IPv6 – A 10,000m Perspective

Geoff Huston

Chief Scientist, APNIC

Why?

Because we've run out of addresses!

again!

A bit of history...

The original ARPAnet design of 1969 used the NCP protocol, which used 8 bit addresses

- Maximum network size of 256 nodes
- Enough, yes?

ARPAnet IMP - 1970’s

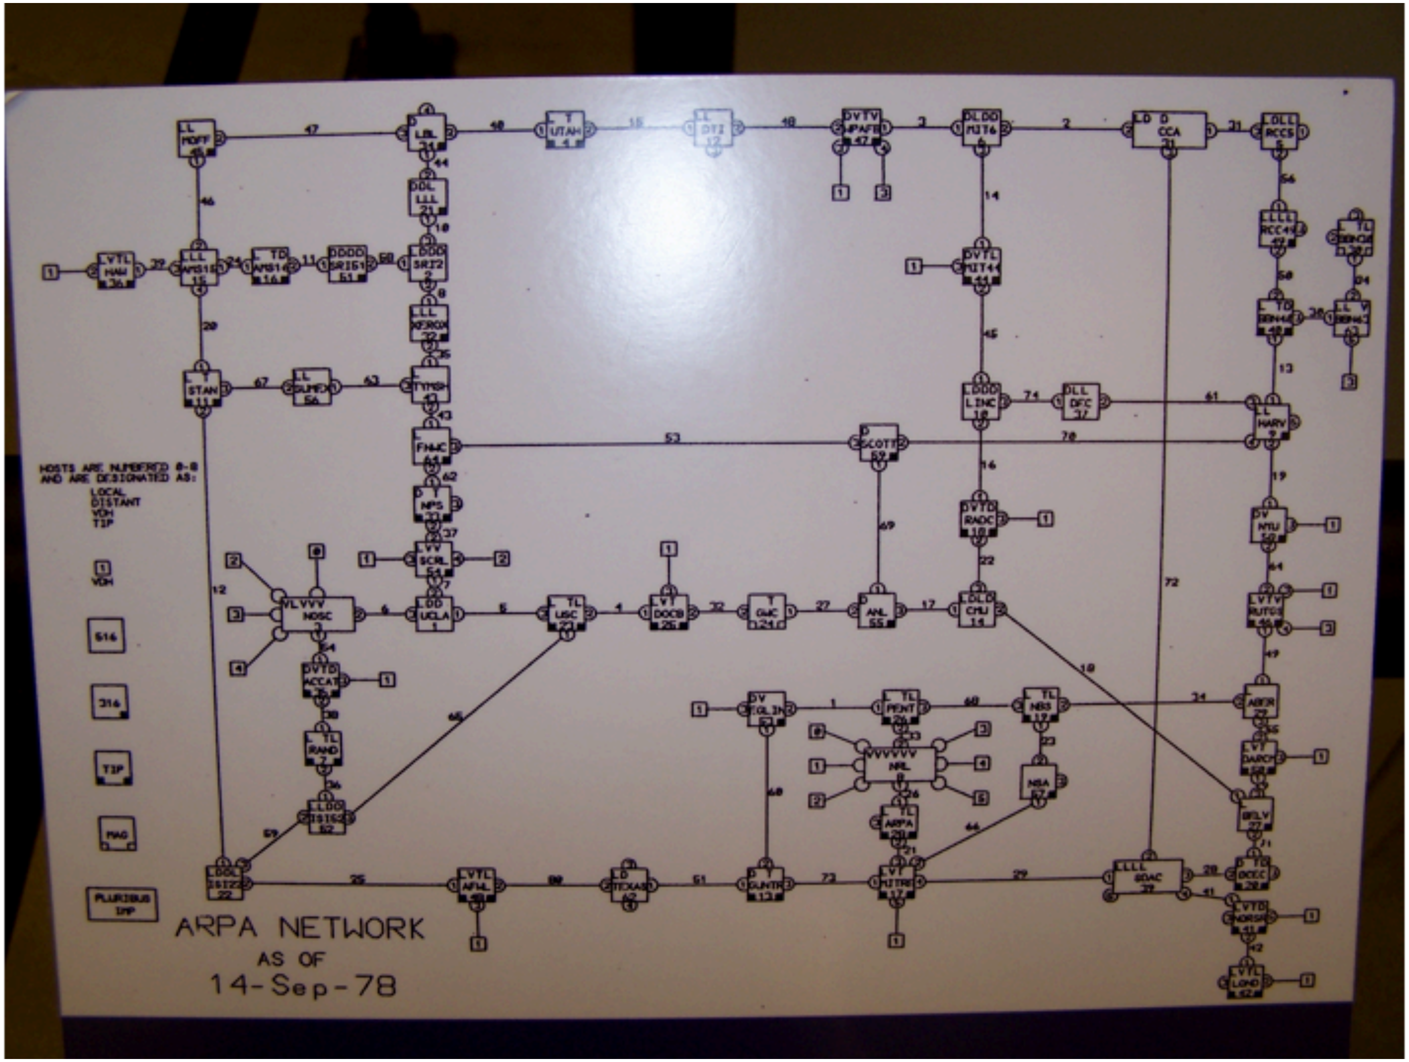

ARPAnet - September 1978

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

Vint Cerf, APRICOT, Feb 2011

"This time, for sure!" *

* Actually Vint didn't say this!

IP Version 4

- 32 bit address field
  - That's 4,294,967,296 addresses
- A triumph of minimalism
  - Basic datagram architecture
  - Stateless network with admission control and without active resource management
  - Variable packet size with fragmentation on the fly
  - Basic header set: Source, Destination, Fragmentation Control, checksums, all in 20 octets
  - Decoupled framework of related functions

It Worked!

- The minimal approach allowed for more efficient use of the common network
  - It was cheap
  - It was easy
  - It scaled

It Worked too well!

- Back in 1983 noone ever truly believed that IP would be the single communications protocol for the 21st century
  - And you would be mad for thinking that
  - OSI was meant to be the answer
  - And we understood so little about computing and communications that it was equally possible that we would find something better than packet switching pretty soon
  - So 32 bits of address space was looking like a decent engineering tradeoff

The cloud on the horizon...

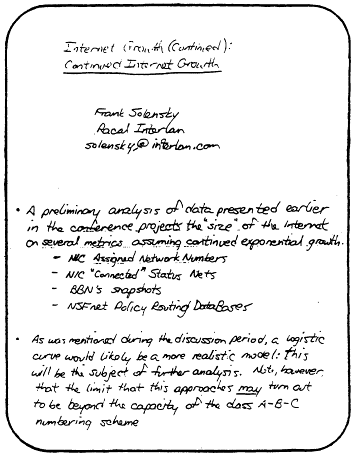

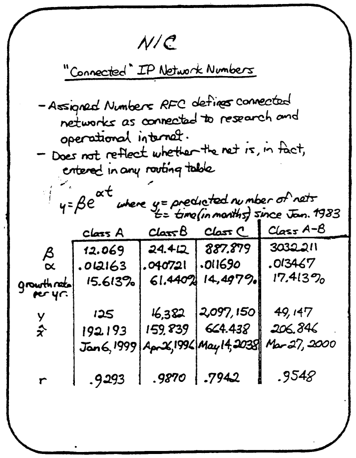

	Class A	Class B	Class C	Class A-B
β	12.069	24.412	887.879	3032.211
α	.012163	.040721	.011690	.013467
growth rate per yr.	15.613%	61.440%	14.497%	17.413%
y	125	16,382	2,097,150	49,147
x̂	192.193	159.839	664.438	206.846
	Jan 6, 1999	Apr 26, 1996	May 14, 2038	Mar 27, 2000
r	.9293	.9870	.7942	.9548

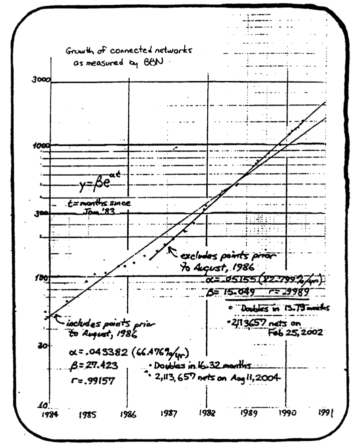

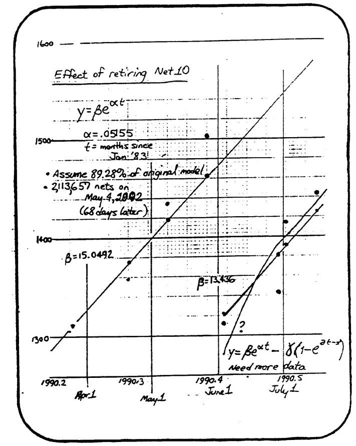

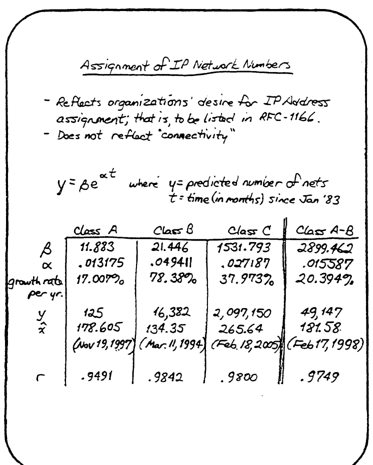

	Class A	Class B	Class C	Class A-B
β	11.883	21.446	1531.793	2899.462
α	.013175	.049411	.027187	.015587
growth rate per yr.	17.007%	78.38%	37.973%	20.394%
y	125	16,382	2,097,150	49,147
x̂	178.605	134.35	265.64	181.58
	(Nov 19, 1997)	(Mar. 11, 1994)	(Feb. 18, 2005)	(Feb 17, 1998)
r	.9491	.9842	.9800	.9749

59

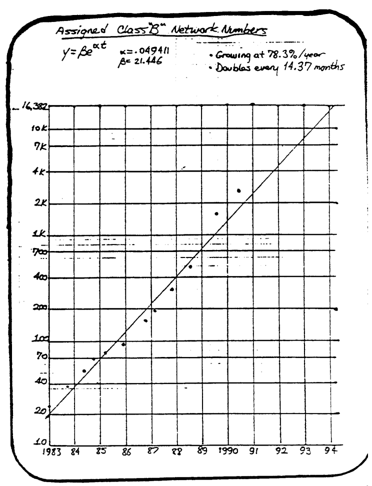

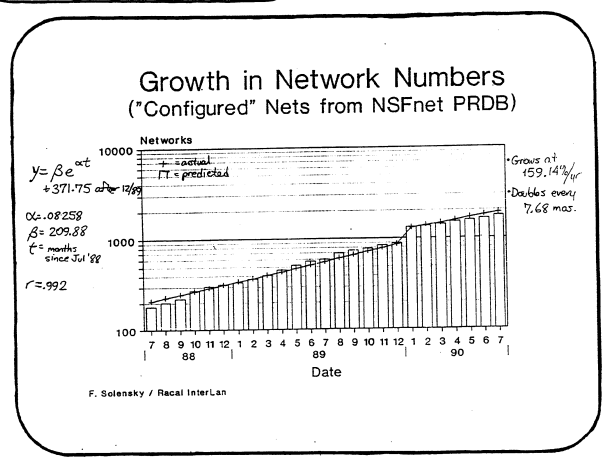

60

Frank Solensky, IETF, August 1990

Doomsday – Mk 1

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.

Frank Solensky, IETF, August 1990

IPv4 Address Allocations

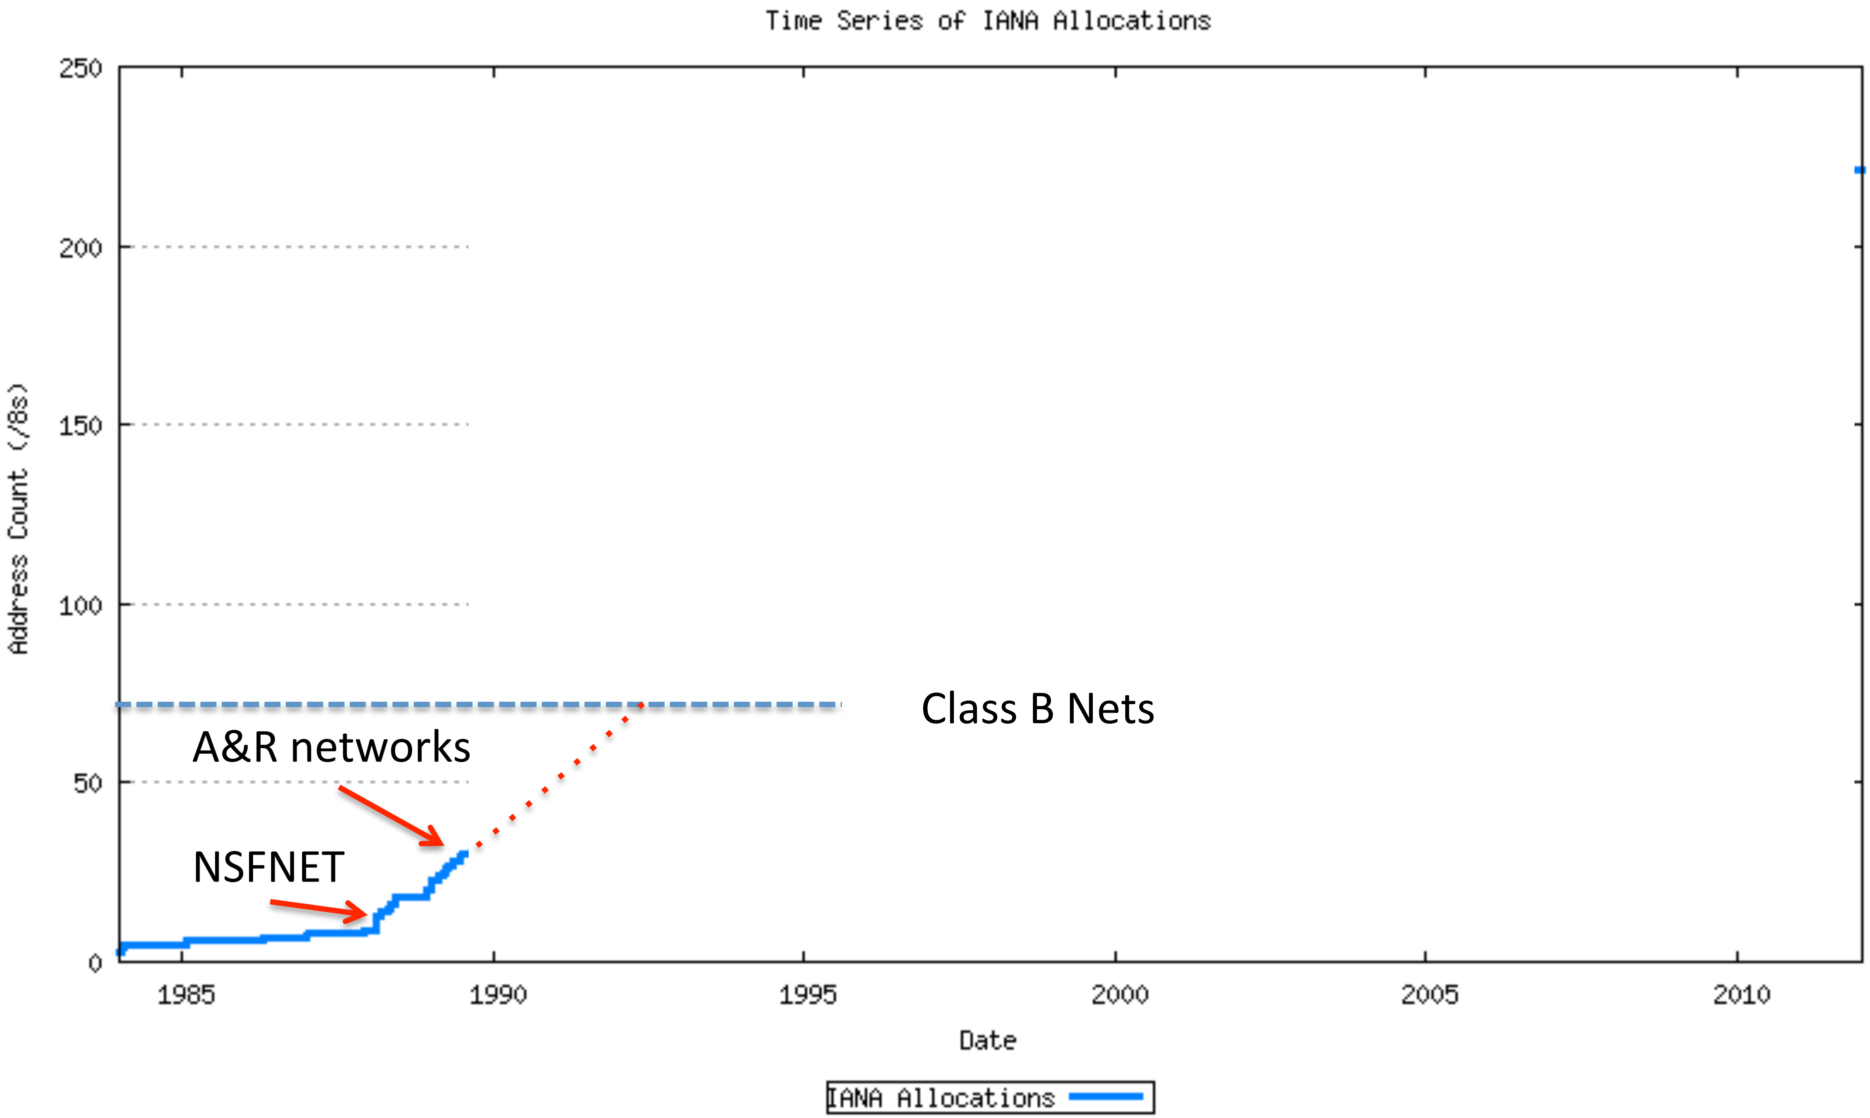

The CIDR Band-Aid

- It was clear by 1991 that we needed a new protocol
  - There was just no way we could hack extra bits of address space into the IPv4 header
  - And maybe we should think about what should/and should not be in the packet header at the same time as we enlarged the address size
- So we needed to buy a few years of breathing space
  - We did this by removing the fixed network/host boundary points
  - Classless Inter-Domain Routing was rushed in as a quick fix

The CIDR Fix

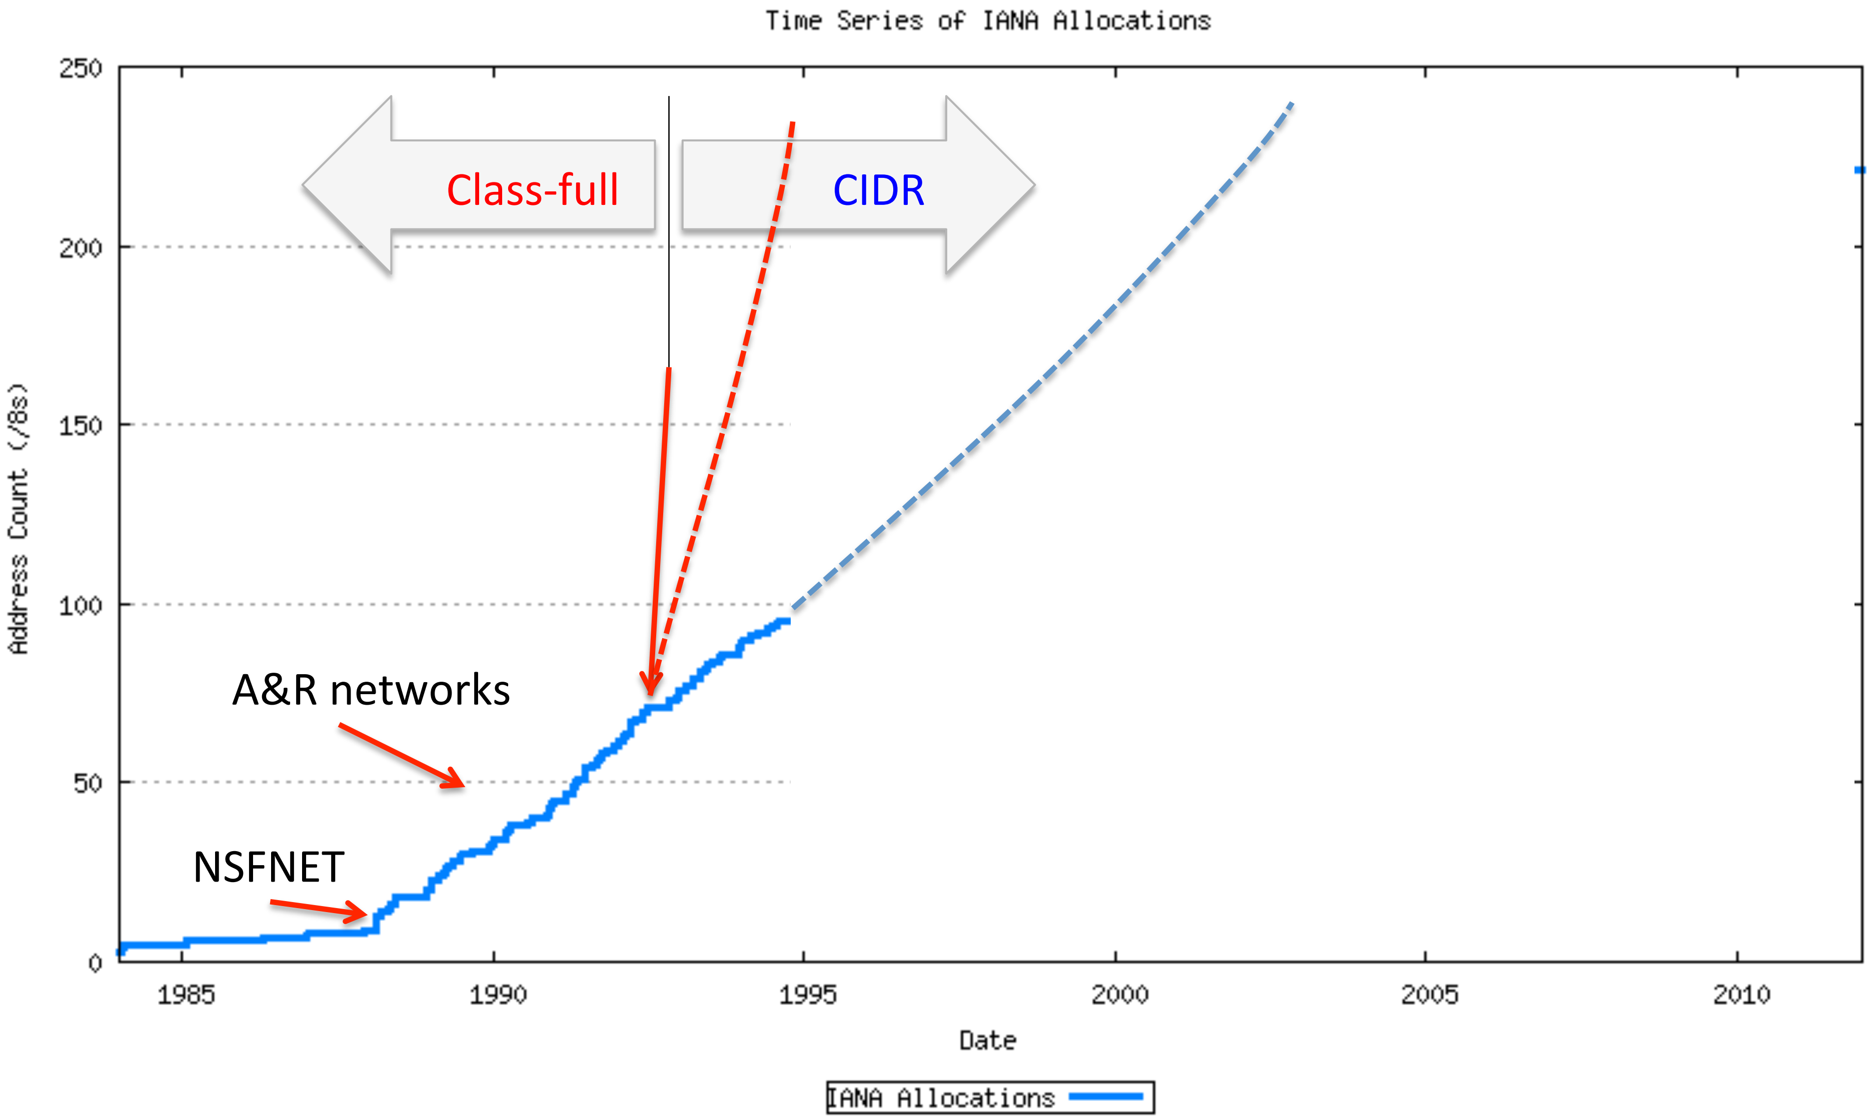
Time Series of IANA Allocations
250
200
150
100
50
0
Address Count (/8s)
1985
1990
1995
2000
2005
2010
Date
IANA Allocations
Class-full
CIDR
A&R networks
NSFNET

And the long term plan?

IPng

- There was no OSI any more, so this had to be the one and only protocol in the eyes of the protocol designers
- It was envisaged to have a lifespan of 30 – 100 years
- And encompass ubiquitous deployment to the order of trillions of connected nodes

Problem Solved!

- We set the protocol designers onto the problem
  - we were naive enough to think that a committee could engineer a better architecture
- We planned to worry about transition on a later day once the protocol design had been worked out
- And we turned back to building the network
  - and making money
  - a LOT of money, as it turned out

zzzzzz

For a while the problem of the need for a new protocol became LESS urgent

- The network grew at ever faster rates
- But CIDR allowed us to use vastly fewer addresses
- And then consumer NATS allowed us to use even fewer addresses
- So IPv6 became a perennial "sometime" issue that never quite became a "now" item

IPv4 + CIDR + NATs worked!

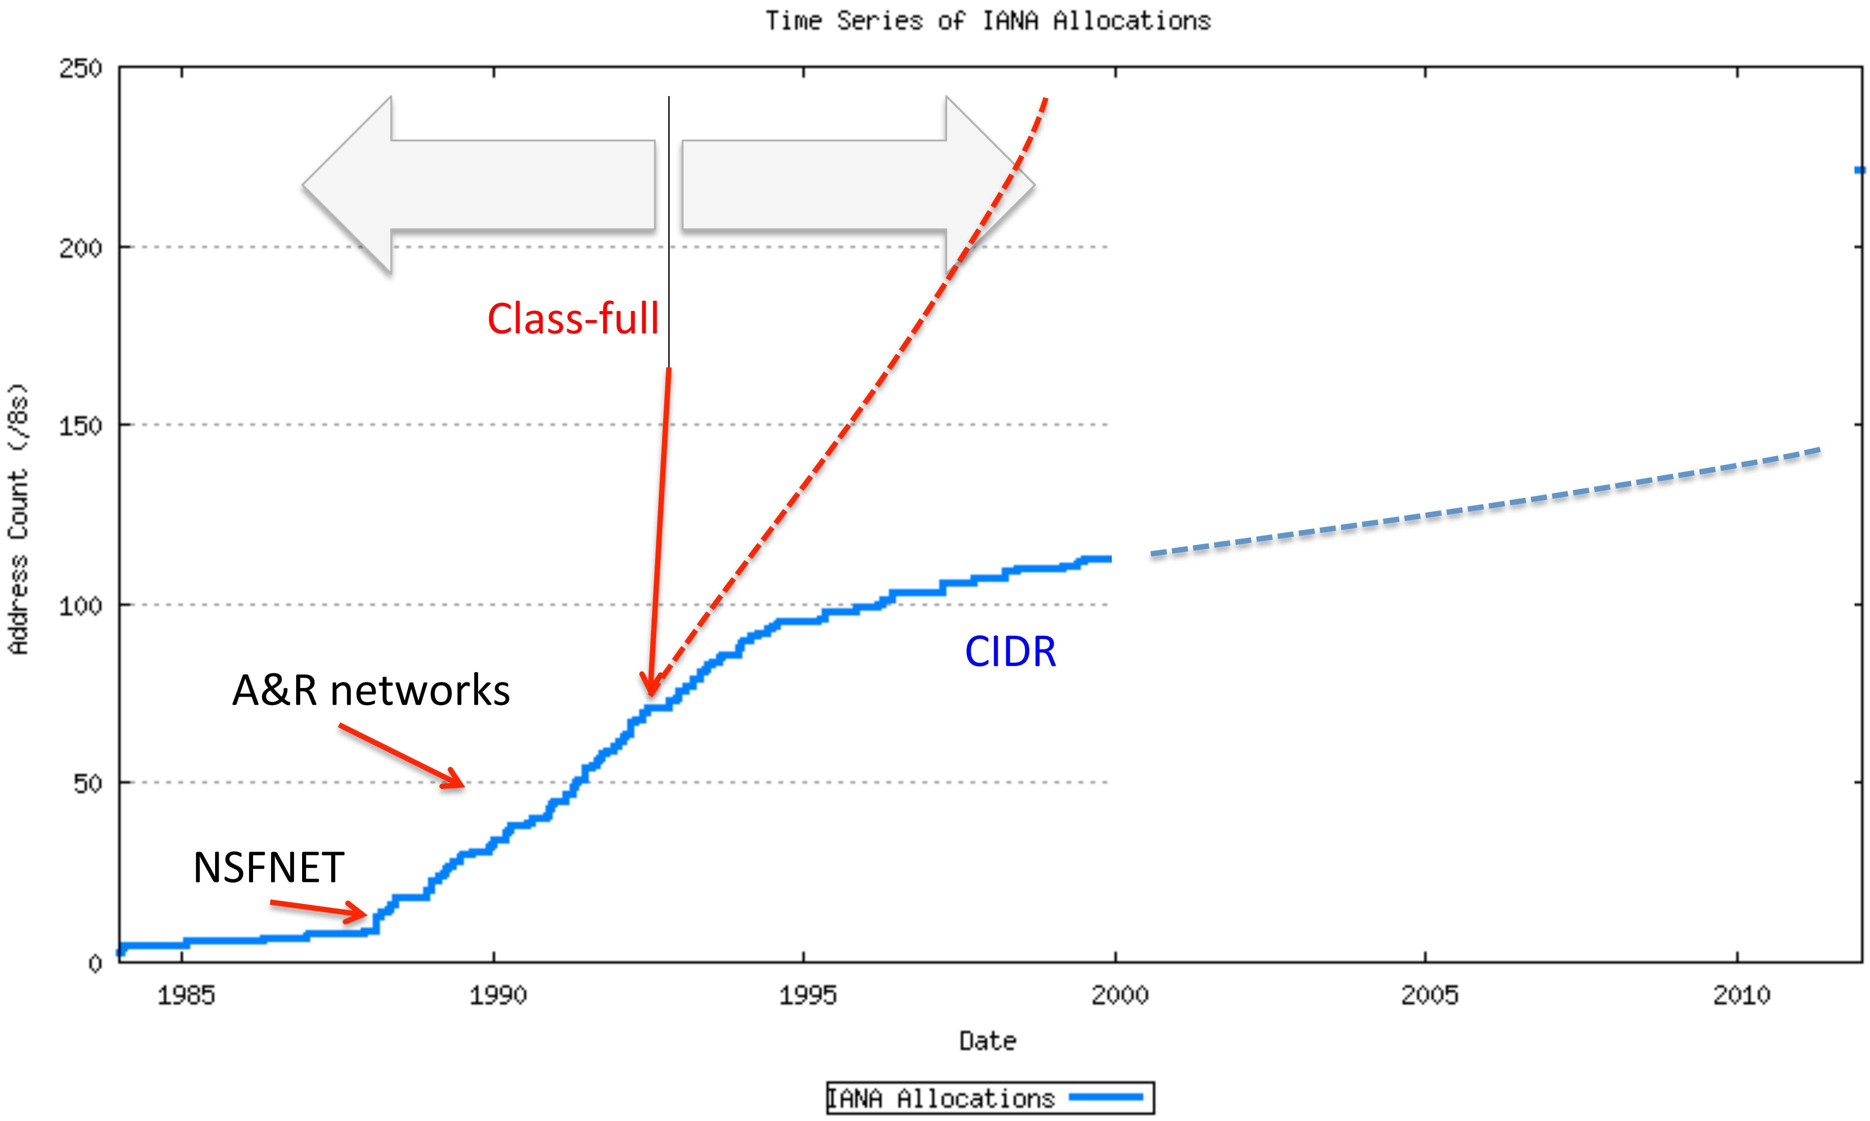

And worked...

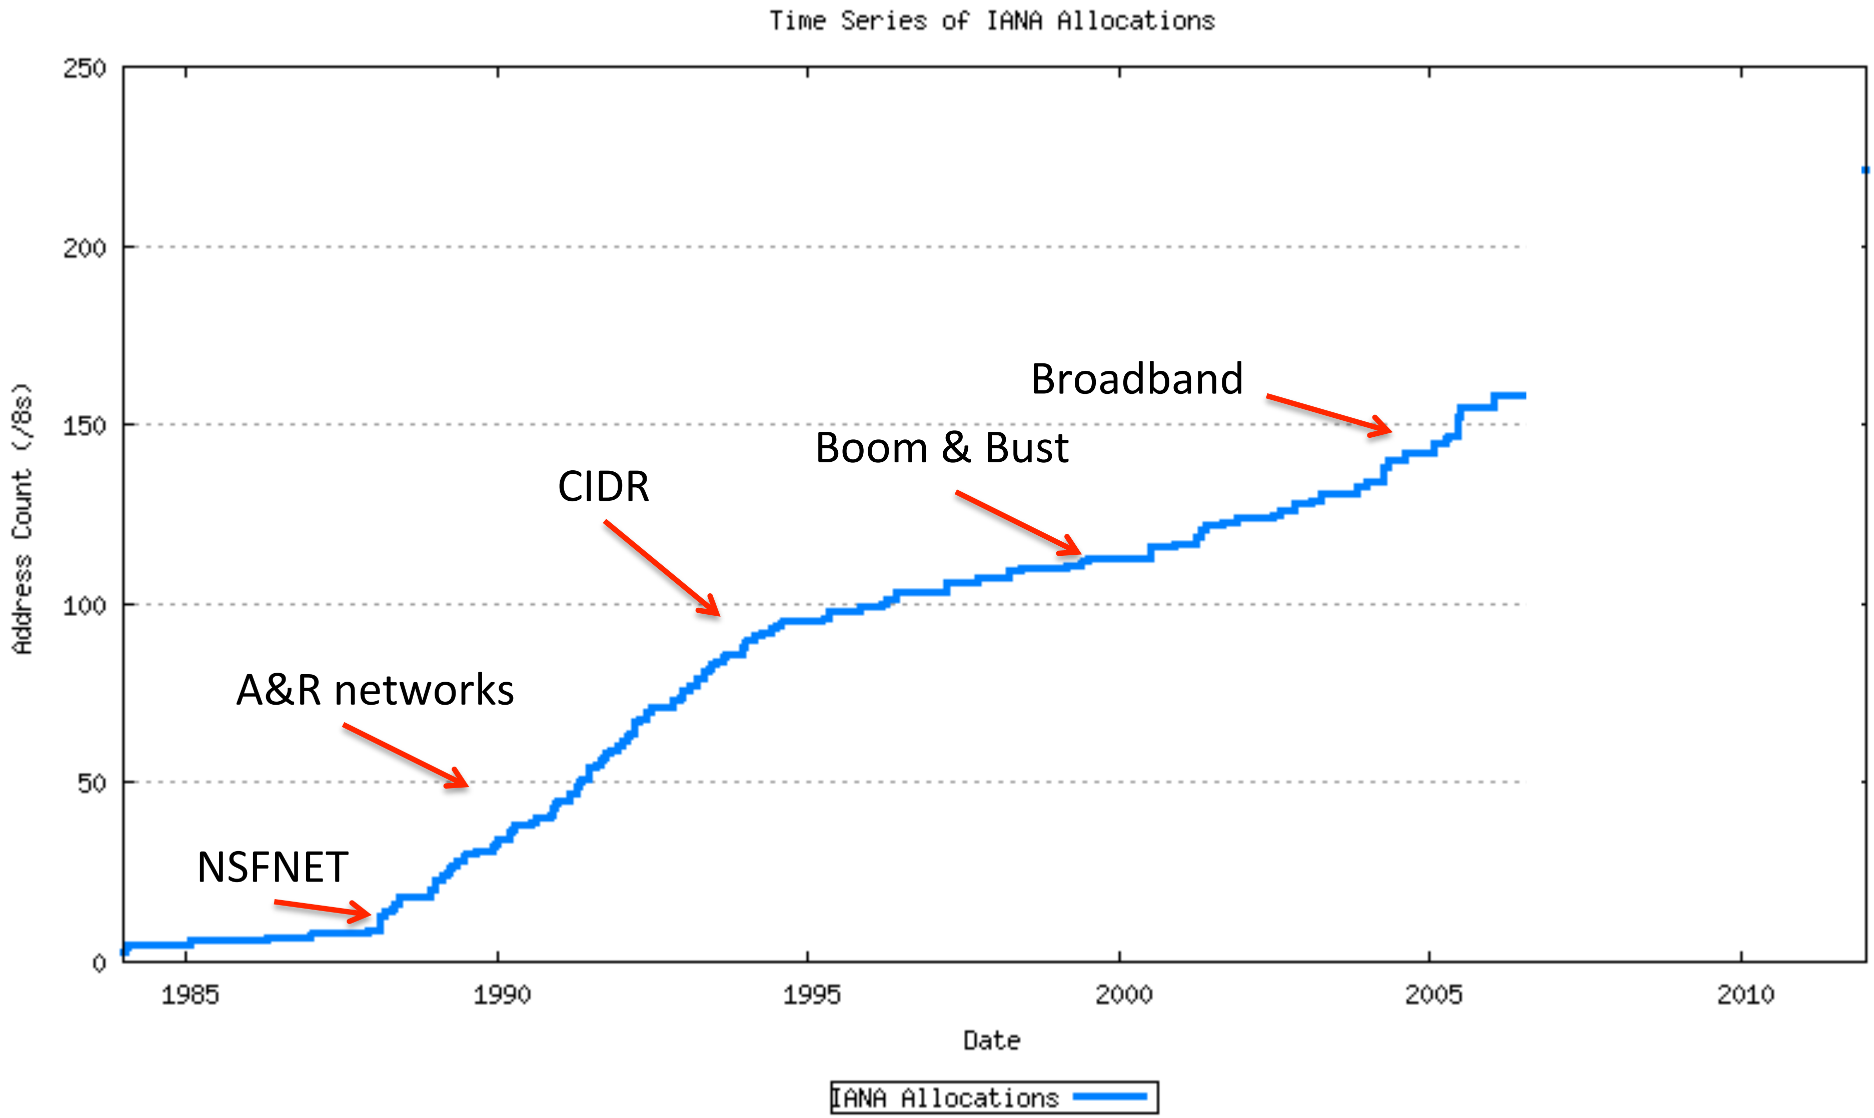

Time Series of IANA Allocations
Address Count (/8s)
250
200
150
100
50
0
1985
1990
1995
2000
2005
2010
Date
IANA Allocations
NSFNET
A&R networks
CIDR
Boom & Bust
Broadband

And worked!

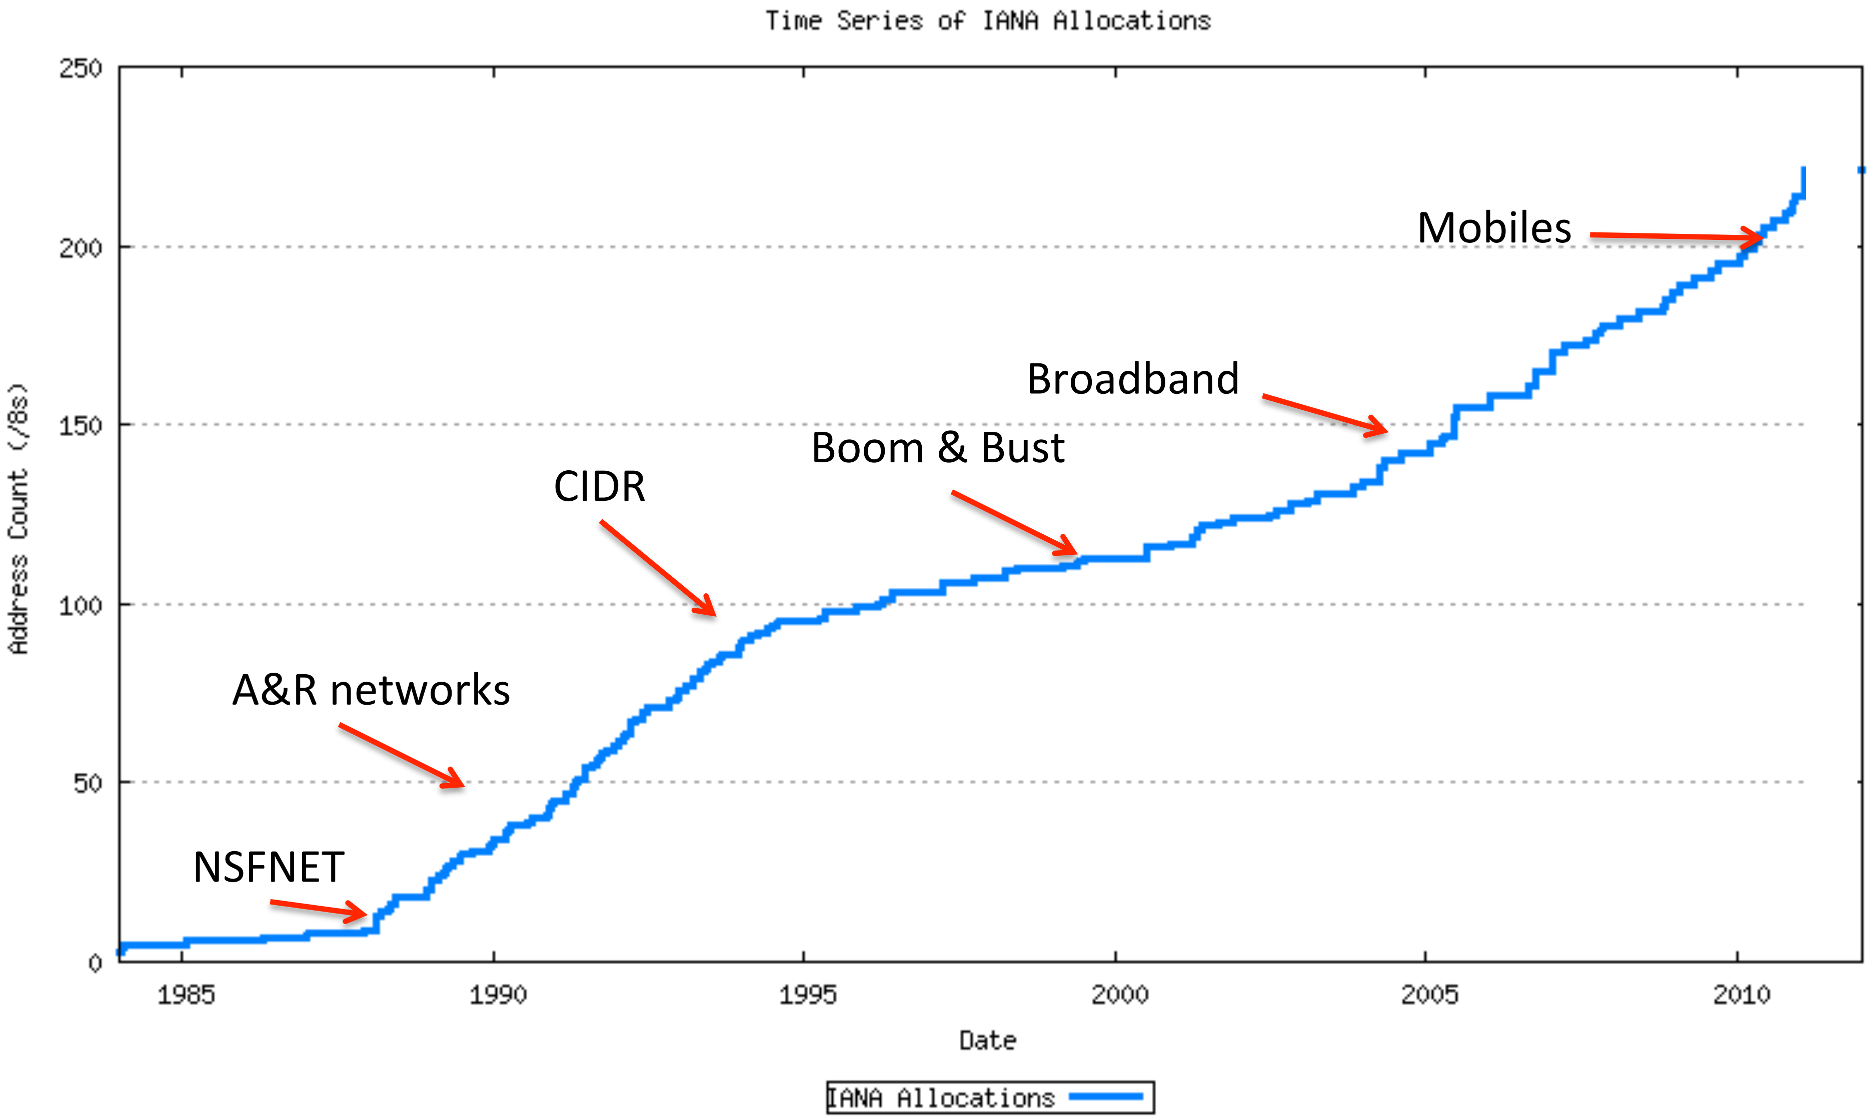

Time Series of IANA Allocations
Address Count (/8s)
0
50
100
150
200
250
1985
1990
1995
2000
2005
2010
Date
IANA Allocations
NSFNET
A&R networks
CIDR
Boom & Bust
Broadband
Mobiles

Until it didn't work any more!

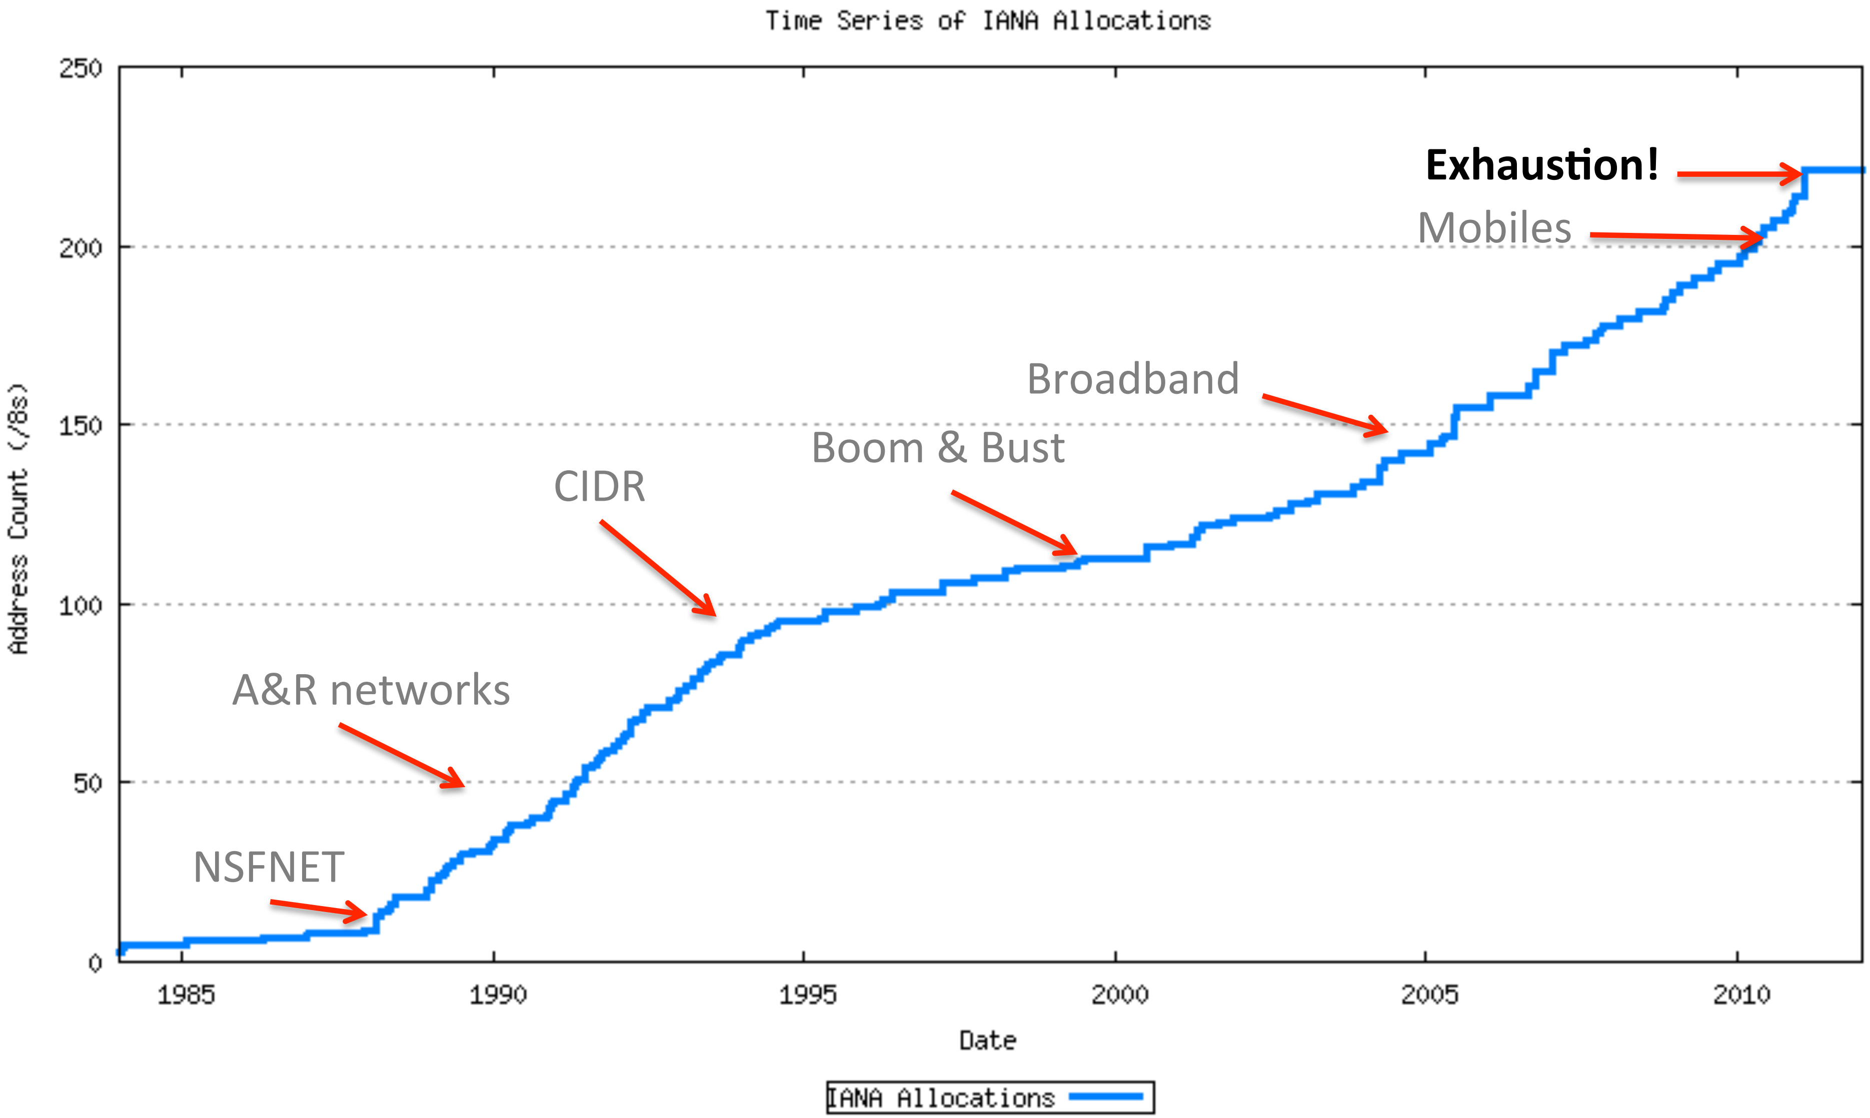

3 February 2011

Free Pool of IPv4 Address Space Depleted

IPv6 adoption at critical phase

Montevideo, 3 February 2011 – The Number Resource Organization (NRO) announced today that the free pool of available IPv4 addresses is now fully depleted. On Monday, January 31, the Internet Assigned Numbers Authority (IANA) allocated two blocks of IPv4 address space to APNIC, the Regional Internet Registry (RIR) for the Asia Pacific region, which triggered a global policy to allocate the remaining IANA pool equally between the five RIRs. Today IANA allocated those blocks. This means that there are no longer any IPv4 addresses available for allocation from the IANA to the five RIRs.

IANA assigns IPv4 addresses to the RIRs in blocks that equate to 1/256th of the entire IPv4 address space. Each block is referred to as a "/8" or "slash-8". A global policy agreed on by all five RIR communities and ratified in 2009 by ICANN, the international body responsible for the IANA function, dictated that when the IANA IPv4 free pool reached five remaining /8 blocks, these blocks were to be simultaneously and equally distributed to the five RIRs.

"This is an historic day in the history of the Internet, and one we have been anticipating for quite some time," states Raúl Echeberría, Chairman of the Number Resource Organization (NRO), the official representative of the five RIRs. "The future of the Internet is in IPv6. All Internet stakeholders must now take definitive action to deploy IPv6."

"This is truly a major turning point in the on-going development of the Internet," said Rod Beckstrom, ICANN's President and Chief Executive Officer. "Nobody was caught off guard by this, the Internet technical community has been planning for IPv4 depletion for quite some time. But it means the adoption of IPv6 is now of paramount importance, since it will allow the Internet to continue its amazing growth and foster the global innovation we've all come to expect."

Ooops!

IPv6 is the "next generation" of the Internet Protocol, providing a hugely expanded address space and allowing the Internet to grow into the future. "Billions of people world wide use the Internet for everything from sending tweets to paying bills. The transition to IPv6 from IPv4 represents an opportunity for even more innovative applications without the fear of running out of essential Internet IP addresses," said Vice President of IANA Elise Gerich.

Adoption of IPv6 is now vital for all Internet stakeholders. The RIRs have been working with network operators at the local, regional, and global level for more than a decade to offer training and advice on IPv6 adoption and ensure that everyone is prepared for the exhaustion of IPv4.

"Each RIR will have its final full /8 from IANA, plus any existing IP address holdings to distribute. Depending on address space requests received, this could last each RIR anywhere from a few weeks to many months. It's only a matter of time before the RIRs and Internet Service Providers (ISPs) must start denying requests for IPv4 address space. Deploying IPv6 is now a requirement, not an option," added Echeberría. IPv6 address space has been available since 1999. Visit http://www.nro.net/ipv6/ for more information on IPv6, or

IPocalypse?

Maybe not

- Its a massive industry
- And exhaustion is not a sudden state change

- But the network grew by more than 280 million services in 2011
  - Which was the largest year so far for the Internet

It’s more like this!

Transition is hard!

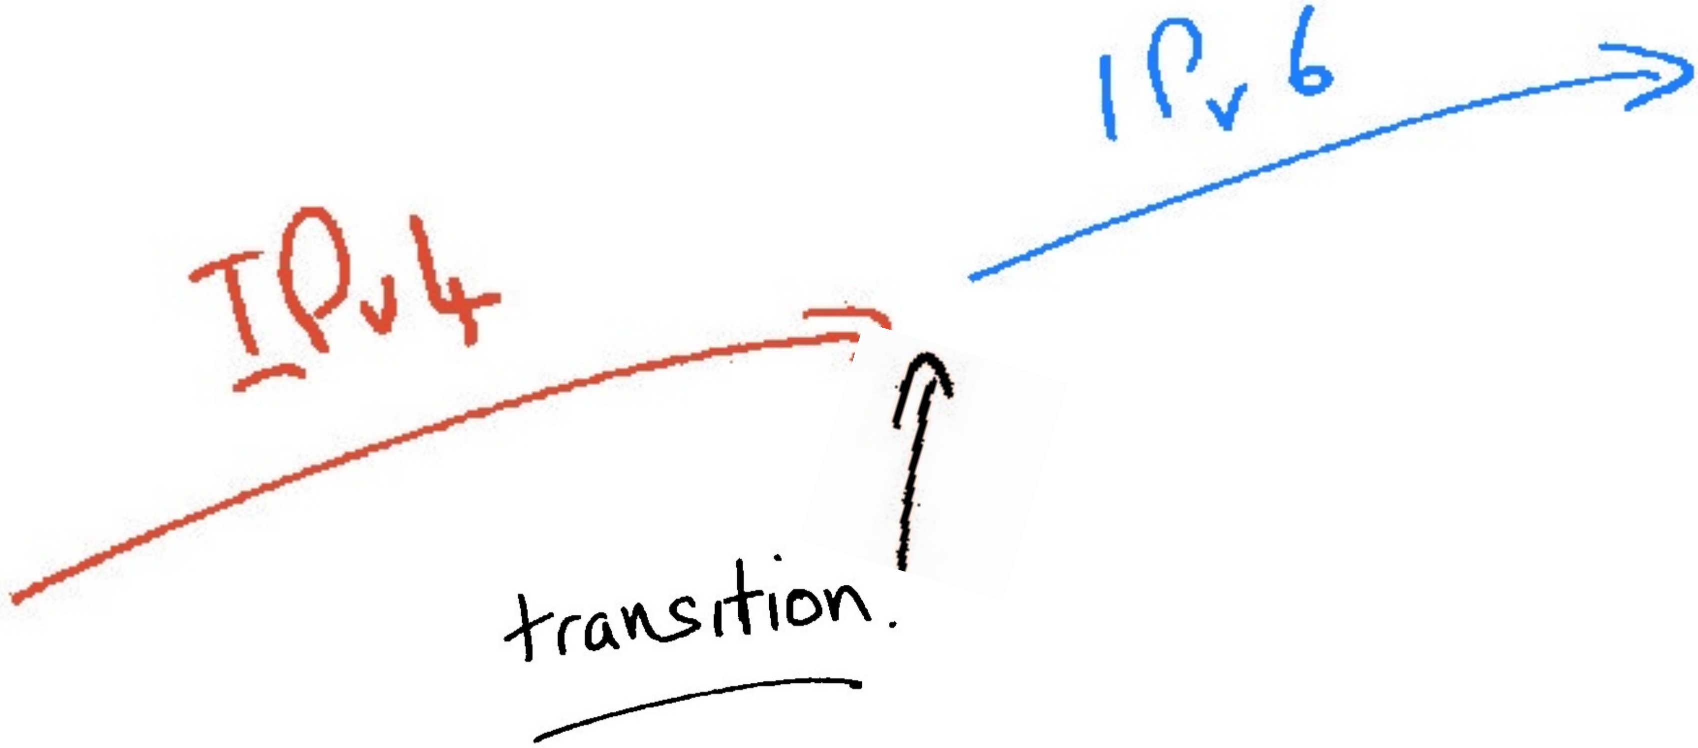

Switchover?

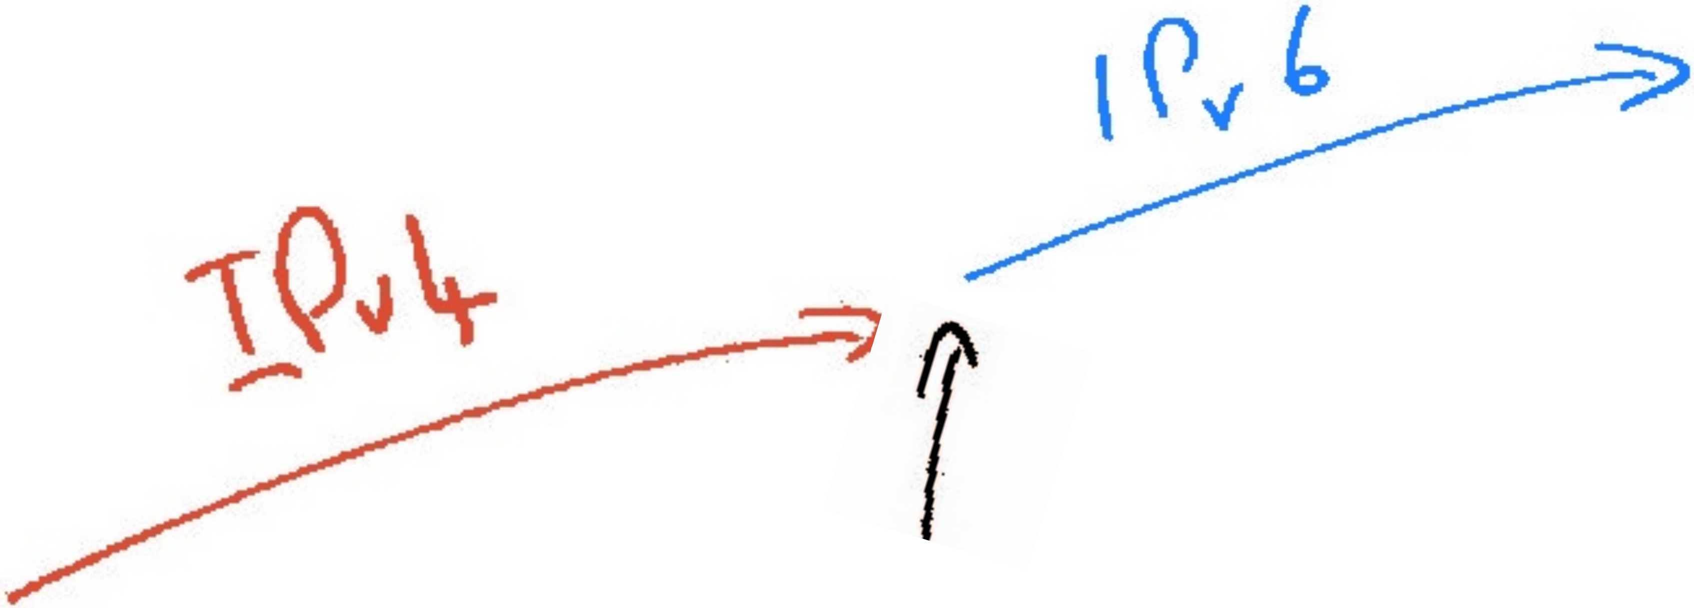

We all agree to turn off IPv4 and turn on IPv6 EVERYWHERE

All at the same time! All over the Internet!

Switchover?

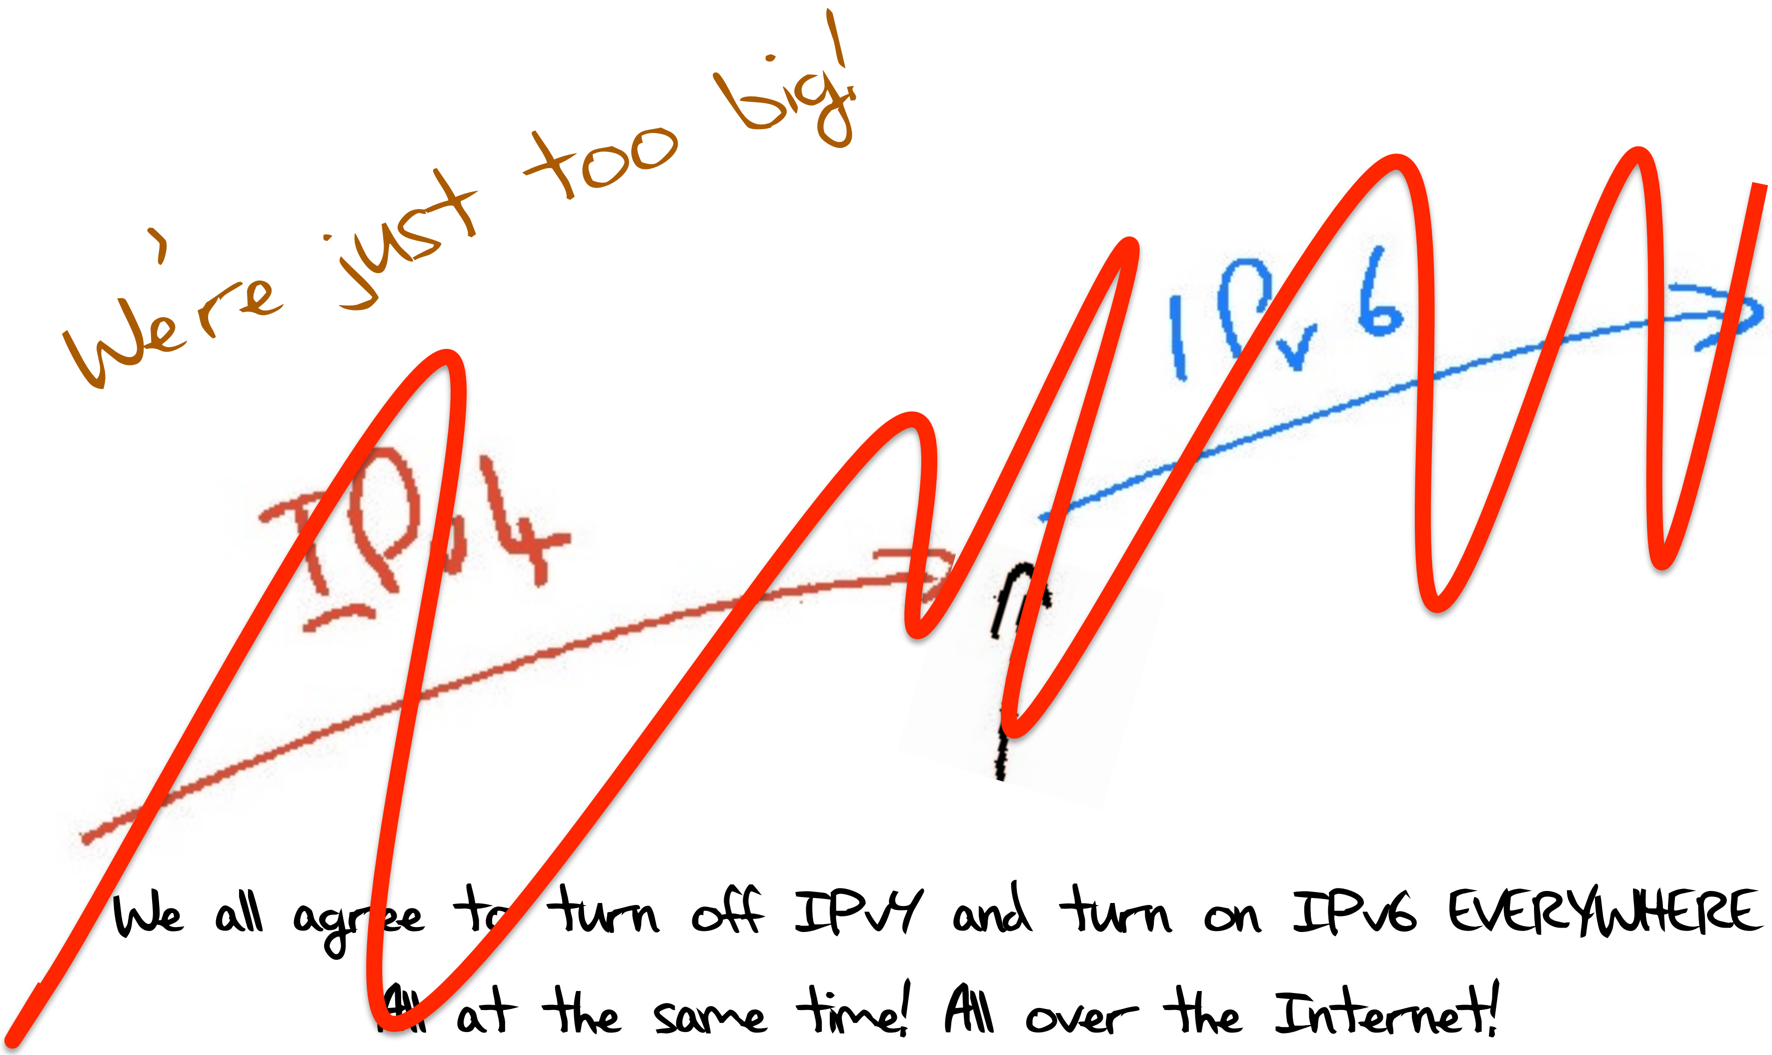

We all agree to turn off IPv4 and turn on IPv6 EVERYWHERE

All at the same time! All over the Internet!

Piecemeal Switchover?

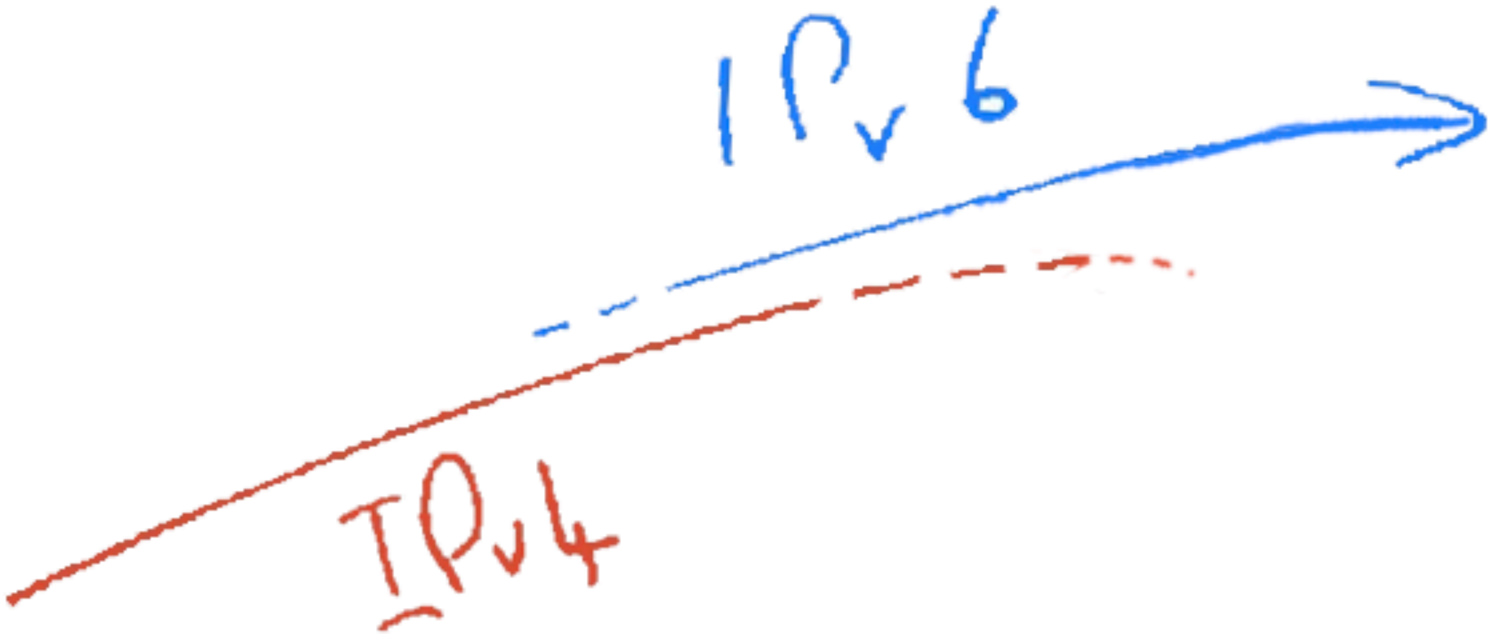

One-by one networks switchover to Ipv6,

Switching off ipv4 when they complete their transition

Piecemeal Switchover?

IPv6 is NOT backwards compatible!

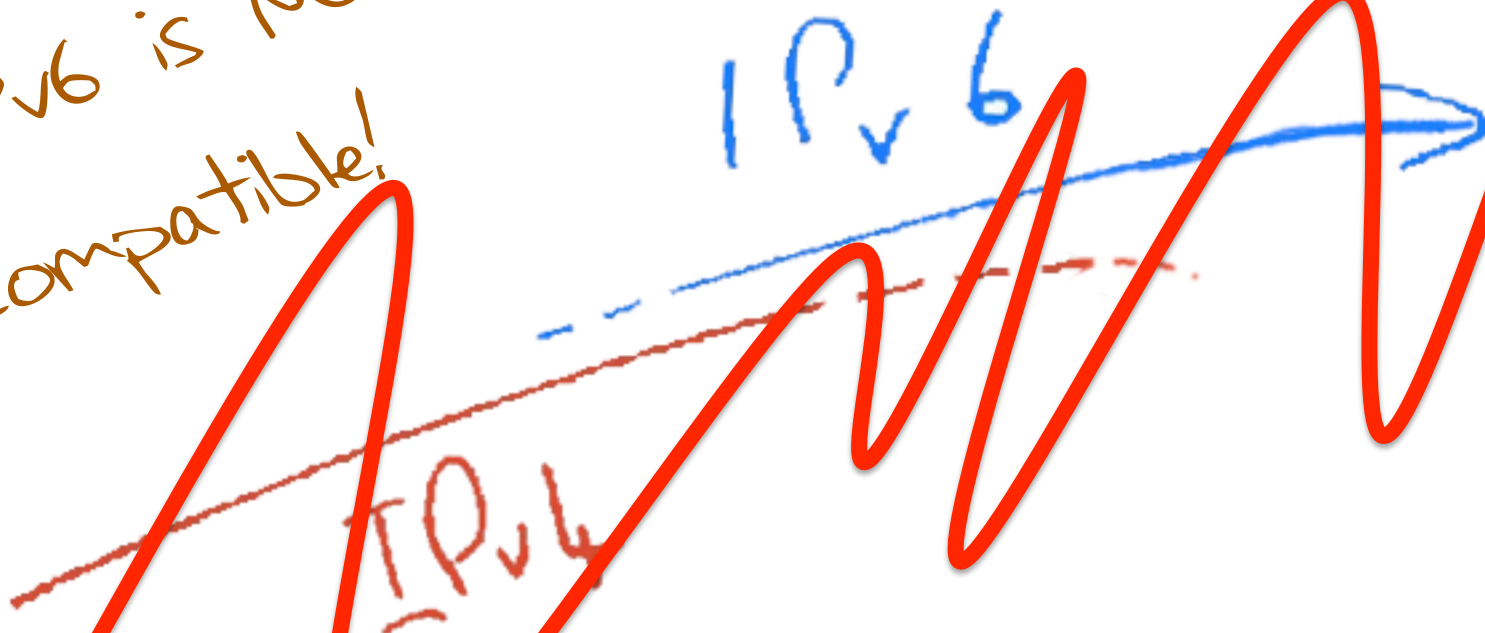

One-by one networks switchover to IPv6,
Switching off ipv4 when they complete their transition

Dual Stack Transition

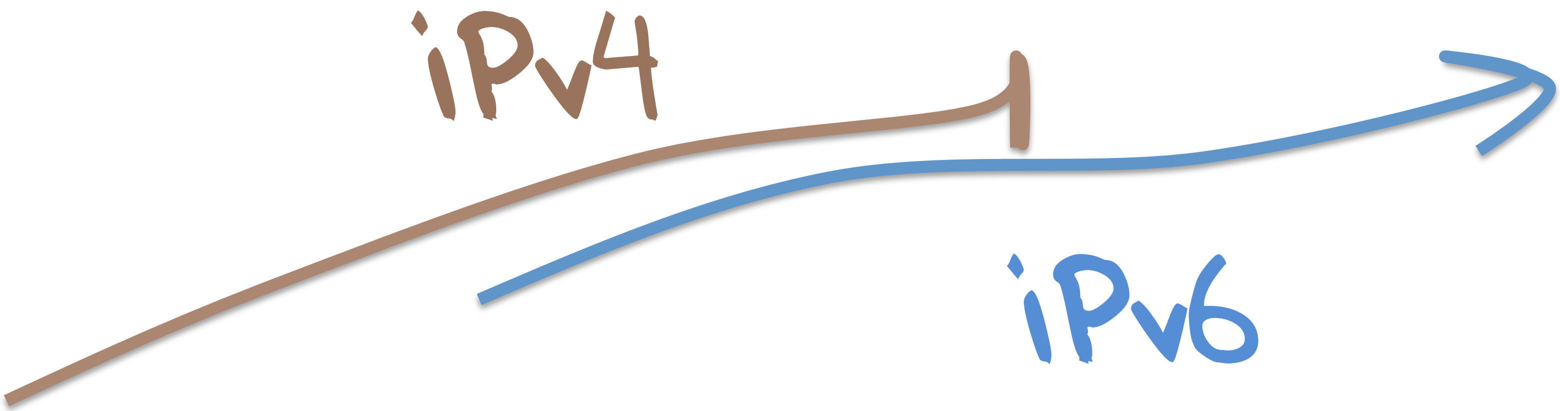

One-by one networks and hosts have IPv6 added to IPv4

Switching off ipv4 when every element has both IPv4 and IPv6

Dual Stack Transition

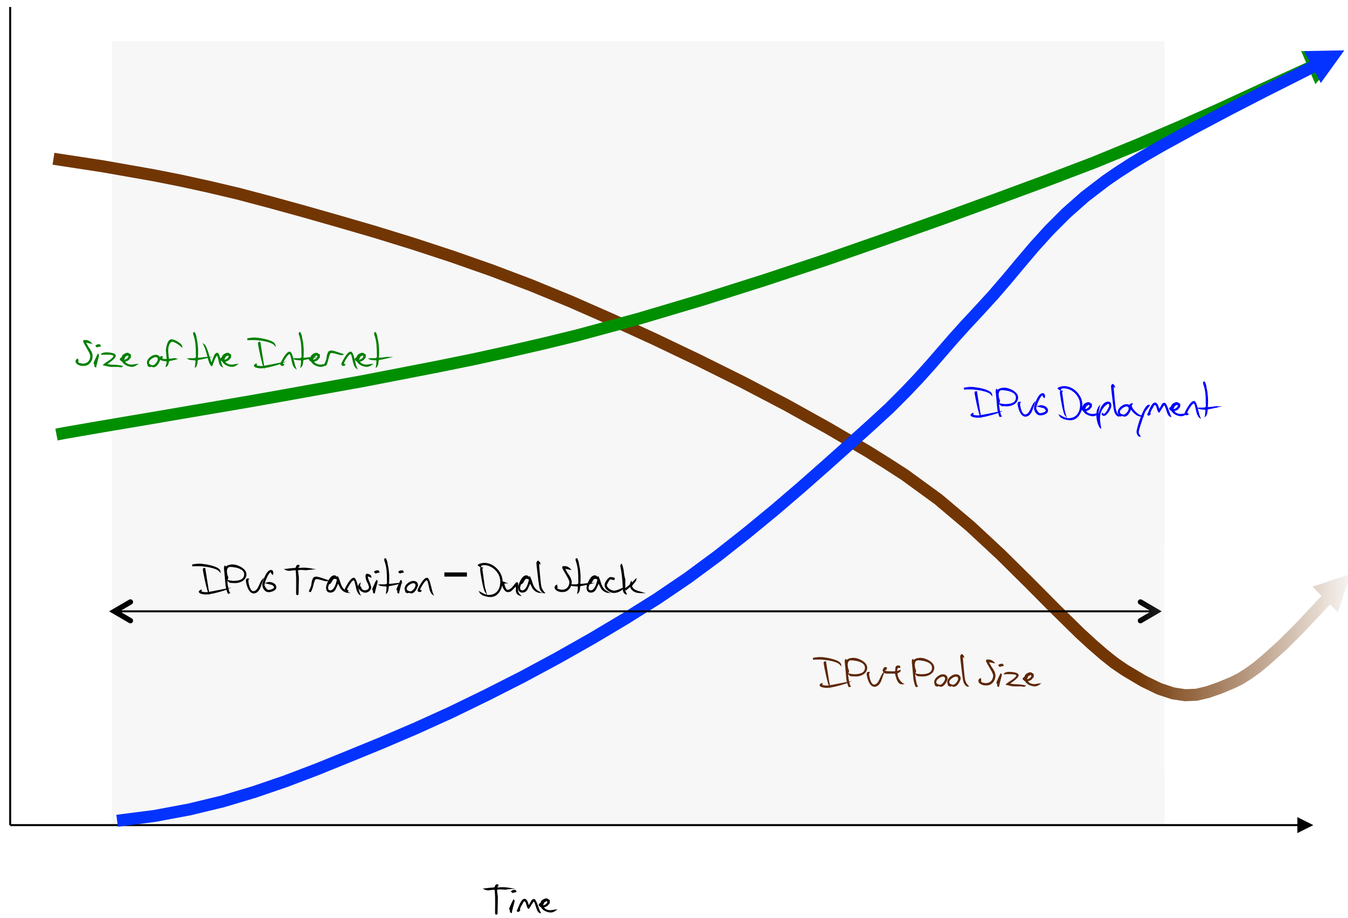

For this to work we have to start early and finish BEFORE IPv4 address pool exhaustion

Dual Stack Transition

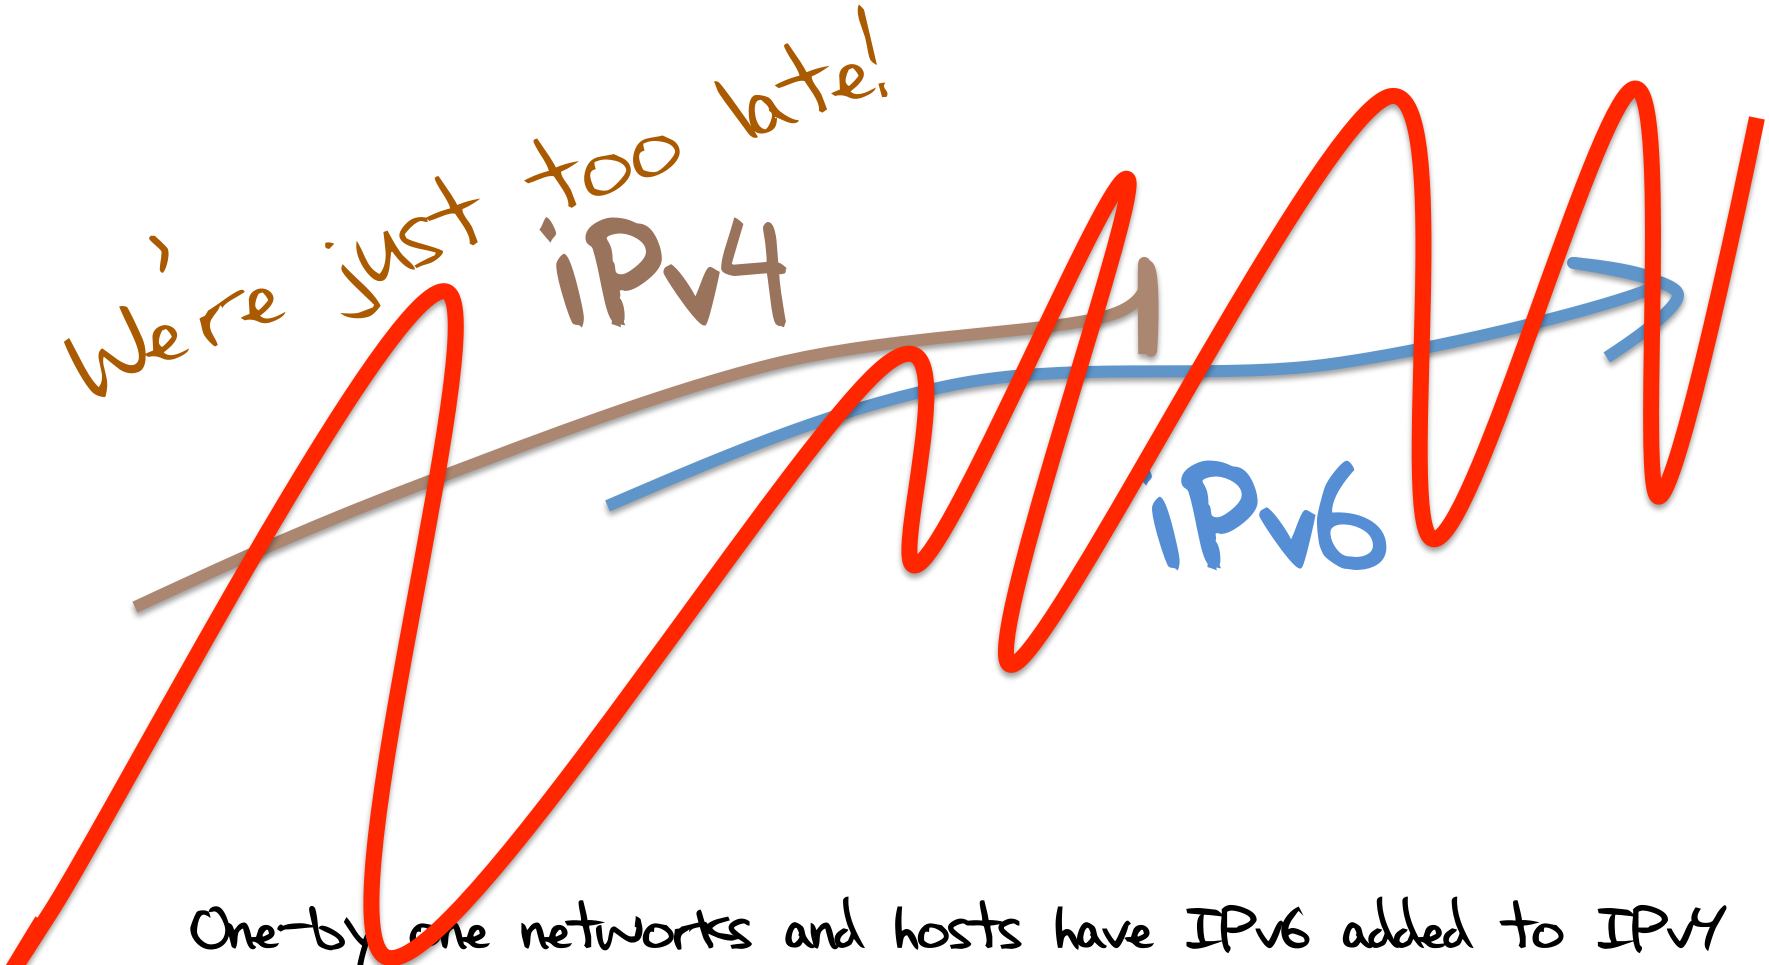

One-by-one networks and hosts have IPv6 added to IPv4

Switching off ipv4 when every element has both IPv4 and IPv6

Dual Stack Transition

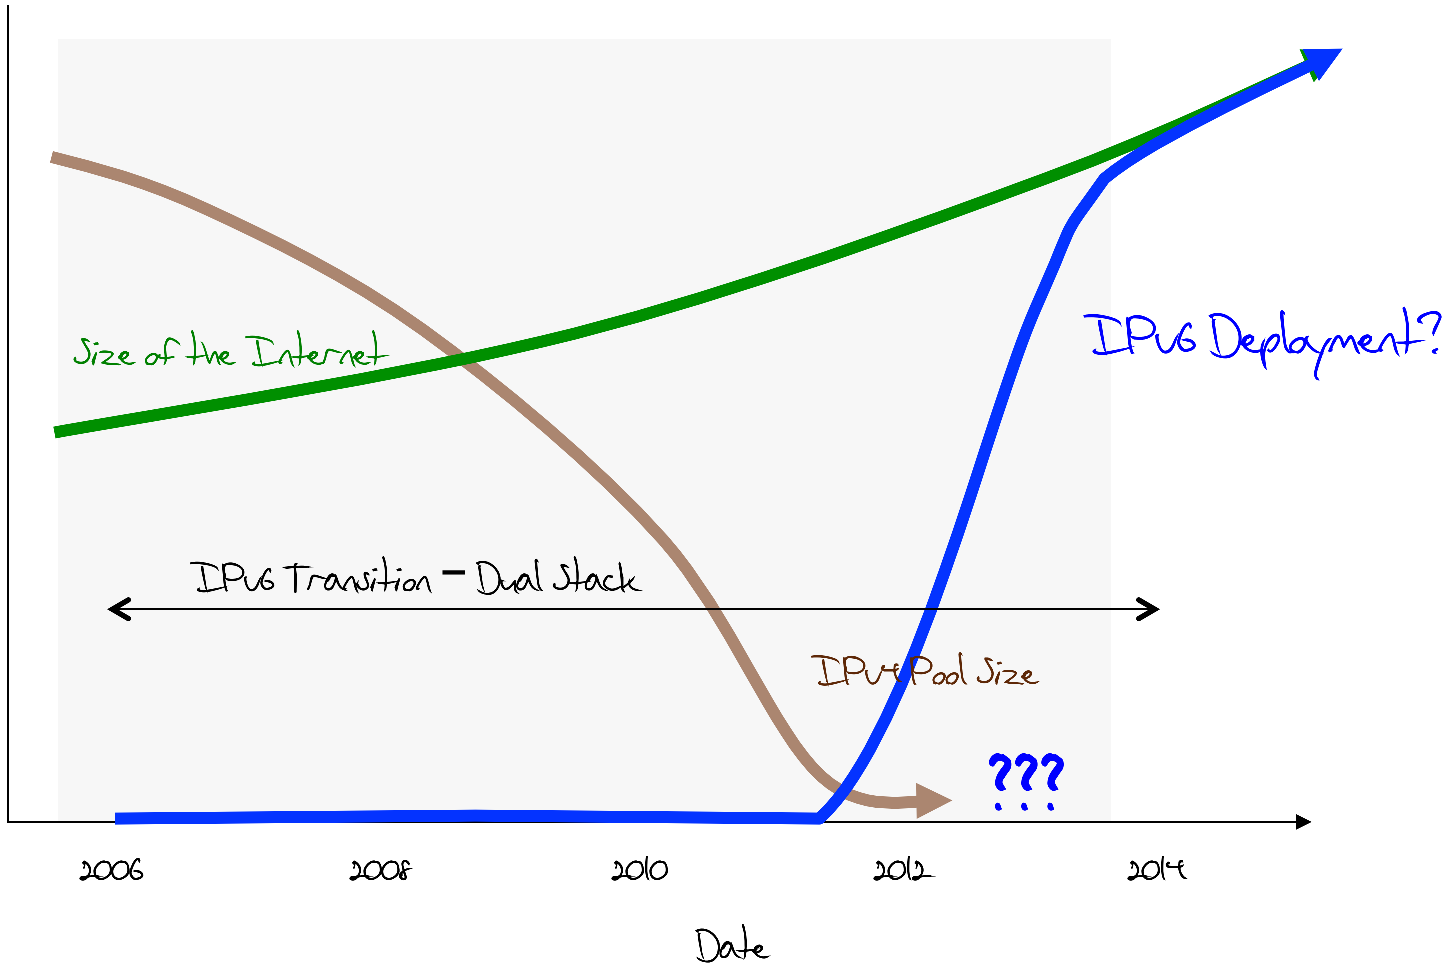

More Band Aids!

The increasing scarcity of IPv4 will force carriage providers to add address sharing mechanisms into the IPv4 network

Plan F – The Kitchen Sink Approach!

To get from "here" to "there" requires an excursion through an environment of CGNs, CDNs, ALGs and similar middleware 'solutions' to IPv4 address exhaustion

Where are we with IPv6?

It's a mixed story

- Some components of the Internet have had IPv6 for many years
  - There is far more IPv6 out there in the Internet if you know where to look
  - About one half of today's Internet devices show that they have an active IPv6 stack
- But some critical parts of the Internet are still determined not to make any shift away from IPv4
  - While one half of the Internet's devices have IPv6, less that 1 in a hundred devices can actually use IPv6 on the Internet

Where are we with IPv6?

- Every Host?
  - 50% do, 50% do not
  - Strong Points:
    - Microsoft Windows Vista and 7 with IPv6 on by default
    - Mac OSX with IPv6 on by default
    - Unix servers with IPv6 on by hand
  - Weak points:
    - Mobile devices do NOT have IPv6 in their radio systems
      - Any many do not have it at all so far
    - Waiting for the world to turn off XP

Where are we with IPv6?

- Every Network?
  - 50% of the transit networks do, 50% do not
  - 4% of the access networks (or less) do
    - Weak points:
      - DSL deployments with customer-owned CPE are a major impediment to transition
      - BRAS / BFLETs IPv4 only
      - CPE IPv6 story is patchy to bad
      - 3G networks are a problematical in GGSN services
      - 4G networks – still early days
  - Server / Data centre infrastructure weak
    - Not many of the load management products support IPv6
    - And dual stack in a data centre is messy
    - IPv6 internals with a dual stack external presentation is an efficient approach for a data centre – but few centres are willing to make the call and transition to Ipv6 yet

Where are we with IPv6?

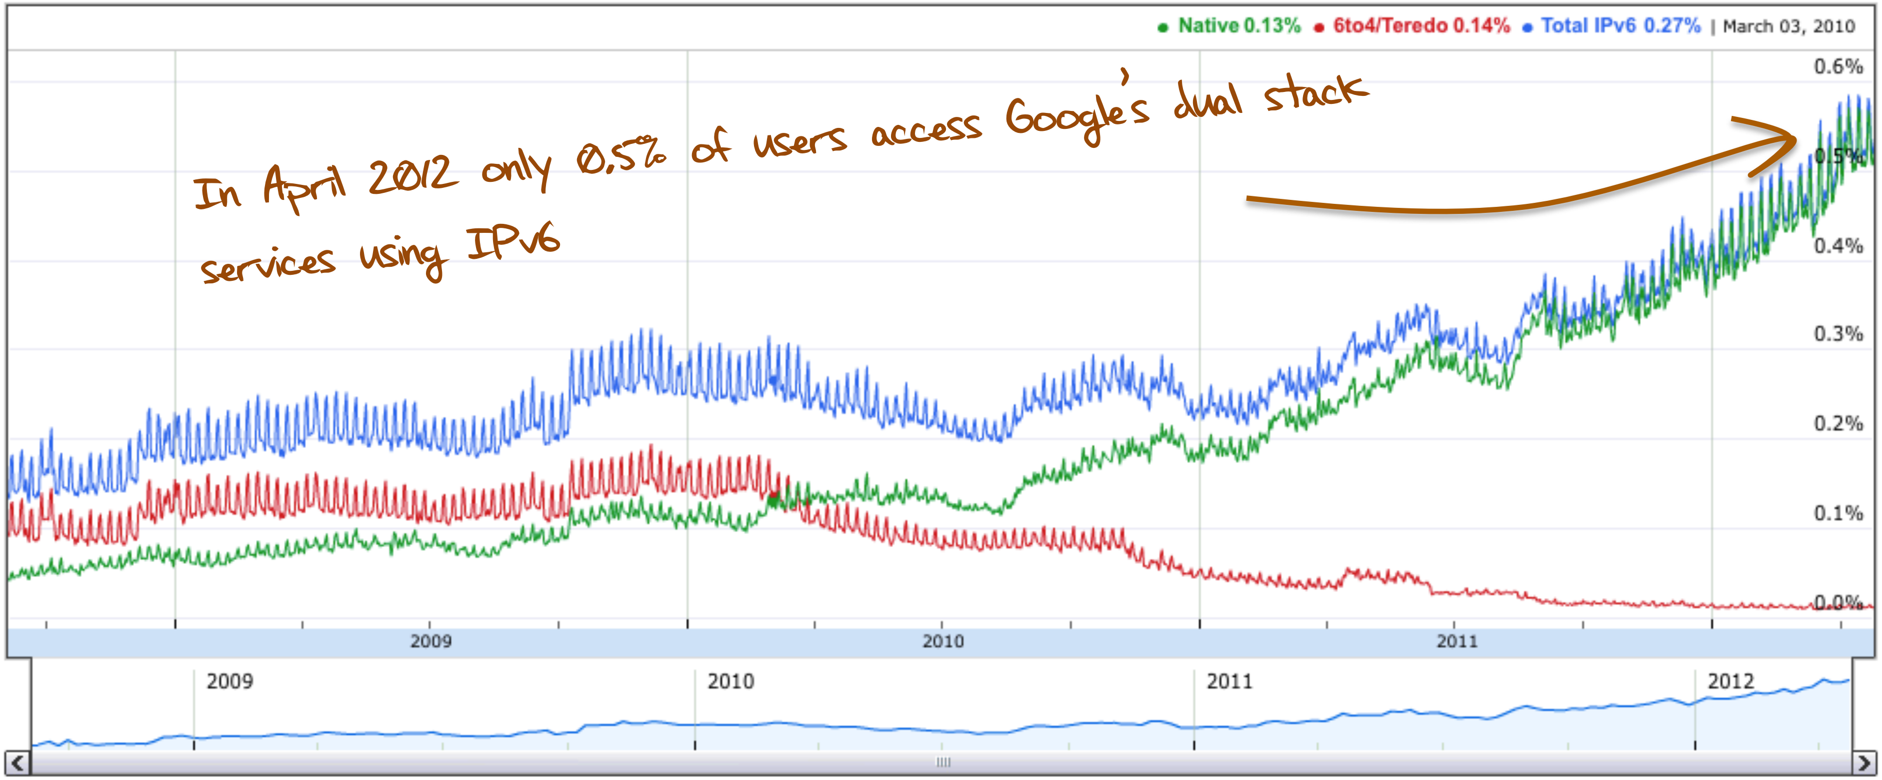

http://www.google.com/intl/en/ipv6/statistics/

Where is Australia with IPv6?

IPv6 Users by Country

Date: 01 Jul 2012

Index	ISO-3166 Code	Internet Users	V6 Use ratio	V6 Users (Est)	Population	Country
1	RO	8667240	7.52%	651776	22110307	Romania
2	FR	49976813	4.07%	2034056	64736805	France
3	LU	465300	2.83%	13167	509081	Luxembourg
4	EU	0	2.41%	0	0	European Union
5	JP	100956617	1.83%	1847506	126195772	Japan
6	US	247624831	1.06%	2624823	316251381	United States of America
7	CN	515779845	0.96%	4951486	1343176680	China
8	SI	1418244	0.90%	12764	1997527	Slovenia
9	HR	2653046	0.89%	23612	4481497	Croatia
10	CH	6447178	0.69%	44485	7656981	Switzerland
11	CZ	7217097	0.64%	46189	10179263	Czech Republic
12	LT	2097651	0.57%	11956	3525465	Lithuania
13	NO	4575368	0.54%	24706	4707169	Norway
14	SE	8456772	0.51%	43129	9103092	Sweden
15	NC	80153	0.43%	344	235053	New Caledonia
16	SK	4343746	0.41%	17809	5484529	Slovakia
17	RU	61169412	0.38%	232443	138079937	Russian Federation
18	DE	67966373	0.36%	244678	82184248	Germany
19	NL	15138030	0.32%	48441	16914000	Netherlands
20	HU	6418129	0.31%	19896	9828683	Hungary
21	AU	19773262	0.30%	59319	22019223	Australia
22	TZ	5013280	0.30%	15039	43593747	United Republic of Tanzania
23	PT	5469159	0.29%	15860	10787297	Portugal
24	TW	16185162	0.29%	46936	23121661	Taiwan
25	SV	1261045	0.22%	2774	6092008	El Salvador
26	FI	4663439	0.21%	9793	5263476	Finland
27	MK	1933362	0.21%	4060	3897909	The former Yugoslav Republic of Macedonia
28	UA	15204978	0.21%	31930	44852443	Ukraine
29	MY	16689667	0.21%	35048	27049705	Malaysia
30	ID	55622256	0.17%	94557	248313646	Indonesia

These are low numbers

- Less than 1% of the Internet's user base with IPv6 active in 2012 is a very weak position
- And time is running out
- This cannot be an extended transition
  - Either we all move and move the entire Internet to supporting IPv6 by around 2015
  - Or we'll lose focus and momentum and turn our collective attention to engineering insane adornments for CDNs, ALGs and similar active middleware in an all-IPv4 network

This was always going to be tough

- Deregulated industry structure
- Commoditization of carriage provision
- Dominance of content services
- Disjoint cost and benefit in V6 deployment for access provider industry
- Significant resistance from the carrier sector
- No clear consumer benefits in cost or utility

Why should you care?

- No perceived need - already have IPv4 for enterprise
- Few IPv6 enterprise products available
- No IPv6 expertise in IT management and operations units

- Difficult, costly, and it addresses no perceived need

BUT

Why did you deploy IPv4 in the first place?

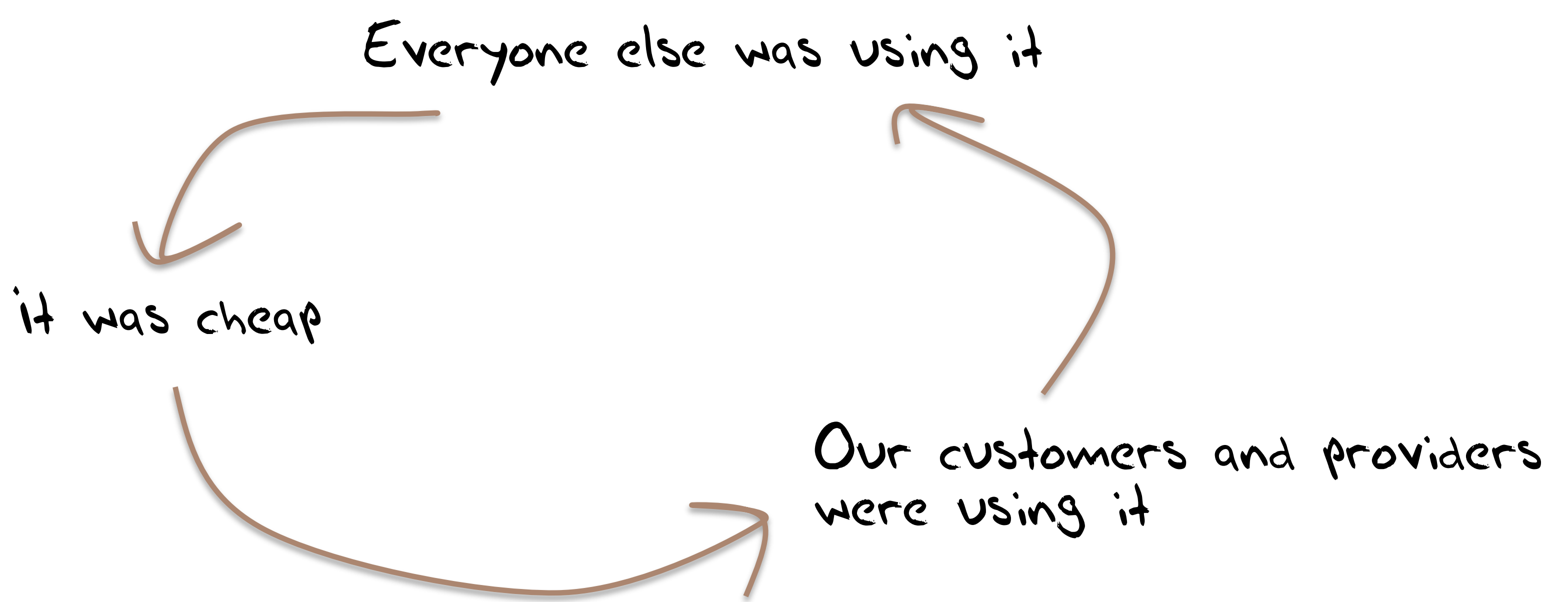

BUT

Why did you deploy IPv4 in the first place?

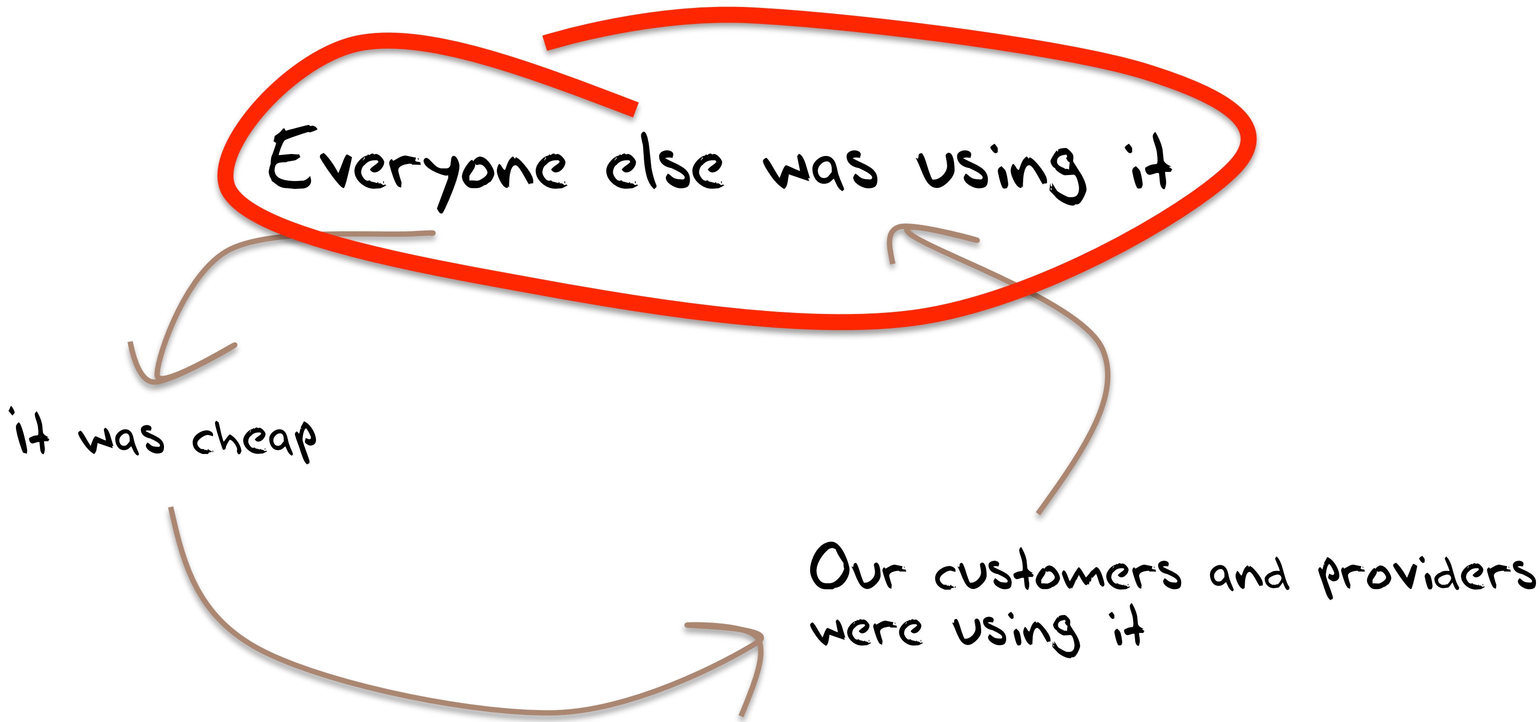

If not IPv6 ...

There is no other plan!

The path with IPv4 leads to Carrier NATS, Application Level Gateways and ultimately to fragmentation and piecemeal networks

The major benefits of IPv4 lie in its openness and universality

- neither of these are sustainable attributes for more than 2 – 3 years at most!
- After that expect to see IPv4 segment itself into a set of carrier-limited islands

Timing is everything

- What's your threshold for IPv6?
- There are no clear first adopter advantages
- So "wait and see" is a pervasive attitude
- But there are clear long term common risks of inaction in terms of cost, efficiency, openness and utility of the common network platform
- "when" is a big question here if we all want to avoid these risks

A Modest Suggestion

- Start **small**, but start **now**
  - Dual Stack the external front of house
    - Contract with your data centre / service provider for front-side IPv6 access
    - Put IPv6 on your front of house service platforms
    - Enable Dual stack your server application
    - Add AAAA records to your DNS
    - Include IPv6 monitoring in your operational monitoring
    - Measure the results

And while you are at it...

- No more purchasing “IPv4 only” products
- Dual stack should be a mandatory purchase criteria
  – Its a simple case of ensuring a reasonable service life for your equipment
  – And it also exposes your operational environment to introducing dual stack services internally
  – And allows you to flexible in adjusting to the moves of your suppliers, peers and customers

What changes with IPv6?

What changes with IPv6?

- Not much
  - Most users never notice when they connect to a dual stack access network
  - It's still the Web, its still the Internet, and things work much the same as ever
- But that's what we intended – if nothing changes then we've succeeded!

However, it's not perfect

- We learned a huge amount as we deployed IPv4
- And its clear that IPv6 still has its wrinkles
- And doubtless many network operators and their customers will encounter new issues in this deployment
- But that's no excuse to wait...

We need to move quickly with IPv6

as there is just no more IPv4!

Thank You