Just what are we doing about this IPv6 transition?

Geoff Huston
APNIC

research@APNIC.net

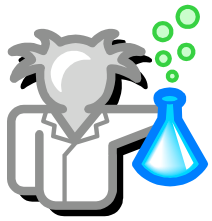

The story so far…

The story so far…

In case you hadn't heard by now, we appear to be running quite low on IPv4 addresses!

Maybe you've had enough of the train wreck analogy for IPv4 exhaustion, despite some truly excellent wrecks that were especially prepared for your enjoyment.

So if you like your visual analogies to be a little more catastrophic in nature ...

IPv4 Address Exhaustion

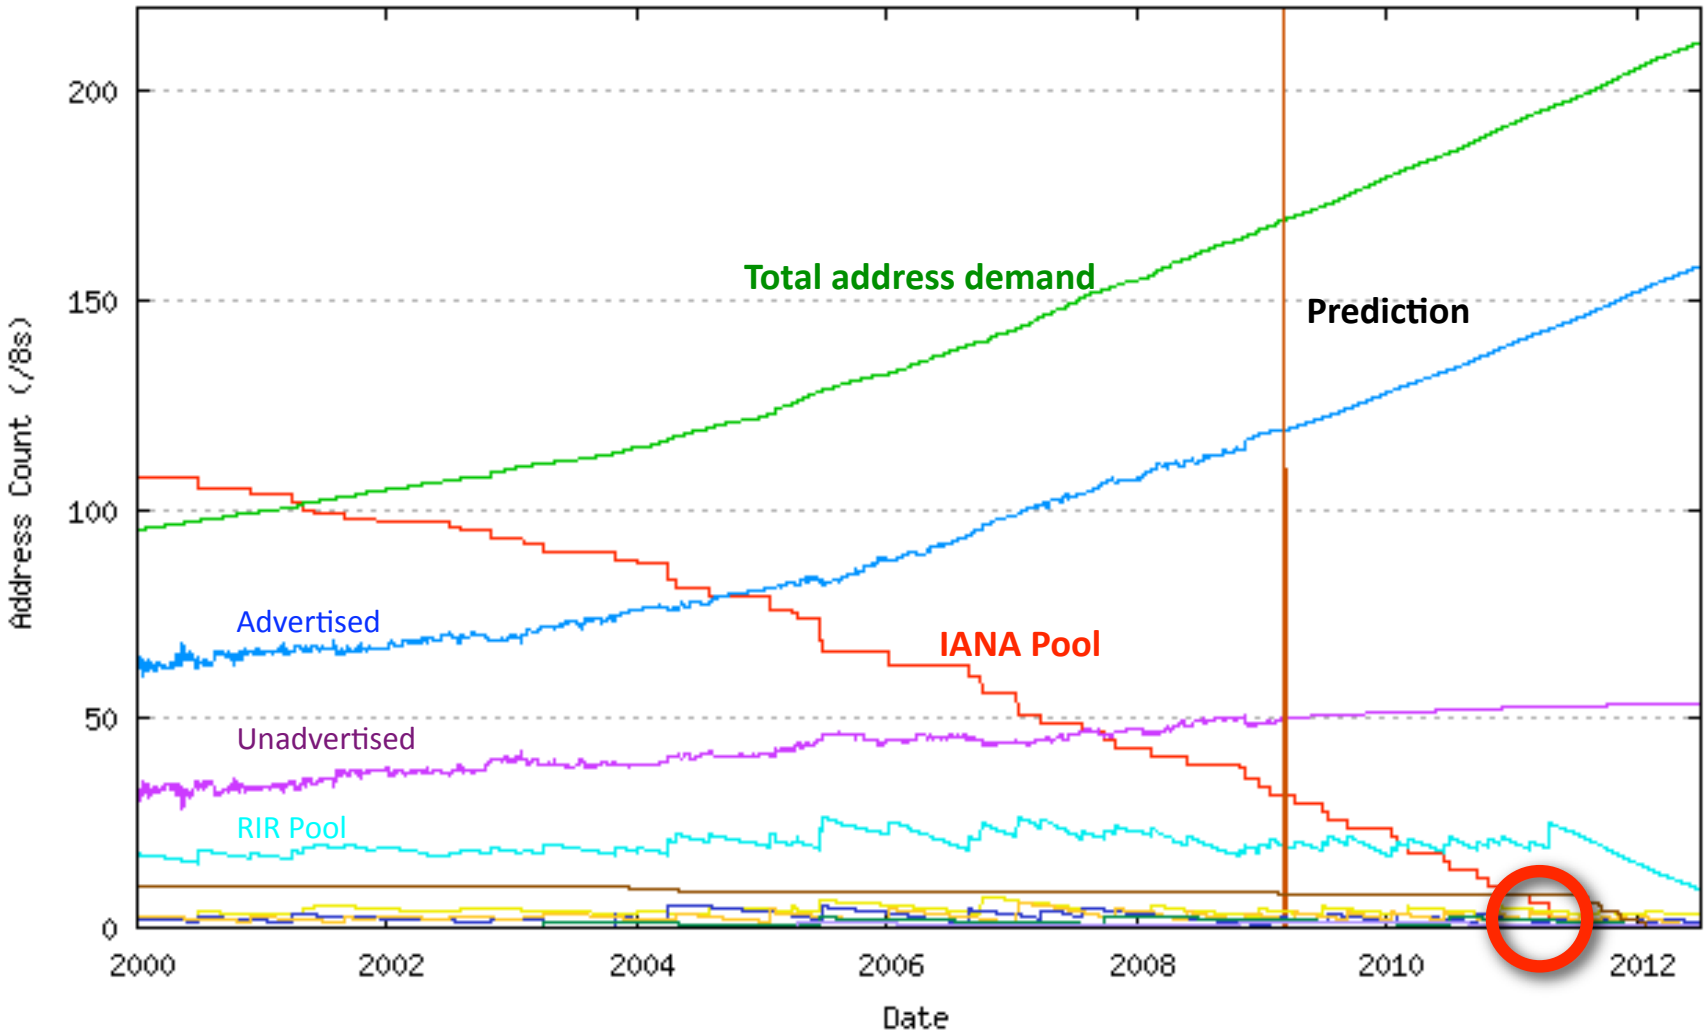

In this model, IANA allocates its last IPv4 /8 to an RIR on the 13th July 2011

This is the model's predicted exhaustion date as of the 23rd August 2009. The predictive model is updated daily at:
http://ipv4.potaroo.net

Ten years ago we had a plan …

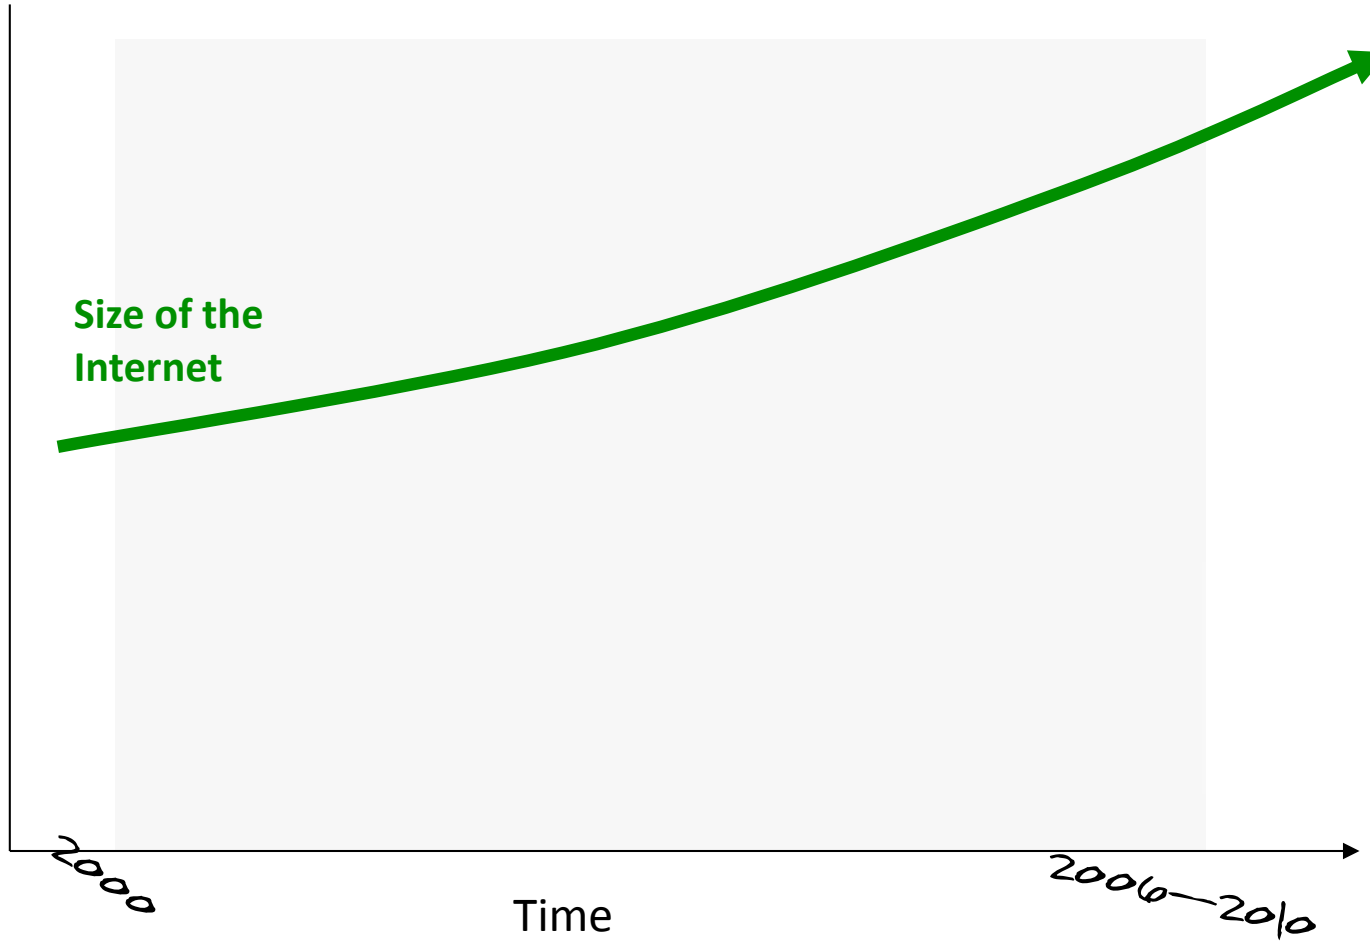

Ten years ago we had a plan …

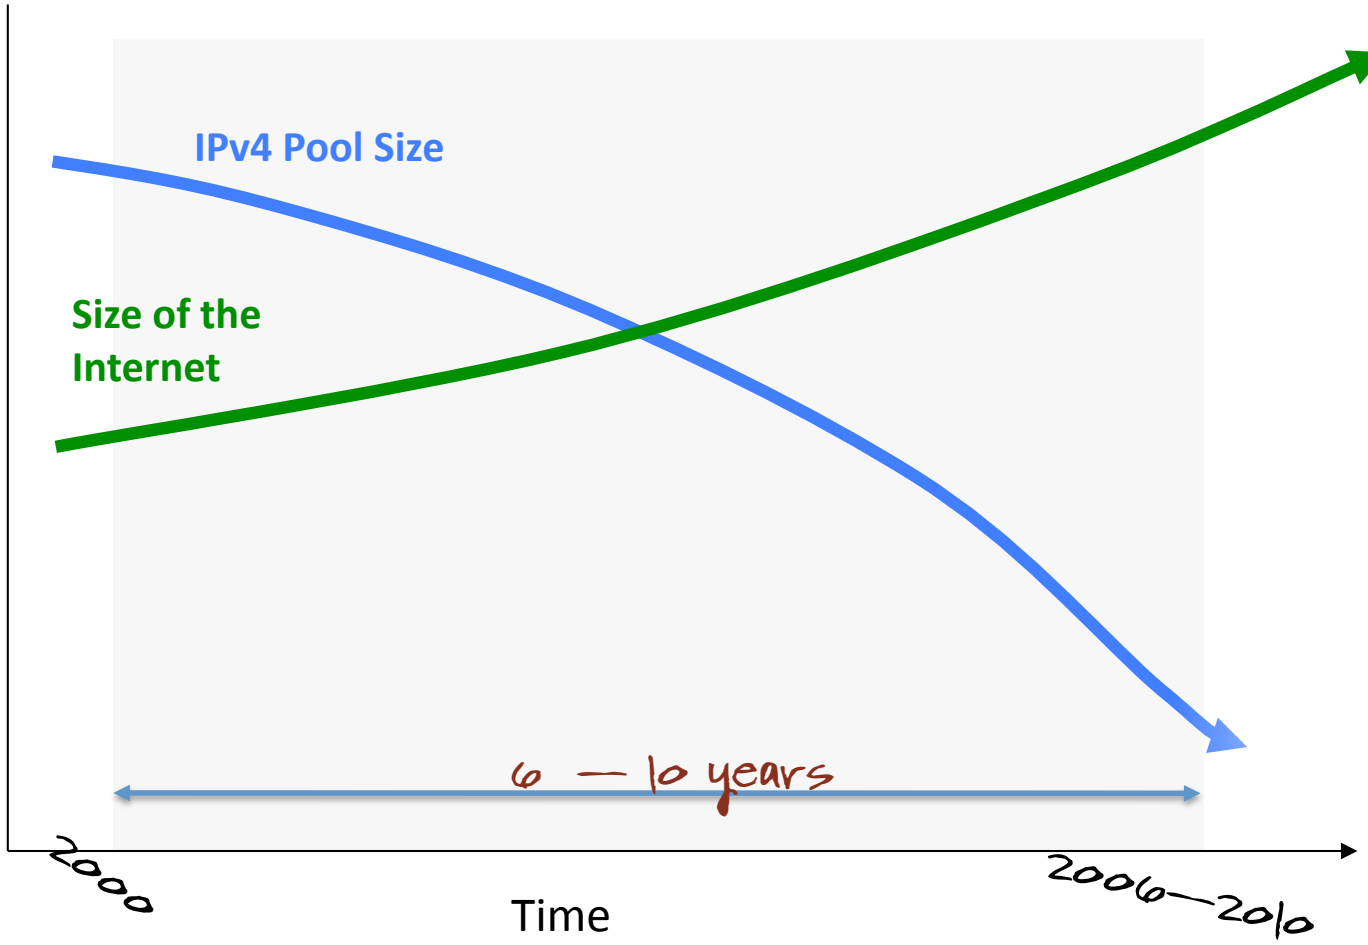

Ten years ago we had a plan …

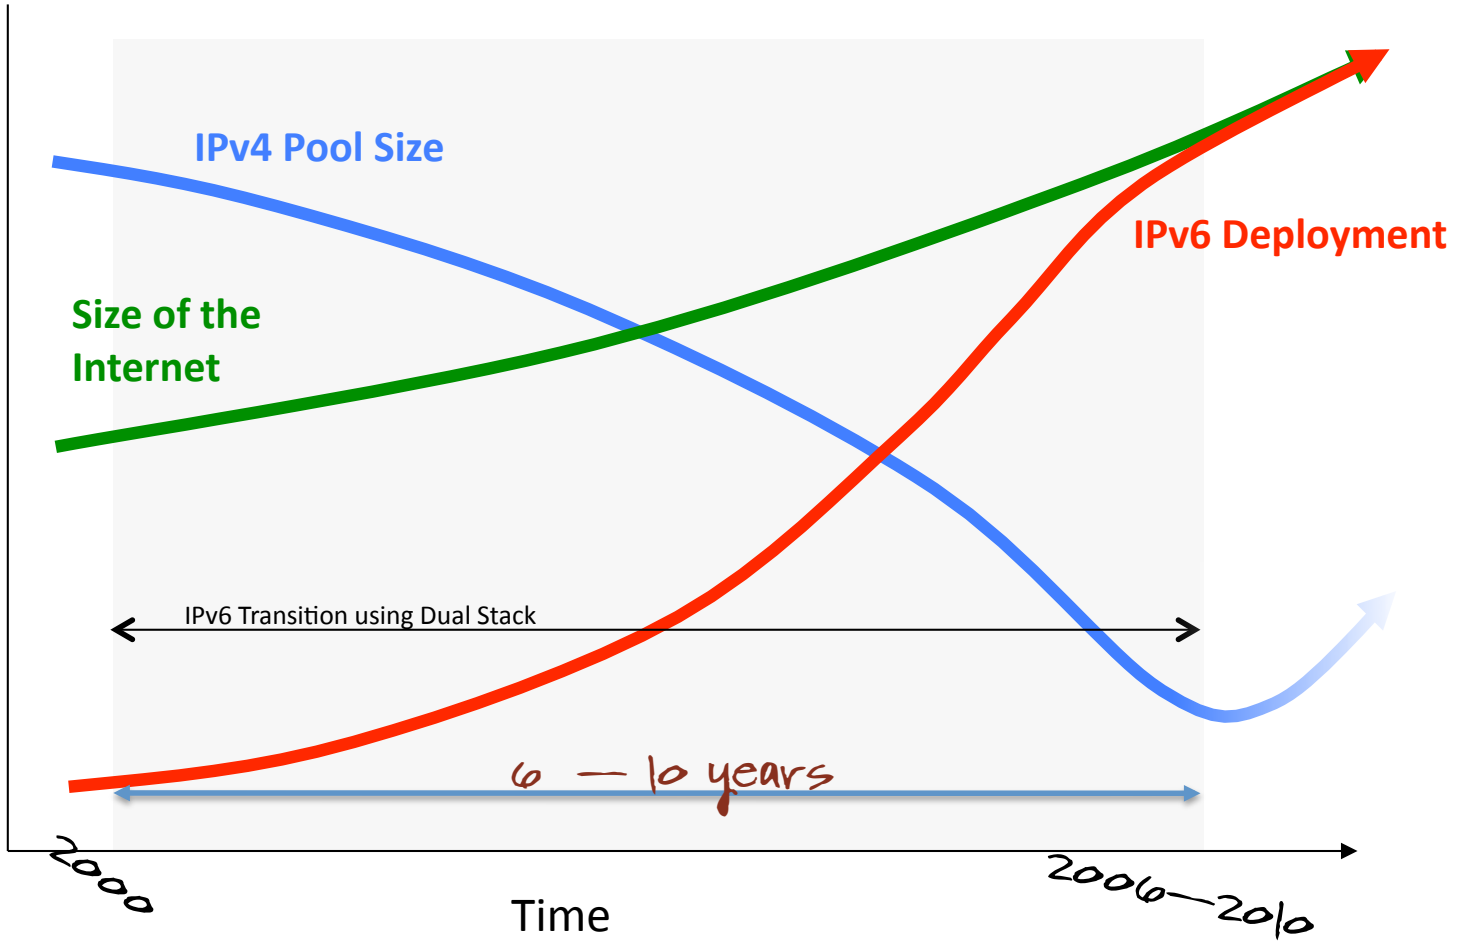

Ten years ago we had a plan …

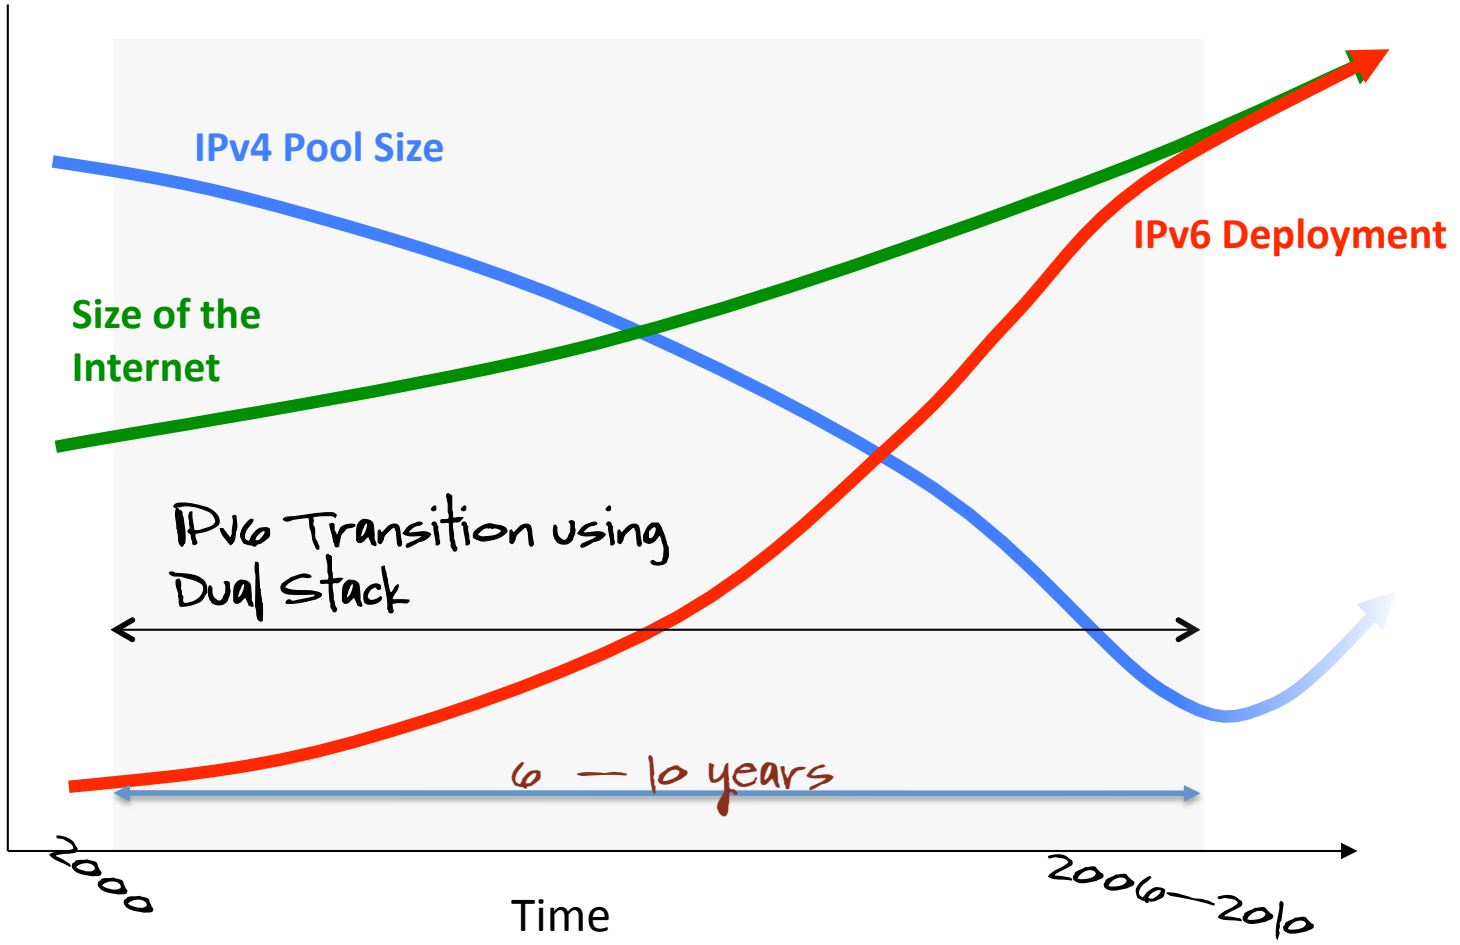

How are we going today with this plan?

OR: How much IPv6 is being used today?

Web-based IPv6 Stats

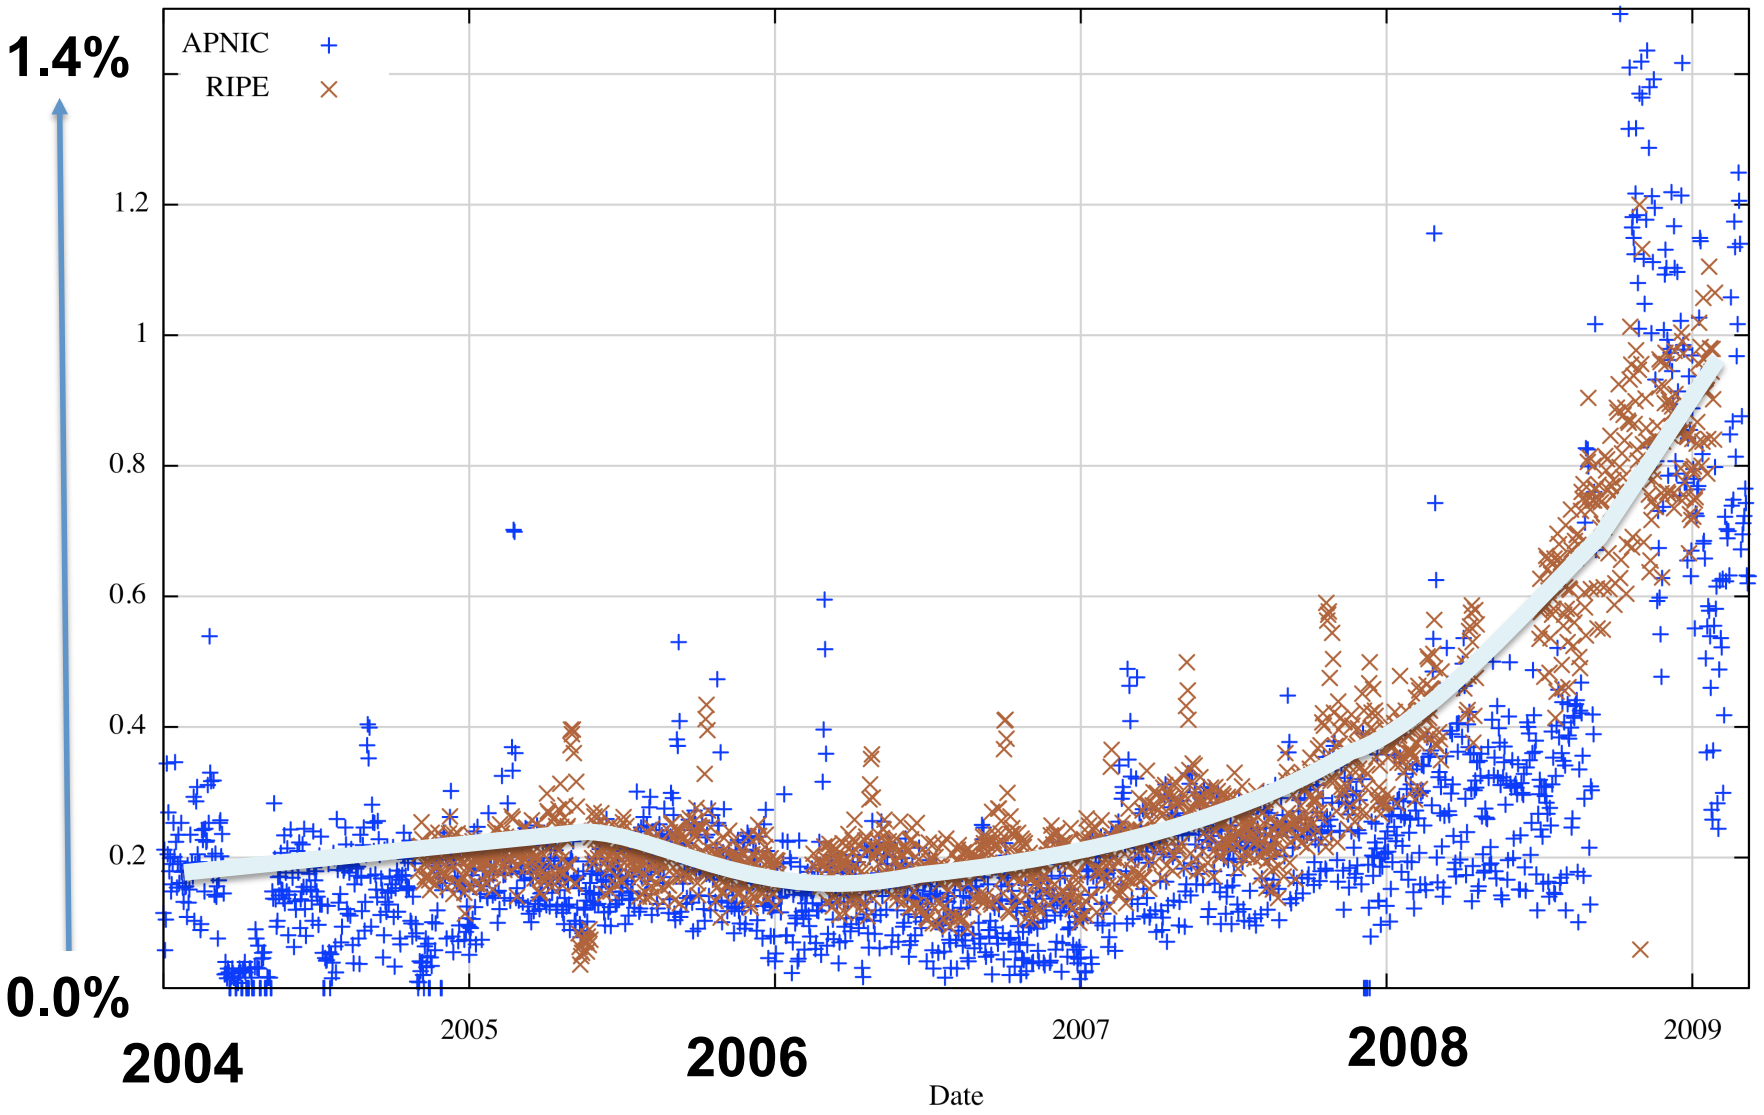

Where are we today with IPv6?

Compared with the size of the IPv4 network, the IPv6 network is around one hundred times smaller (or 1%)

This figure is based on end—to—end capability measurements from a small sample of dual stack web sites. The bias in the data set means that the figure may well be very much smaller than 1% for the larger Internet

What's the revised plan?

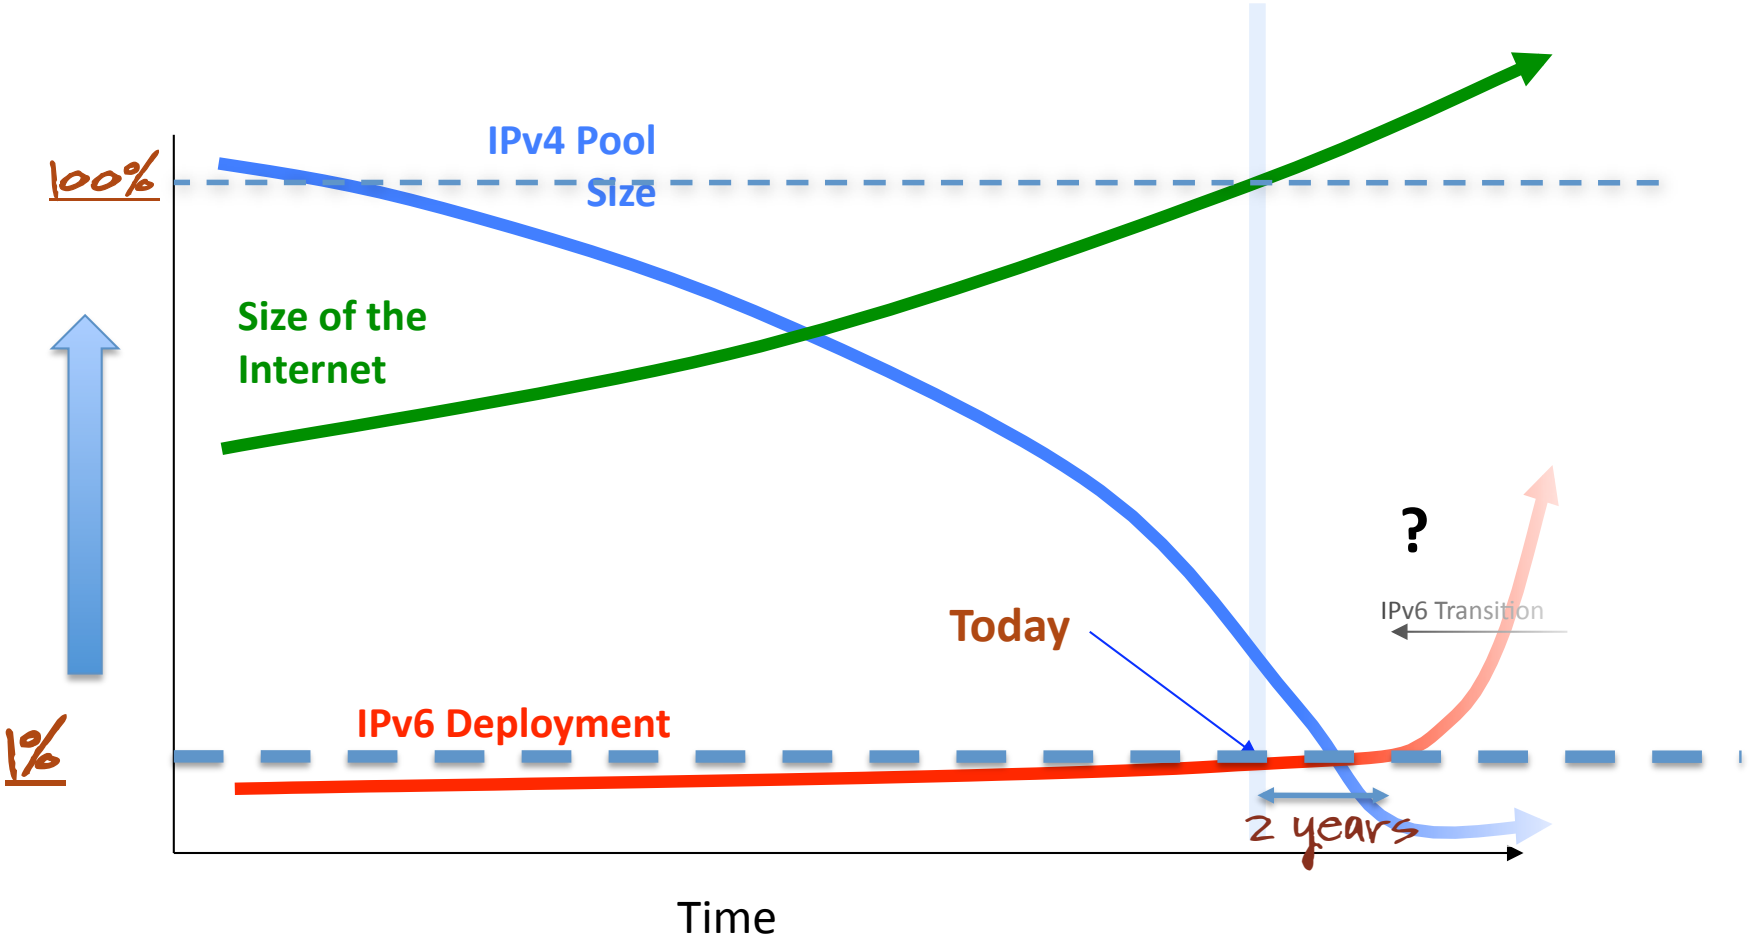

Its just not looking good is it?

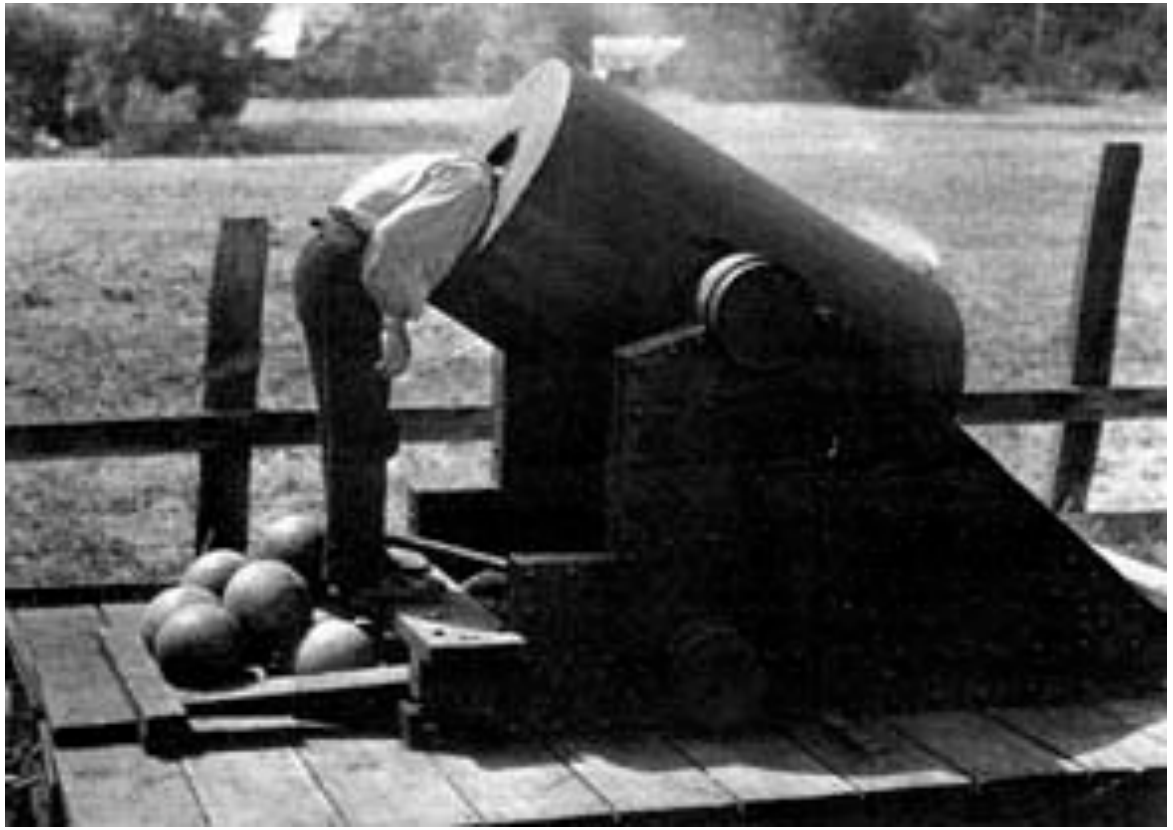

The Grand Dual-Stack Transition Plan

IPv6 is not "backward compatible' with IPv4

The Grand Dual-Stack Transition Plan

IPv6 is not "backward compatible' with IPv4

So the plan was to undertake the transition at the edges, progressively equipping end hosts and apps with IPv6 as well as IPv4

The Grand Dual-Stack Transition Plan

IPv6 is not "backward compatible' with IPv4

So the plan was to undertake the transition at the edges, progressively equipping end hosts and apps with IPv6 as well as IPv4

When the overall majority Internet host population and Internet applications were dual-stack equipped we could then shut down IPv4 support

Dual Stack Transition

Dual Stack Transition

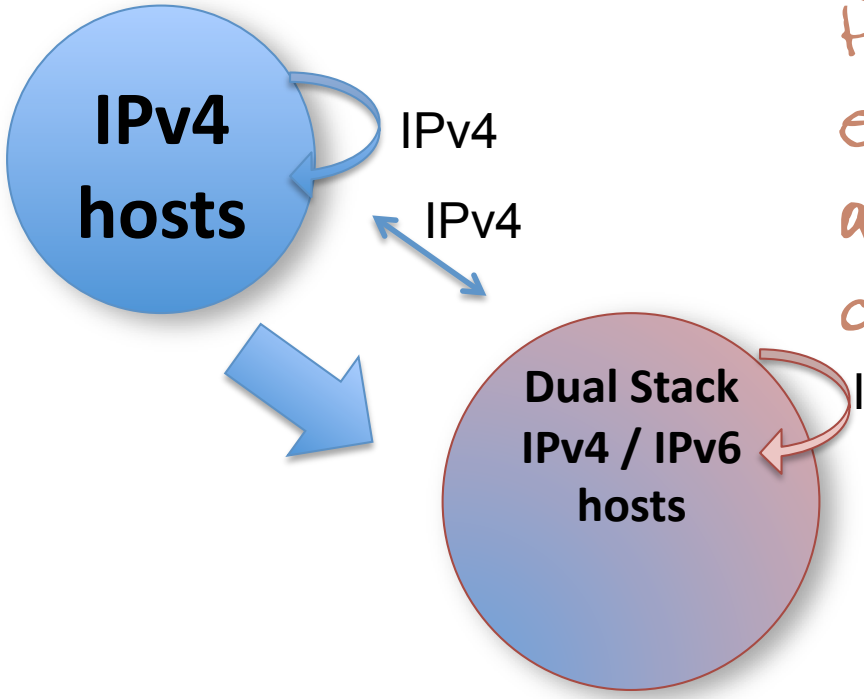

Phase 1 –

Progressively equip all end host systems and apps with Dual stack capability

Dual Stack Transition

Phase 2 – Phase out Dual Stack in favour of IPv6

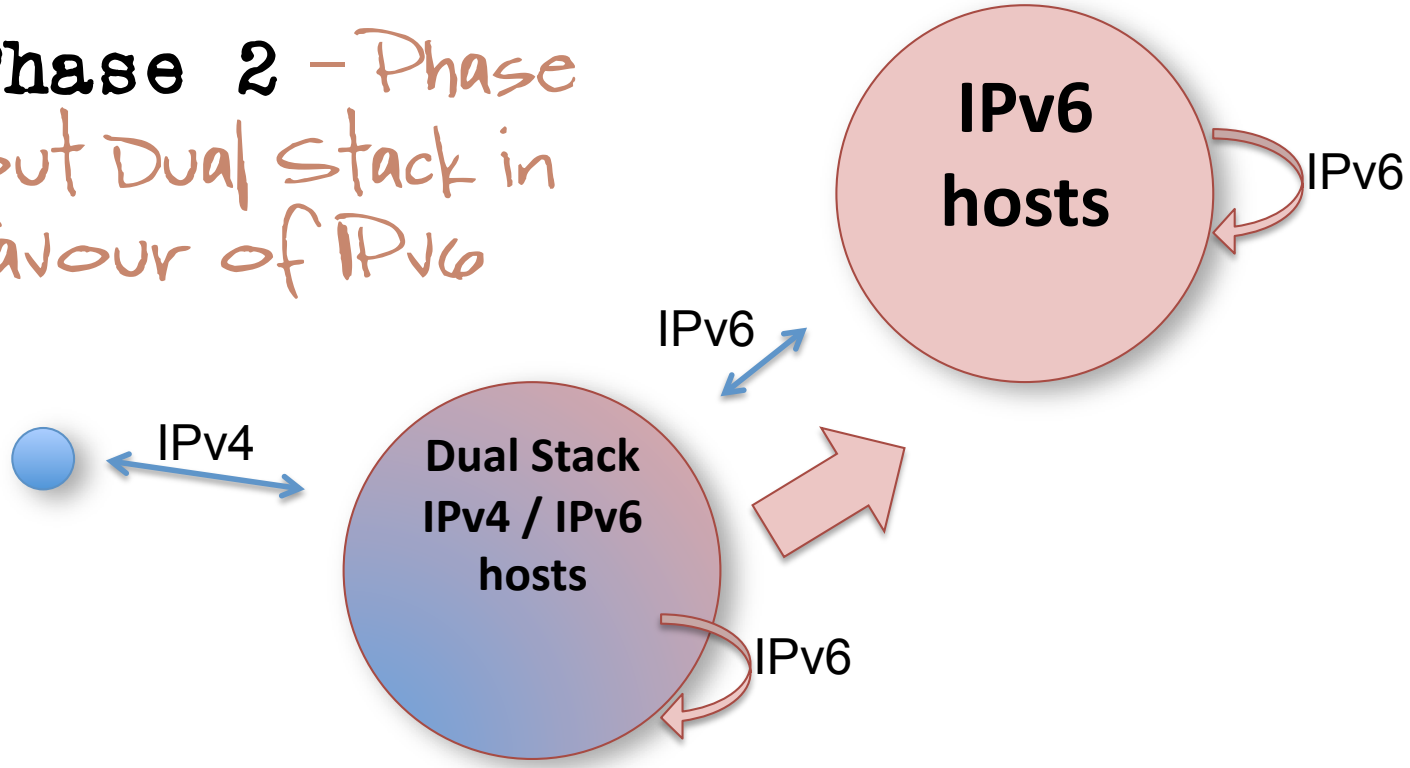

Dual Stack Transition

Phase 2 – Phase out Dual Stack in favour of IPv6

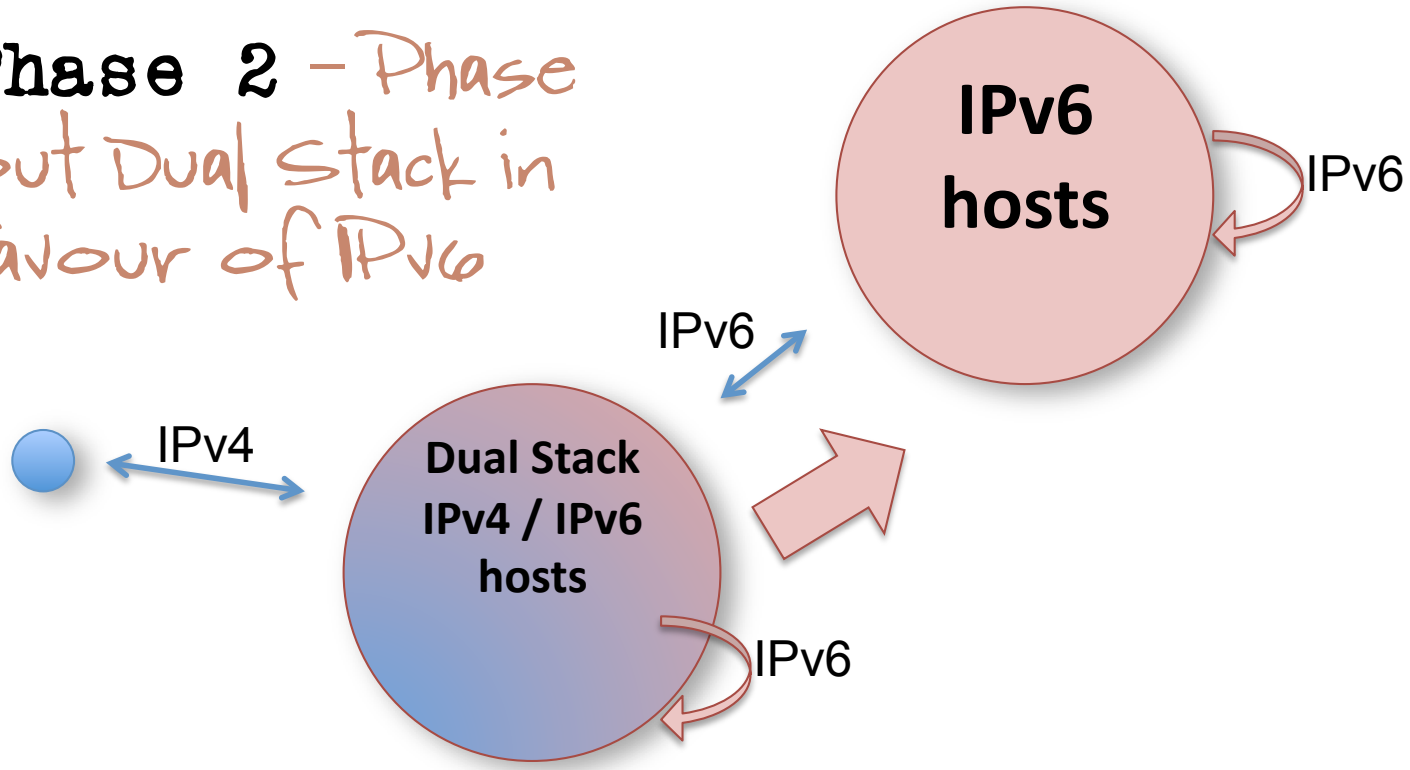

If we ever get to phase 2, the execution of phase 2 will be quick – once all(*) hosts are IPv6 capable, then there is no need to continue support for ipv4

Dual Stack Transition

How long will Phase 1 take?

For how many years from now will we need to keep on providing IPv4 addresses to every host?

Phase 1 - Option A

We perform a miracle!

The global Internet, with more than 1.7 billion users, a similar population of end hosts and devices, and hundreds of millions of routers, firewalls, and billions of lines of configuration codes, and hundreds of millions of ancillary support systems, where only a very small proportion are IPv6 aware today, are all upgraded and fielded to work with IPv6 in the next **500 days**, and then completely quits all use of IPv4 in **30 days** later.

Phase 1 - Option A

We perform a miracle!

The global Internet, with more than 1.7 billion users, a similar population of end hosts and devices, and hundreds of millions of routers, firewalls, and billions of lines of configuration codes, and hundreds of millions of ancillary support systems, where only a very small proportion are IPv6 aware today, are all upgraded and fielded to work with IPv6 in the next **500 days**, and then completely quits all use of IPv4 in **30 days** later.

Yeah - right!

Phase 1 - Option B

We go so slowly that it stalls!

- Transition extends for more than a decade
- The Internet grows to 4 - 10 times its current size using intense IPv4 NATs and a shift to universal adoption of client/server architectures and translation gateways

Phase 1 - Option B

We go so slowly that it stalls!

VERY UGLY!

At what point in time is IPv6 dropped as a common objective and the networked environment shift to large scale disjoint network realms with application level gateways with content capture and provider lock—in?

Phase 1 - Option C

We have at most about 4-5 years:

To get to the point where so much of the host population is dual-stack capable that what's left on IPv4 is not a stalling factor

How can this happen?

Deploy IPv6/IPv4 Dual Stack on EVERYTHING!

and clean up the IPv6 infrastructure as we do so!

And increase NAT density in V4

How can this happen?

Deploy IPv6/IP[illegible] Stack on EVERYTH[illegible]

We have an idea how to do this

and clean up the IPv6 infrastructure as we do so!

And increase NAT density in V4

How can this happen?

Deploy IPv6/IP[illegible] Stack on EVERYTH[illegible]

We have an idea how to do this

and clean up the IPv6 infras[illegible]ture as we do so!

And increase [illegible]cusity in V4

This one could be tricky...

Today

NATs exist in the CPE

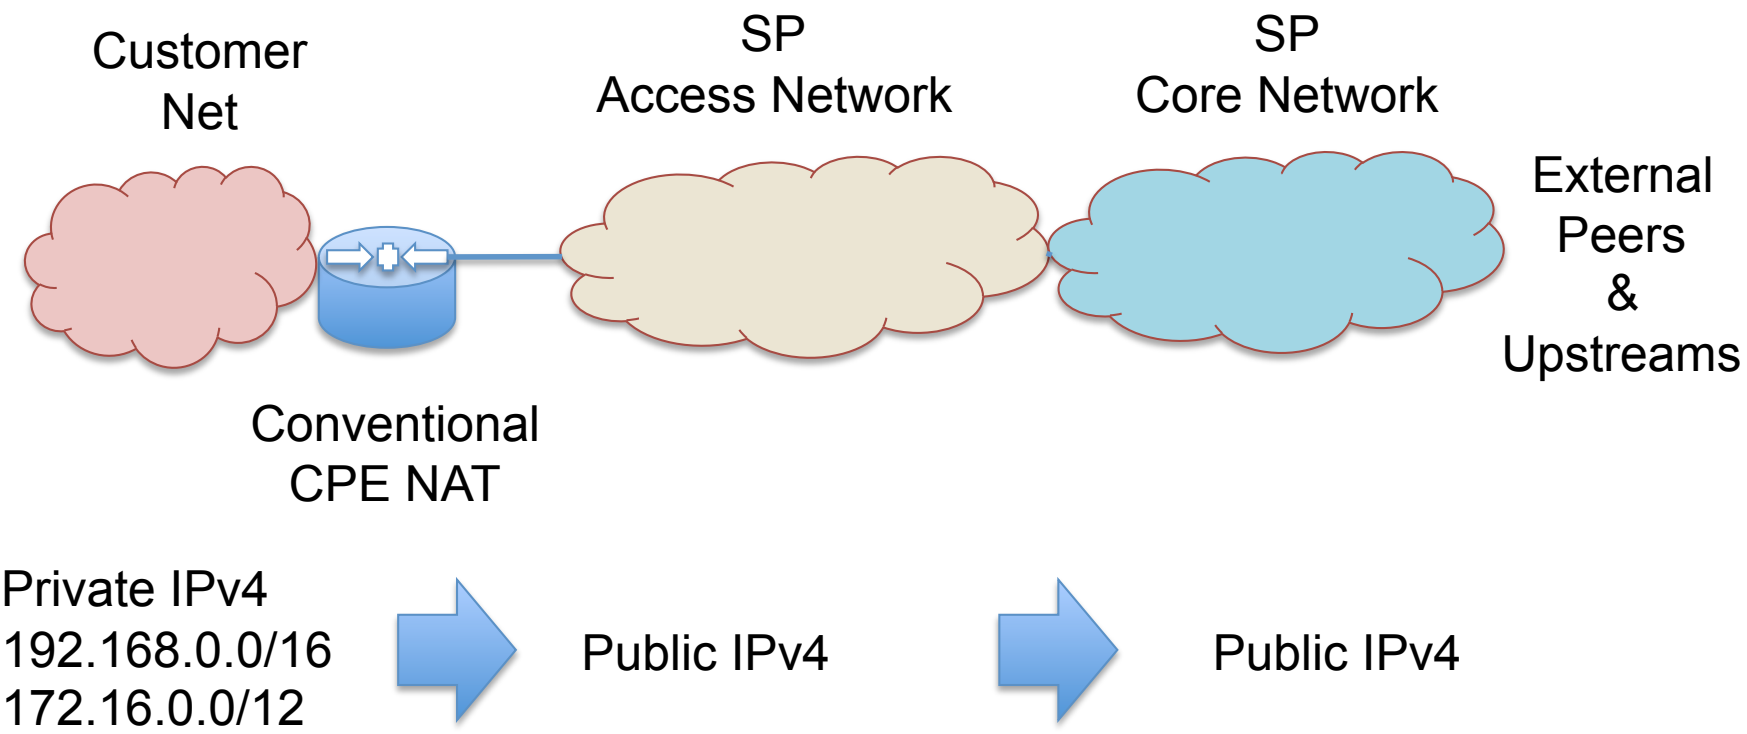

Carrier Grade NAT

Add another NAT in the path

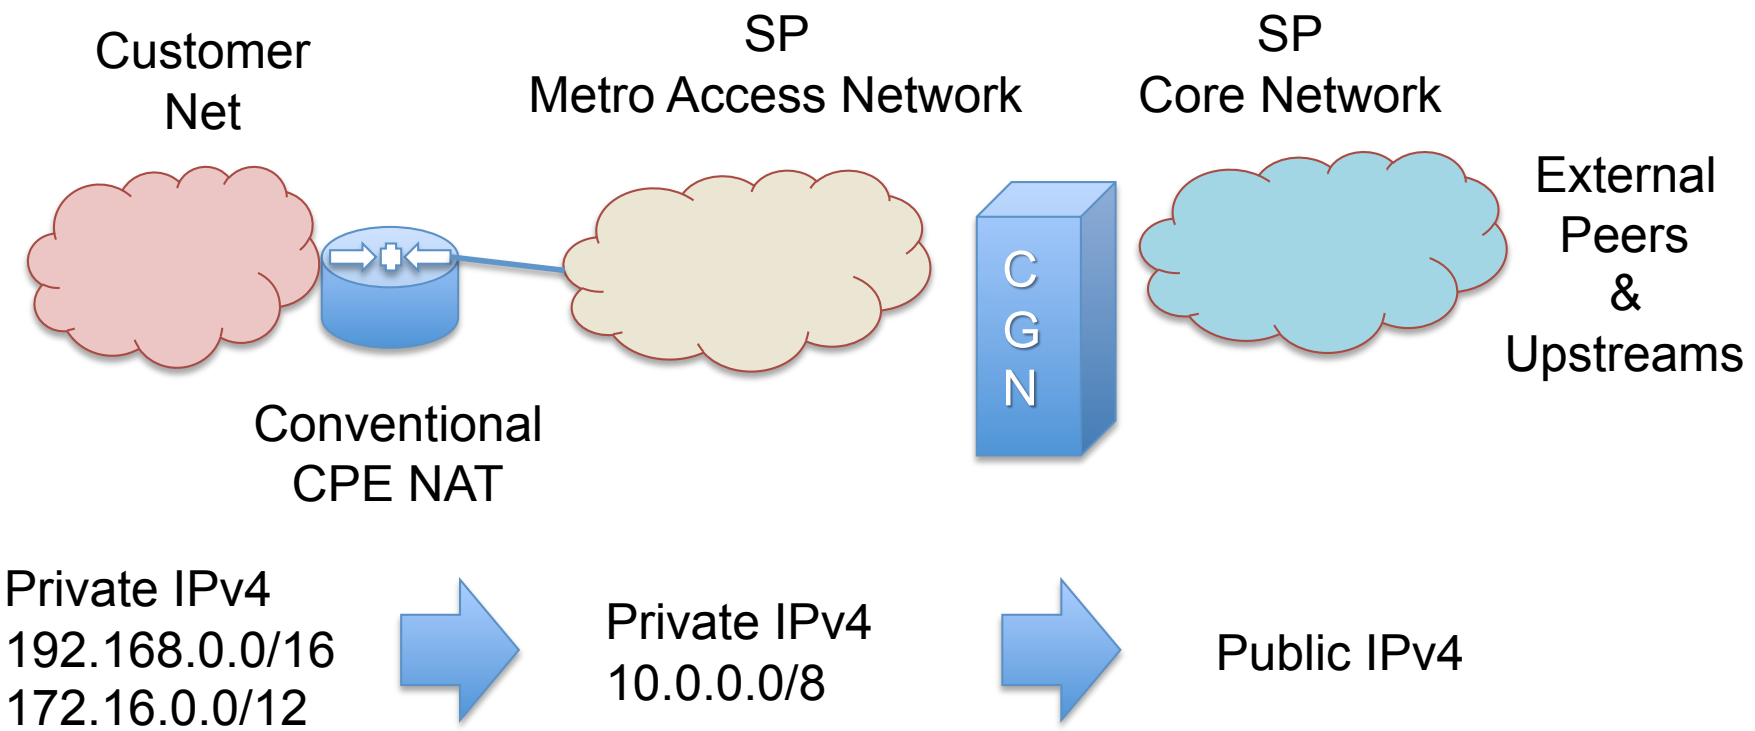

Variations

- Dual Stack Lite

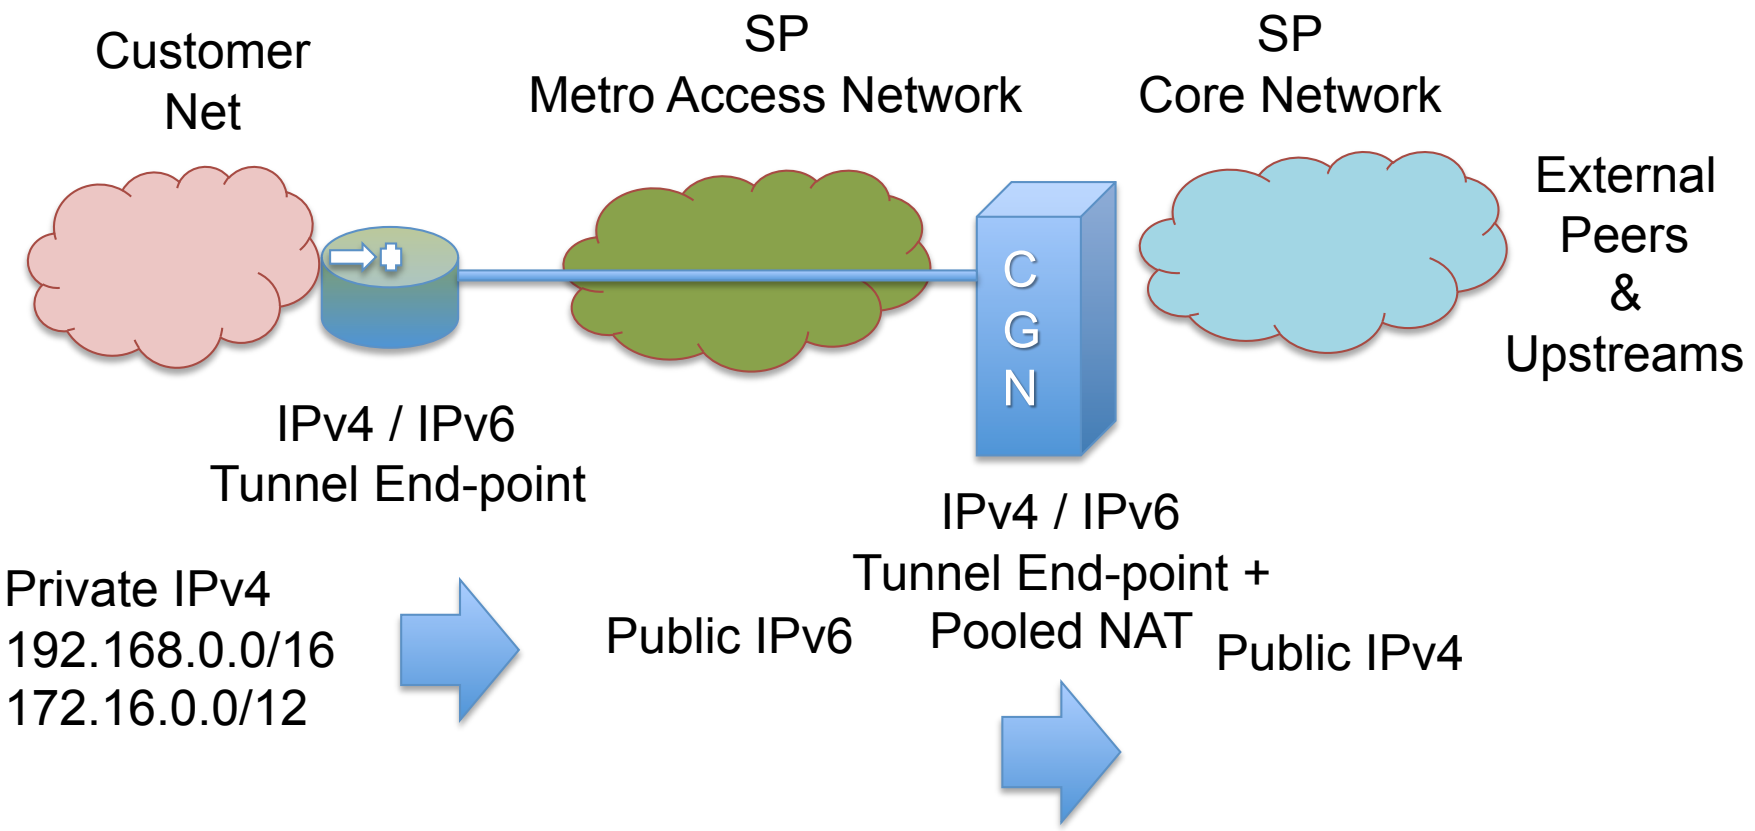

Variations

- Address + Port

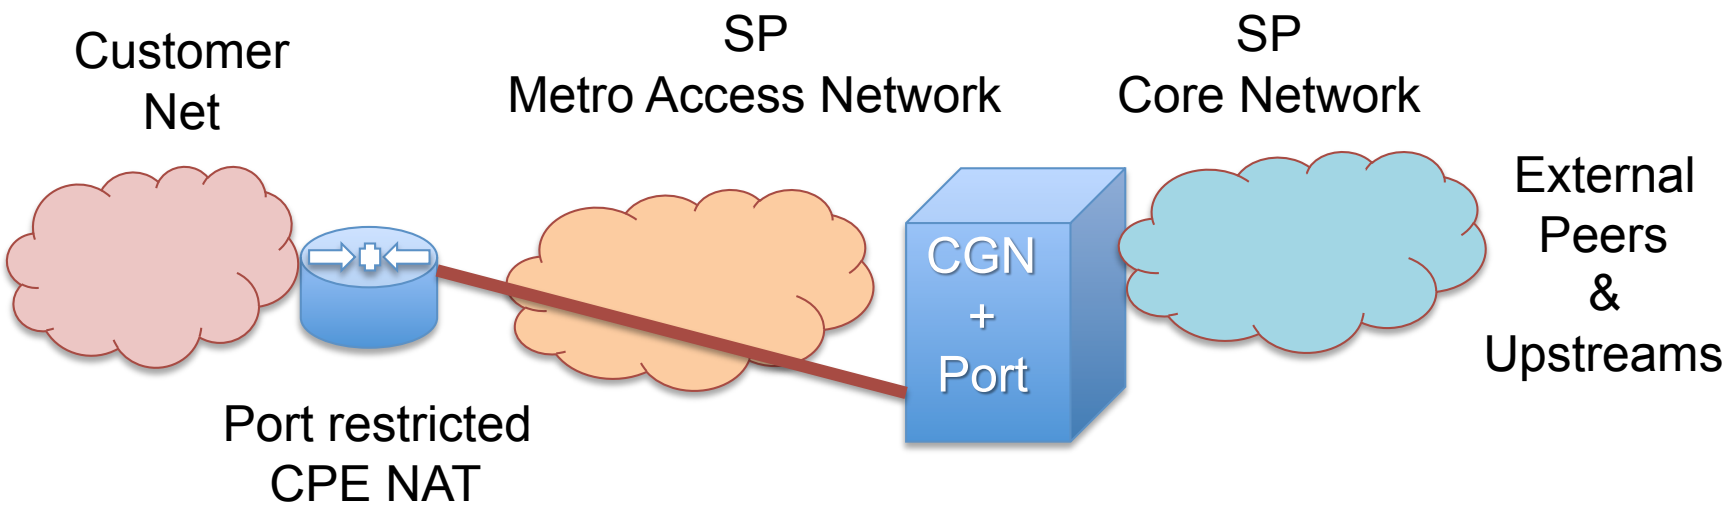

Private IPv4
192.168.0.0/16
172.16.0.0/12

Shared Private IPv4
Port Forwarding
/ Tunnelling

Public IPv4

NATs, NATs and NATs

- Use the port address in the TCP / UDP header to distinguish between CPE end points

  i.e. share an SP's IPv4 address across multiple CPE endpoints

  - CGN: dynamic port pool operation, but with complications of dual NAT traversal
  - D-S Lite: shift the NAT to the SP and eliminate the CPE NAT
  - A+P: explicit port rationing at the CPE and eliminate the SP's CGN

But…

None of these are commercial products as yet ..

- **CGN** requires equipment to be deployed in the SP network (and will probably break some existing applications)
- **D-S Lite** requires CPE change plus CGN equipment plus IPv6 SP deployment in the access net
- **A+P** requires CPE change plus CGN equipment plus SP change to permit port forwarding

What won't work

NAT-PT

- at a packet-to-packet, statically mapped translation level you can make it fly
  - and there are implementations out there
- but when you add the DNS and various application level behaviours into the mix, then lying about destination addresses, even for Good, is a Bad Thing in a packet datagram architecture

What won't work

Assuming that this industry is ill-informed and stupid

- the impediments to rapid dual stack deployment across all products and services are not based on ignorance of IPv6 within the industry.
- more outrageous exhortations and overblown hype about IPv6 serves no useful purpose other than providing mild amusement!
- it may be better to look to the business model and public policy framework of today's Internet

What's missing?

- Transition appears to be a necessary activity, and we will have to make Dual Stack last well beyond exhaustion, including IPv4
- So one way or another we are facing some form of carrier NAT solution, and possibly a number of approaches
- If this is a necessary future, then what's missing from what we have now in order to make this happen?

1. No Money

~~Good~~, ~~Fast~~, and Cheap?

- Cheap is what drives the economics of the internet
- For an ISP, address scarcity has, so been a cost imposed on **customers**, not the ISP up until now…
- BUT all this is changing with address sharing proposals
- All these address sharing models impose new roles (and costs) on ISPs
- These models do not generate commensurate additional revenue
- Leading to a situation of displaced costs and benefits - the major benefits of this investment appear to be realized at the services and application layer rather than by existing large scale infrastructure incumbents, yet the major costs of such address sharing measures will be borne by the large scale incumbent operators of low layer access services

Sound Familiar?

2. No Time

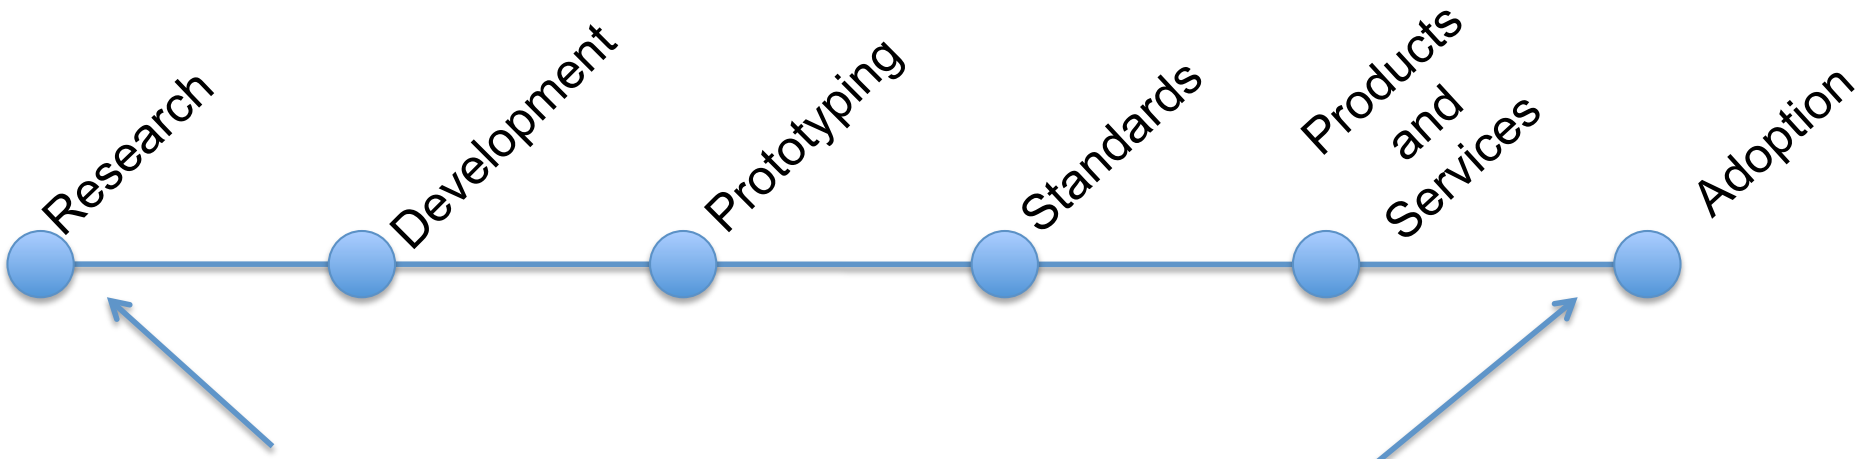

- We appear to be at the initial steps of this process of novel NAT technology to underpin IPv4 networks post-exhaustion
- We would like to be at the final stages of this process in a month or three from today
- Is this scale of development and deployment over the entire Internet likely? Possible? Plausible? Implausible? Impossible?

3. No Common Consensus

Confusion and Chaos

- Given that available effort is finite, where should we invest to effect the greatest leverage?
  - Port rationing in IPv4 ?
  - IPv6?
  - IMS and Application Level Gateways?
  - Application Level Peer networks
- Or will each or us make our own individual decisions and create chaotic and unviable outcomes for the network as a whole?
- No commonality of purpose or direction
- What's a "natural" evolution here?

Where Next?

- Do we need to address EVERYTHING with shared addressing models?
- Or do we just need to allow web access to work? (The "everything over http" model of Internet services)
- How will the next generation of application models react to this situation?

Or...

when all else fails, there is always denial


Supertramp
CRISIS? WHAT CRISIS?
394 560-2