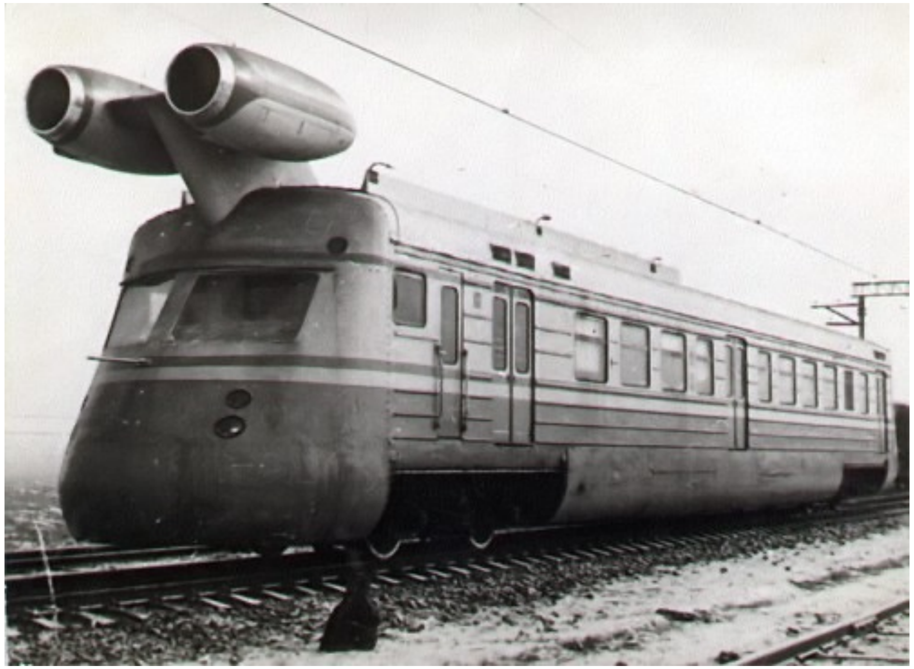

IPv6!

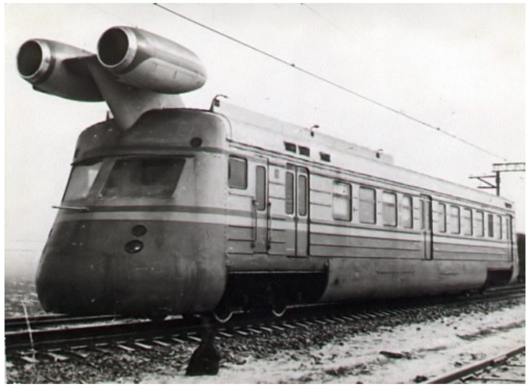

Minimal changes to IP:

- 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

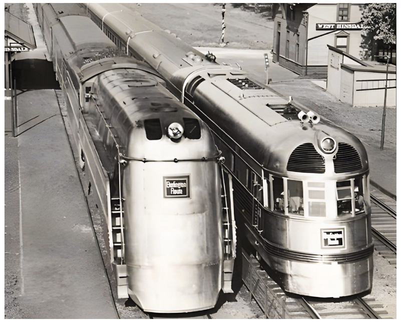

- But IPv6 was not backward compatible with IPv4 on the wire
- The plan was that we needed to run some form of a “dual stack” transition process
  - Network-level proxies/translators were deemed to be too insecure
- Which meant that we needed to equip hosts and networks with two protocol stacks
- But the network was too big to ”just do it” so we needed to devise a “transition plan” that allowed for piecemeal adoption

Dual Stack Assumptions

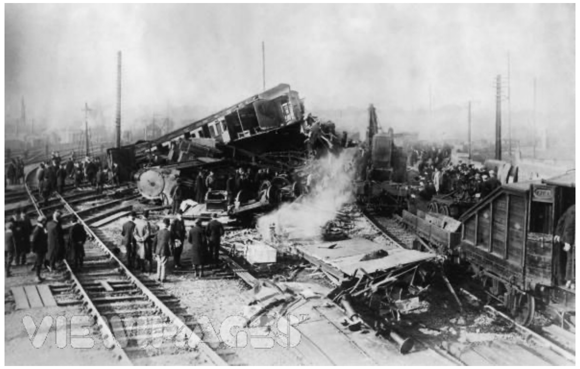

- That we could drive the entire transition to IPv6 while there were still ample IPv4 addresses to sustain the entire network and its growth
- Transition would be driven by individual local decisions to deploy dual stack support
- The *entire* transition would complete *before* the IPv4 unallocated pool was exhausted

Dual Stack
Transition to IPv6

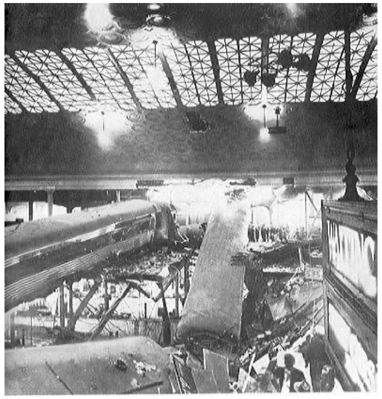

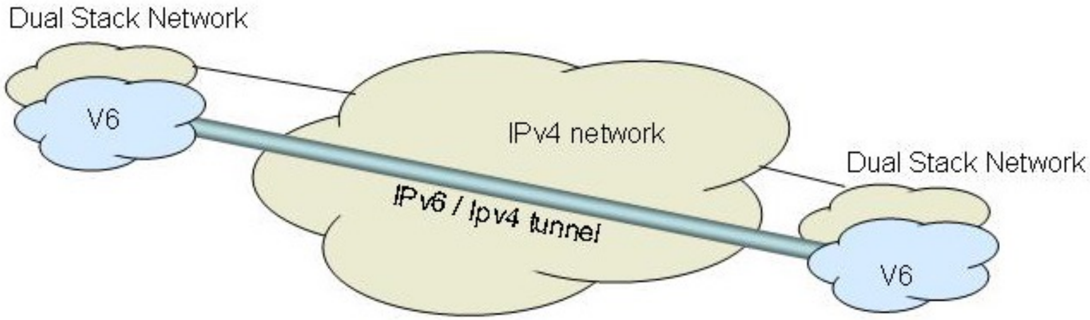

Phase 1

- "Initial" Dual Stack deployment:

  Dual stack networks with V6 / V4 connectivity

  Dual Stack hosts attempt V6 connection, and use V4 as a fallback

Dual Stack
Transition to IPv6

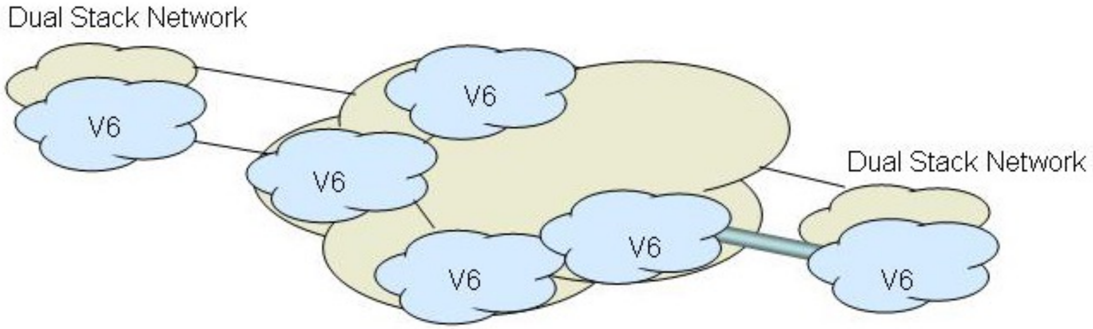

Phase 2

- "Intermediate"
  - Older V4 only networks are retro-fitted with dual stack V6 support

Dual Stack Transition to IPv6

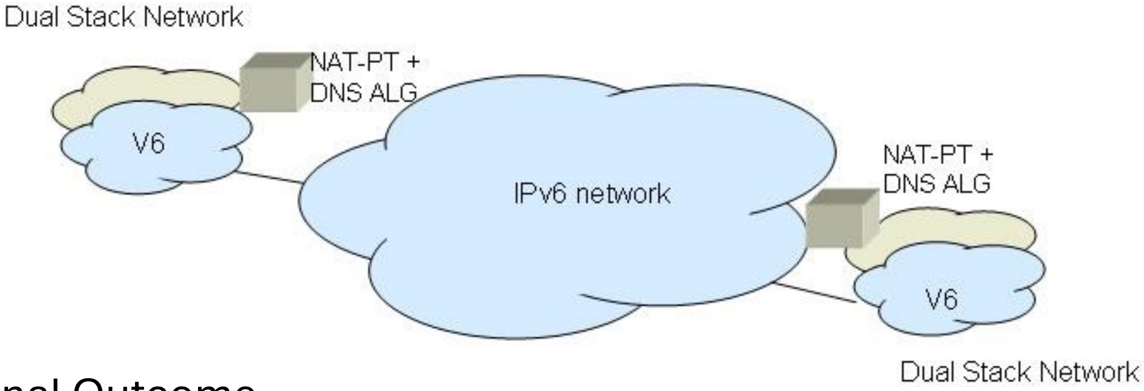

The Final Outcome

- "Completion"
    - V4 shutdown occurs in a number of networks
    - Connectivity with the residual V4 islands via DNS ALG + NAT-Protocol Translation
    - Outside the residual legacy deployments, the network is single protocol V6

Problem solved!

- We had a technology solution to address depletion
- Because 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…

2009 - Inexorable Growth

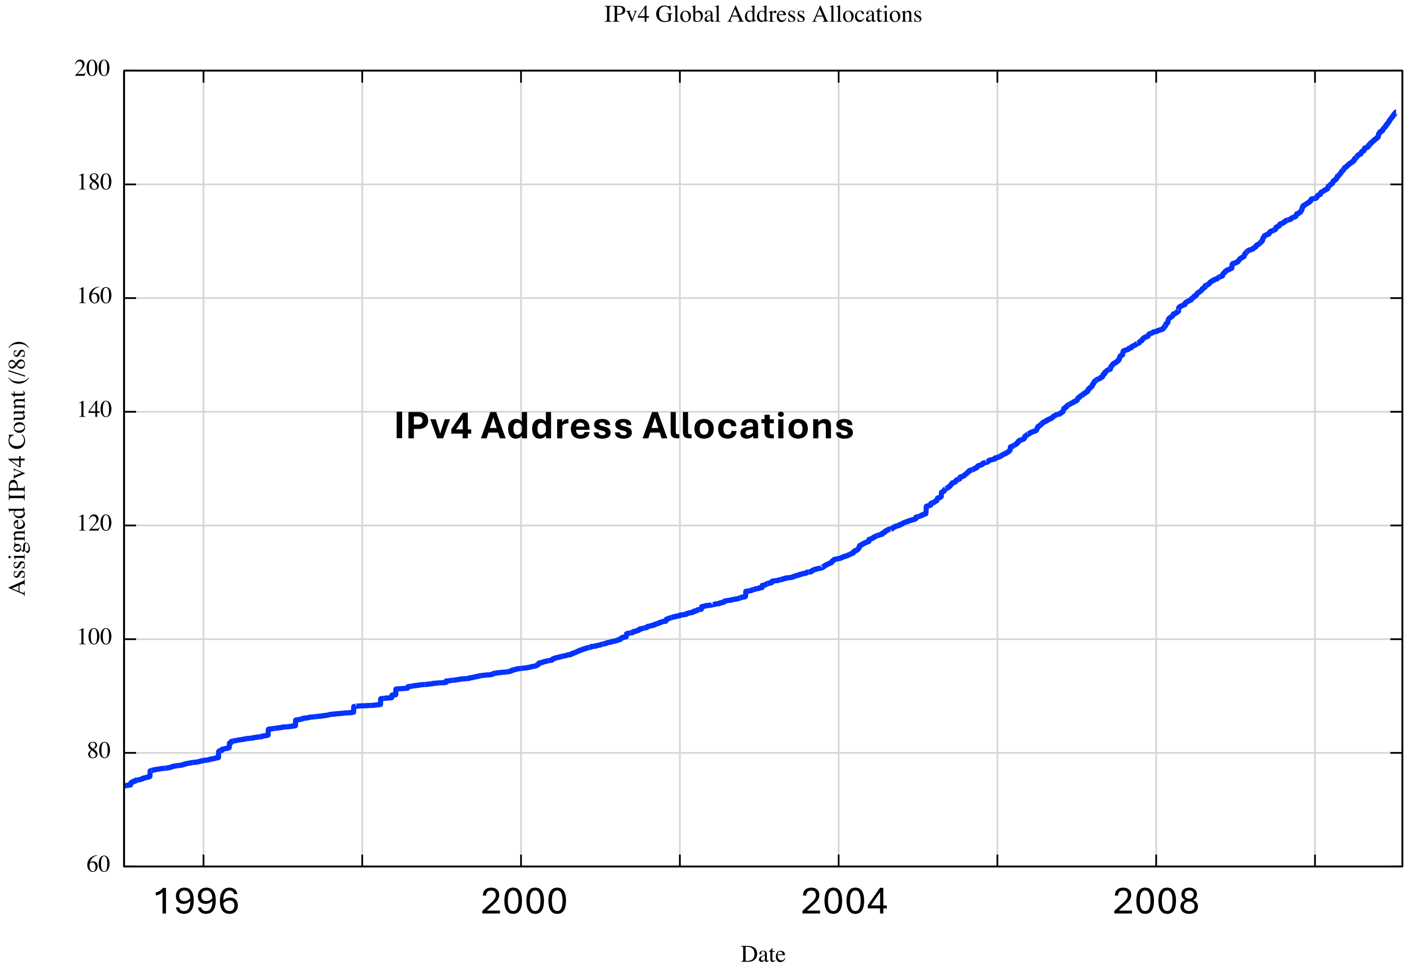

2010 - Accelerating Growth

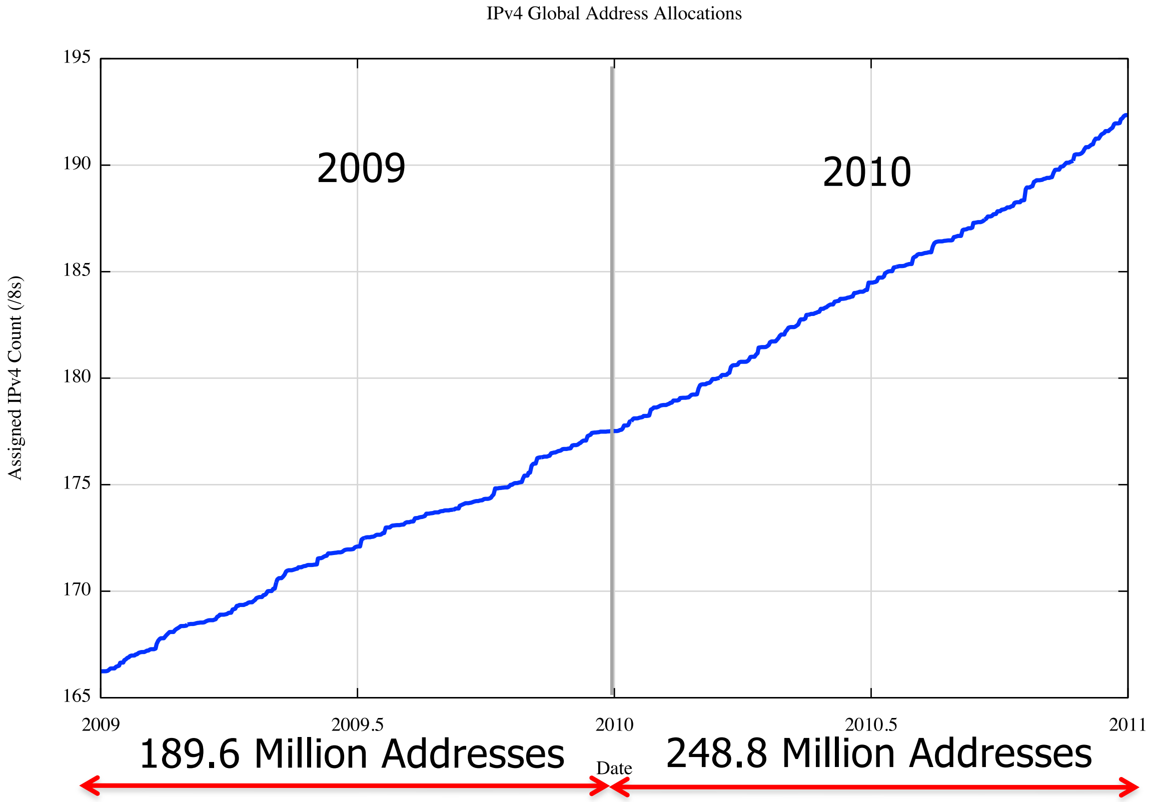

2009

2010

189.6 Million Addresses

248.8 Million Addresses

Accelerating Growth

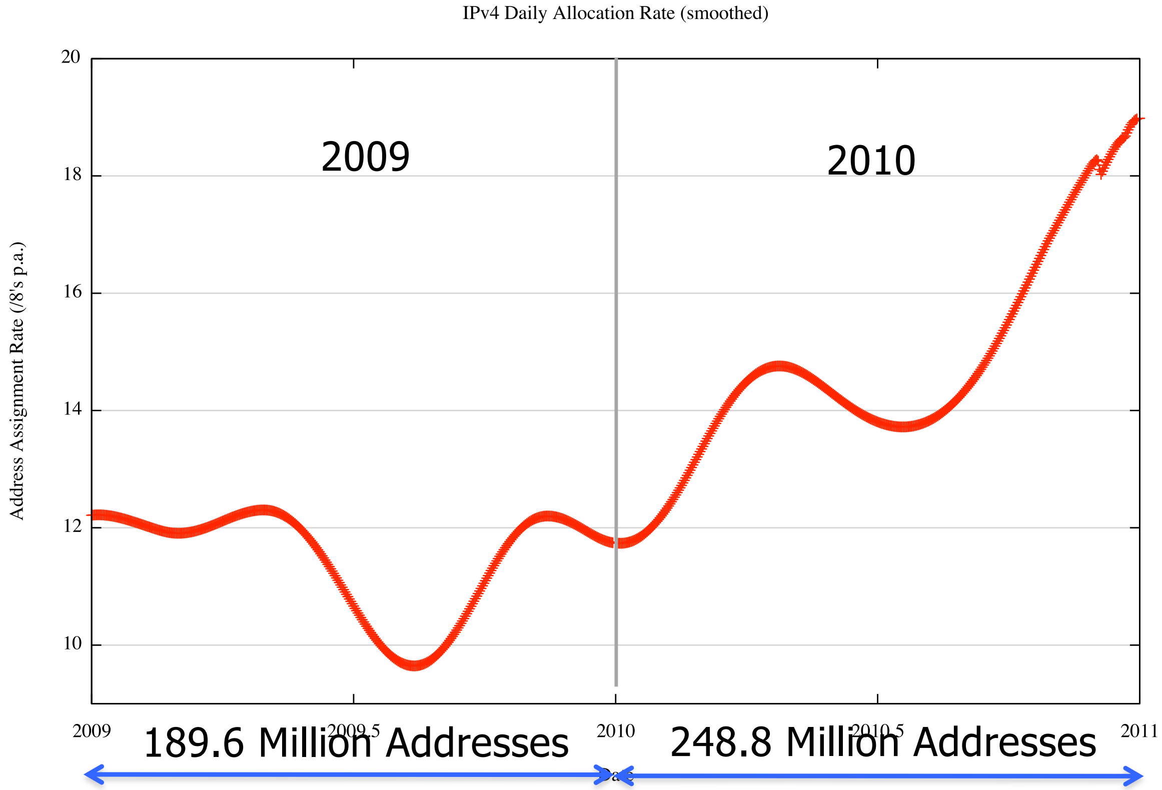

IPv4 Daily Allocation Rate (smoothed)
Address Assignment Rate (/8's p.a.)
20
18
16
14
12
10
2009
2009.5
2010
2010.5
2011

2009

2010

189.6 Million Addresses

248.8 Million Addresses

We had this plan …

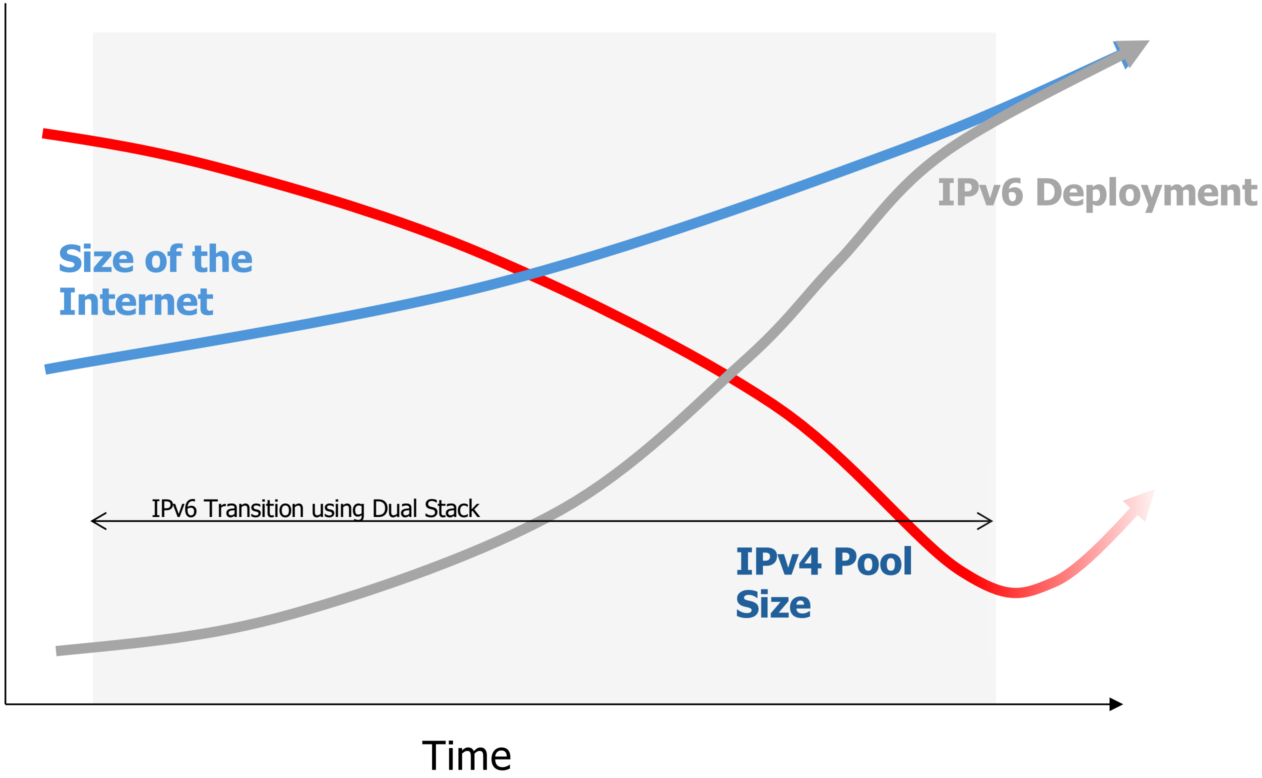

But we had strayed off-plan!

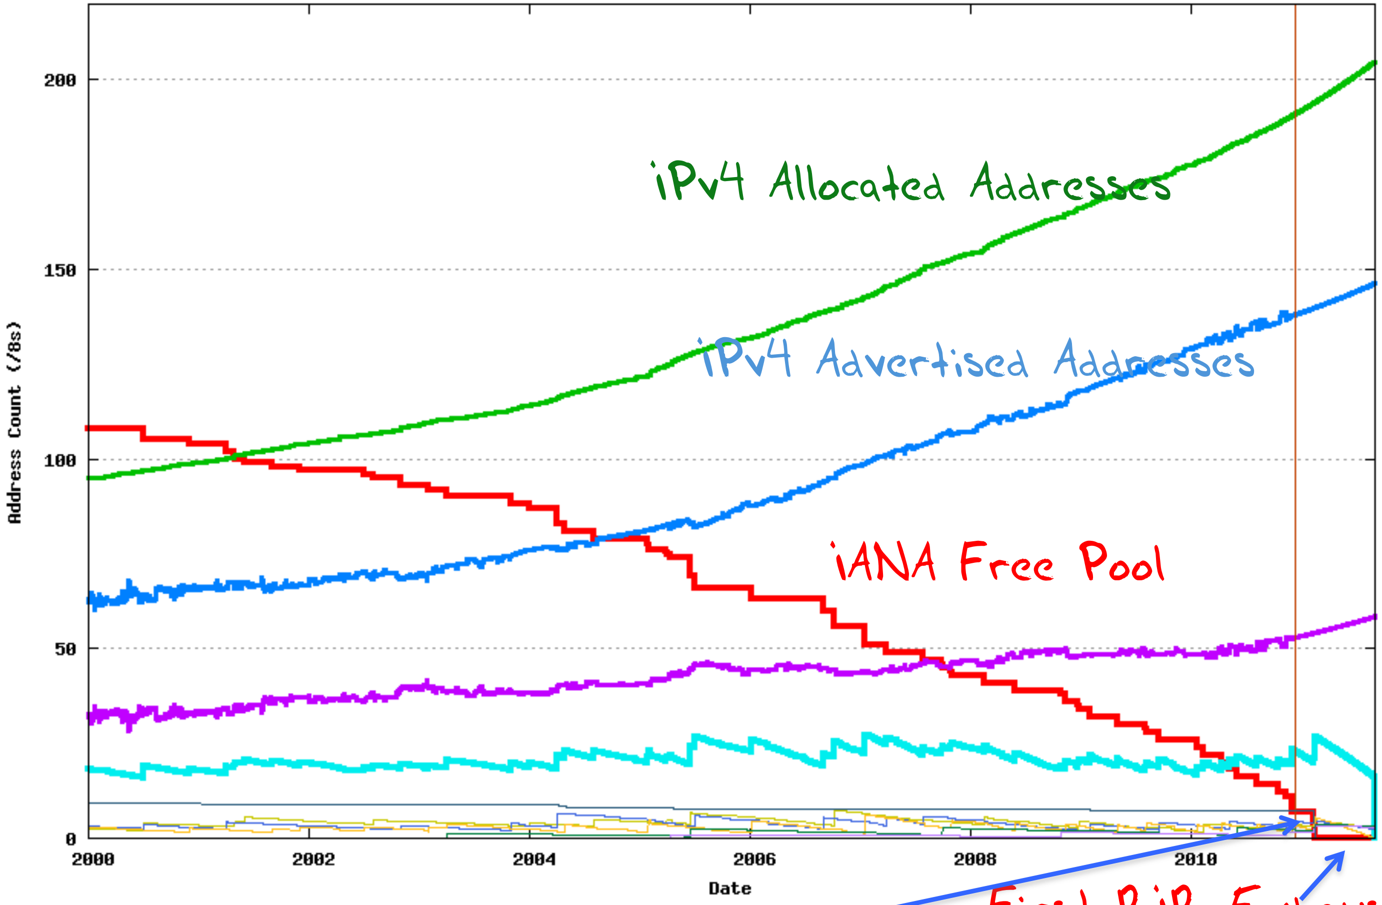

IANA Exhaustion February 2011

First RIR Exhaustion July 2011

Where were we with IPv6 deployment?

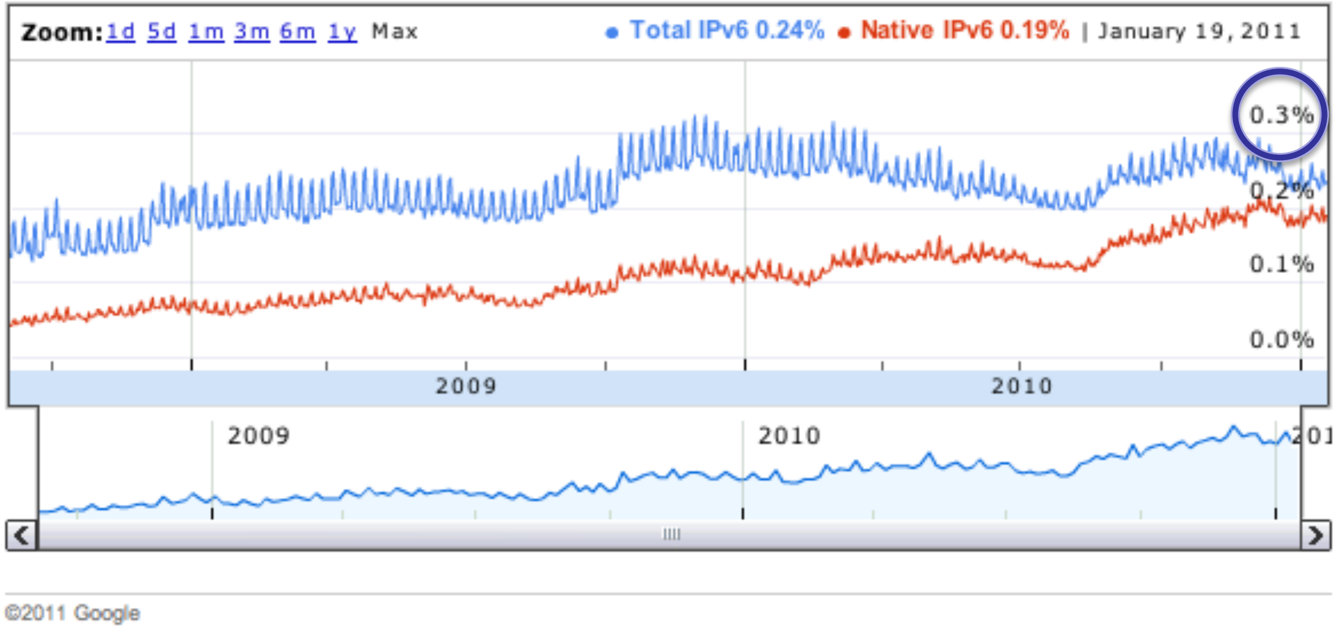
Zoom: 1d 5d 1m 3m 6m 1y Max
● Total IPv6 0.24% ● Native IPv6 0.19% | January 19, 2011
0.3%
0.2%
0.1%
0.0%
2009
2010
2009
2010
201
©2011 Google

http://www.google.com/intl/en/ipv6/statistics/

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 2011 IPv6 Transition Plan

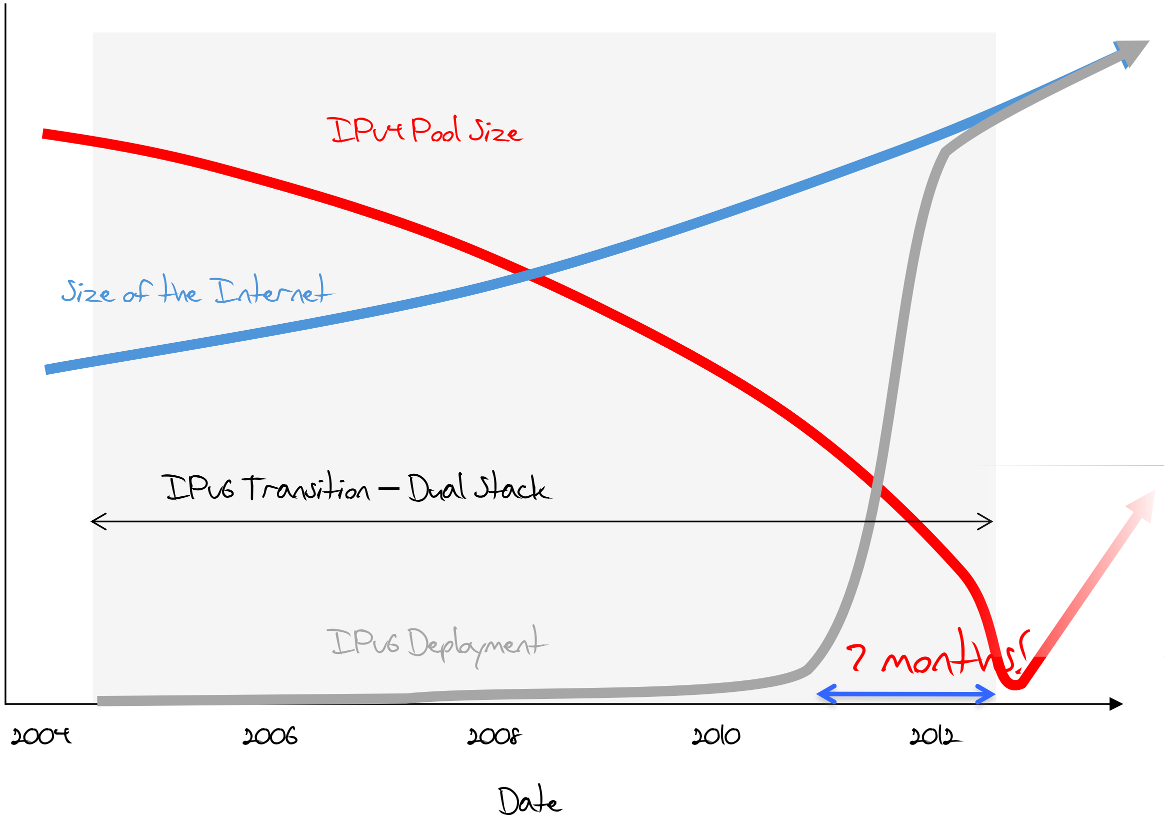

Was this Plan Feasible?

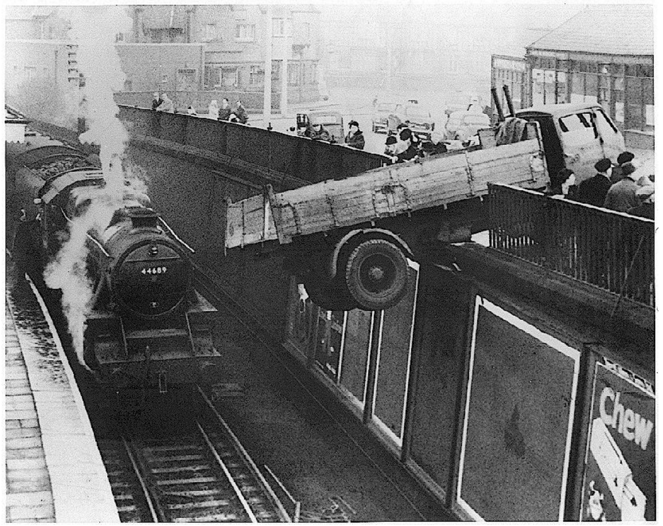

Deploy IPv6 across some 1.8 billion users, with more than a billion end hosts, and upgrade hundreds of millions of routers, firewalls, middleware units and CPEs, and audit billions of lines of configuration codes and filters, and audit hundreds of millions of ancillary support systems -

all within 200 days.

The 2012 IPv6 Transition Plan

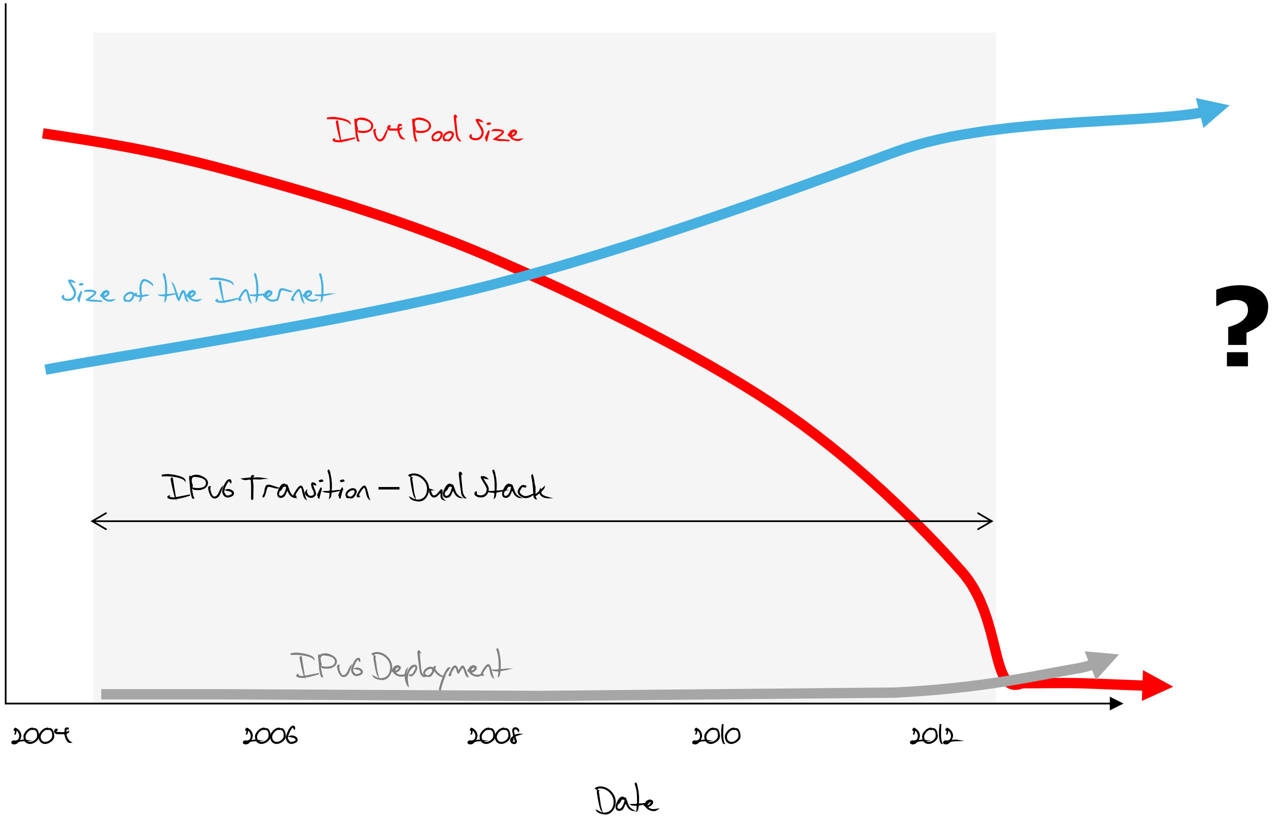

What next?

- 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

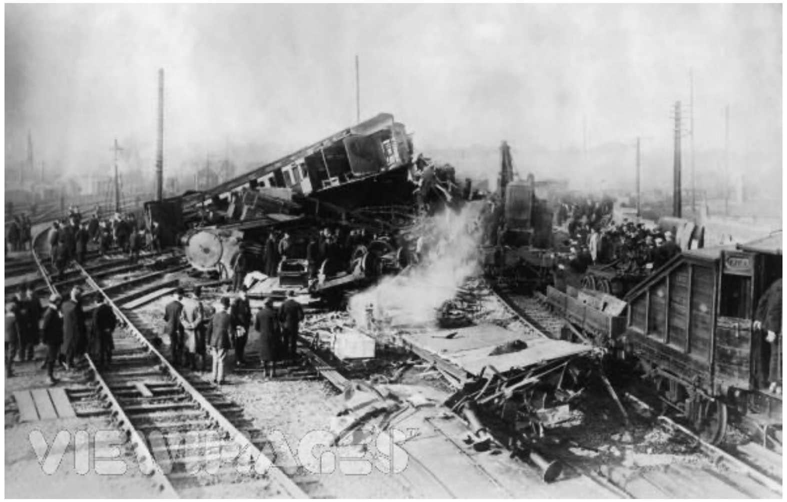

- This low friction response to IPv4 address depletion had been used for more than a decade in client/server network architectures
- **Clients** initiate a service transaction and only need an external address/port binding for the duration of the transaction
- **Servers** sit in central data centres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

Implications

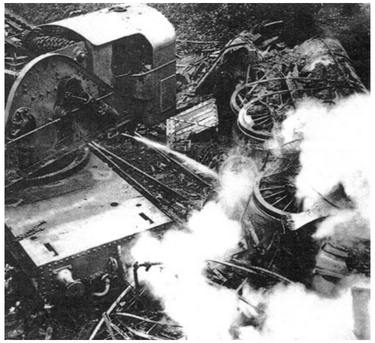

- IPv4 addresses continue to be in demand far beyond the date of exhaustion of the unallocated pool
  - In the transition environment, all new and expanding network deployments will need IPv4 service access and addresses for as long as we are in this dual track transition

- But the process is no longer directly controlled through today’ 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

Making IPv4 Last Longer

The Wreck of Maine Central Train No. 13, Oakland
(W.H. Bunting)

- For how long?
- For what cumulative address demand?
- For what level of fairness of access?
- At what cost?
- For whom?
- To what end?
- What if we actually achieve something different?
  - How would the Law of Unintended Consequences apply here?
  - Would this negate the entire “IPv6 is the solution” philosophy?

Why is the Internet wedged on IPv6 transition?

Not everyone is feeling the pressure to move to Dual Stack

Use of IPv6 for Northern America (XQ)

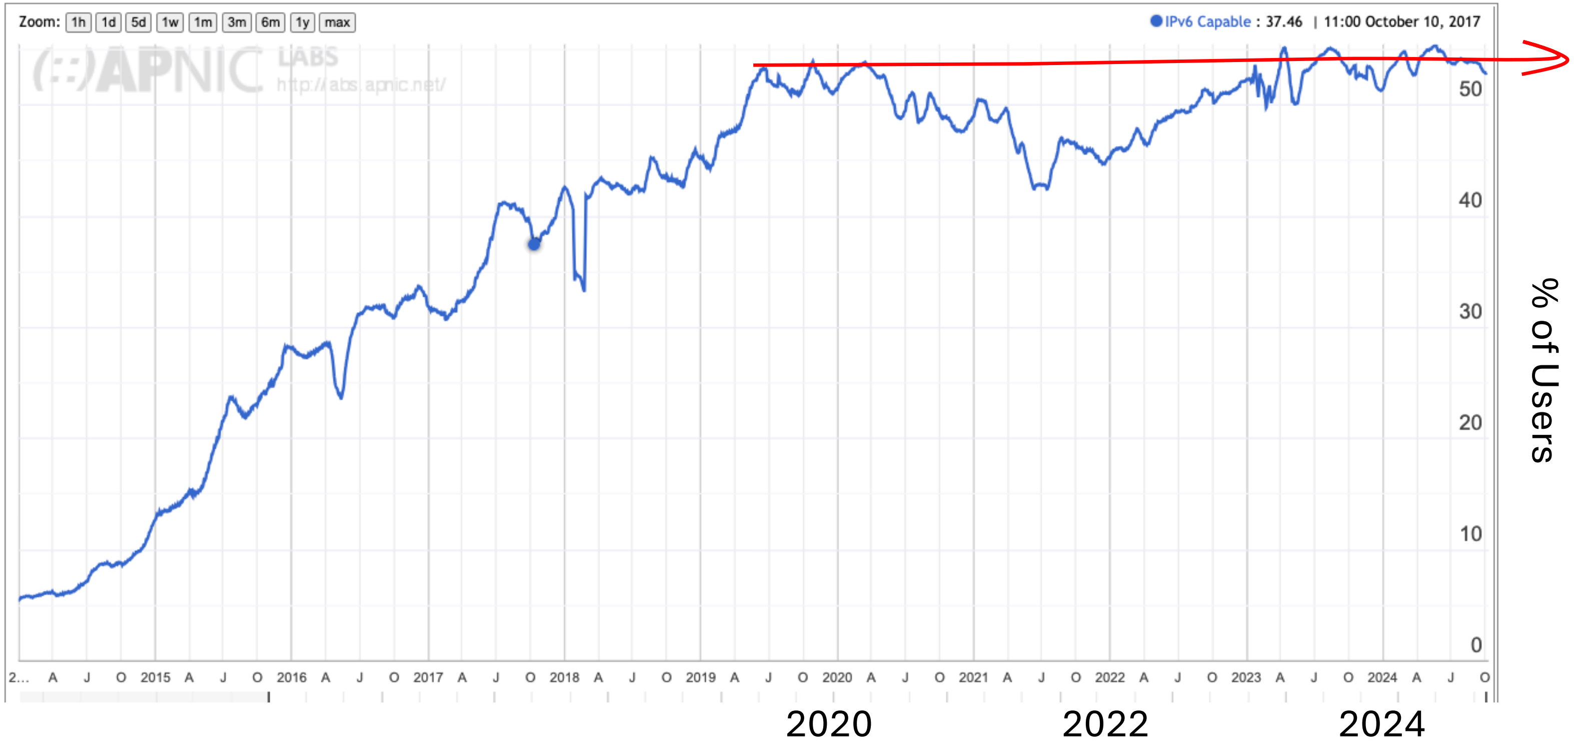

https://stats.labs.apnic.net/ipv6/XQ

Maybe this really is the IPv6 trajectory that we're now locked into

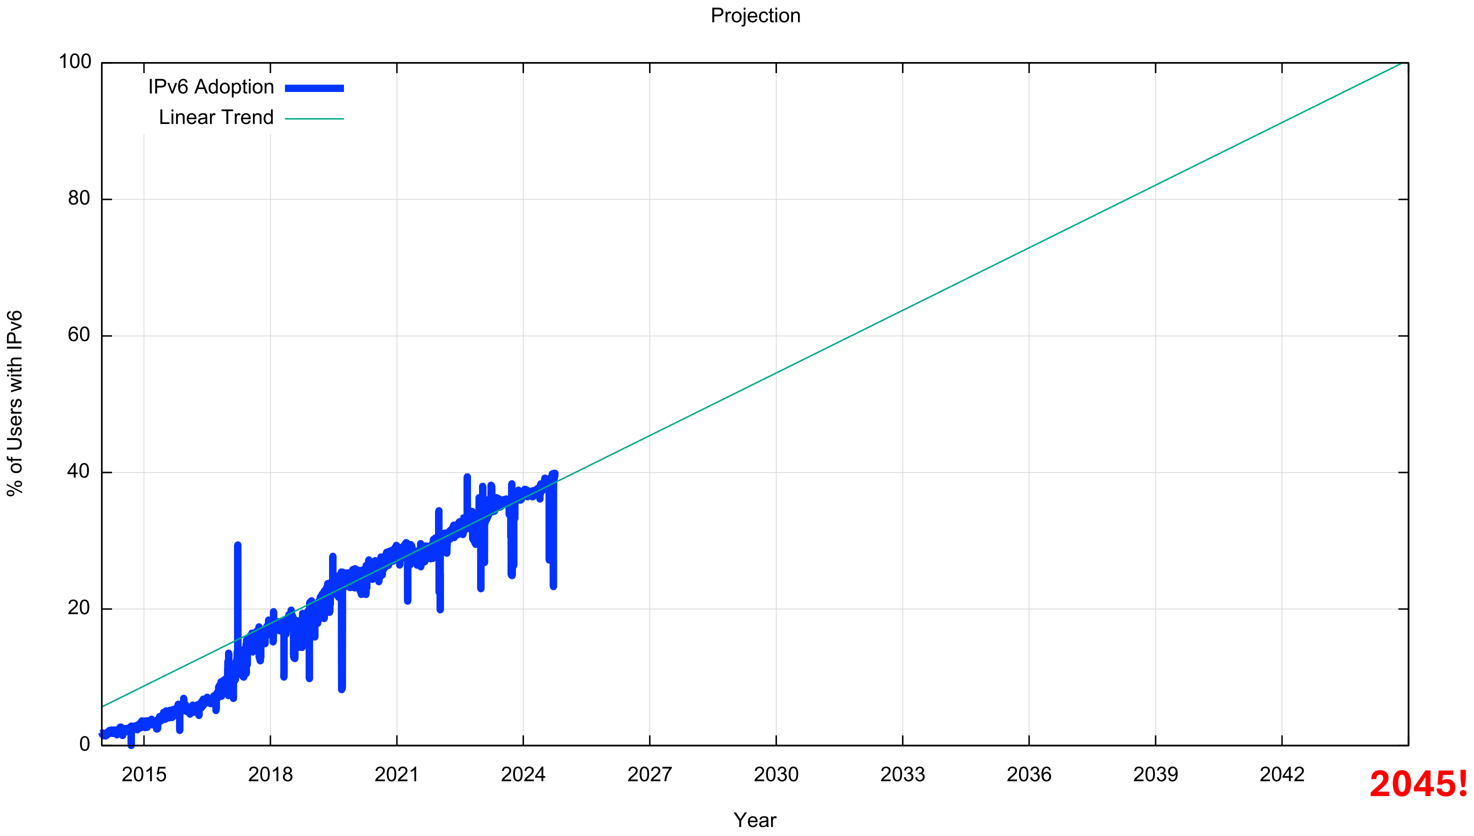

2045!

Why is the Internet wedged on IPv6 transition?

- 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

Because it wasn't just an IPv4 to IPv6 transition

Follow the money…

The "Classical" Internet

- IP was a ***network*** *protocol* that provided services to attached devices
- It was the role of Network Providers to allow clients to consume content and access services
  - The costs of operating the network dominated the entire cost of the Internet
  - In networking distance dominates all cost models
  - In the Internet ecosystem the role of transit providers was paramount
    - We used to spend all our time talking about peering and transit!
  - ISPs were the brokers of rationing the scarce resource of distance capacity

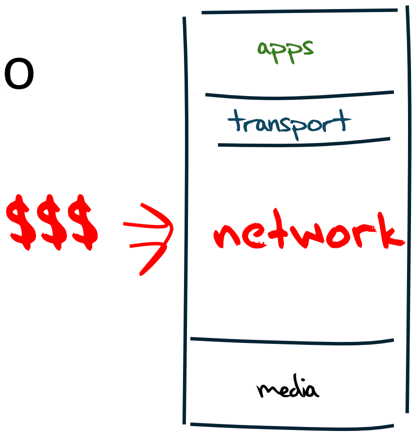

What's driving change today?

- **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

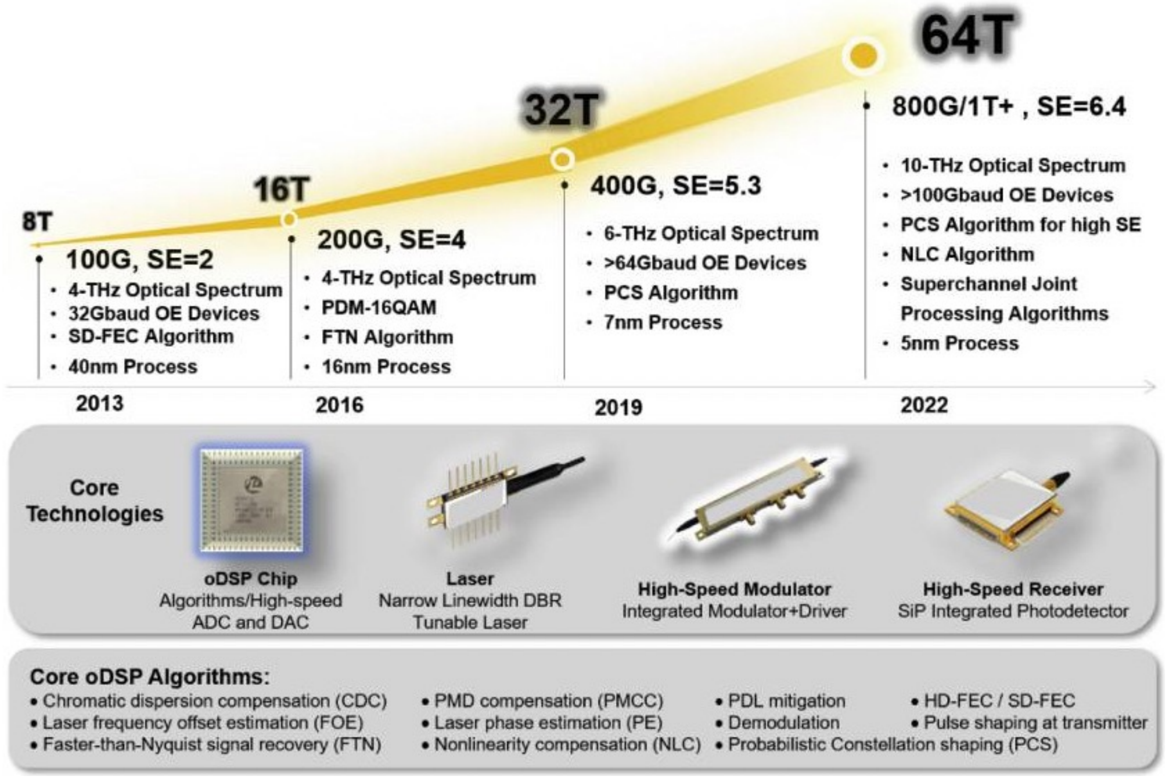

100G, SE=2
- 4-THz Optical Spectrum
- 32Gbaud OE Devices
- SD-FEC Algorithm
- 40nm Process

2013

200G, SE=4
- 4-THz Optical Spectrum
- PDM-16QAM
- FTN Algorithm
- 16nm Process

2016

400G, SE=5.3
- 6-THz Optical Spectrum
- >64Gbaud OE Devices
- PCS Algorithm
- 7nm Process

2019

800G/1T+ , SE=6.4
- 10-THz Optical Spectrum
- >100Gbaud OE Devices
- PCS Algorithm for high SE
- NLC Algorithm
- Superchannel Joint Processing Algorithms
- 5nm Process

2022

Core Technologies

oDSP Chip
Algorithms/High-speed ADC and DAC

Laser
Narrow Linewidth DBR Tunable Laser

High-Speed Modulator
Integrated Modulator+Driver

High-Speed Receiver
SiP Integrated Photodetector

Core oDSP Algorithms:
- Chromatic dispersion compensation (CDC)
- Laser frequency offset estimation (FOE)
- Faster-than-Nyquist signal recovery (FTN)
- PMD compensation (PMCC)
- Laser phase estimation (PE)
- Nonlinearity compensation (NLC)
- PDL mitigation
- Demodulation
- Probabilistic Constellation shaping (PCS)
- HD-FEC / SD-FEC
- Pulse shaping at transmitter

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

Moore's Law: The number of transistors on microchips doubles every two years

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important for other aspects of technological progress in computing – such as processing speed or the price of computers.

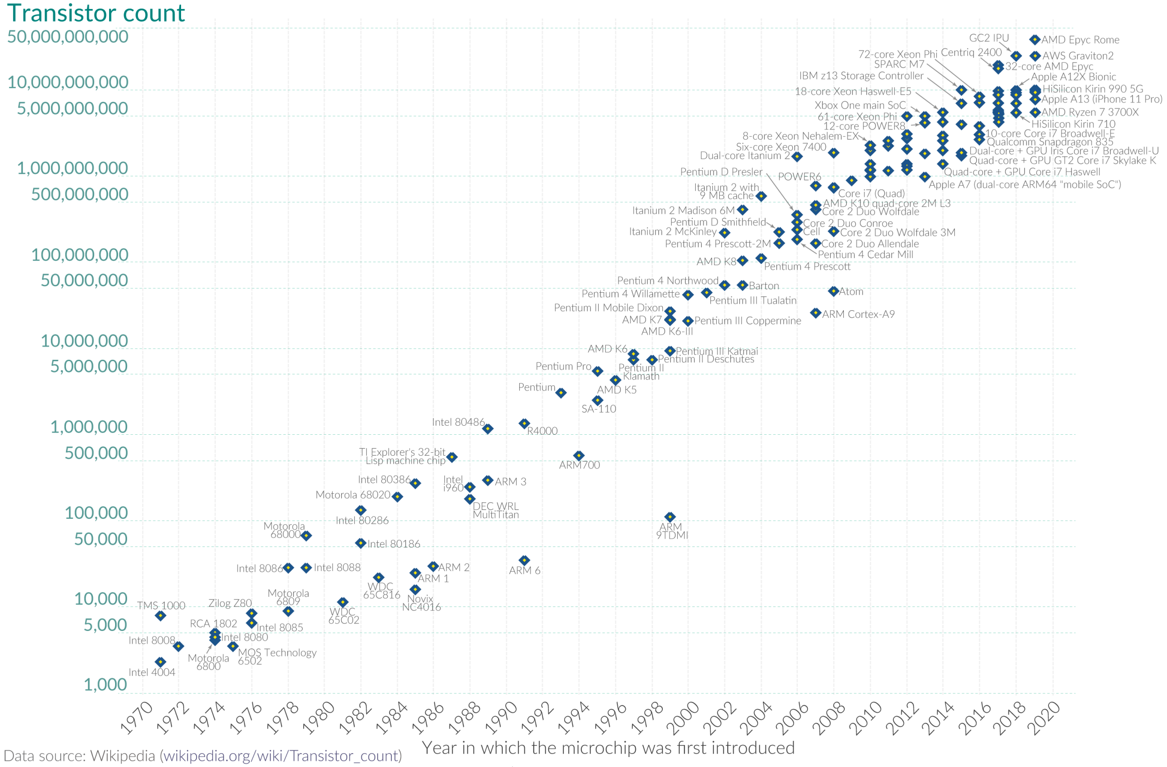

Data source: Wikipedia (wikipedia.org/wiki/Transistor_count)

OurWorldinData.org – Research and data to make progress against the world's largest problems.

Licensed under CC-BY by the authors Hannah Ritchie and Max Roser.

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

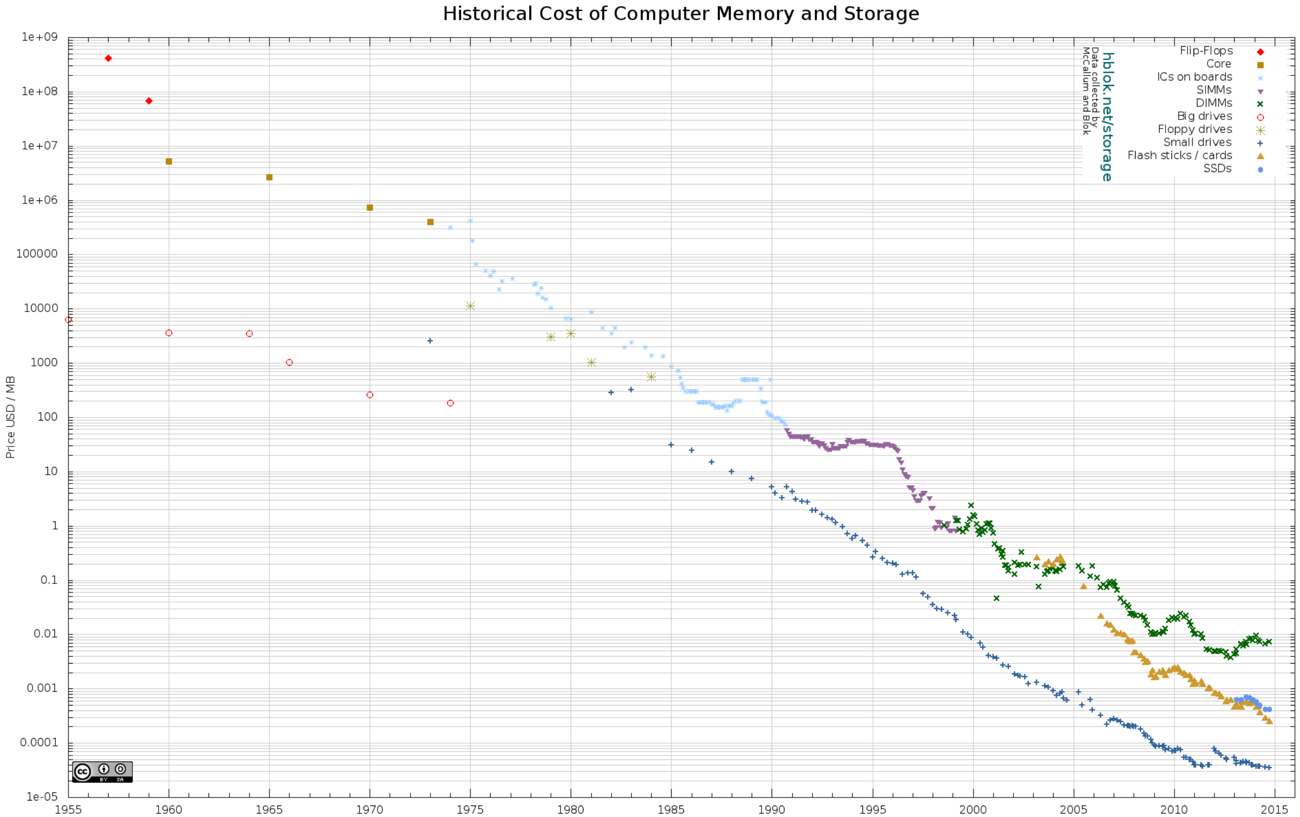

http://aiimpacts.org/wp-content/uploads/2015/07/storage_memory_prices_large-_hblok.net_.png

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

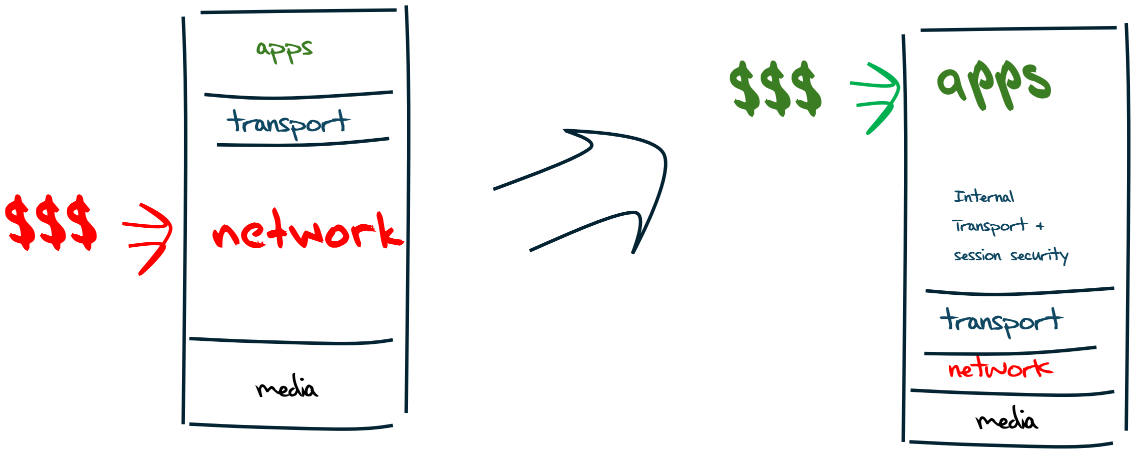

So, who needs to pay?

- Networks need to make the investment to switch to a dual stack mode that includes IPv6
- But neither the user base not 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 pull 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

It's Names that drive today's Internet

- TLS is the only underpinning of authenticity in the network
  - DNSSEC is largely a market failure
  - RPKI really does not matter in a routing-free network!
- TLS is a name-based framework for validating authenticity
- These days its DNS outages can be globally catastrophic
  - While address and routing outages are generally just annoying!
- “Evilness” and “DNS abuse” are homonyms these days!

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 Pv6 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 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

Today's Internet:

- Names Matter
- The DNS Matters

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

Longer Term Trends?

Pushing EVERYTHING out of the network and over to applications

- Transmission infrastructure is becoming an abundant commodity
  - Network sharing technology (multiplexing) is decreasingly relevant
- We have so much network and computing that we no longer have to bring consumers to service delivery points - instead, we are bringing services towards consumers and using the content frameworks to replicate servers and services
- With so much computing and storage **the application is becoming the service**, rather than just a window to a remotely operated service

Do Networks matter any more?

- We have increasingly stripped out network-centric functionality in our search for lower cost, higher speed, and better agility
- We are pushing functions out to the edge and ultimately off “the network” altogether and what is left is just dumb pipes?
- What defines “the Internet”?
    - A common shared transmission fabric, a common suite of protocols and a common protocol address pool?

  or

    - A disparate collection of services that share common referential mechanisms using a common name space?

Some issues to think about

What matters today?

- **End Point Addressing** – IPv4 / IPv6 / IPv? Absolute? Relative?
  - Is universal unique end-point addressing a 1980’s concept whose time has come and gone?
  - If network transactions are localised, then what is the residual role of unique global end point addressing for clients or services?
  - And if we cannot find a role then why should we bother?
  - Who decides when to drop it?
    - Is this a market function, so that a network that uses local addressing can operate from an even lower cost base gains a competitive market edge?
    - Or are carriage services so cheap already that the relative benefit in discarding the last vestiges of unique global addresses are so small that it’s just not worth bothering about?

Some issues to think about

What matters today?

- **Naming and Name Spaces** – DNS evolution?
  - Are ”names” a common attribute of the network, or an attribute of a service environment or application realm?
  - Should names be persistent over time?
  - Is the resolution of a name absolute or relative to the content of the resolution query?
  - In a world of densely replicated service delivery points how does a client rendezvous with the “best” service point? Does the client work it out? Or the network? Or the service?
  - If names are an attribute of applications, then why do we need a single name domain? Surely each application realm can define its own name space? How can we associate a referential name space with a given name?
  - If both names and addresses are ephemeral and unstable then what defines the Internet?

Some issues to think about

What matters today?

- **Referential Frameworks?**
  - Without a common referential space then how do we usefully communicate?
  - What do we mean by “common” when we think about referential frameworks?
  - How can we join the ‘fuzzy’ human language spaces with the tightly constrained deterministic computer-based symbol spaces?

Thank You!