IPv6

Geoff Huston
APNIC

Why?

Because we've run out of addresses

Again.

We've been here before ...

The original ARPAnet design from 1969 used the NCP protocol, which used 8 bit addresses in the NCP packet header

– Maximum network span of 256 nodes
– Enough?

– Well yes, because at the time computers were the size of entire rooms, cost many millions of dollars and there were only a few hundred in the entire world.
– At the time shrinking the computer to something you could hold in one hand and trying to connect billions of them together was just too far into the future to worry about

ARPAnet iMP

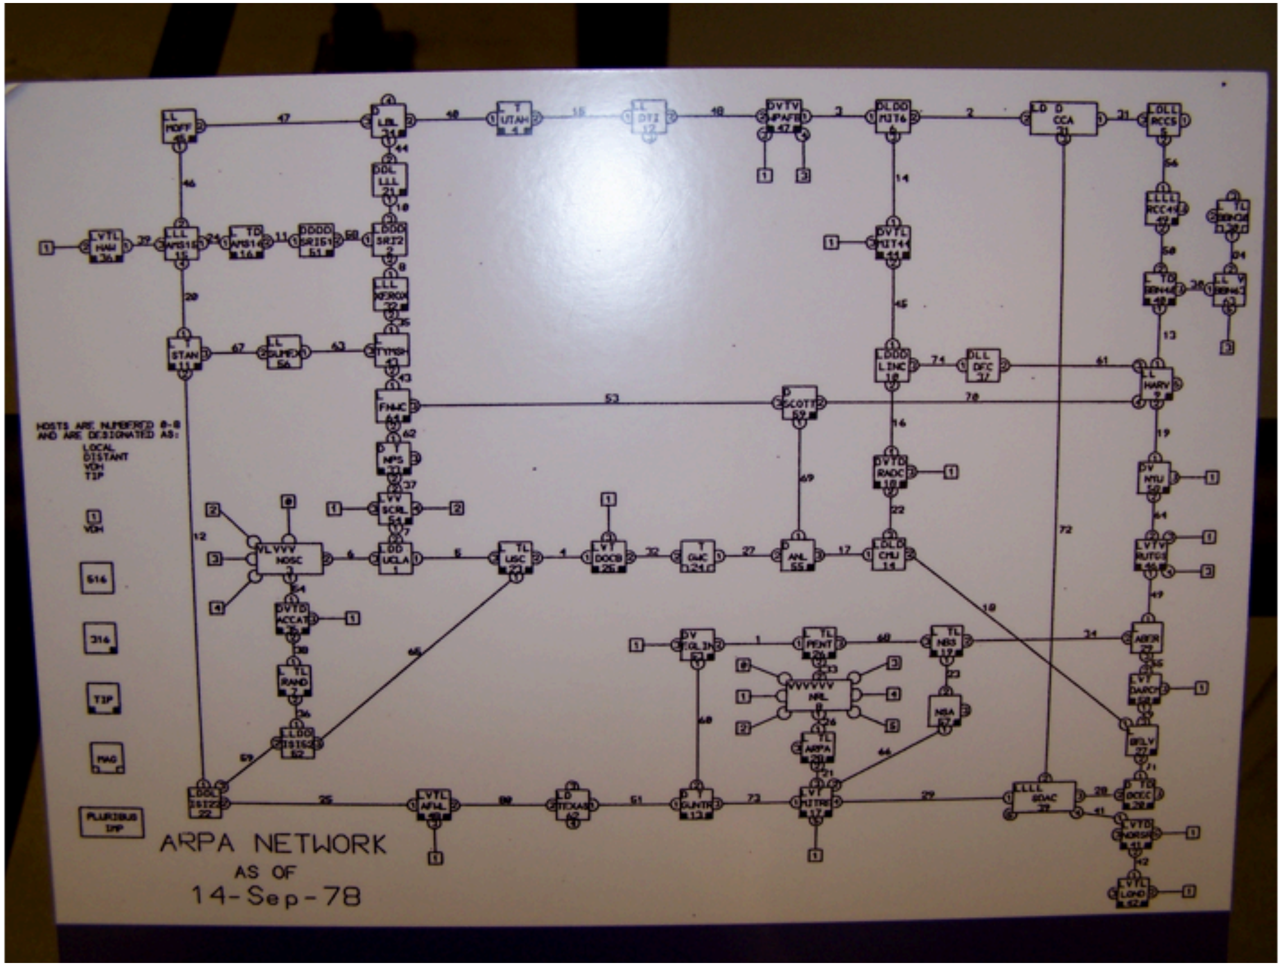

ARPAnet - September 1978

Transition V1.0

- Turns out that 8 bits of addresses was not enough for the next generation of mini computers
- ARPAnet undertook a transition from NCP to a new protocol: TCP/IP
  - Expansion from 8 to 32 bit addresses
  - Flag Day: 1 January 1983
  - Shutdown and reboot every node into the new protocol

"This time, for sure!" *

* Actually Vint didn't say this!

IP Version 4

- 32 bit address field
  - That's 4,294,967,296 addresses
- We've used this to build today's Internet:
  - Some 400,000 networks
  - Around 1.5 billion connected devices
- Some 29 years later, we've run out of addresses - again!

IPv4 Address Allocations

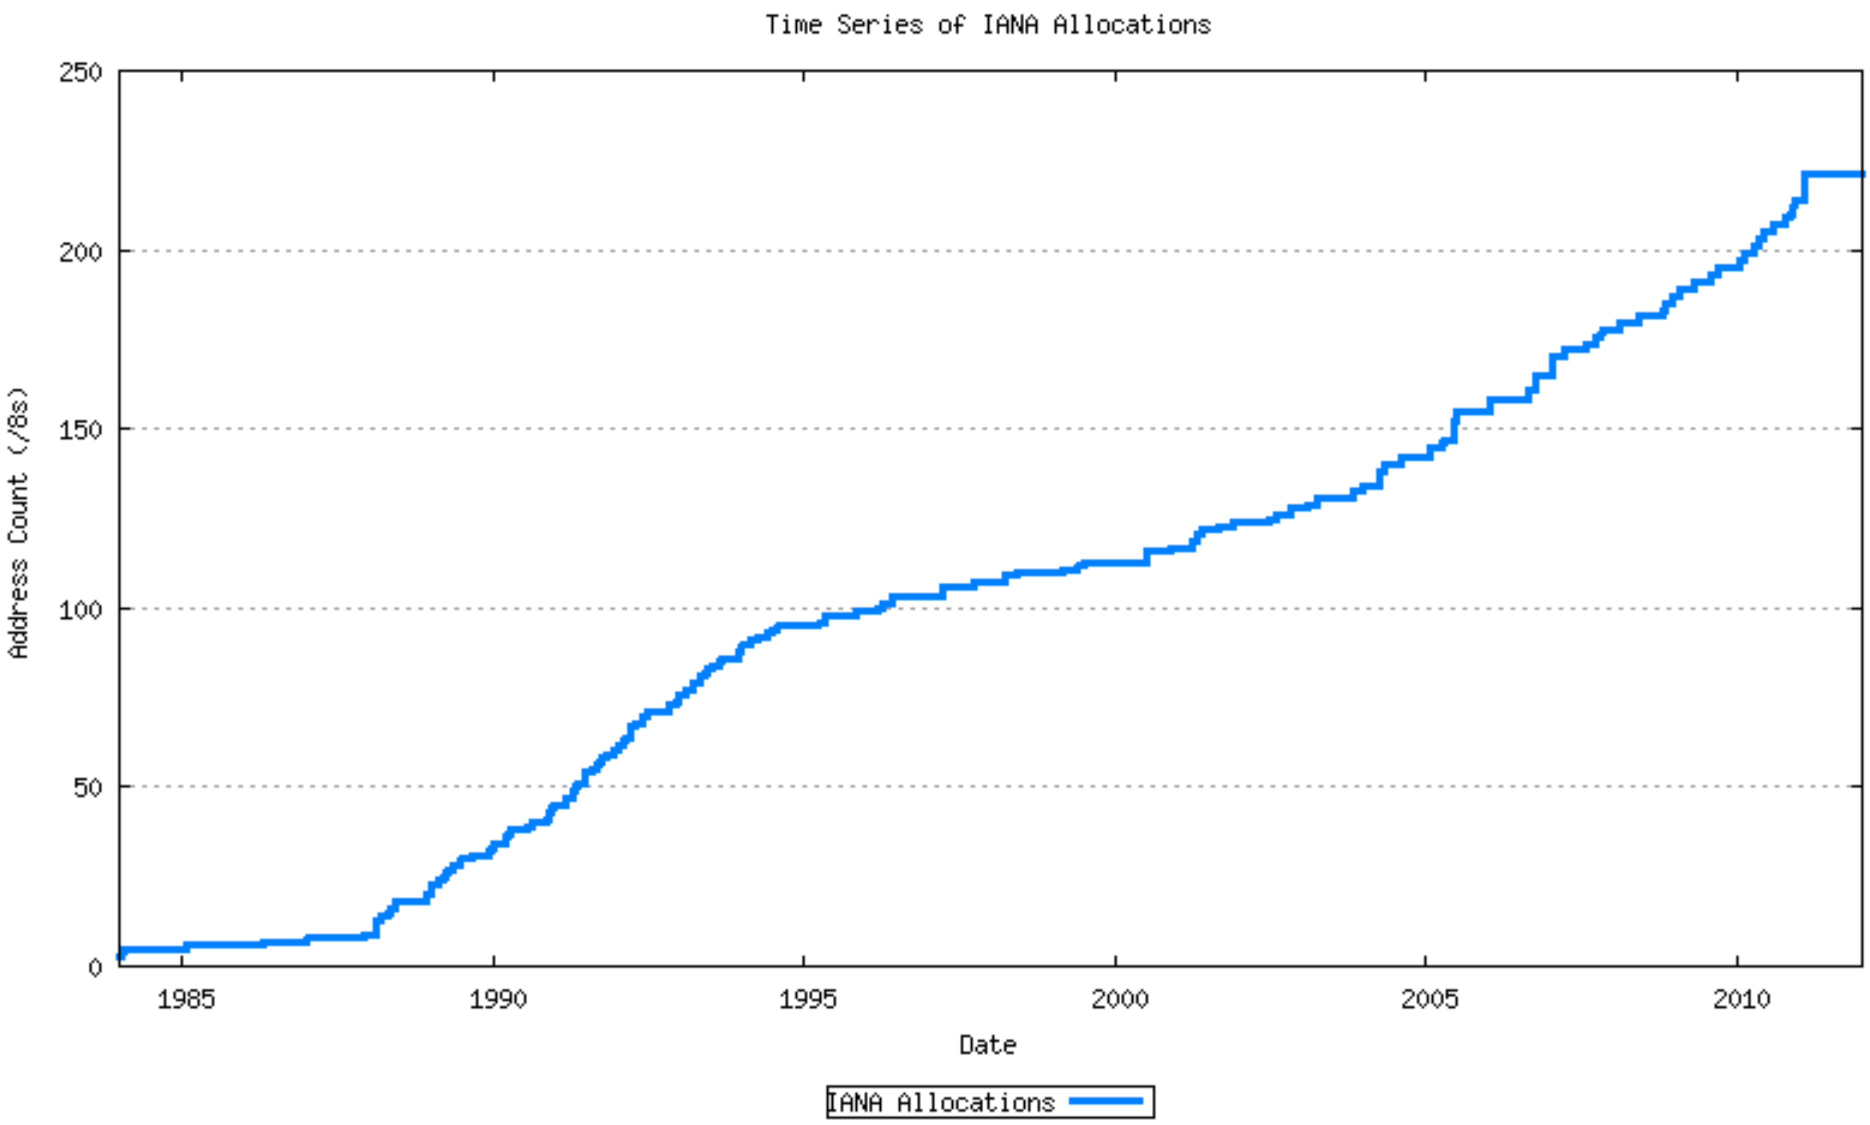

IPv4 Address Allocations

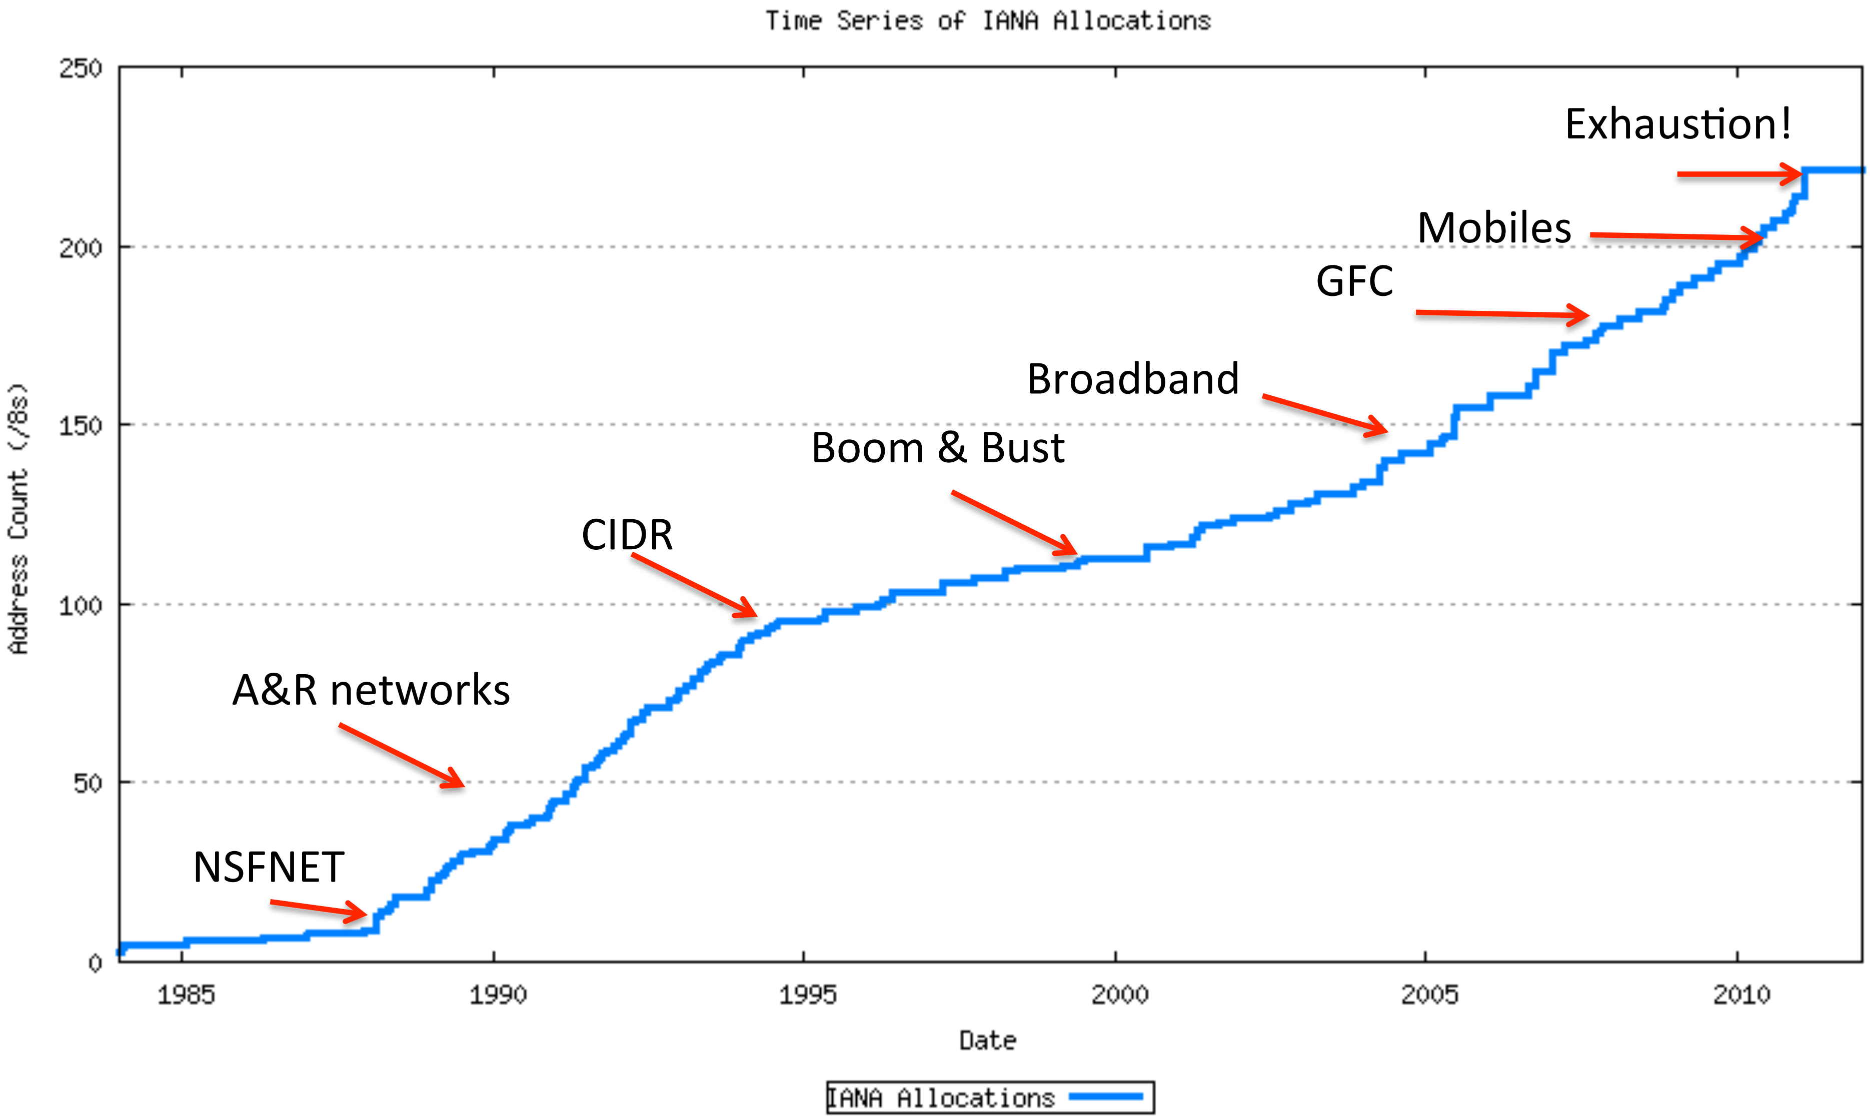

iPocalypse?

So we'll hit the wall – right?

Maybe not...

Many ISPs have been stockpiling IPv4 addresses in anticipation of this run out event

In many networks address "recovery" programs are underway

So it's not a sudden halt to the Internet

But the addressed part of the network grew by more than 250 million services in 2010

- Which was the largest year so far for the Internet

it's more like this!

What are we transitioning to?

IPv6!

Layer-3 Protocol Surgery

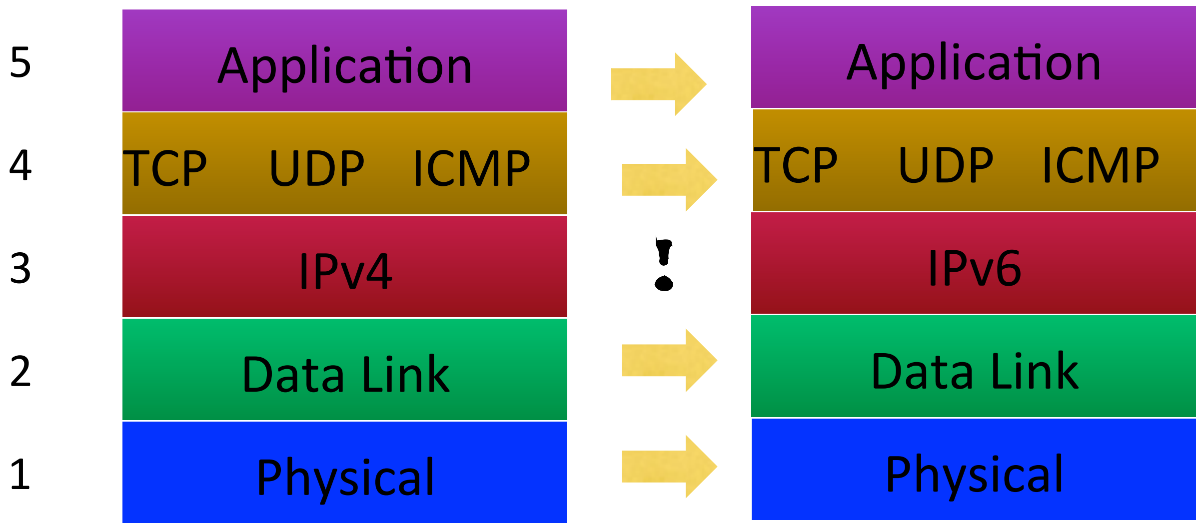

Only the IP layer changes – nothing else!

What changes with IPv6?

What changes with IPv6?

- 128 bit address fields
- Fixed host/network boundary
- Replace Broadcast and ARP with Multicast and SLAAC
- Removed on-the-fly fragmentation with ICMP6 notification to source
- No NATS!
- Multi-Addressing
- Scoped Addresses

What's giving us grief?

- 128 bit address fields
- Fixed host/network boundary
- Replaced Broadcast and ARP with Multicast and SLAAC
  - But we need to retain DHCP for DNS auto-config
- Removed on-the-fly fragmentation with ICMP6 notification to source
- No NATS!
- Multi-Addressing
- Scoped Addresses
- No Backwards Compatibility

Technology Considerations

- For simple LANs it is possible to “just turn it on”
  - Although the lack of a NAT can be an issue in terms of shifting from default to explicit security rules in firewalls
- For more complex networks IPv6 requires careful engineering
  - Particularly around prefix delegation
  - And firewall configuration
- And the dual stack environment introduces a whole new set of application problems

Transition, the second time around

- A "Flag Day" switchover is impossible
- Piecemeal replacement won't work either as IPv6 is not backward compatible with IPv4
- So we need to run both protocols in tandem "for a while"
- But bear in mind that one protocol has already run out of addresses
- And network growth continues at record levels

Transition, the second time around

We need to :

- deploy IPv6 in parallel with IPv4
- deploy ever more stringent IPv4 address conservation measures within the network
- allow the network to expand at an ever increasing rate

All at the same time!

Maybe it's like this!

Why is this so hard?

The IPv6 Transition Plan

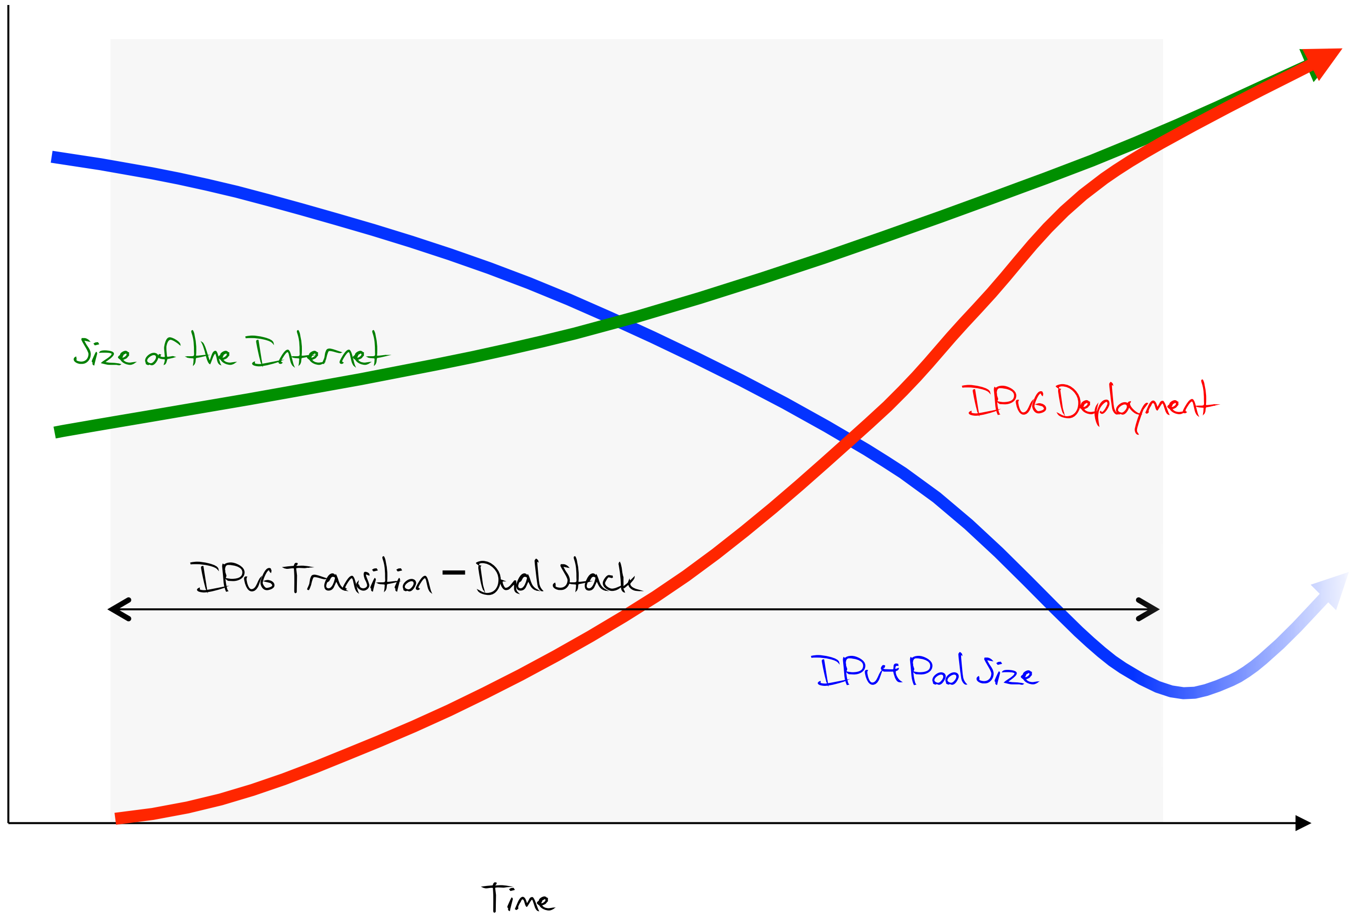

IPv4 Depletion

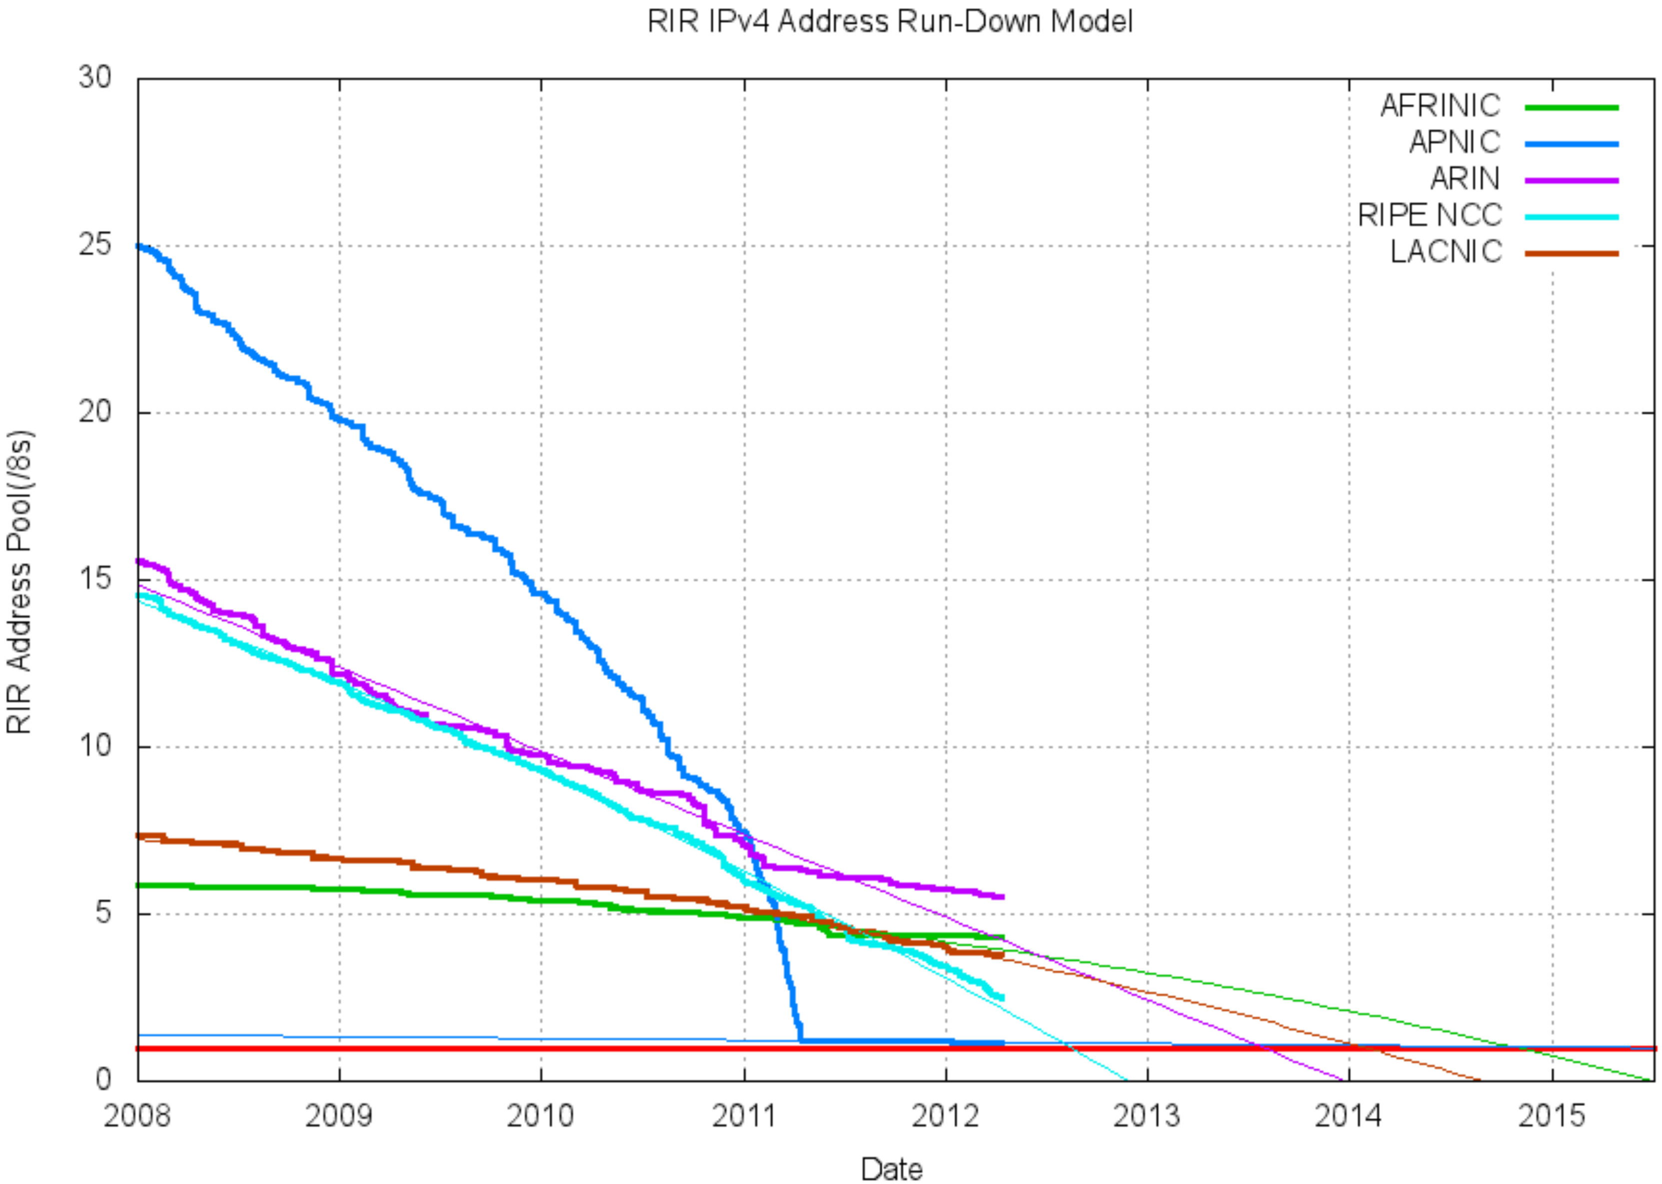

A Census of the Network Edge

- Counting IPv6 in client devices:
  - Some **45**% of devices run Windows Vista or Windows 7 - with IPv6 turned on
  - Some 8% of devices run Mac OS X - with IPv6 turned on
  - Some 35% of devices run Windows XP
- About half of the devices out there have IPv6 installed and active
  - And a large proportion of the other half are probably running Windows XP

http://en.wikipedia.org/wiki/Usage_share_of_operating_systems

A Census of the Network Core

- 4,882 ASNs originate IPv6 prefixes (out of a total of 39,535 ASNs in the IPv4 routing table)
- But 33,909 ASNs are stubs and 5,626 ASNs are transit
- So 49% of the IPv4 transit ASNs in routing space originate IPv6 prefixes

http://bgp.potaroo.net/v6/as2.0/

IPv6 capability, as seen by Google

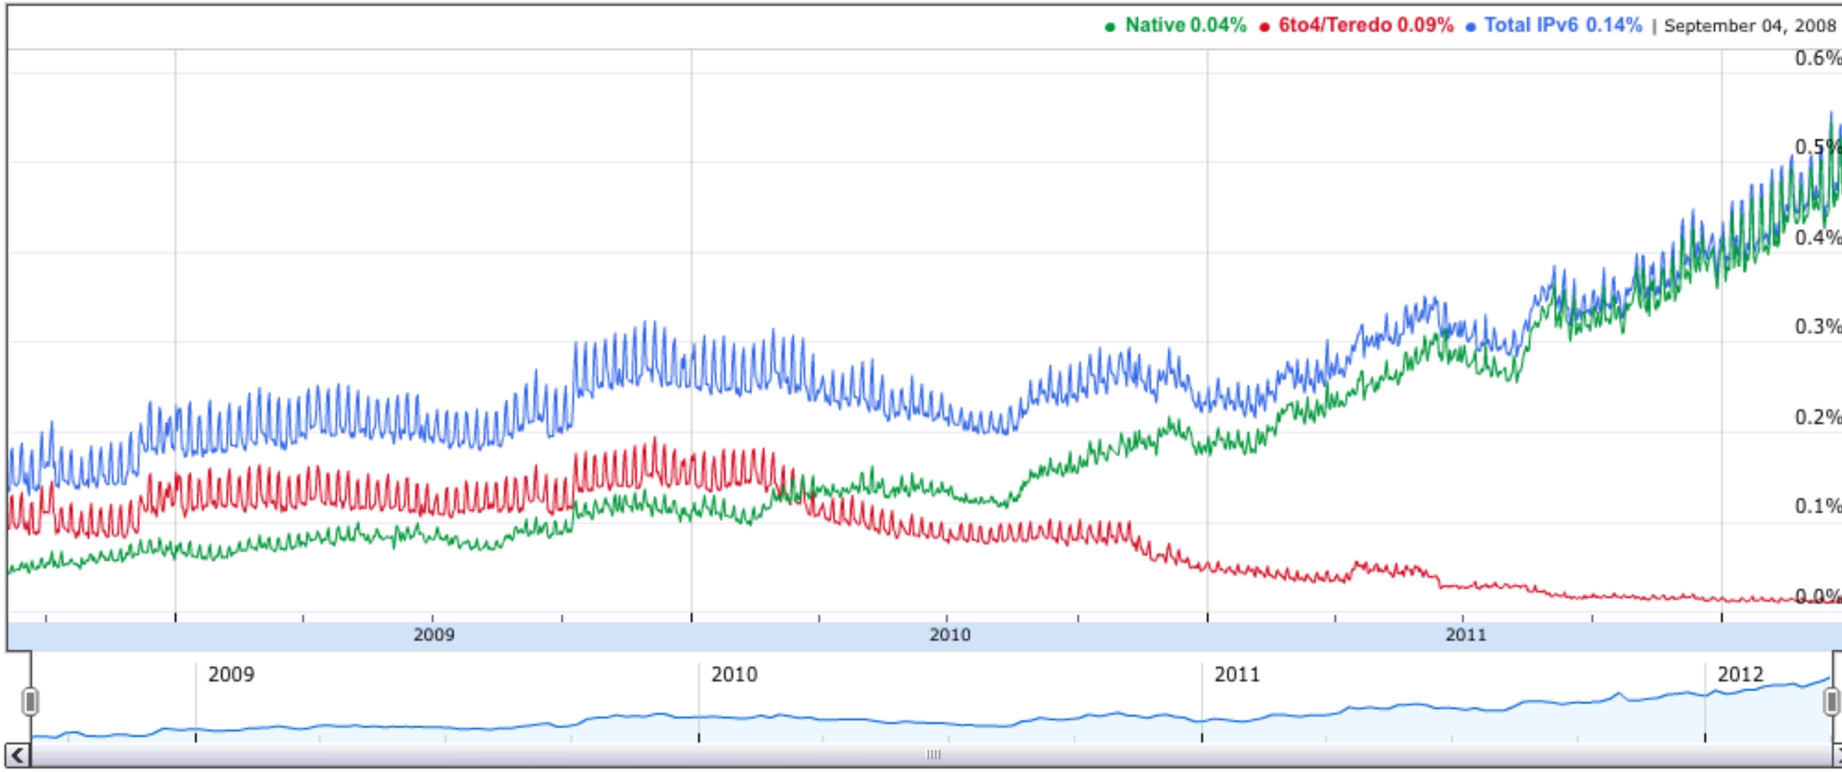

http://www.google.com/intl/en/ipv6/statistics/

IPv6 capability, as seen by APNIC

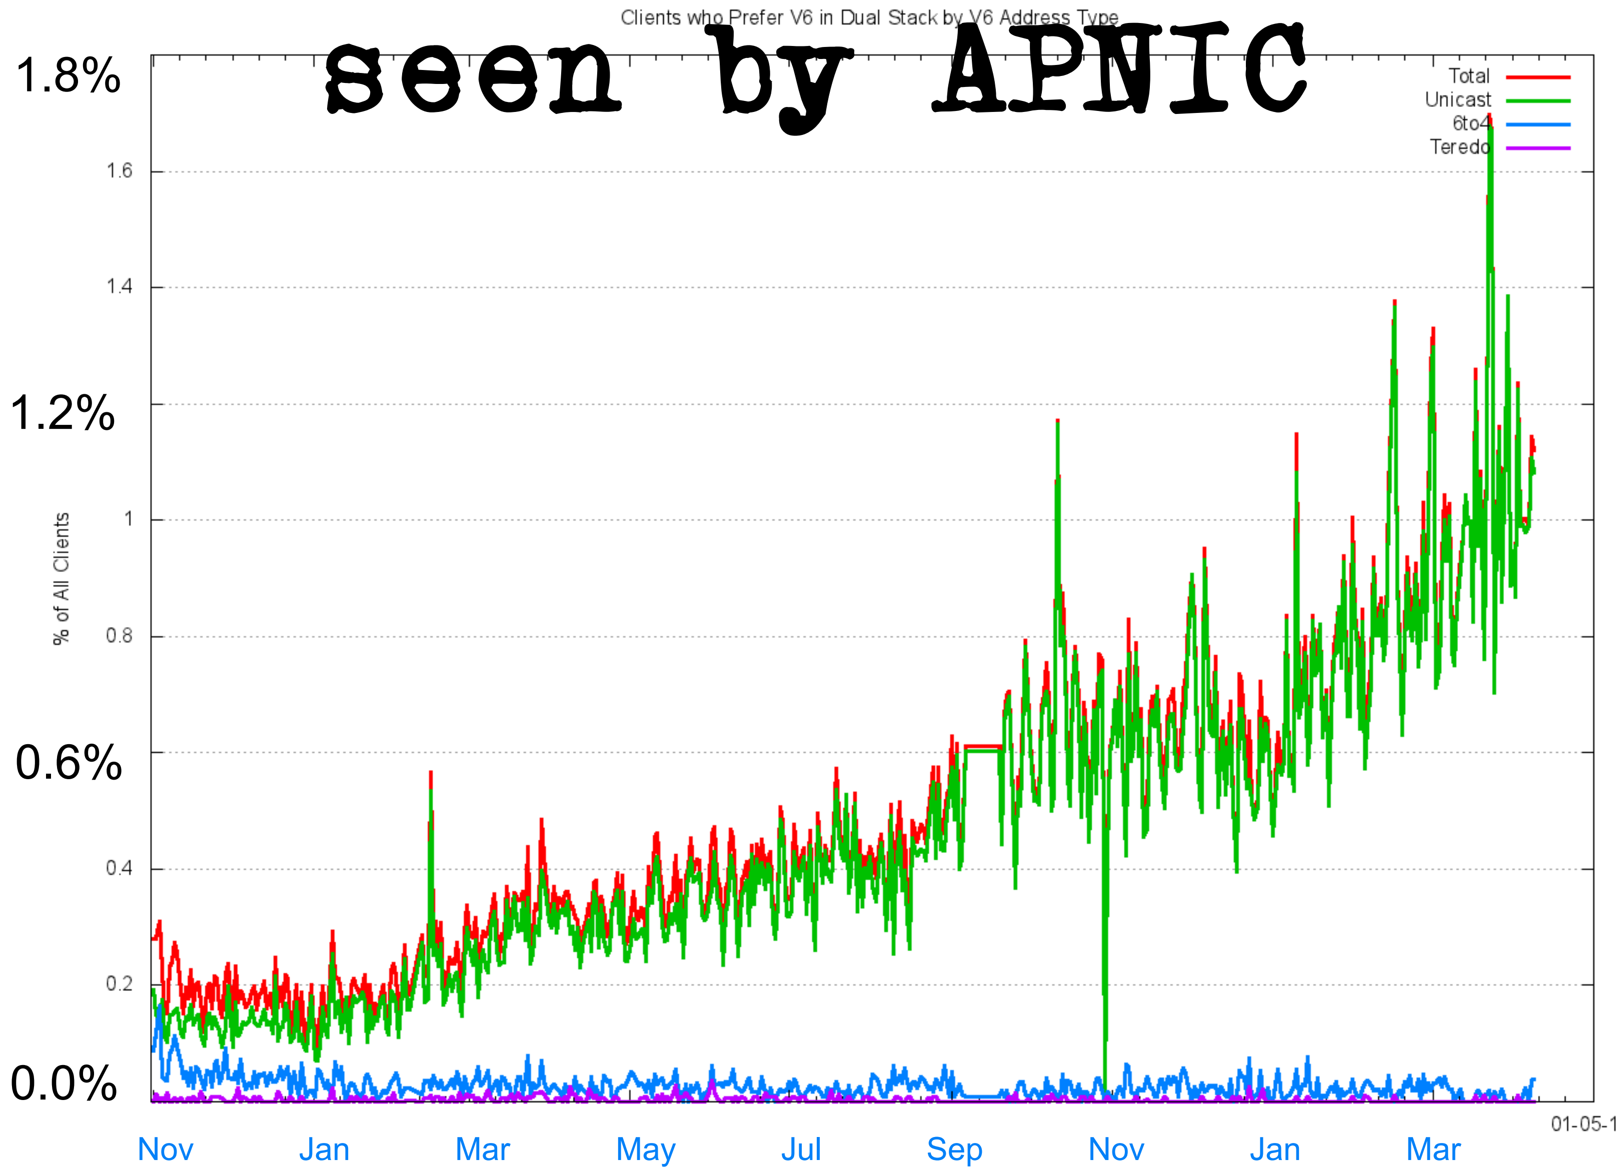

http://www.potaroo.net/stats/1x1/sitec/

Ooops!

- Access – 0.5% of end clients are served with an IPv6 access service that provides the client with a native IPv6 unicast address
- Services – 0.7% of the Alexa top 1M web sites have AAAA records

The IPv6 Transition Plan - V2.0

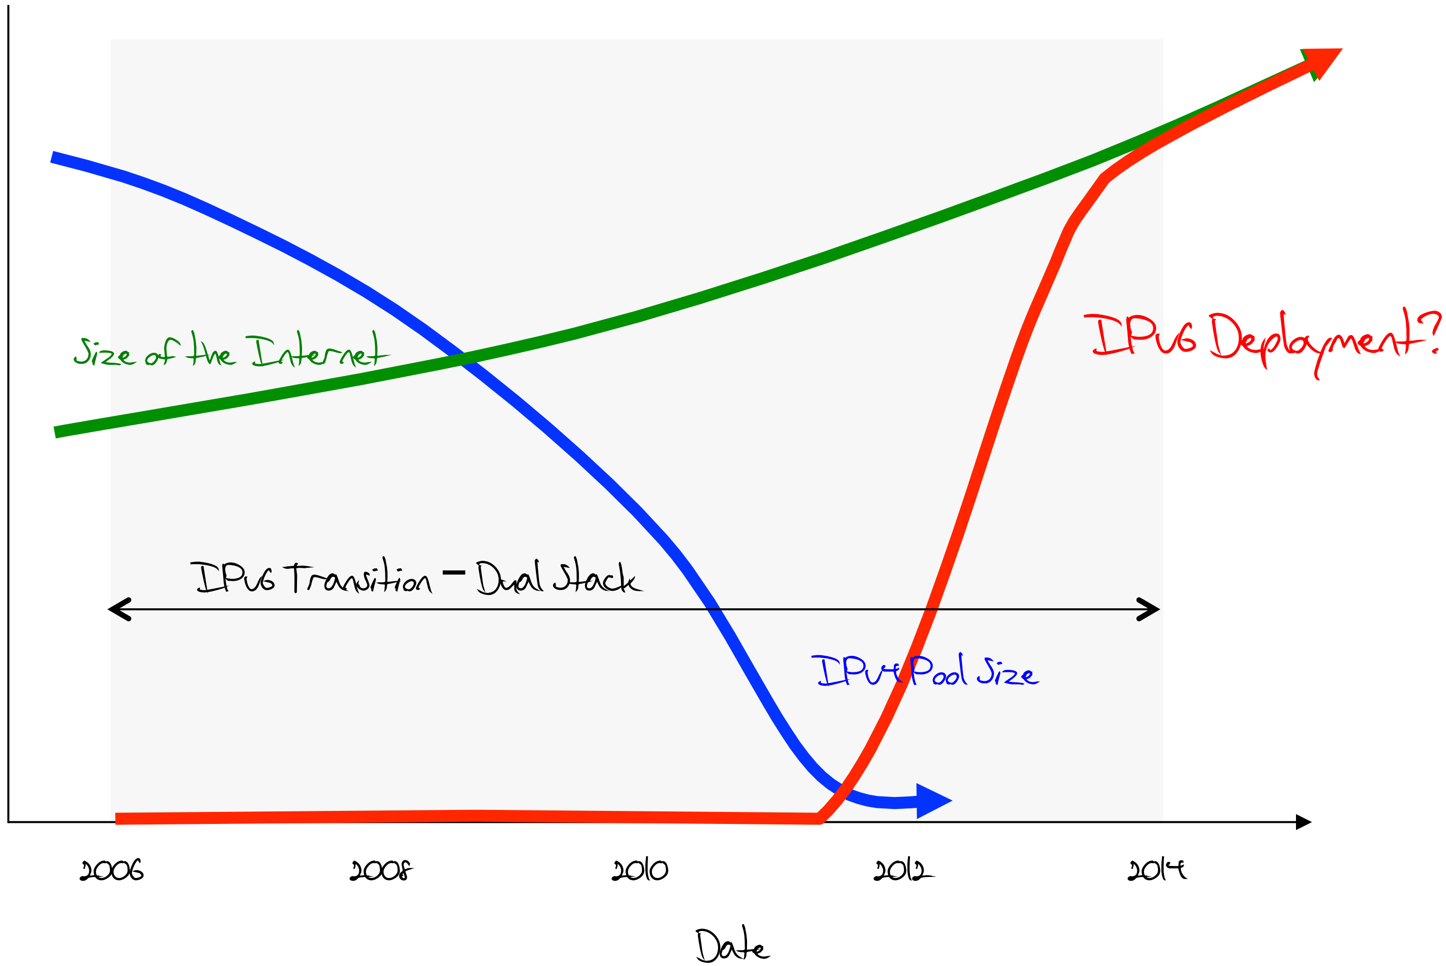

What's gone wrong?

- It seems that we've managed to achieve only 2 out of 3 objectives for IPv6 deployment
- And now the access industry has to deploy (and fund) IPv4 address extension mechanisms in addition to funding an IPv6 deployment
- What's going wrong in this gap between core and edge?
  - Why has the access service sector been disinterested in any meaningful levels of IPv6 deployment so far?
  - Why is the content industry lagging on IPv6 deployment?

Lessons from the Past

If this transition to IPv6 is proving challenging, then how did we ever get the IPv4 Internet up and running in the first place?

IPv4 Deployment Lessons

Technology: packet switching vs circuit switching

- lower network costs though pushing of functionality and cost to end systems exposed a new demand schedule for communications services

i.e. packet switching was far cheaper than circuit switching. This drop in cost exposed new market opportunities for emergent ISPs

The Demand Schedule

Higher prices reduce demand, while the higher margins incent higher volumes of supply

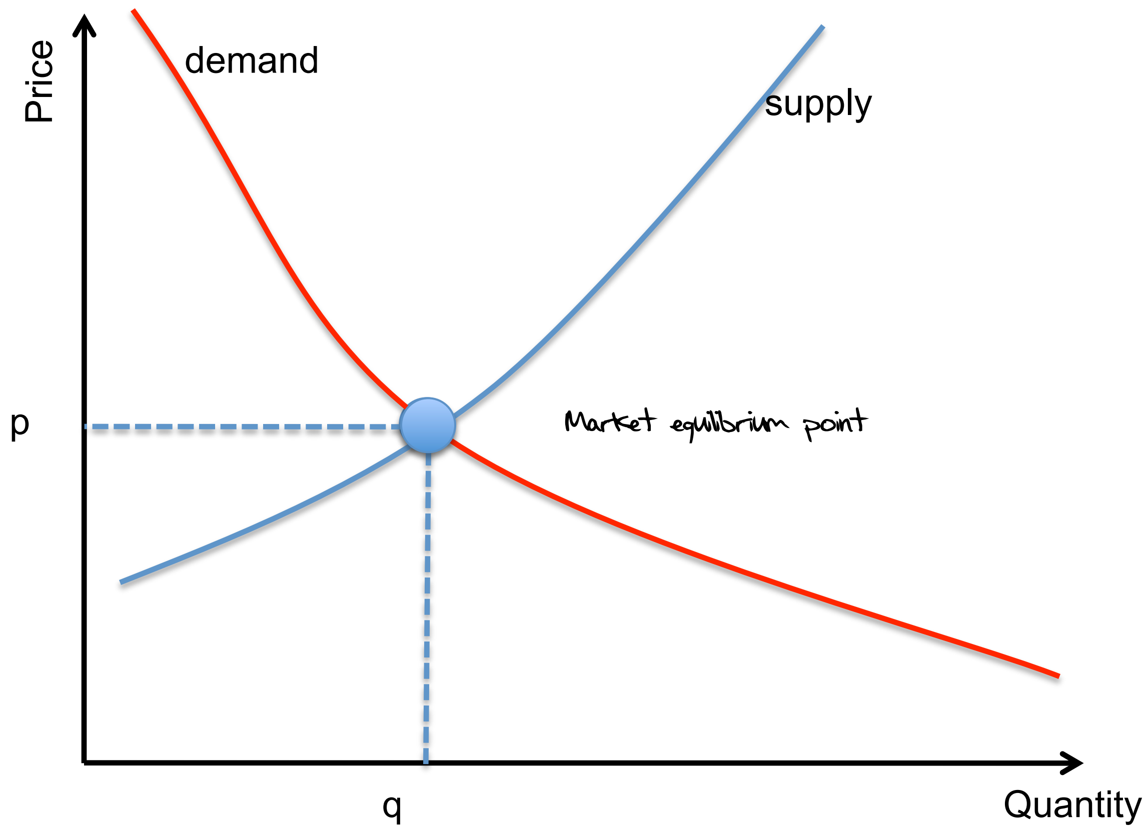

Lower prices increase Demand, but reduce the incentive to increase supply

Circuits to Packets: The Demand Schedule Shift

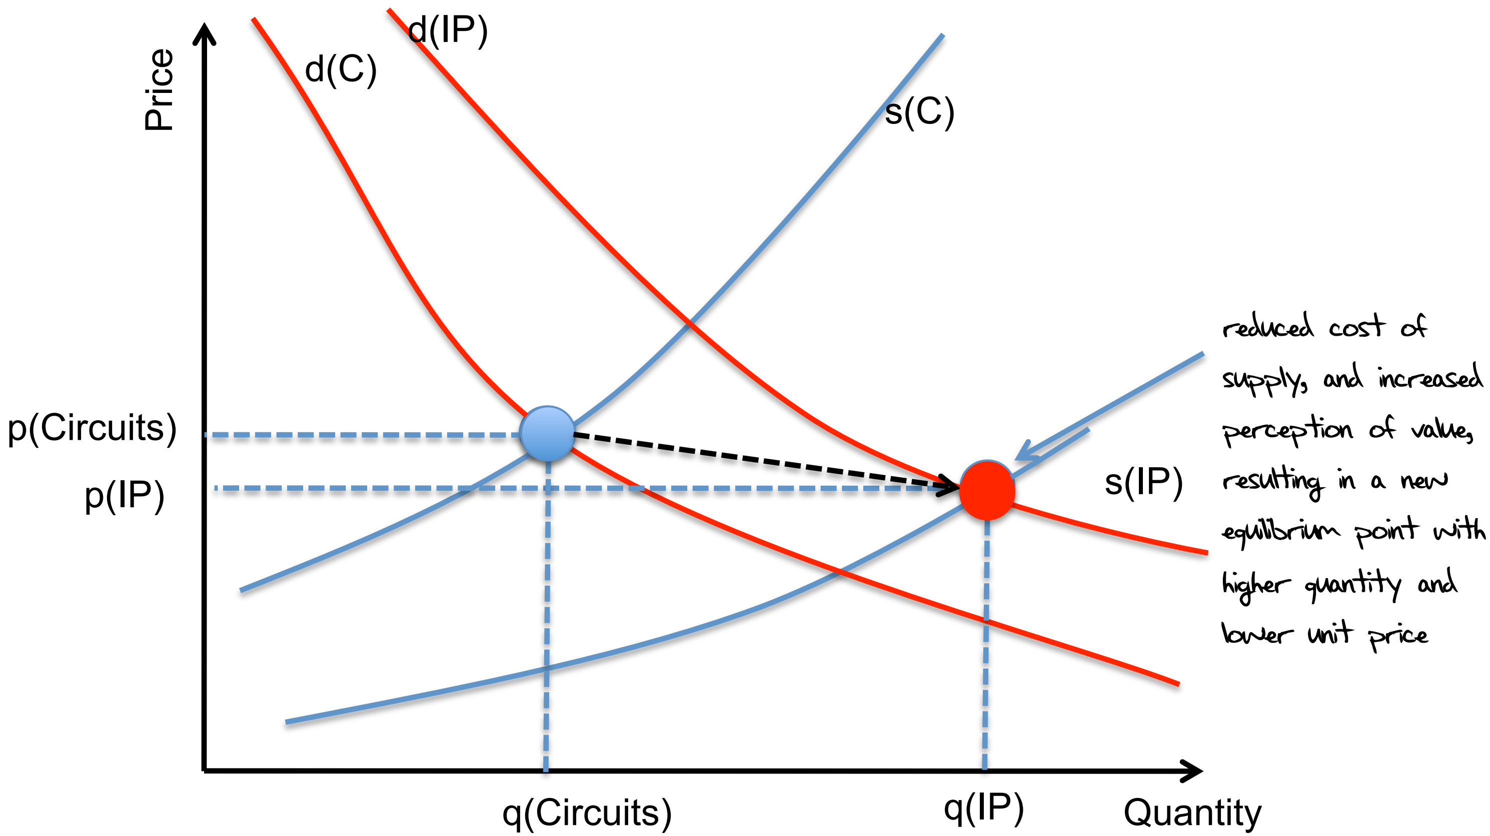

IPv4 Deployment

<u>Business</u>: exposed new market opportunity in a market that was actively shedding many regulatory constraints

- exposed new market opportunities via arbitrage of circuits
  - buy a circuit, resell it as packets
- presence of agile high-risk entrepreneur capital willing to exploit short term market opportunities exposed through this form of arbitrage
- volume-based suppliers initially unable to redeploy capital and process to meet new demand
  - unable to cannibalize existing markets
  - unwilling to make high risk investments

IPv4 Deployment

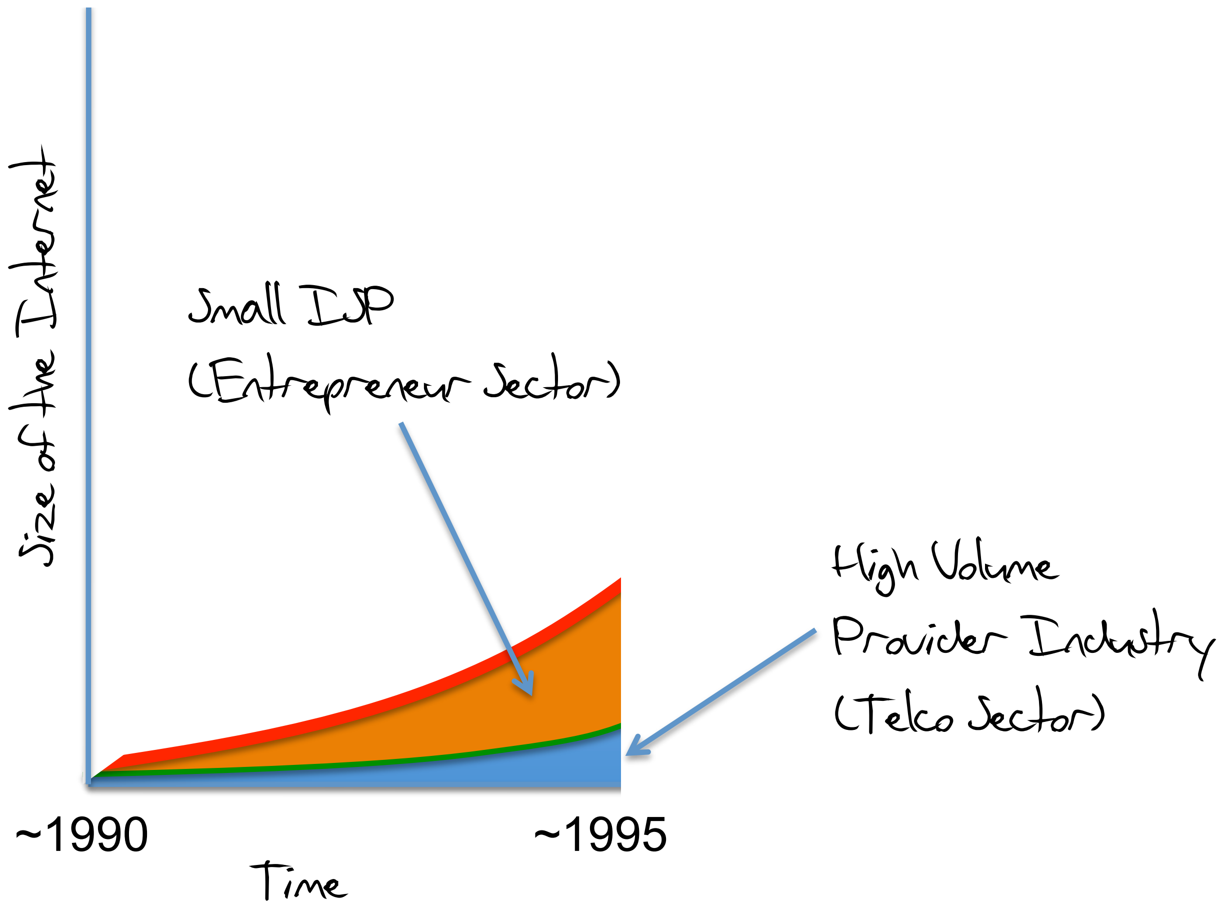

IPv4 Deployment

Business: exposed new market opportunity in a market that was actively shedding many regulatory constraints

- exposed new market opportunities via arbitrage of circuits
  - buy a circuit, resell it as packets
- presence of agile high-risk entrepreneur capital willing to exploit short term market opportunities exposed through this form of arbitrage
- volume-based suppliers initially unable to redeploy capital and process to meet new demand
  - unable to cannibalize existing markets
  - unwilling to make high risk investments
- the maturing market represented an opportunity for large scale investment that could operate on even lower cost bases through economies of scale

IPv4 Deployment

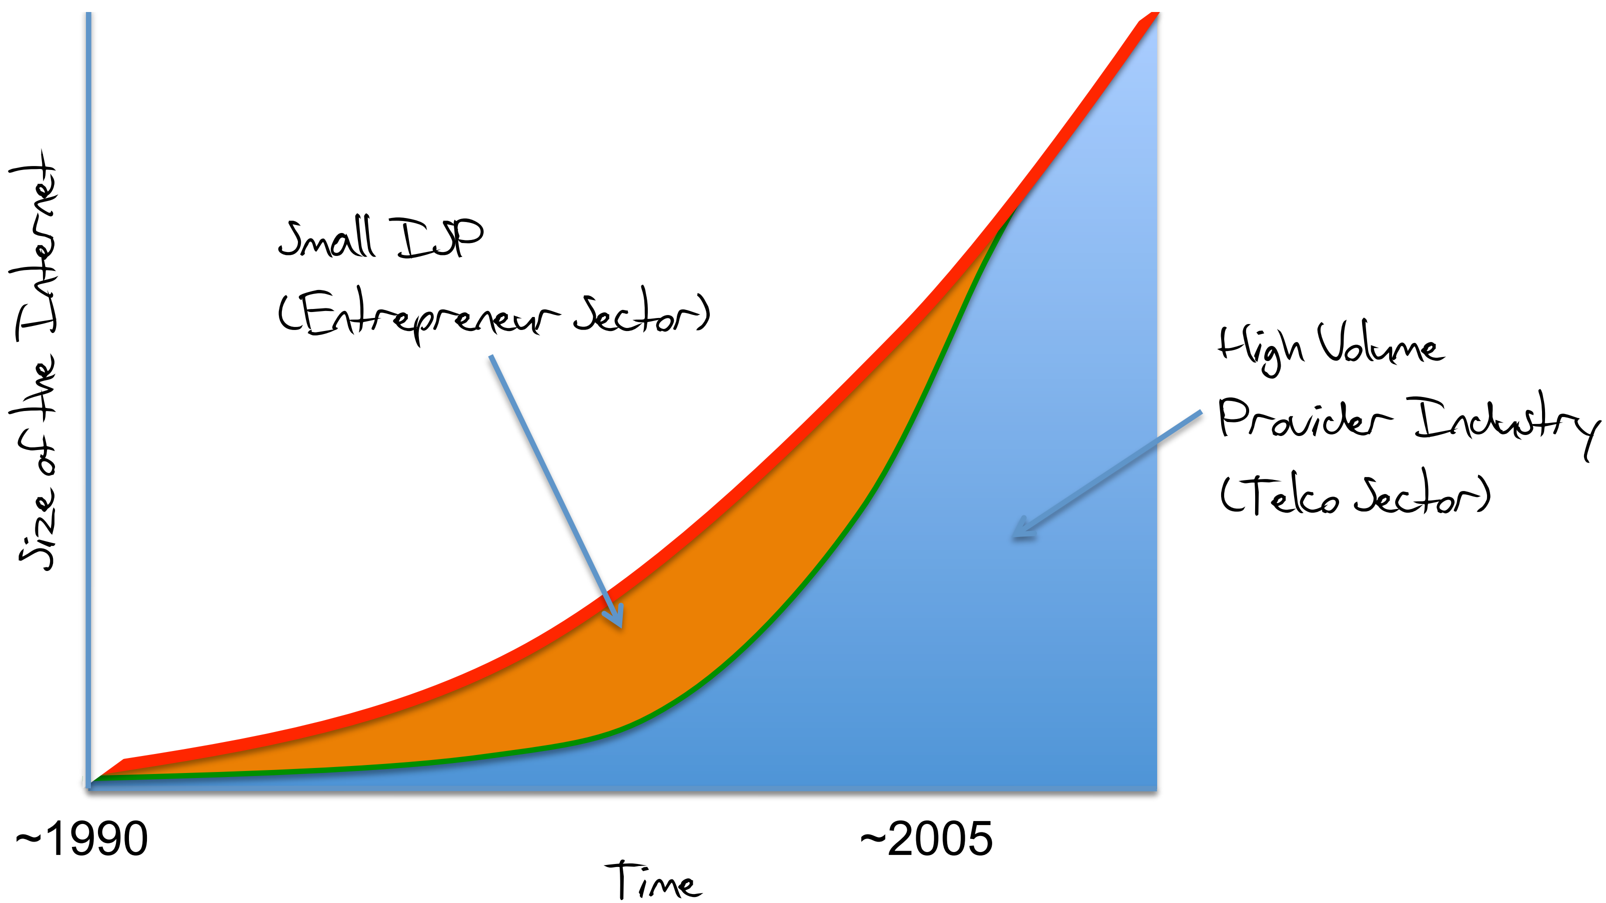

What about IPv6 Transition?

Will the same technology, cost and regulatory factors that drove the deployment of the IPv4 Internet also drive this industry through the transition from IPv4 to IPv6?

IPv6 vs IPv4

Are there *competitive differentiators*?

- no cost differential
- no functionality differential
- no inherent consumer-visible difference
- no killer app
- no visible consumer demand
- no visible competitive differentiators other than *future risk*

IPv4 to Dual Stack: The Demand Schedule Shift

Supply side cost increase due to Dual Stack operation

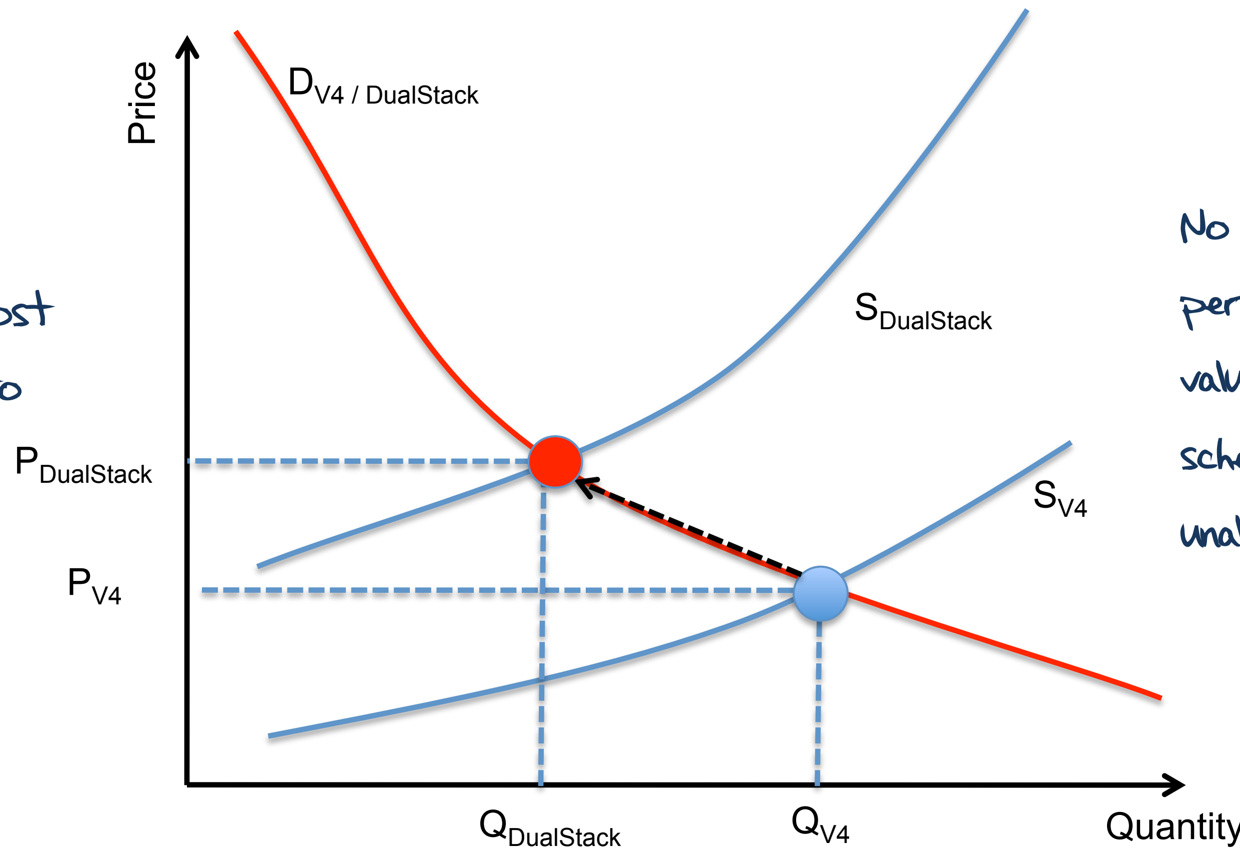

No change in perception of value, so demand schedule is unaltered

Equilibrium point is at a lower quantity if Dual Stack supply costs are passed on to customers

The Transition to IPv6

Given that we've left it so late in terms of the scale of the transition and the degree of difficultly with IPv4 exhaustion, and given that there appears to be little motivation from some critical industry segments to embark on this transition ...

will it happen at all?

The Transition to IPv6

Alternatively, is this transition an instance of a *market failure*?

“Market Failure”

Wikinomics:

“In economics, a market failure exists when the production or use of goods and services by the market is not efficient. That is, there exists another outcome where market participants' overall gains from the new outcome outweigh their losses (even if some participants lose under the new arrangement). Market failures can be viewed as scenarios where individuals' pursuit of pure self-interest leads to results that are not efficient – that can be improved upon from the societal point-of-view. The first known use of the term by economists was in 1958, but the concept has been traced back to the Victorian philosopher Henry Sidgwick.”

http://en.wikipedia.org/wiki/Market_failure

The Transition to IPv6

Alternatively, is this transition an instance of a *market failure*?

Individual self-interest leads to inefficient supply outcomes, as self-interest does not lead the installed based of consumers and suppliers to underwrite the cost of dual stack operation within the transition

IPv6 Transition as a "Public Good?"

- Is the transition to IPv6 is *non-excludable* and *non-rivalrous?*
- In which case this transition issue parallels that of a *public good*
  - With an implication that conventional market dynamics in a deregulated environment will not lead to this transition being undertaken
  - And a corollary that if this transition is considered to be necessary or essential then some form of public good solution needs to be considered

Public Good “solutions”

There are a number of conventional approaches to the distribution of a *public good:*

- Assurance contracts
- Coasian solutions
- Government enterprise provisioning
- Tariffs
- Subsidies
- Taxation remedies
- Regulatory impost

Regulatory Impost

- A regulatory constraint is placed on the ISP carrier licence holders that IPv6 services are to be provided by a given deadline
  - as has happened with digital television in many regulatory regimes.
- This regulatory constraint acts a form of a *assurance contract*, where all providers are in effect bound to produce a particular solution

Government Purchase Contracts

- Where the public sector collectively require the provision in IPv6 in all their service contracts.
- This is a form of a *coasian solution* where a group of potential beneficiaries pool together their willingness to pay for the public good.
  - We have seen this approach in the past with the Government OSI Profiles (GOSIP) of the late 1980's when the approach proved ineffectual.
  - There is no assurance that such collective actions on the part of the public sector have sufficient mass and momentum to create a broader sustainable market that will impel the private sector to undertake the transition.

Subsidies and Incentives

- Where the production of the good is subsidised in some fashion by public funds
  - This can be in the form of direct payments to service providers, or in the form of vouchers to consumers which can be redeemed only in exchange for the supply of a specified service.
- Related incentive measures include the use of taxation incentives related to infrastructure investment, where the investment in a certain class of infrastructure or in a certain sector can be provided with advantaged taxation treatment.

Public Provision

- Where the service is provided by a publically-owned enterprise.
- The funding for such an enterprise can be provided by government-backed investment bonds, or directly from public revenues, and operating losses are underwritten by the public purse.
  - This measure was used for most national telephone service providers for a significant part of the twentieth century, so it is not exactly a completely foreign concept for this industry.

What About IPv4 Exhaustion?

- Does IPv4 address exhaustion change this picture?
- What are the economic implications of service providers adding CGNs to the current service offering based on IPv4?
- Are CGNs and IPv6 mutually exclusive investment options for access providers?

Adding CGNs to IPv4: The Demand Schedule Shift

Supply side cost increase due to Dual Stack operation

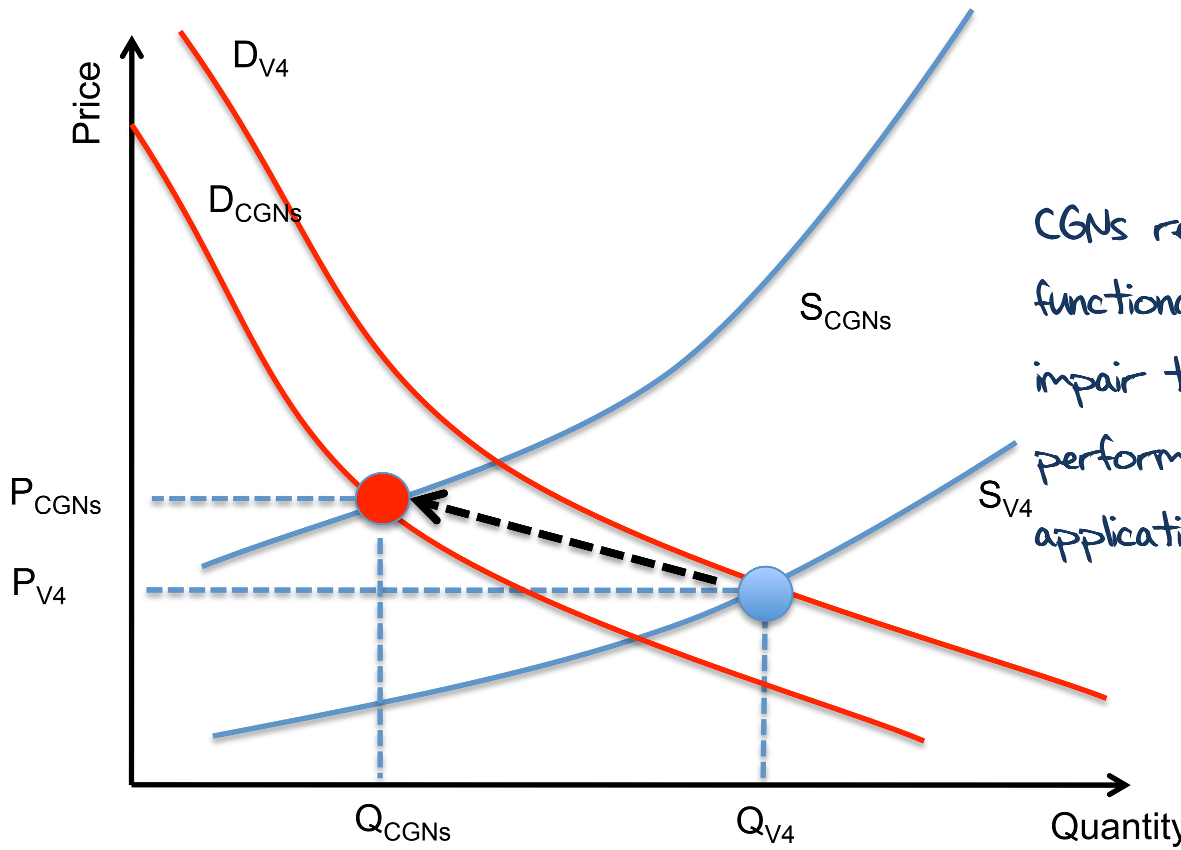

CGNs reduce functionality and impair the performance of some applications

CGNs represent higher cost and lower value for customers

But is this all there is to CGNs?

- Will CGN's alter the user's experience of services?
- Does this alter the role (and location) of CDNs?
  - Or has the CDN model already evolved to accommodate this evolution?
- Do CGNs alter the leverage of the access provider with respect to service deployment?
  - Is this an instance of a forced carriage toll gate that allows the carriage sector to renegotiate their relationship with the content access model

Further musing...

Do we really understand the dynamics and inter-relationships of the components of this industry?

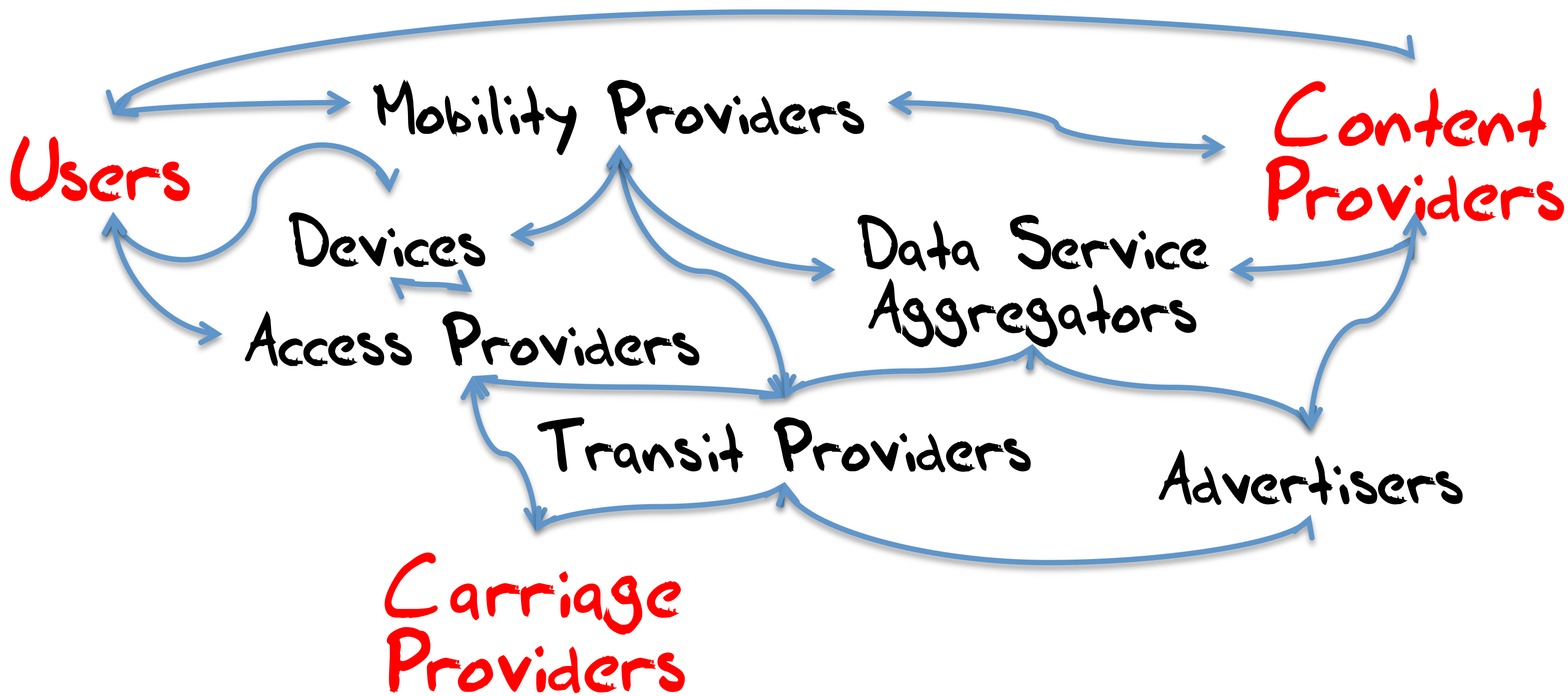

Further musing...

- What drives the carriage sector?
- What drives the content sector?
- Who's winning today?

- Is the open network architecture being offered by IPv6 fundamental to the objectives of either of these sectors?
- Will the access part of the carriage industry invest in IPv6 infrastructure and service provision?
  - If so, then why?
  - If not, then why not?

Your Thoughts?

- Carriage vs Content
  - Currently advances in IT and the Internet has allowed content to shed carriage mediation and negotiate directly with the end consumer
  - Will scarcity in the carriage activity enable carriage players to re-enter the content distribution function in a mediation role and extort toll revenues from content providers?

Thank You!