Open Life in a Post iPocalyptic Network

Scribblings by Geoff Huston, APNIC

The mainstream telecommunications industry has a rich history

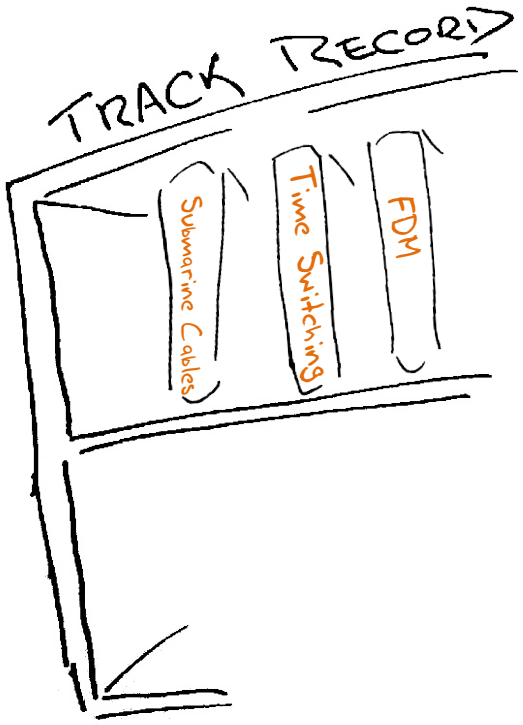

The mainstream telecommunications industry has a rich history

…of making some really poor technology choices

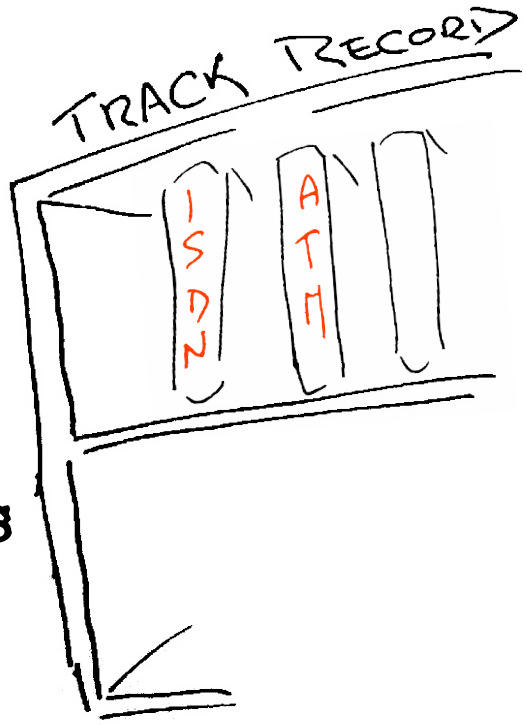

The mainstream telecommunications industry has a rich history

…of making very poor technology guesses

and regularly being taken by surprise!

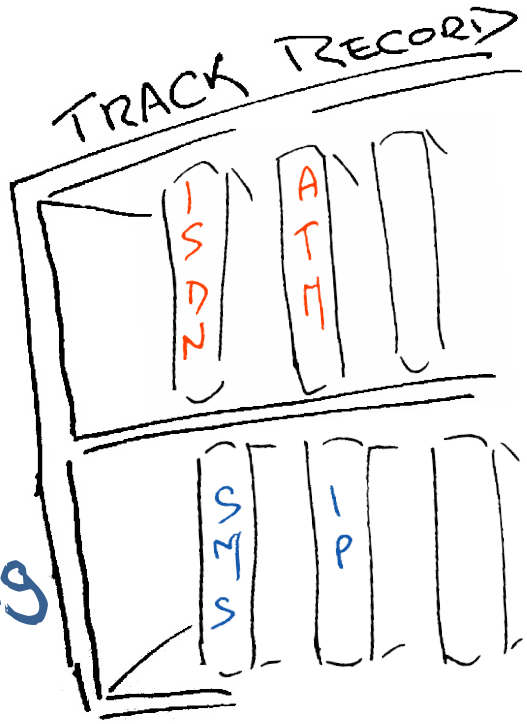

The Internet...

Has been a runaway success that has transformed not just the telecommunications sector, but entire social structures are being altered by the Internet!

The Internet...

Has been a runaway success that has transformed not just the telecommunications sector, but entire social structures are being altered by the Internet!

And now we've used up most of the original pool of IP addresses!

The Internet...

Has been a runaway succ[illegible]t
has transformed not [illegible]e
telecommunicati[illegible]or, but
entire socia[illegible]tures are
being al[illegible] the Internet!

This is should not be news - we've known about this looming iPocalypse for the past twenty years!

[illegible] now we've used up most of the original pool of iP adresses!

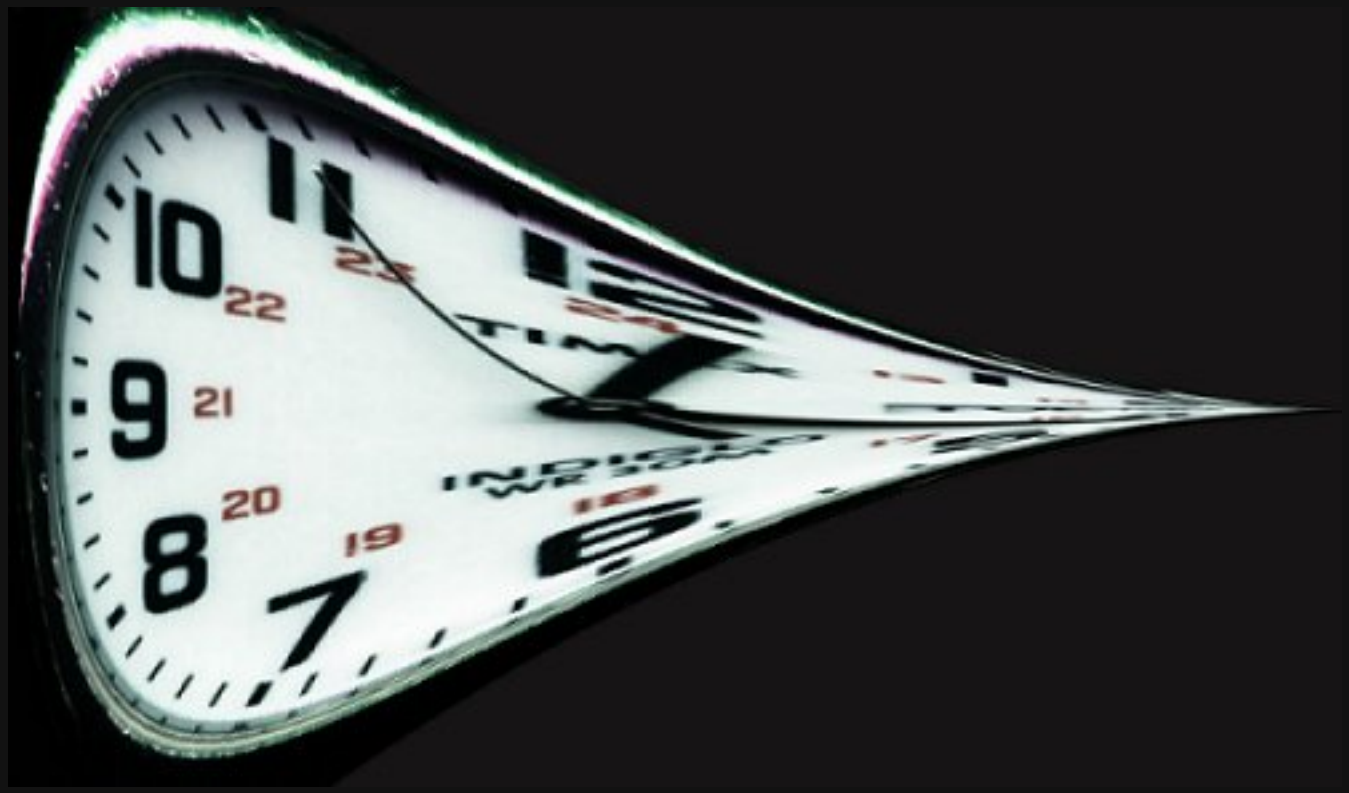
11
10
22
9
21
20
8
19
7

IETF Meeting - August 1990

Internet Growth (Continued):
Continued Internet Growth

Frank Solensky
Racal InterLan
solensky@interlan.com

- A preliminary analysis of data presented earlier in the conference projects the "size" of the Internet on several metrics, assuming continued exponential growth.
  - NIC Assigned Network Numbers
  - NIC "Connected" Status Nets
  - BBN's snapshots
  - NSFnet Policy Routing DataBases
- As was mentioned during the discussion period, a logistic curve would likely be a more realistic model: This will be the subject of further analysis. Note, however, that the limit that this approaches may turn out to be beyond the capacity of the class A-B-C numbering scheme

NIC

"Connected" IP Network Numbers

- Assigned Numbers RFC defines connected networks as connected to research and operational internet.
- Does not reflect whether the net is, in fact, entered in any routing table.

$y = \beta e^{\alpha t}$ where y = predicted number of nets, t = time (in months) since Jan. 1983

	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
$\hat{x}$	192.193	159.839	624.438	206.346
	Jan 6, 1999	Apr 26, 1996	May 14, 2038	Mar 27, 2000
r	.9293	.9870	.7942	.9548

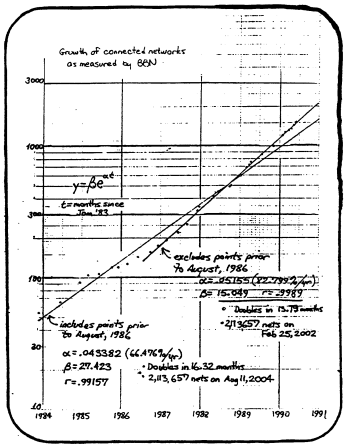

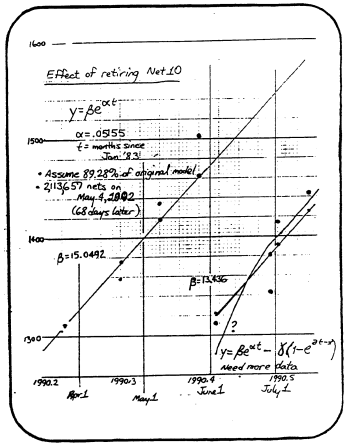

Assignment of IP Network Numbers

- Reflects organizations' desire for IP Address assignment; that is, to be listed in RFC-1166.
- Does not reflect "connectivity"

$y = \beta e^{\alpha t}$ where y = predicted number of nets, t = time (in months) since Jan '83

	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.	17.008%	78.389%	37.973%	20.394%
y	125	16,382	2,097,150	49,147
$\hat{x}$	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

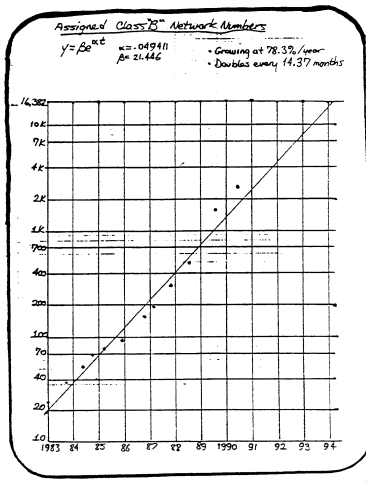

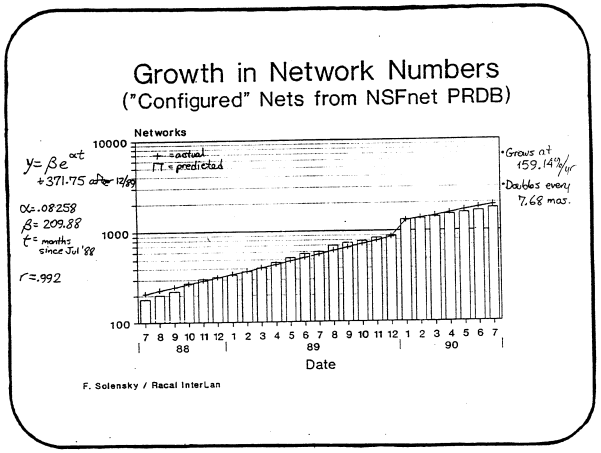

IETF Meeting - August 1990

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.

What did we do back in 1992?

What did we do back in 1992?

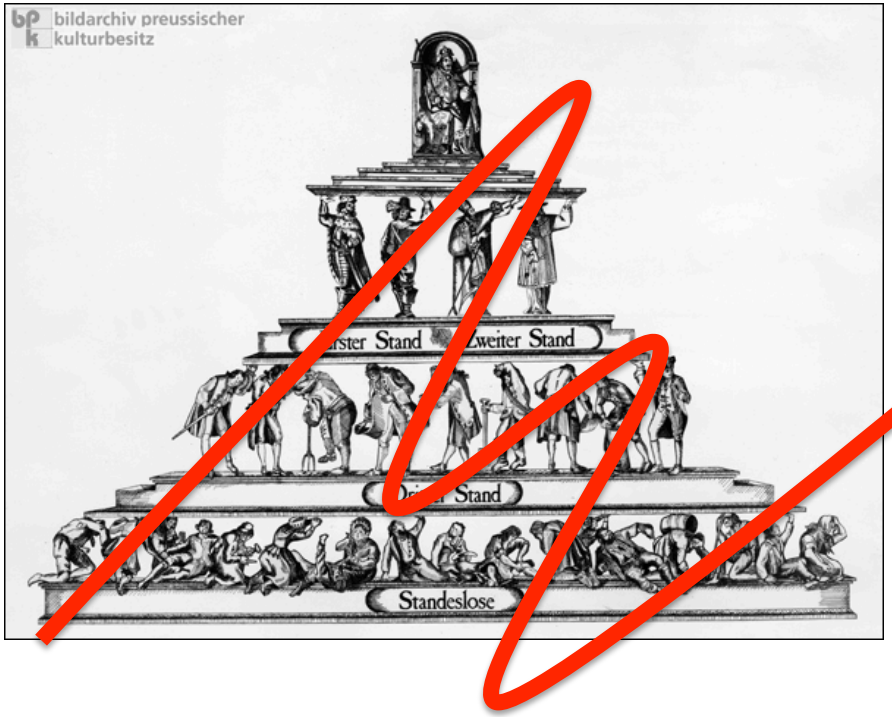

We bought some time by removing the CLASS A, B, C address structure from IP addresses

What did we do back in 1992?

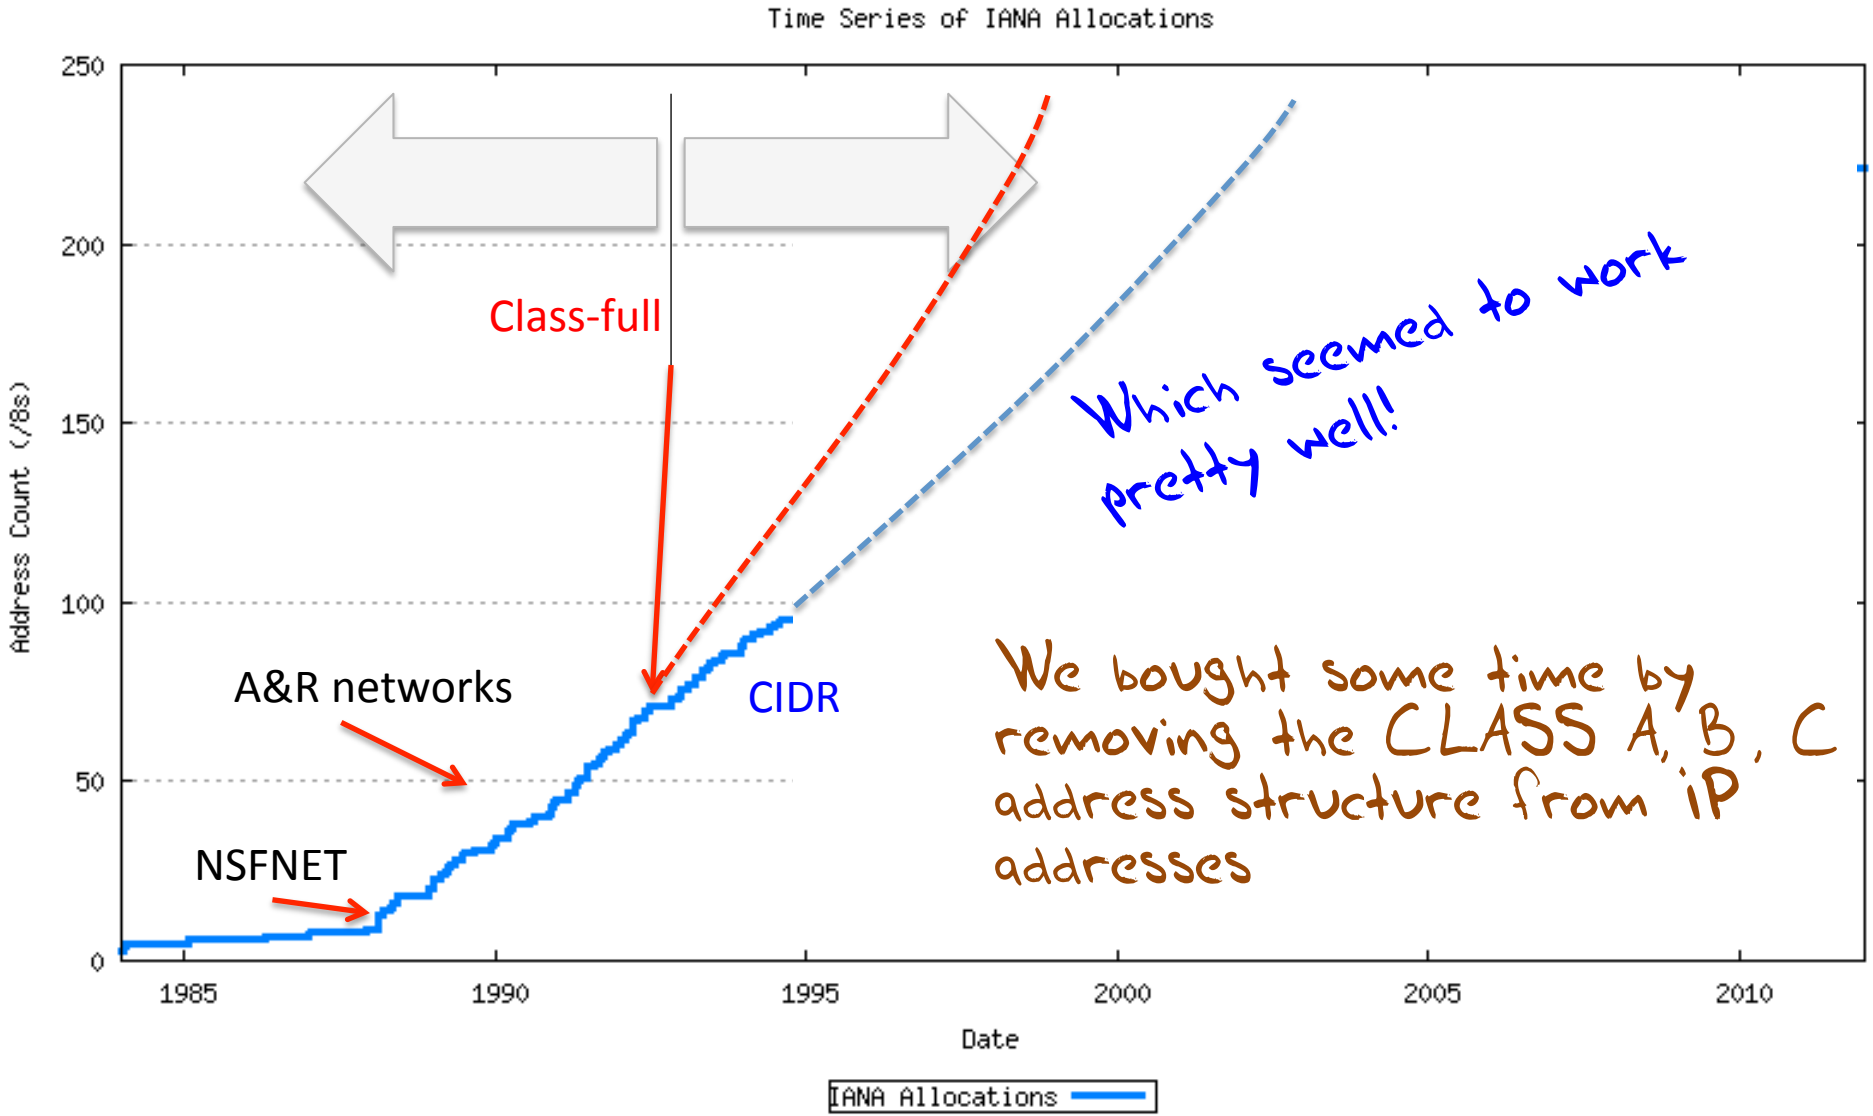

What else did we do back in 1992?

And we started working on a new Internet Protocol - to become IPv6 - to replace IPv4

What else did we do back in 1992?

And we started working on a new Internet Protocol - to become IPv6 - to replace IPv4

We left the task of transition until after we had figured out what this new protocol would look like

What else did we do back in 1992?

We developed some new middleware - an address sharing protocol that worked for TCP and UDP: NAT (RFC 1631)

"It is possible that CIDR will not be adequate to maintain the IP Internet until the long-term solutions are in place. This memo proposes another short-term solution, address reuse, that complements CIDR or even makes it unnecessary."

zzzzzz . . .

For a long while this did not look to be an urgent problem...

CIDR + NATs just worked!

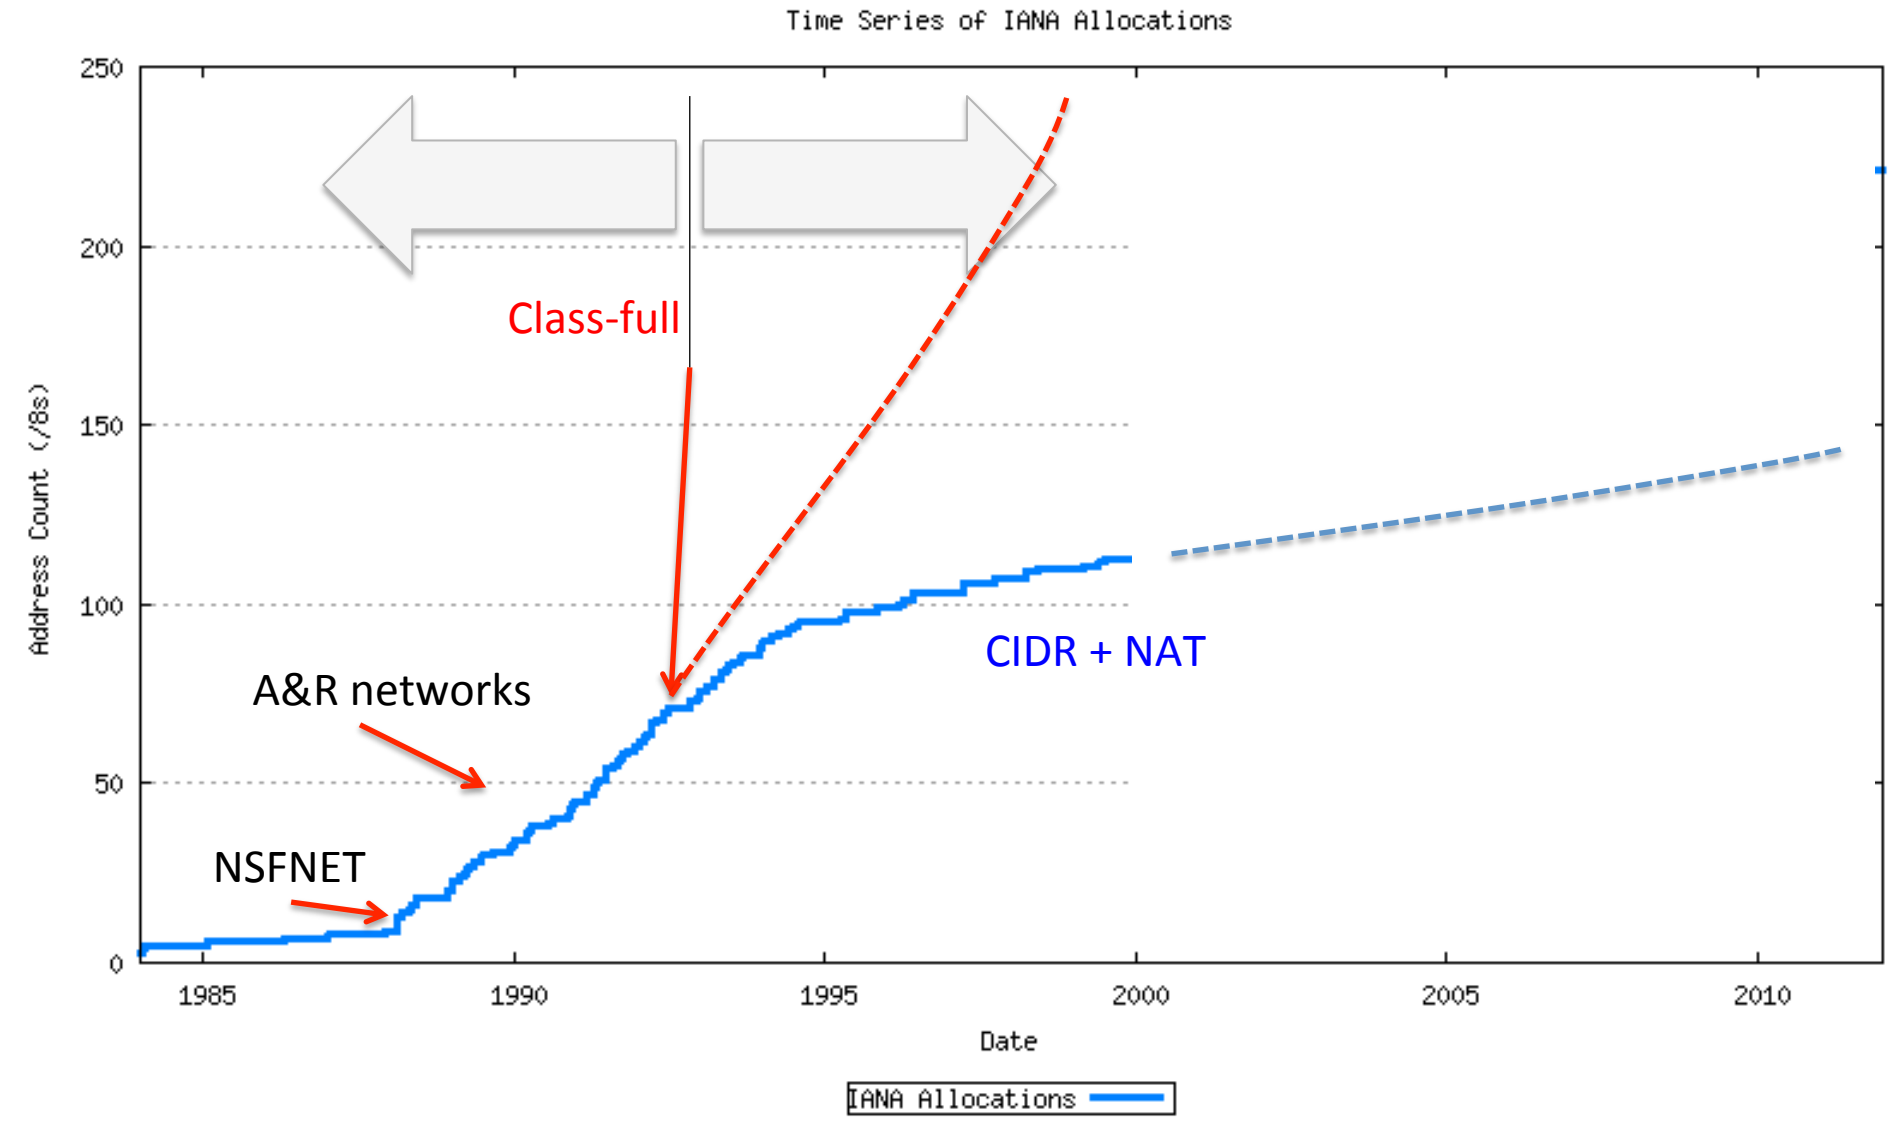

Meanwhile, we continued to build (IPv4) networks

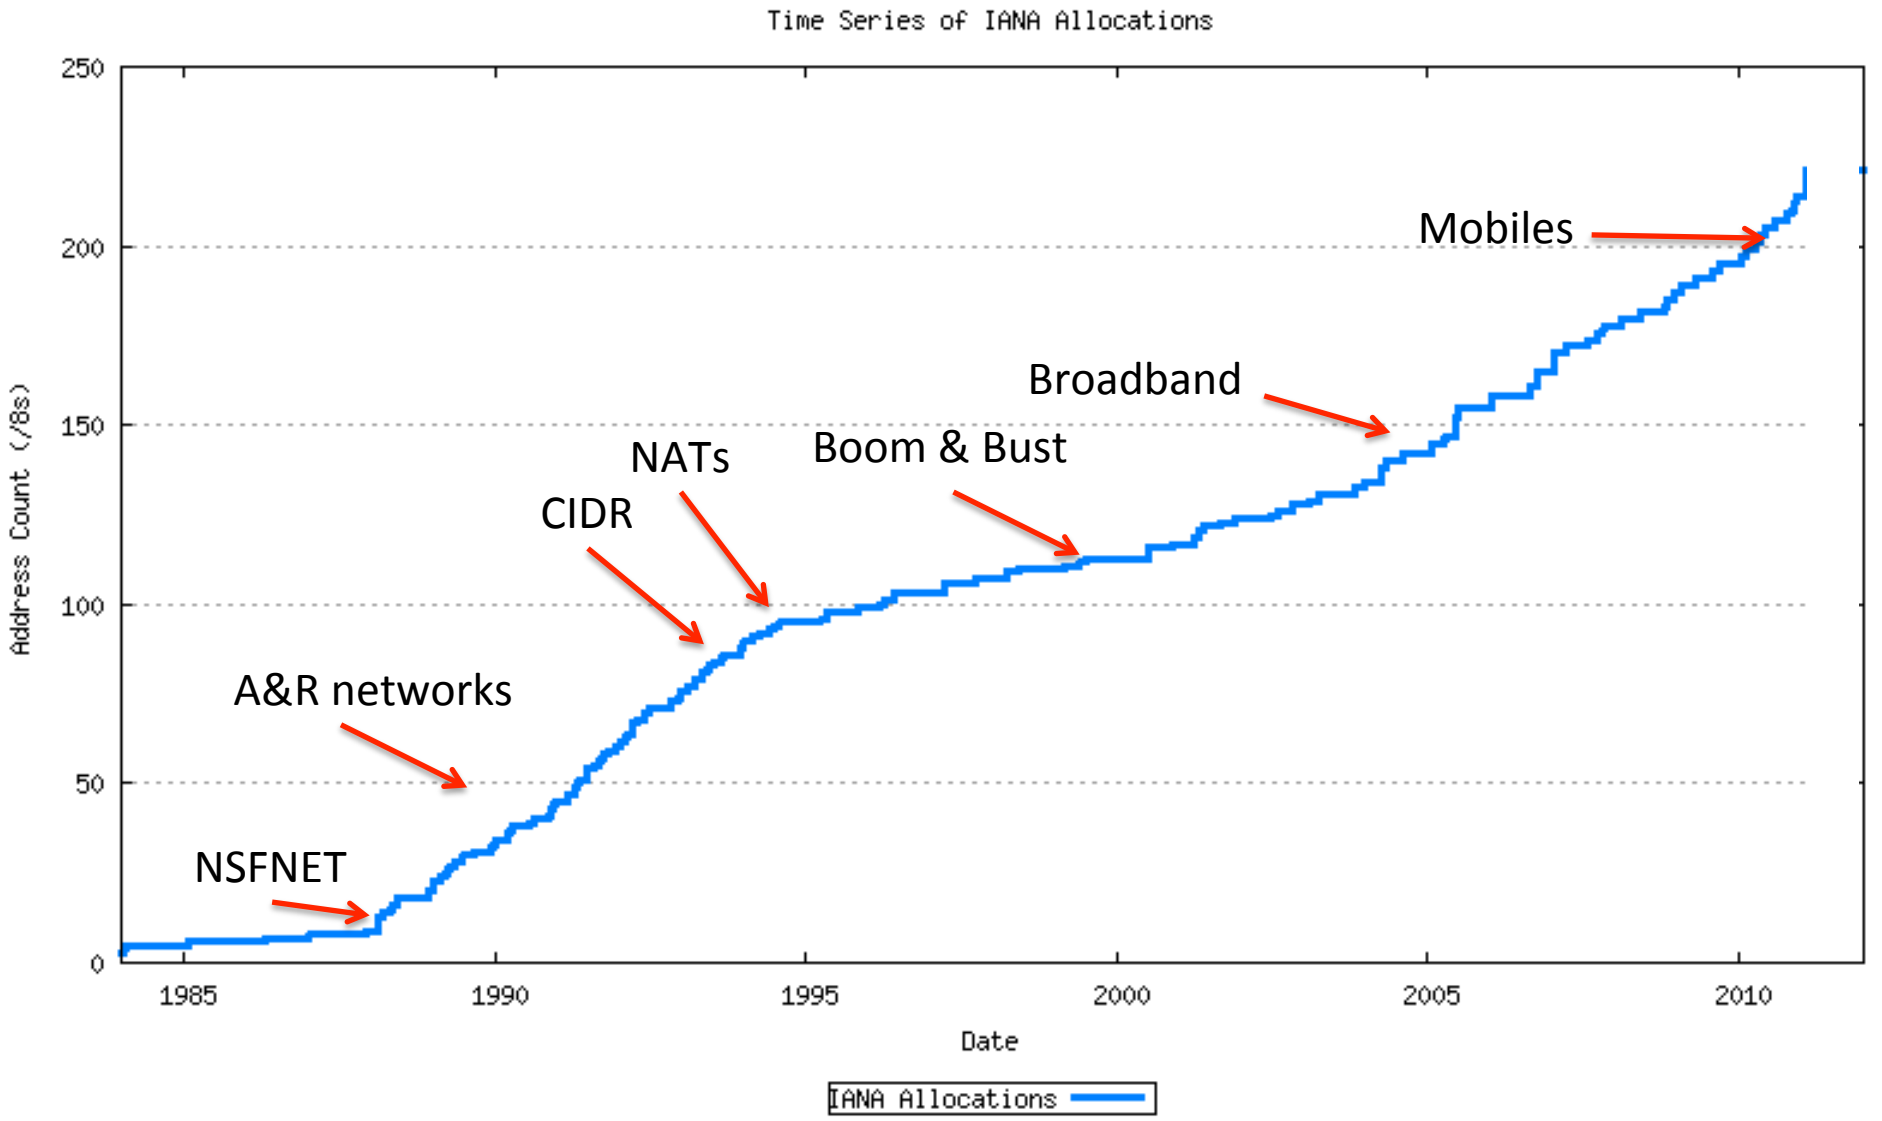

The rude awakening

Until all of a sudden the IPv4 address piggy bank was looking extremely empty...

IPv4 Address Allocations

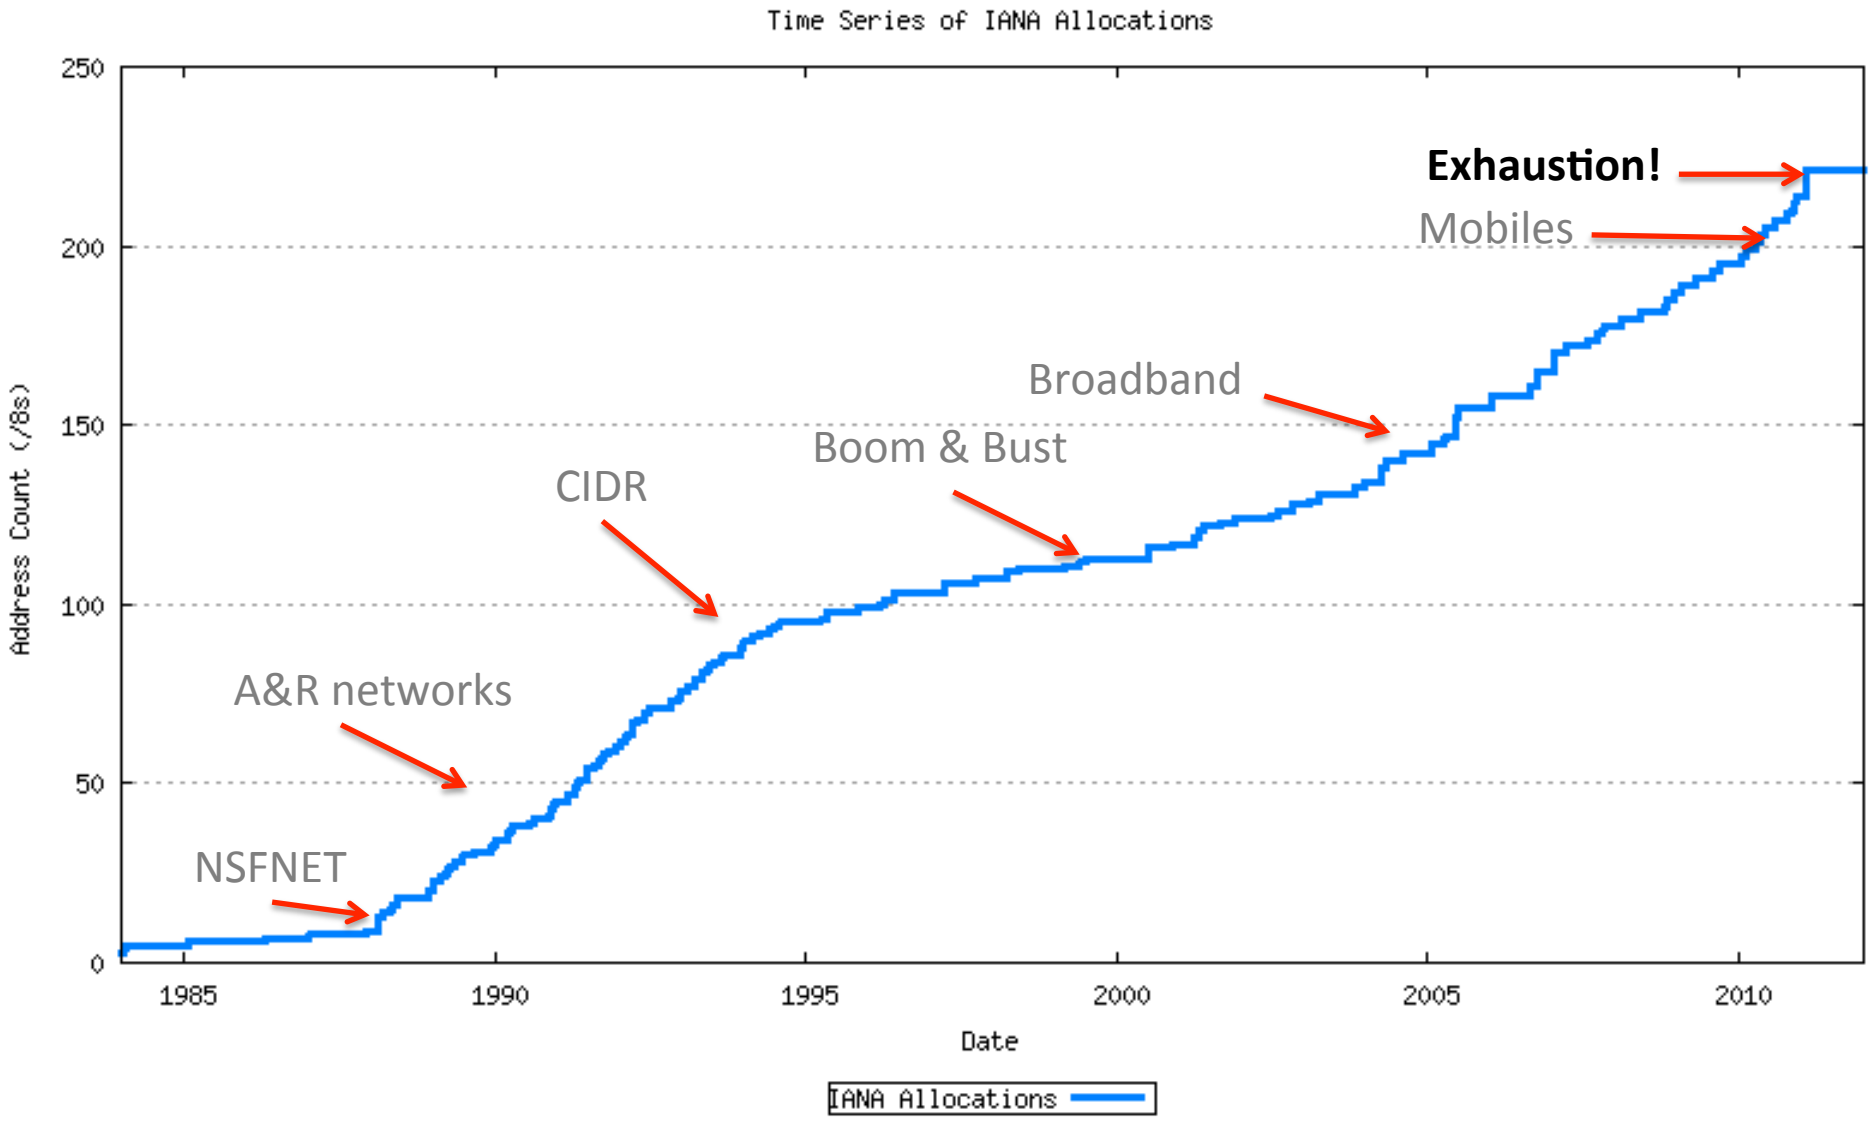

Going...

28 January 2011 Last updated at 11:01 GMT

Share

Net approaches address exhaustion

By Mark Ward
Technology correspondent, BBC News

The last big blocks of the net's dwindling stock of addresses are about to be handed out.

The event that triggers their distribution is widely expected to take place in the next few days.

When that happens each of the five regional agencies that hand out net addresses will get one of the remaining blocks of 16 million addresses.

Time is running out for the internet's current addressing scheme

The addresses in those last five blocks are expected to be completely exhausted by September 2011.

Final five

The trigger event will likely come from the agency that oversees net addresses in the Asia-Pacific region, a body known as Apnic.

When Apnic's store of addresses falls below a key threshold, said Geoff Huston, chief scientist at the agency, it will ask for more from the central repository - the Internet Assigned Numbers Authority (IANA).

Related Stories
Internet faces 'turbulent times'
Test looms for net address scheme
Tuning in to the background hum of the net

Top Stories

Pakistan explosions kill scores
Afghan mission in 'last chapter'
Silver Linings surprise at Oscars
American Express cuts 5,400 jobs
Klaus Kinski 'was child abuser'

Going...

Going...

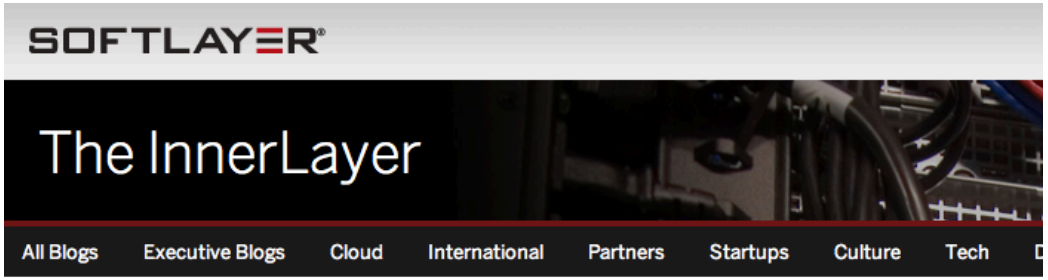

January 21, 2011

What Does IPv4 Exhaustion Mean for You?

Posted by Kevin Hazard in SoftLayer

THE SKY IS FALLING! EVERYBODY MOOOOOOVVVVEEEE! WWWHHHYYY??!! OH THE HUMANITY!!!

Are those your reactions to the depletion of IPv4 space? Probably not. If you haven't seen the IPv4 Exhaustion Rate countdown in the sidebar of SoftLayer.com, head over there and check it out ... At the current rate, there will be ZERO unallocated IPv4 blocks by the middle of February 2011, and that's not a good thing for the Internet as we know it.

Will you need to move your servers into a bomb shelter to protect your now-even-more-valuable IP addresses? Will Google stop Googling? Will there be riots in the streets as over-caffeinated sysadmins flip cars and topple dilapidated buildings in pursuit of lost 32-bit addresses? What does it really mean for you as a hosting customer and web surfer?

The sky won't fall. Your servers are safe in their data centers. Google will still Google. Sysadmins will still be working hard at their desks. But the belt is going to start tightening, and after a while, it might get pretty uncomfortable.

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 [illegible] action to deploy IPv6."

Oh shit!

"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."

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

Panic?

AsiaPac: Panic.

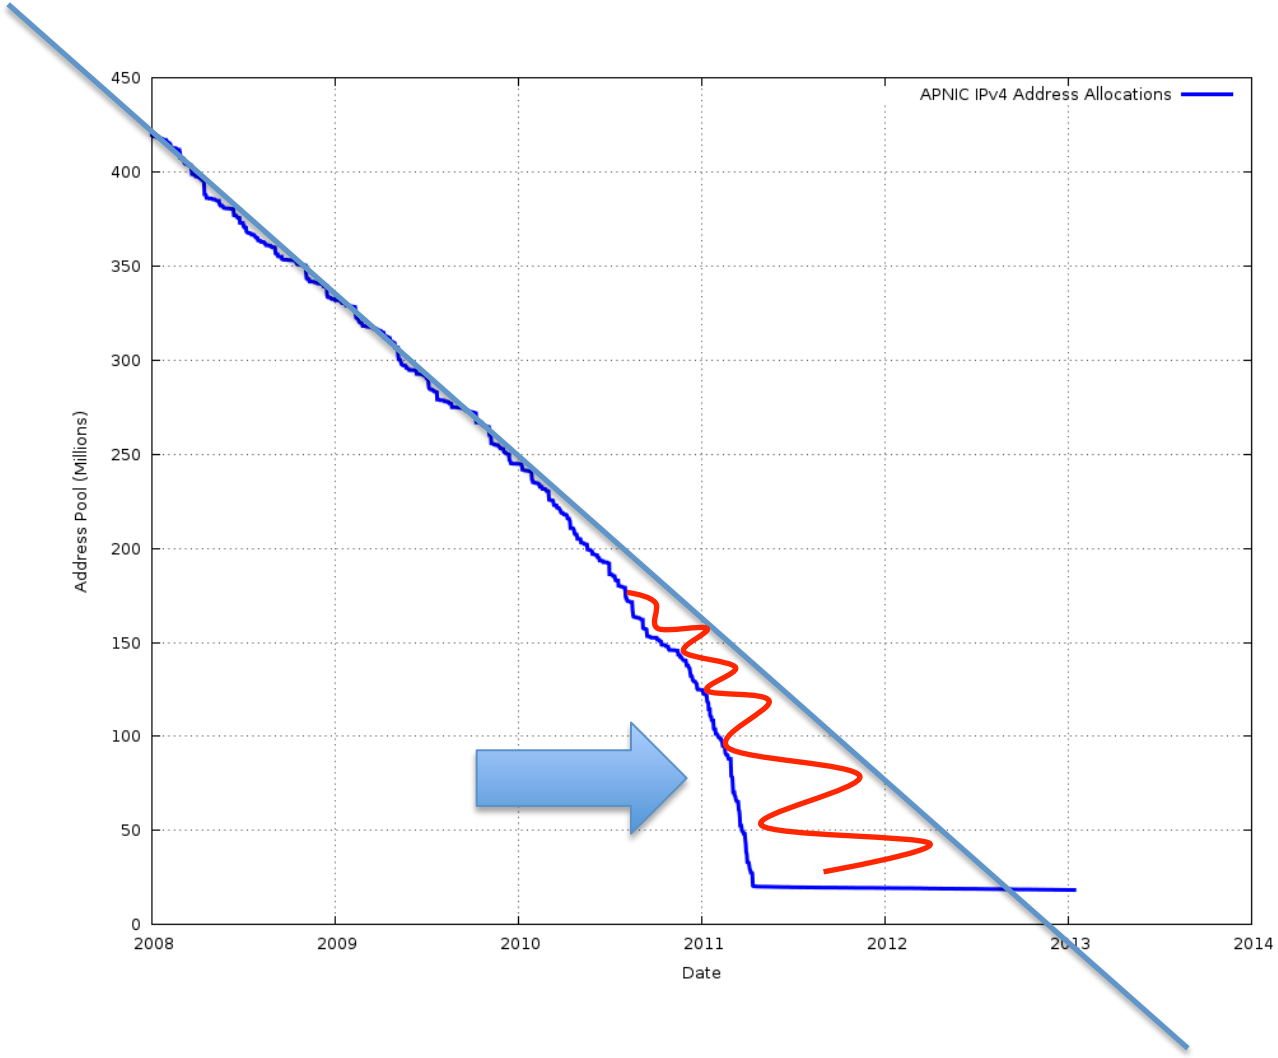

Europe: Distracted.

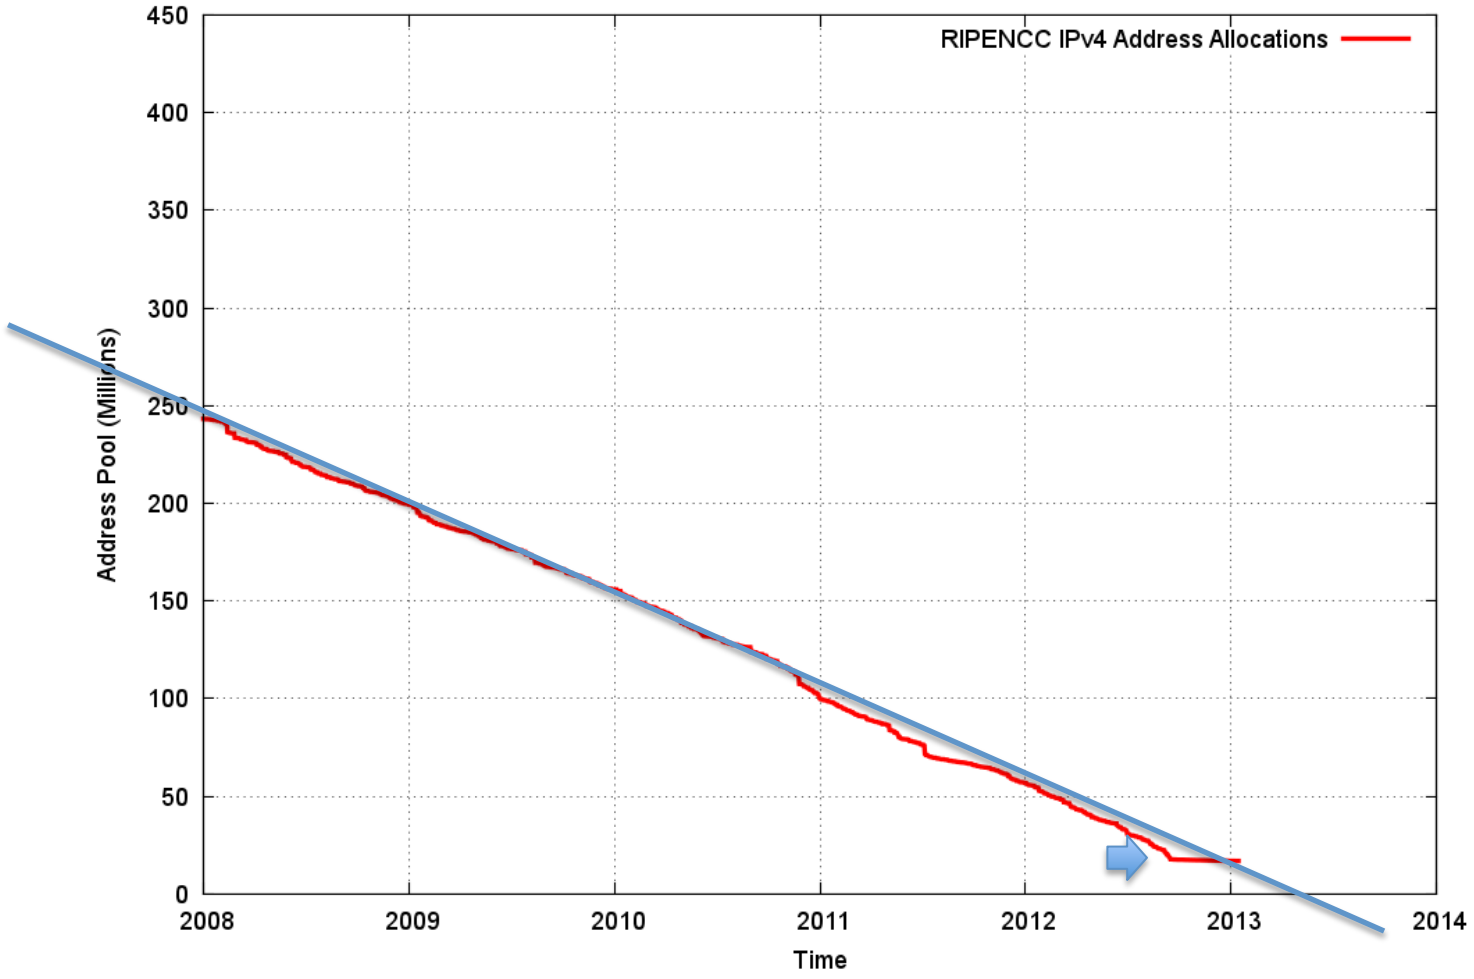

America: Confused.

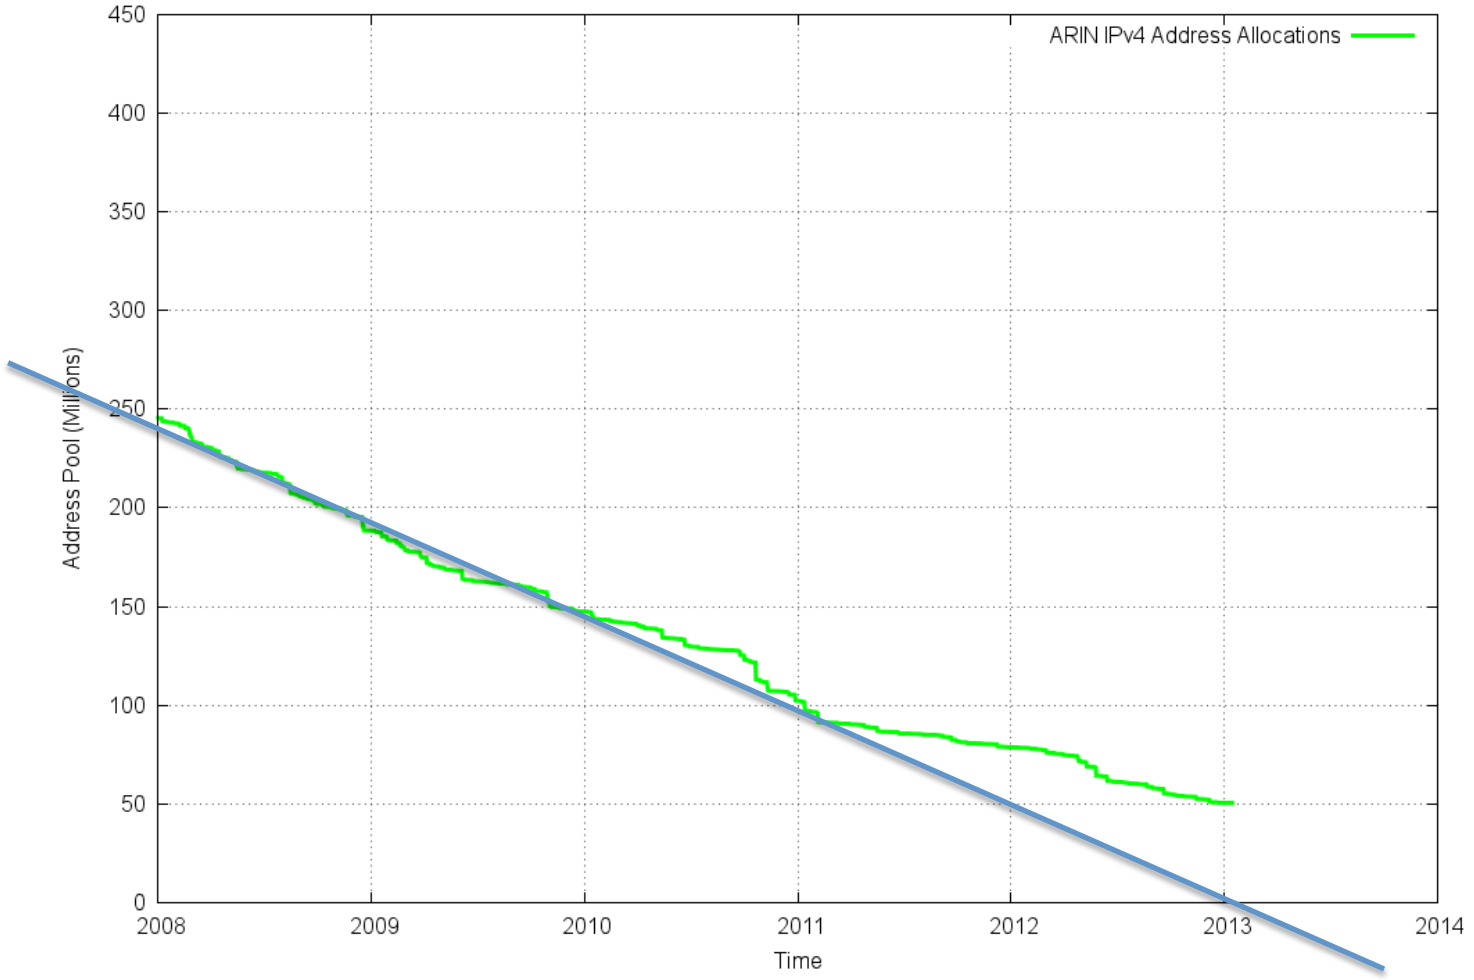

Numerology

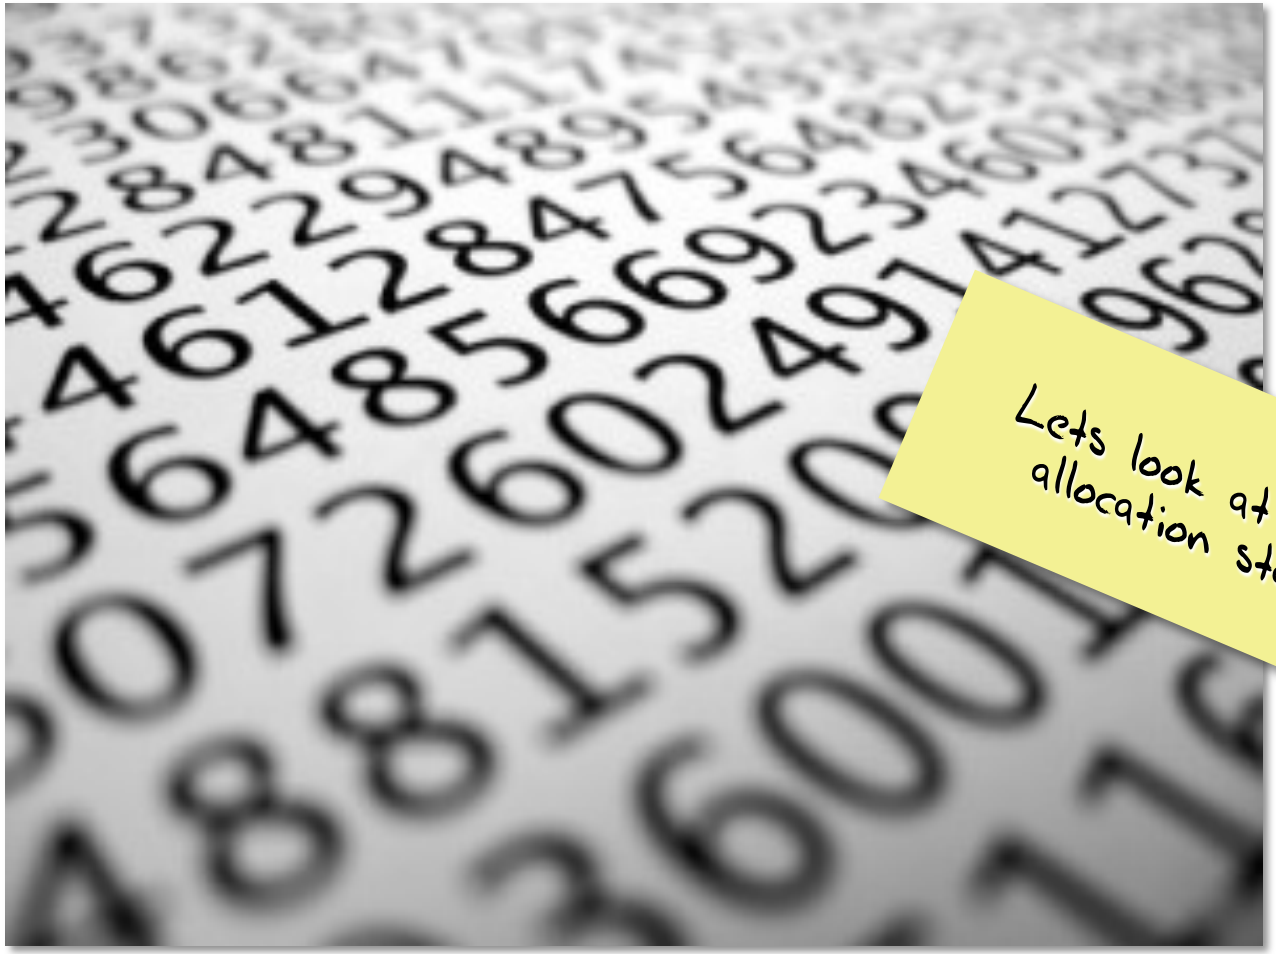

Lets look at the IPv4 allocation statistics

IPv4 Allocations in APNIC

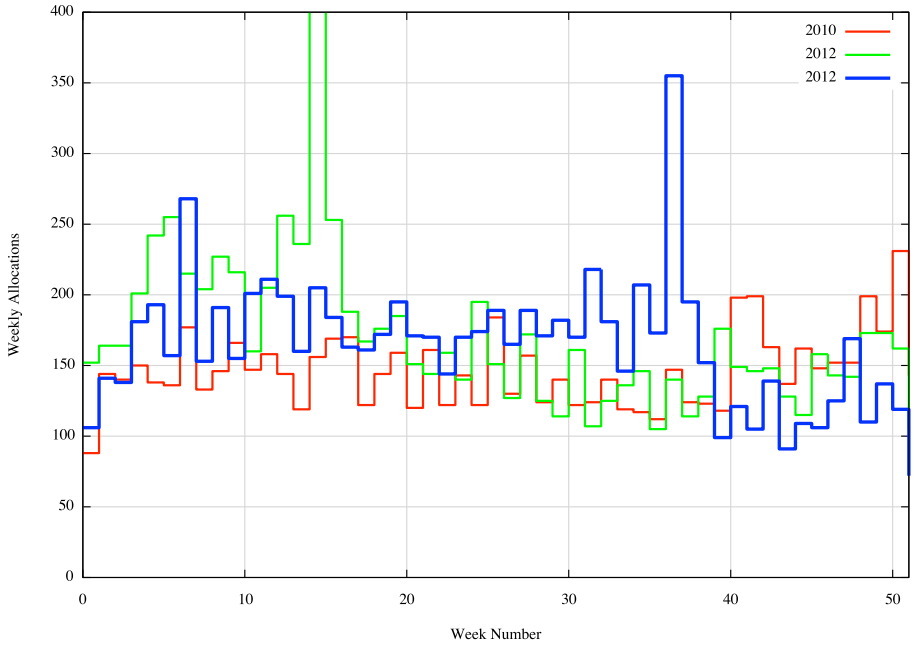

2010: 146 IPv4 allocations / week
2011: 168 IPv4 allocations / week
2012: 166 IPv4 allocations / week

The allocation rate has been pretty constant for the past 3 years!
is this a work to rule at APNIC?
Or some underlying business dynamic?

Why is this production profile steady at 20 allocations per working day?

All RIRs

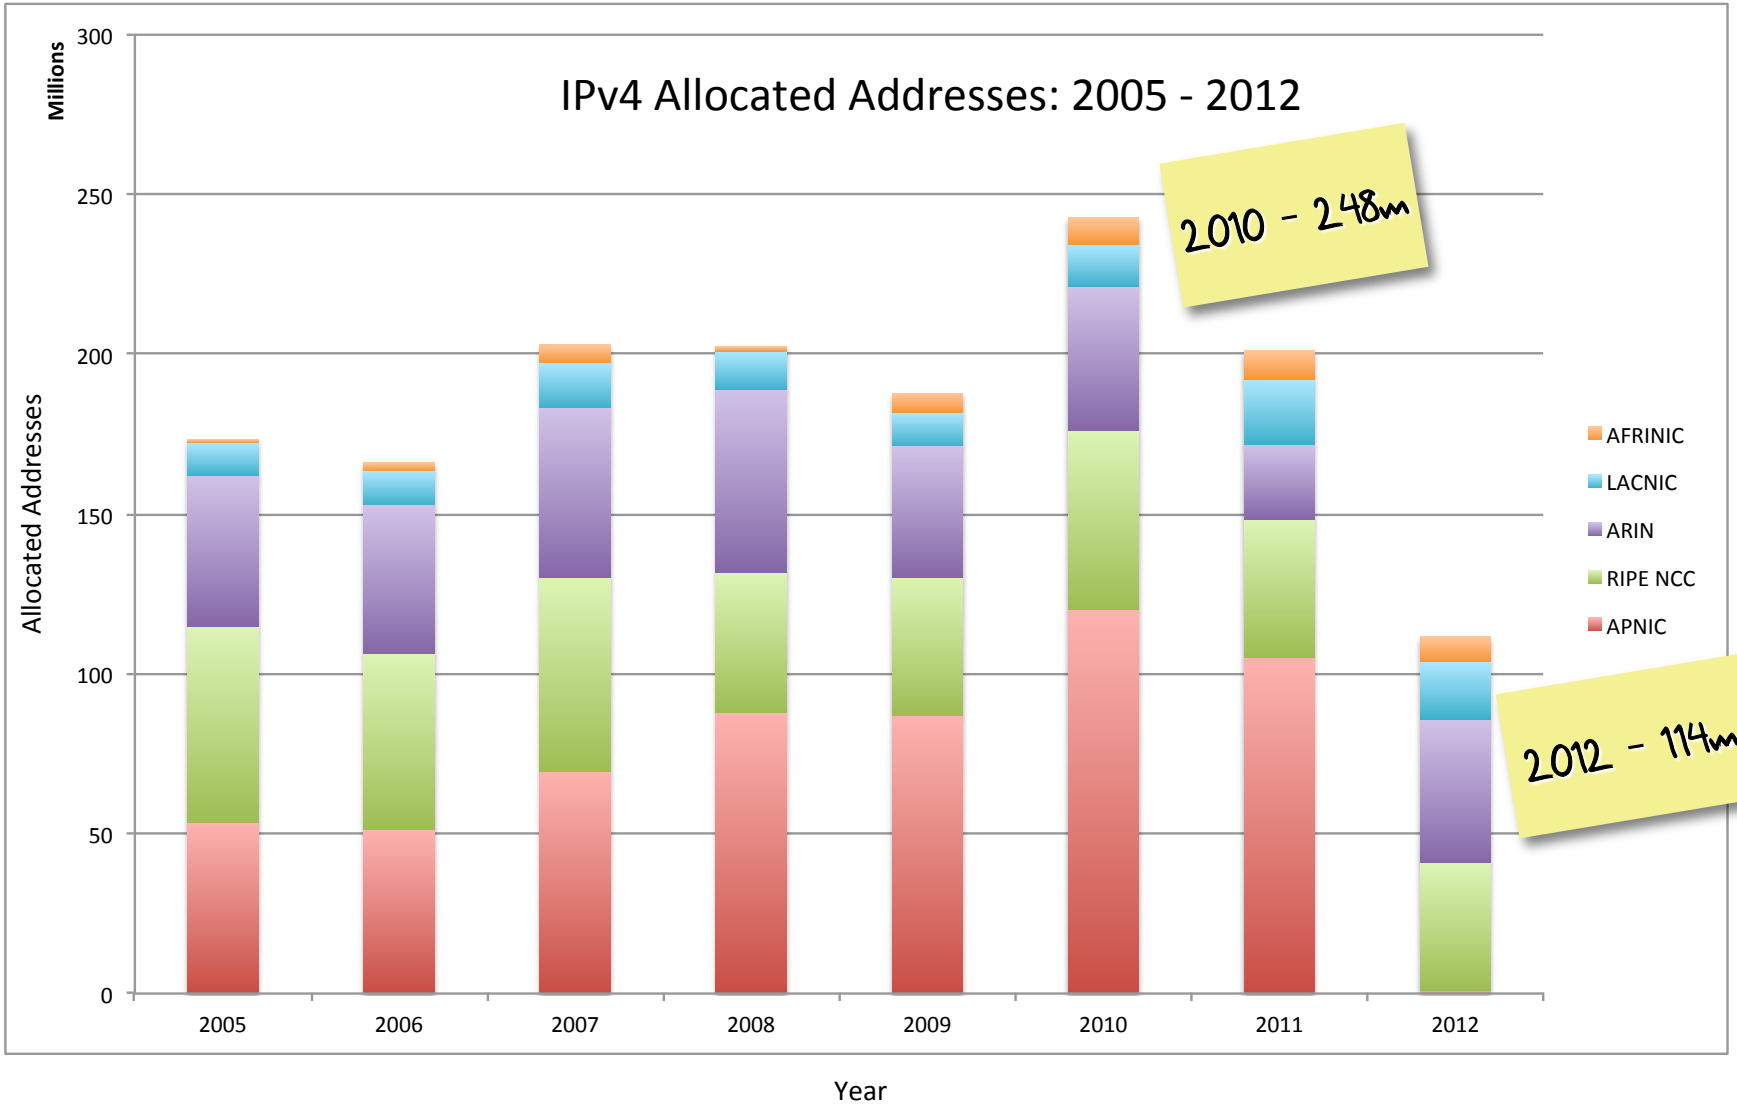
IPv4 Allocated Addresses: 2005 - 2012
Millions
Allocated Addresses
300
250
200
150
100
50
0
2005
2006
2007
2008
2009
2010
2011
2012
Year
AFRINIC
LACNIC
ARIN
RIPE NCC
APNIC
2010 - 248m
2012 - 114m

Top 10 Countries, 2009-2012

Rank	2009		2010		2011		2012	
1	China	50.67	China	45.2	China	53.07	USA	28.2
2	USA	38.55	USA	42.32	USA	21.21	Canada	16.7
3	Japan	11.04	Rep.Korea	25.73	Japan	16.91	Brazil	8.4
4	Rep.Korea	10.95	Japan	10.02	Rep.Korea	7.68	Russia	5.3
5	Russia	5.46	Australia	9.63	Indonesia	7.09	Iran	4.5
6	Brazil	4.19	India	9.43	Brazil	6.29	Germany	3.4
7	UK	4.19	UK	8.13	India	6.01	South Africa	3.4
8	Italy	4.16	Germany	6.97	France	5.39	Italy	3.3
9	France	3.85	Russia	6.46	Russia	5.02	Colombia	2.6
10	Germany	3.6	Brazil	6.29	Germany	4.92	Romania	2.6

Largest Allocations in 2011

Rank	Economy	Organization	Addresses(M)	
1	Japan	NTT Communications Corporation	8.39	*
2	China	China Mobile Communications Corporation	8.39	*
3	Brazil	Comite Gestor da Internet no Brasil (Brasil NIR)	6.29	
4	Indonesia	PT Telekomunikasi Selular Indonesia	6.29	*
5	Japan	KDDI Corporation	4.19	
6	United States	AT&T Mobility LLC	4.19	*
7	United States	AT&T Internet Services	4.19	
8	France	Bouygues Telecom	4.19	*
9	Germany	Telekom Deutschland Mobile	2.1	*
10	China	CHINANET Zhejiang Province Network	2.1	
11	China	China TieTong Telecommunications Corporation	2.1	
12	Pakistan	Pakistan Telecommuication	2.1	*
13	China	China Unicom Shandong province network	2.1	
14	Morocco	Maroc Telecom	2.1	*
15	India	Bharti Airtel Limited	2.1	*
16	Vietnam	Viettel Corporation	2.1	
17	Mexico	Uninet S.A. de C.V., Mexico	2.1	
18	Egypt	TE Data, Egypt	2.1	
		Total	**67.11**	

18 Carriers --> 30÷ of the addresses

Choices, Choices

if you are in AsiaPac, and you need IP addresses...

what are you going to do?

Choices, Choices

1. Apply for a FINAL allocation of 1,024 IPv4 addresses

2. Apply for an IPv6 allocation

3. Buy addresses someone who has some addresses to sell

4. Carrier IPv4 NATs

Choices, Choices

1. Apply for a FINAL allocation of 1,024 IPv4 addresses

That's not a lot of addresses! it is only really useful for small end sites and/or NATs

2. Appl cation

3. Buy addresses someone who has some addresses to sell

4. Carrier IPv4 NATs

Choices, Choices

1. Apply for a [illegible]cation of 1,024 IPv4 addresses

That's not a lot of addresses!

2. Apply for an IPv6 allocation

3. Buy addresses someone who has some addresses to sell

4. Carrier IPv4 NATs

Measuring IPv6

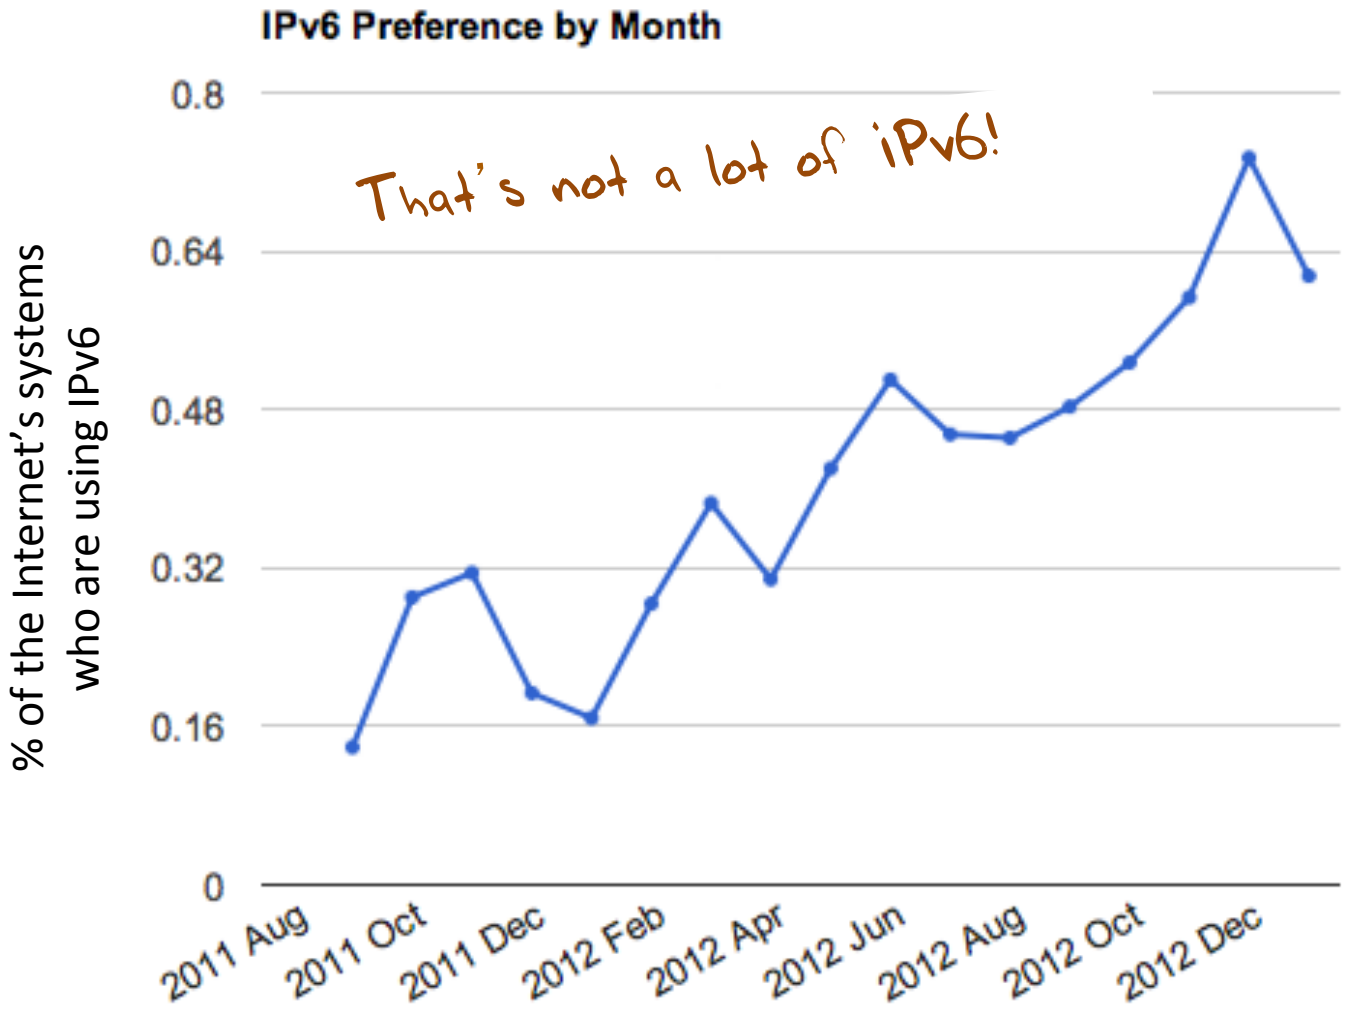

Choices, Choices

1. Apply for a /22 allocation of 1,024 IPv4 addresses

That's not a lot of addresses!

2. Apply for an IPv6 allocation

3. An IPv6-only network without any form of IPv4 mapping or translation capability is a pretty lonely and useless network today!

This won't connect you to the IPv4 internet.

4. Carrier IPv4 NATs

Choices, Choices

1. Apply for an allocation of 1,024 IPv4 addresses

That's not a lot of addresses!

2. Apply for an allocation

This won't connect you to the IPv4 internet

3. Buy addresses someone who has some addresses to sell

4. Carrier IPv4 NATs

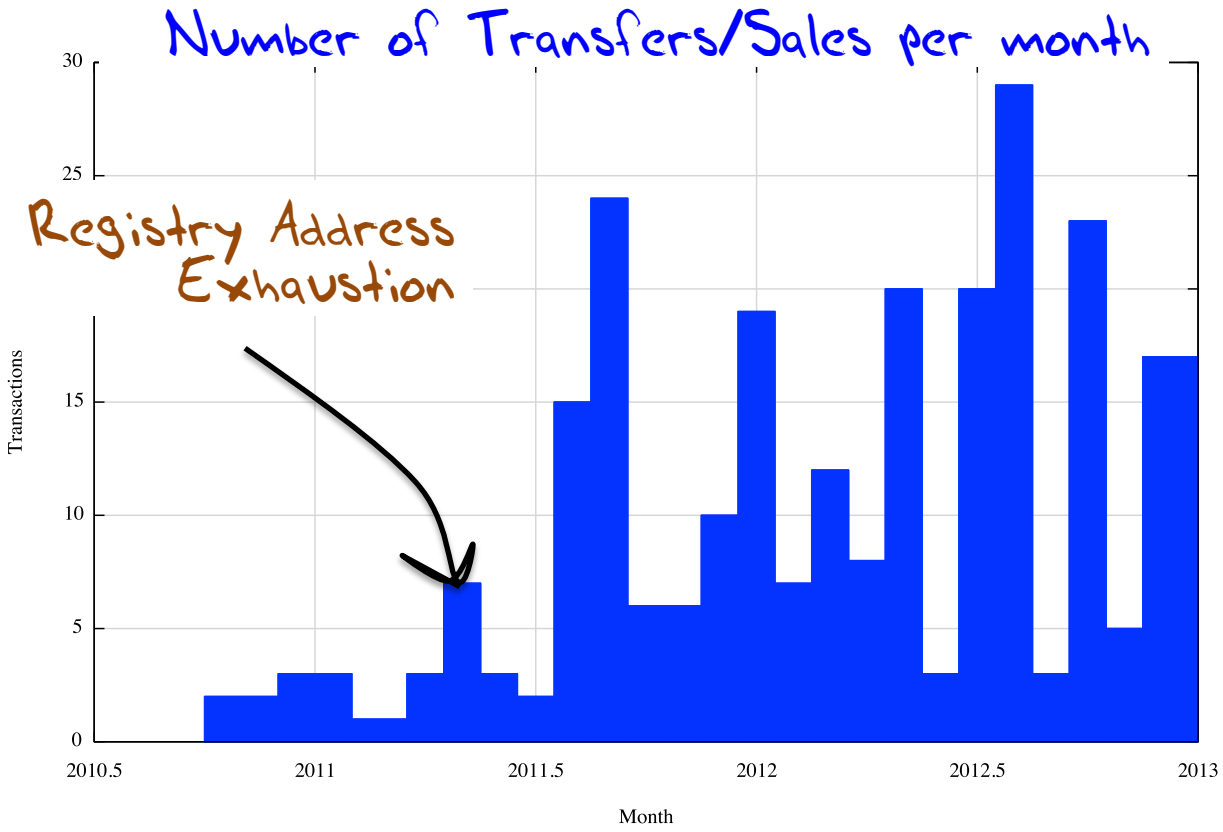

Address Markets Stats

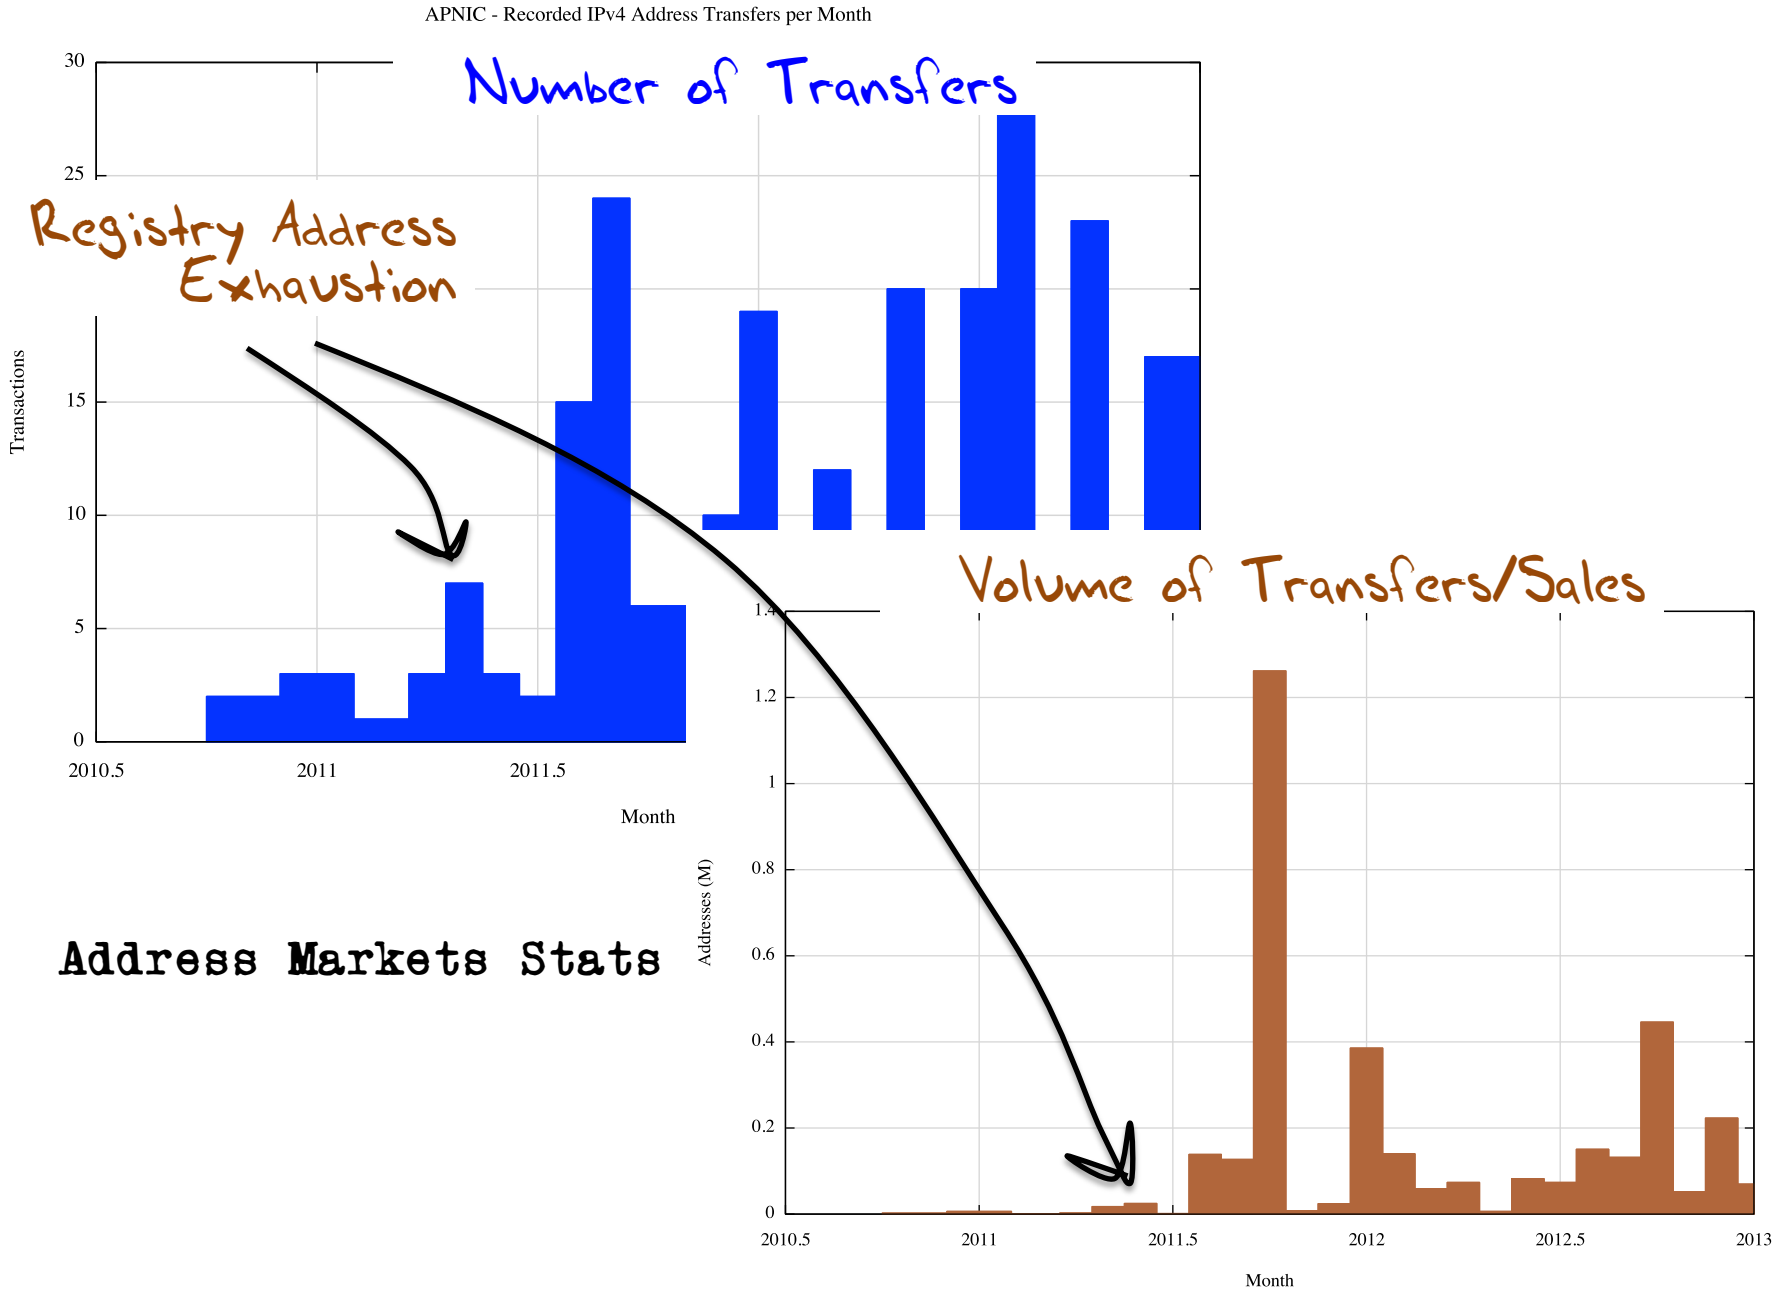
APNIC - Recorded IPv4 Address Transfers per Month
Number of Transfers
Registry Address Exhaustion
Volume of Transfers/Sales
Address Markets Stats
Transactions
Month
Addresses (M)
Month
30
25
15
10
5
0
2010.5
2011
2011.5
1.4
1.2
1
0.8
0.6
0.4
0.2
0
2010.5
2011
2011.5
2012
2012.5
2013

Choices, Choices

1. Apply for [illegible]cation of 1,02[illegible] IPv4 addresses

That's not a lot of addresses!

2. Apply for [illegible]location

This won't connect you to the IPv4 internet

3. Buy addresses someone who has some addresses to sell

4. [illegible]

This is not being widely used. It does not appear to be taken up by ISPs in the region. Supply is limited and costs are volatile

Choices, Choices

1. Apply for a /22 allocation of 1,024 IPv4 addresses

That's not a lot of addresses!

2. Apply for a /22 allocation

This won't connect you to the IPv4 internet

3. Buy addresses — there was some a This is not being widely used addresses to sell

4. Carrier IPv4 NATs

The Goldilocks Option!

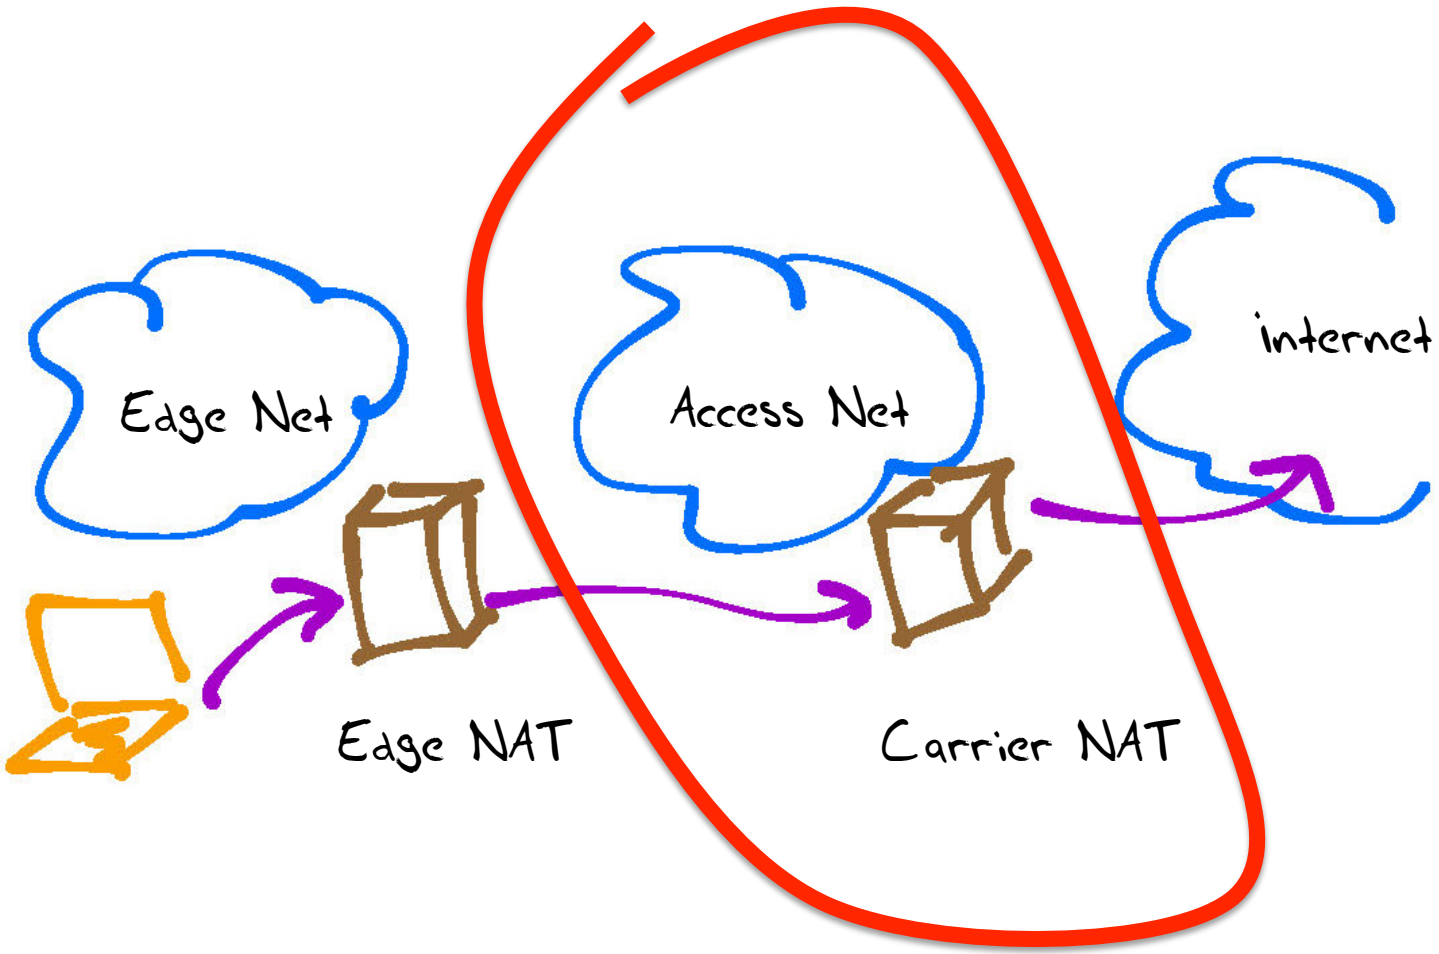
Edge Net
Access Net
internet
Edge NAT
Carrier NAT

Why is this CGN model so attractive?

- incrementally deployable
- No coordination of actions
- Applications are unaltered
- Carrier infrastructure unaltered
- Relatively low marginal cost

Downsides of the CGN model?

- HTTP will work
  - But parallelism will not
- UDP will work
  - But UDP behaviour will be erratic
- Not much else will work
- Which places severe restrictions on how applications operate across the network
- And impacts the current model of network service provision

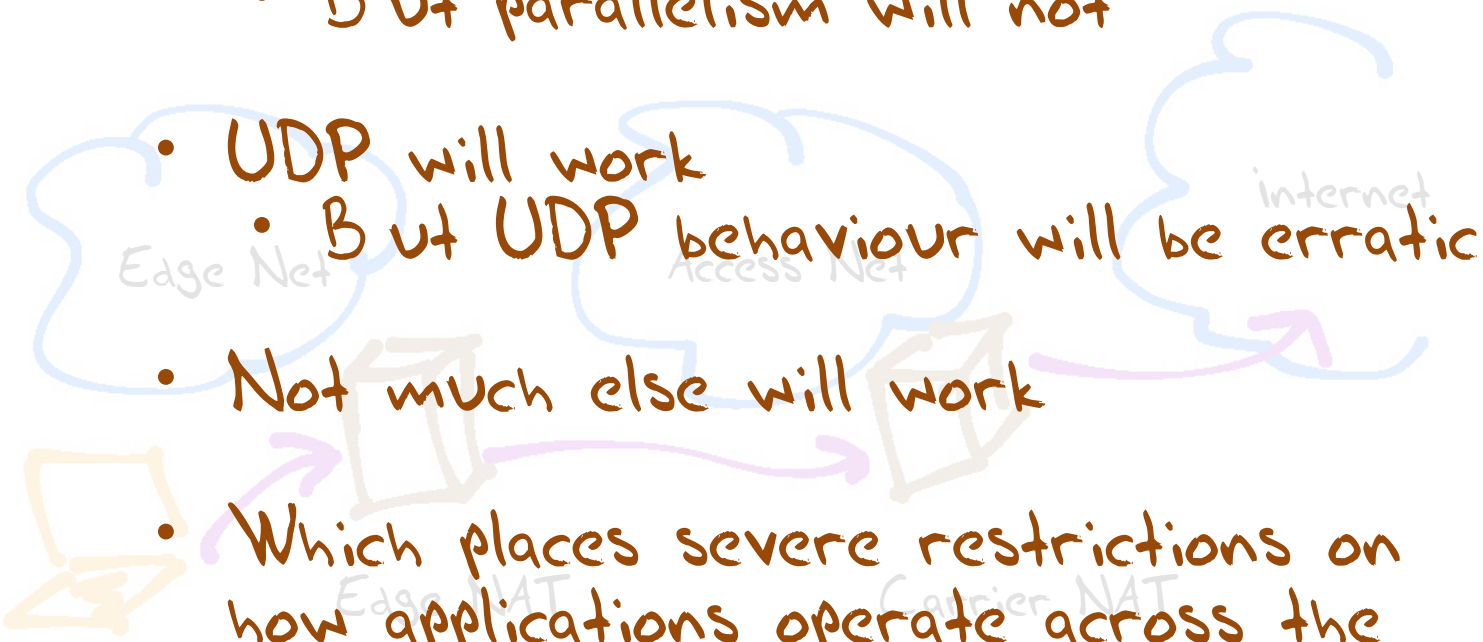

Back to networking basics....

Telco nostalgia...

The historical vertically integrated service architecture

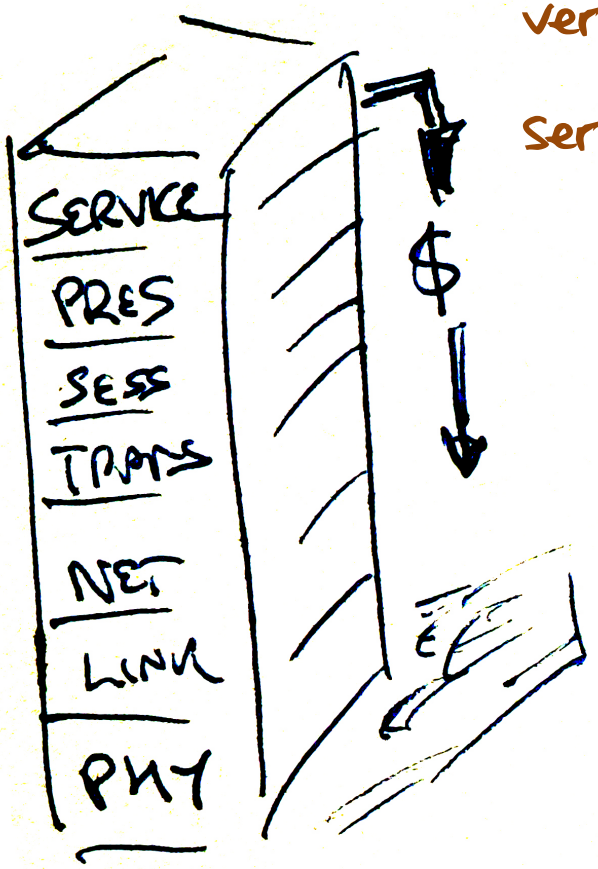

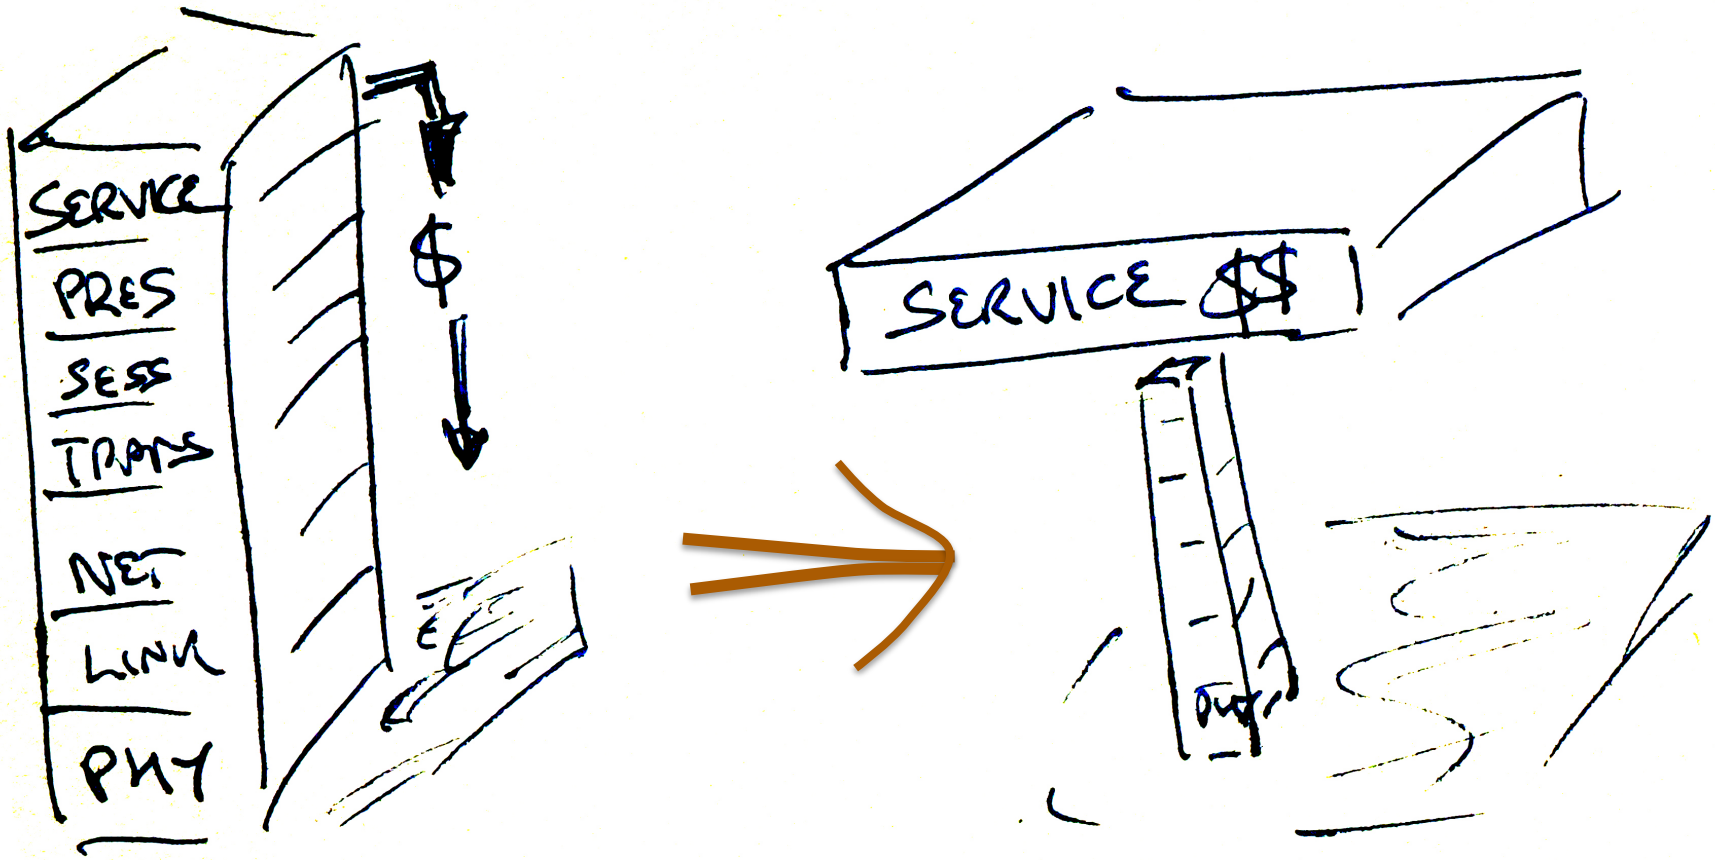

Devolution of the integrated service architecture through an open IP service architecture and deregulation

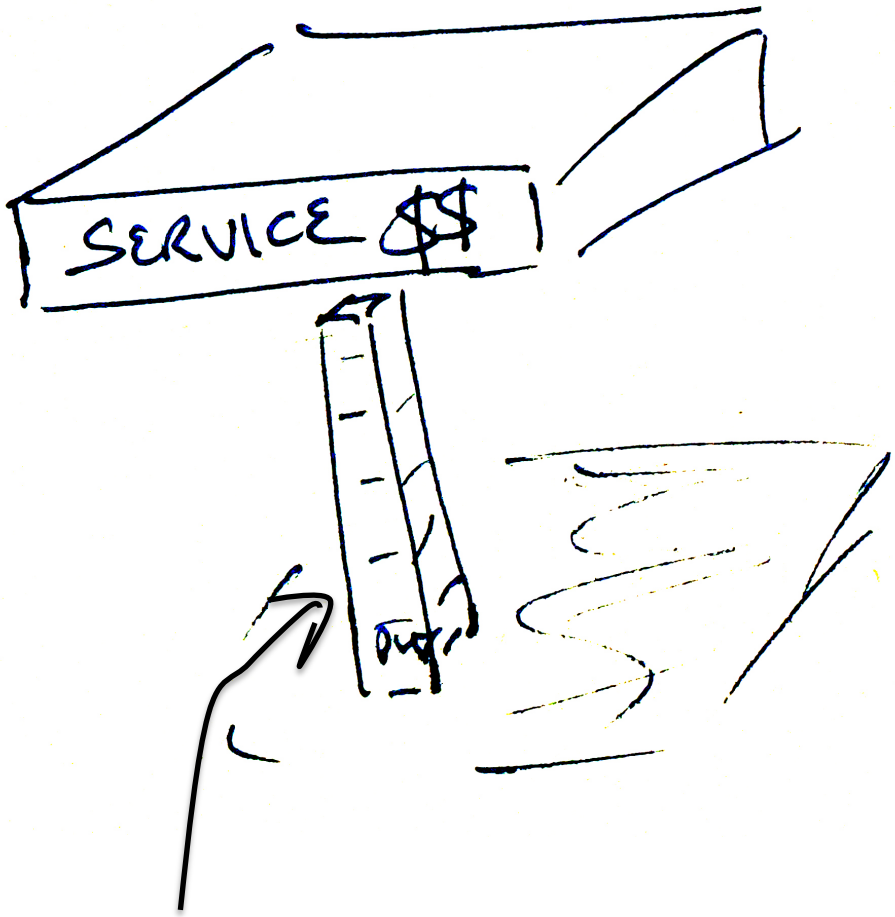

Devolution of the integrated service architecture

Where's the money to invest in new network services?

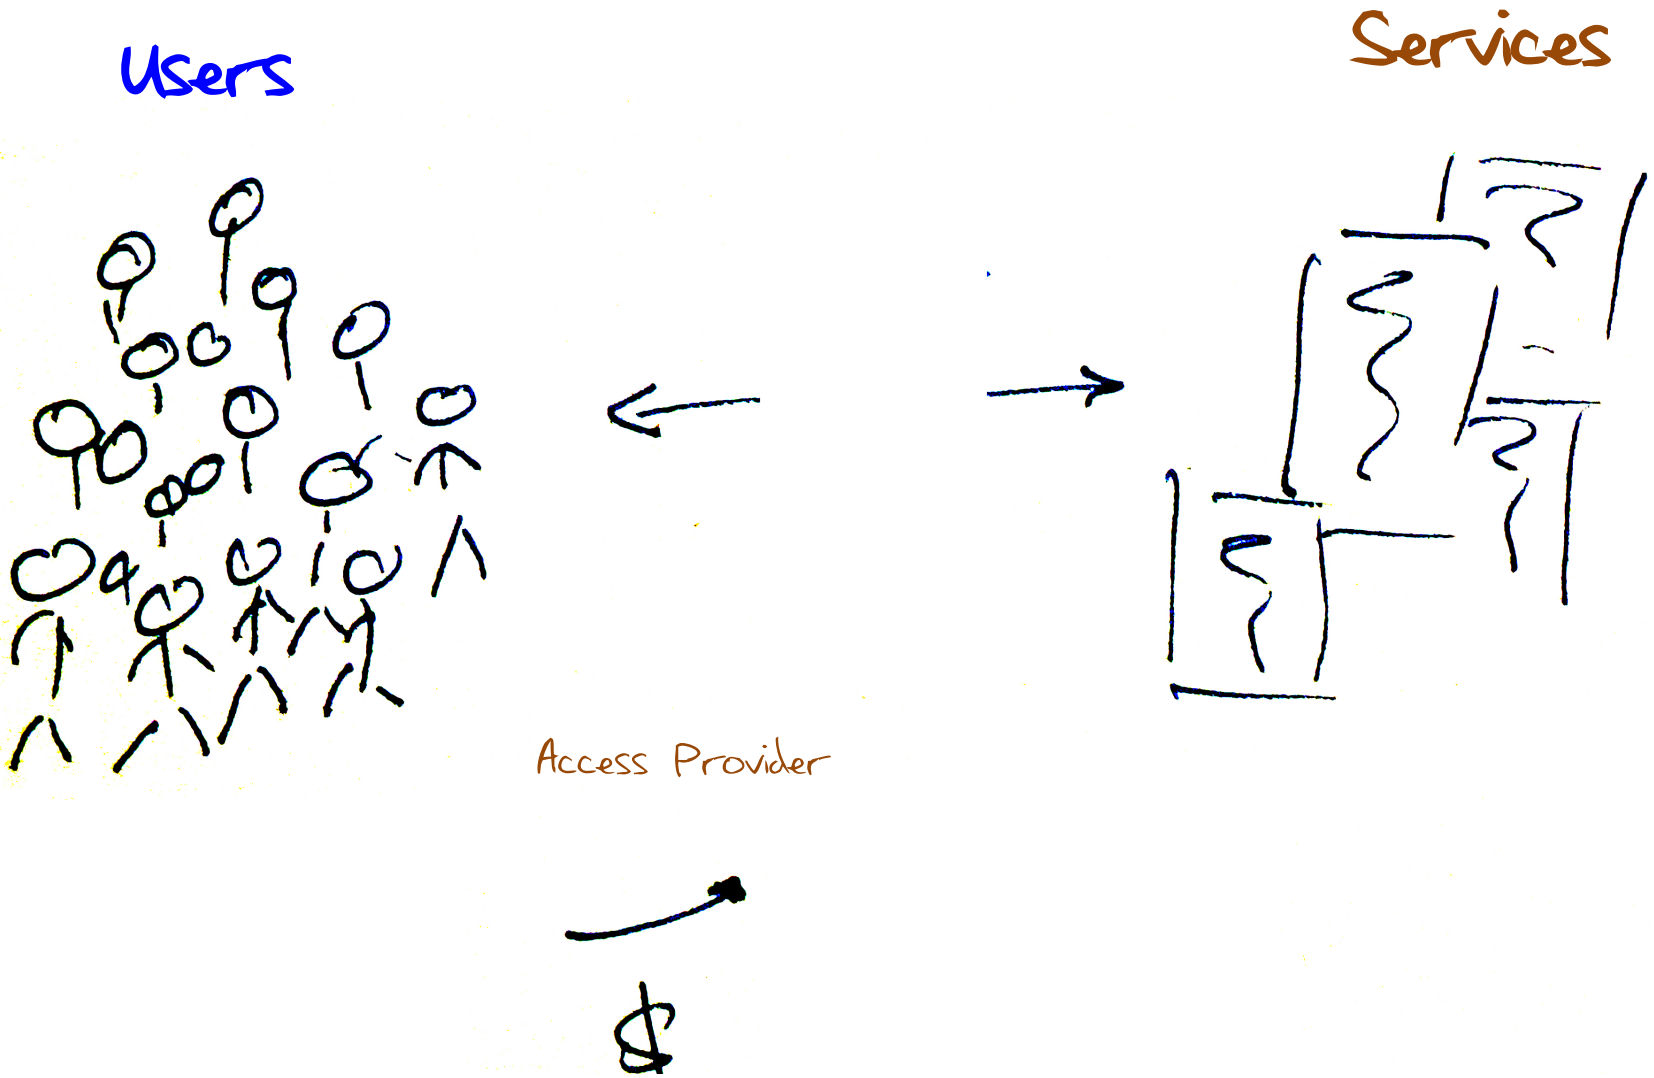
Users
Services
Access Provider
$

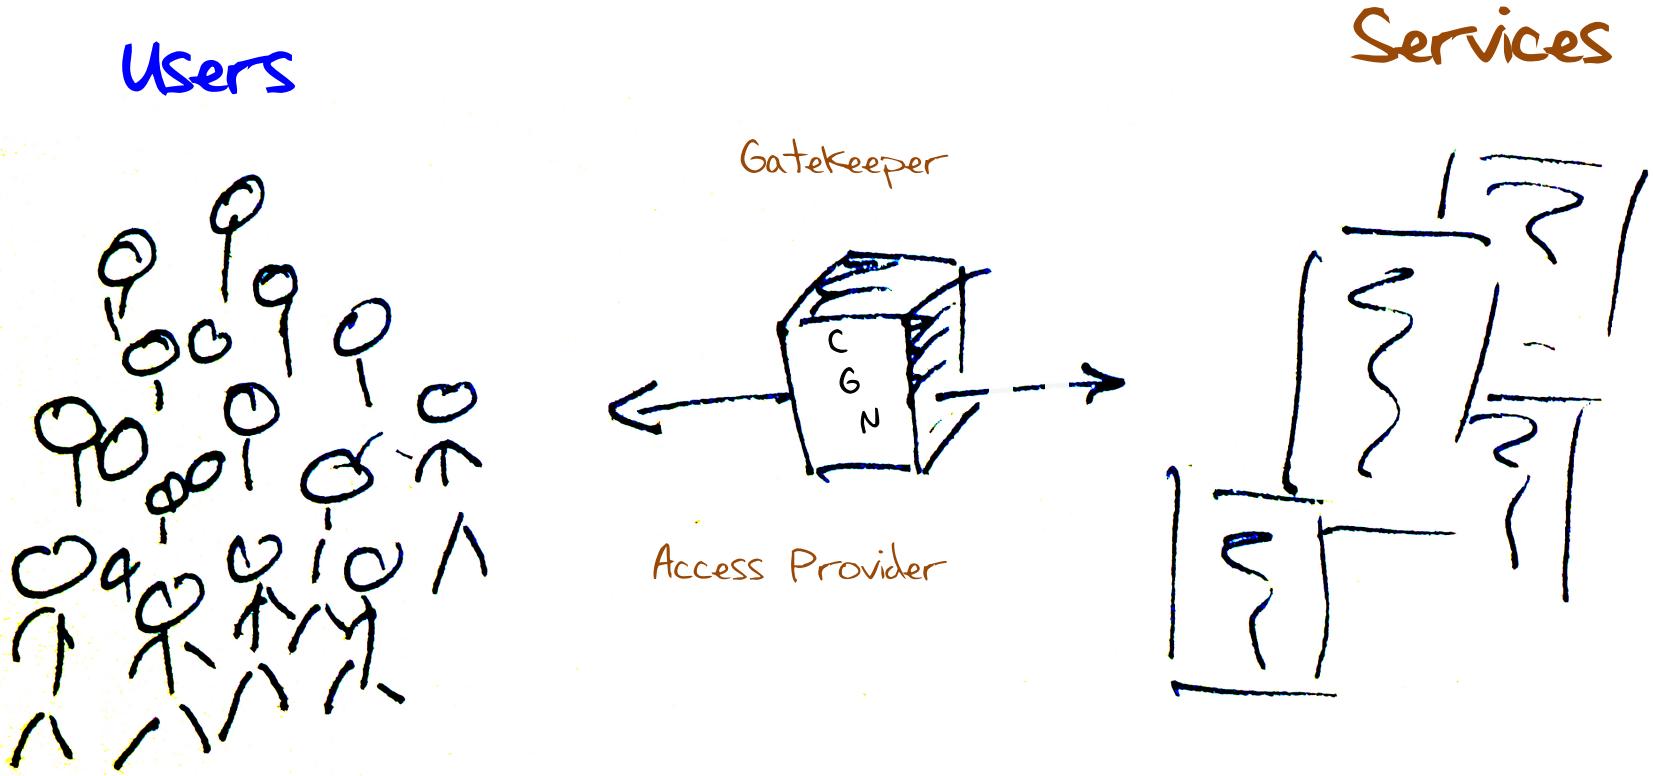
Users
Services
Gatekeeper
C
G
N
Access Provider

CGNs and ALGs and similar IPv4 rationing middleware devices provide control points in the IPv4 network that allow monetary extraction from both consumers and content providers

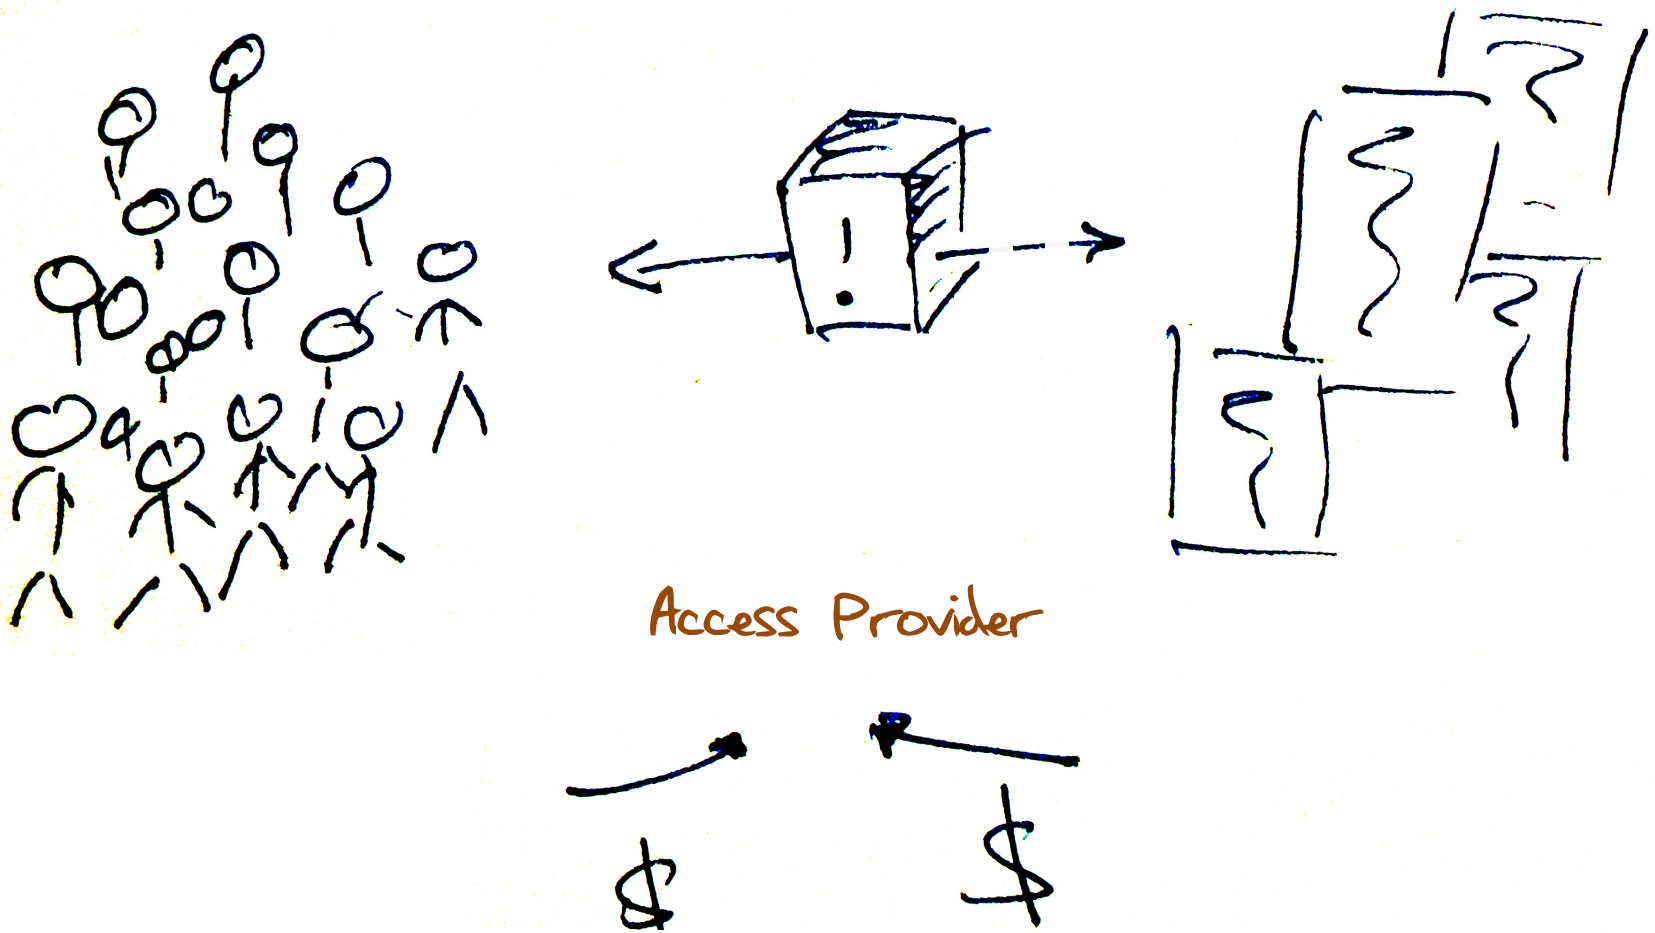

Where are we headed?

Where are we headed?

- The Internet pulled apart content from carriage, and enforced this with concepts of "network neutrality"
- Content folk have thrived
- Carriage folk are claiming otherwise

Where are we headed?

- The Internet pulled apart content from carriage, and enforced [illegible] of [illegible]
- C[illegible] folk have thrived
- Carriage folk are claiming otherwise

But now that we've exhausted V4 addresses this separation of carriage and content roles is being questioned by the ISP carriers

France Rejects Plan by Internet Provider to Block Online Ads

New York Times: 15 January 2013

By ERIC PFANNER
Published: January 7, 2013

PARIS — In a potential test case for Europe, the French government on Monday ordered a big Internet service provider to stop blocking online advertisements, saying the company had no right to edit the contents of the Web for users.

Enlarge This Image

Charles Platiau/Reuters

Fleur Pellerin, left, France's minister for the digital economy, with Maxime Lombardini, chief executive of Iliad, the parent company of the French Internet service provider Free.

Related

Ad Blocking Raises Alarm Among Firms Like Google (January 7, 2013)

The dispute has turned into a gauge of how France, and perhaps the rest of Europe, will mediate a struggle between telecommunications providers against Internet companies like Google, which generate billions of dollars in revenue from traffic that travels freely on their networks.

European telecommunications companies want a share of that money, saying they need it to finance investments in faster broadband networks — and, as the latest incident shows, they are willing to flex their muscles to get it.

Until now, European regulators have taken a laissez-faire approach, in contrast to the U.S. Federal Communications Commission, which has imposed guidelines barring operators of fixed-line broadband networks from blocking access to sites providing lawful content.

On Monday, Fleur Pellerin, the French minister for the digital economy, said she had persuaded the Internet service provider, Free, to restore full access. The company, which has long balked at carrying the huge volume of traffic from sites owned by Google without compensation, had moved last week to block online ads when it introduced a new version of its Internet access software.

WHAT HI-FI?
THE WORLD'S MOST TRUSTED TECH REVIEWS

KOREA: Leading broadband provider throttles 'data-hungry' Samsung Smart TVs

10 Feb 2012

Korea, February 2012

While TV companies are betting on smart TVs being the next big thing, a move by Korea's leading broadband provider is likely to send some shivers through their product-planning departments.

Korea Telecom, or KT, has just announced that from today it will be restricting internet access by Samsung's Smart TVs, removing their ability to download and run apps, and curtailing streaming.

The company says that the TVs, which provide internet access as well as content streaming over home broadband connections 'use internet networks without permission', and that its decision 'was was inevitable in order to protect Internet users and keep market order.'

The move comes after months of negotiations between KT and the makers of smart TVs have failed to reach agreement concerning the huge demands such TVs can put on the country's data networks.

As ever, some of the root of the dispute is financial: Korea's three major broadband providers – KT, SK and LG U+ – think that the makers of smart TVs should be charged according to the amount of network capacity their products use.

There's also the small matter of KT having its own IPTV service: called Olleh, it provides TV and movies to subscribers over broadband networks – services which would be impacted by the predicted huge growth in Smart TV sales.

So What?

Why did the Internet Work?

- Openness
- Agility
- Minimalism
- Efficiency

Why won't ~~did~~ the Internet Work?

- ~~Openness~~ → Proprietary Systems
- ~~Agility~~ → Inflexibility
- ~~Minimalism~~ → Complexity
- ~~Efficiency~~ → Cost

What's Left?

/etc/protocols:

```
#
# Internet protocols
#
# $FreeBSD: src/etc/protocols,v 1.14 2000/09/24 11:20:27 asmodai Exp $
#       from: @(#)protocols     5.1 (Berkeley) 4/17/89
#
# See also http://www.isi.edu/in-notes/iana/assignments/protocol-numbers
#
ip      0       IP              # internet protocol, pseudo protocol number
#hopopt 0       HOPOPT          # hop-by-hop options for ipv6
icmp    1       ICMP            # internet control message protocol
igmp    2       IGMP            # internet group management protocol
ggp     3       GGP             # gateway-gateway protocol
ipencap 4       IP-ENCAP        # IP encapsulated in IP (officially ``IP'')
st2     5       ST2             # ST2 datagram mode (RFC 1819)
tcp     6       TCP             # transmission control protocol
cbt     7       CBT             # CBT, Tony Ballardie <A.Ballardie@cs.ucl.ac.uk>
egp     8       EGP             # exterior gateway protocol
igp     9       IGP             # any private interior gateway (Cisco: for IGRP)
bbn-rcc 10      BBN-RCC-MON     # BBN RCC Monitoring
nvp     11      NVP-II          # Network Voice Protocol
pup     12      PUP             # PARC universal packet protocol
argus   13      ARGUS           # ARGUS
emcon   14      EMCON           # EMCON
xnet    15      XNET            # Cross Net Debugger
chaos   16      CHAOS           # Chaos
udp     17      UDP             # user datagram protocol
mux     18      MUX             # Multiplexing protocol
dcn     19      DCN-MEAS        # DCN Measurement Subsystems
hmp     20      HMP             # host monitoring protocol
prm     21      PRM             # packet radio measurement protocol
xns-idp 22      XNS-IDP         # Xerox NS IDP
trunk-1 23      TRUNK-1         # Trunk-1
trunk-2 24      TRUNK-2         # Trunk-2
leaf-1  25      LEAF-1          # Leaf-1
leaf-2  26      LEAF-2          # Leaf-2
rdp     27      RDP             # "reliable datagram" protocol
irtp    28      IRTP            # Internet Reliable Transaction Protocol
iso-tp4 29      ISO-TP4         # ISO Transport Protocol Class 4
netblt  30      NETBLT          # Bulk Data Transfer Protocol
mfe-nsp 31      MFE-NSP         # MFE Network Services Protocol
merit-inp       32      MERIT-INP       # MERIT Internodal Protocol
sep     33      SEP             # Sequential Exchange Protocol
3pc     34      3PC             # Third Party Connect Protocol
idpr    35      IDPR            # Inter-Domain Policy Routing Protocol
xtp     36      XTP             # Xpress Tranfer Protocol
ddp     37      DDP             # Datagram Delivery Protocol
idpr-cmtp       38      IDPR-CMTP       # IDPR Control Message Transport Proto
tp++    39      TP++            # TP++ Transport Protocol
il      40      IL              # IL Transport Protocol
ipv6    41      IPV6            # ipv6
sdrp    42      SDRP            # Source Demand Routing Protocol
ipv6-route      43      IPV6-ROUTE      # routing header for ipv6
ipv6-frag       44      IPV6-FRAG       # fragment header for ipv6
idrp    45      IDRP            # Inter-Domain Routing Protocol
rsvp    46      RSVP            # Resource ReSerVation Protocol
gre     47      GRE             # Generic Routing Encapsulation
mhrp    48      MHRP            # Mobile Host Routing Protocol
bna     49      BNA             # BNA
esp     50      ESP             # encapsulating security payload
ah      51      AH              # authentication header
i-nlsp  52      I-NLSP          # Integrated Net Layer Security TUBA
swipe   53      SWIPE           # IP with Encryption
narp    54      NARP            # NBMA Address Resolution Protocol
mobile  55      MOBILE          # IP Mobility
tlsp    56      TLSP            # Transport Layer Security Protocol
skip    57      SKIP            # SKIP
ipv6-icmp       58      IPV6-ICMP       # ICMP for IPv6
ipv6-nonxt      59      IPV6-NONXT      # no next header for ipv6
ipv6-opts       60      IPV6-OPTS       # destination options for ipv6
#       61              # any host internal protocol
cftp    62      CFTP            # CFTP
#       63              # any local network
sat-expak       64      SAT-EXPAK       # SATNET and Backroom EXPAK
kryptolan       65      KRYPTOLAN       # Kryptolan
rvd     66      RVD             # MIT Remote Virtual Disk Protocol
ippc    67      IPPC            # Internet Pluribus Packet Core
#       68              # any distributed file system
sat-mon 69      SAT-MON         # SATNET Monitoring
visa    70      VISA            # VISA Protocol
ipcv    71      IPCV            # Internet Packet Core Utility
```

TCP, bits of UDP, and nothing else!

What's Left?

/etc/services:

```
dna-cml         436/udp     # DNA-CML
dna-cml         436/tcp     # DNA-CML
#                           Dan Flowers <flowers@smaug.lkg.dec.com>
comscm          437/udp     # comscm
comscm          437/tcp     # comscm
#                           Jim Teague <teague@zso.dec.com>
dsfgw           438/udp     # dsfgw
dsfgw           438/tcp     # dsfgw
#                           Andy McKeen <mckeen@osf.org>
dasp            439/udp     # dasp      tommy@inlab.m.eunet.de
dasp            439/tcp     # dasp      Thomas Obermair
#                           Thomas Obermair <tommy@inlab.m.eunet.de>
sgcp            440/udp     # sgcp
sgcp            440/tcp     # sgcp
#                           Marshall Rose <mrose@dbc.mtview.ca.us>
decvms-sysmgt   441/udp     # decvms-sysmgt
decvms-sysmgt   441/tcp     # decvms-sysmgt
#                           Lee Barton <barton@star.enet.dec.com>
cvc_hostd       442/udp     # cvc_hostd
cvc_hostd       442/tcp     # cvc_hostd
#                           Bill Davidson <billd@equalizer.cray.com>
https           443/udp     # http protocol over TLS/SSL
https           443/tcp     # http protocol over TLS/SSL
#                           Kipp E.B. Hickman <kipp@mcom.com>
snpp            444/udp     # Simple Network Paging Protocol
snpp            444/tcp     # Simple Network Paging Protocol
#                           [RFC1568]
microsoft-ds    445/udp     # Microsoft-DS
microsoft-ds    445/tcp     # Microsoft-DS
#                           Pradeep Bahl <pradeepb@microsoft.com>
ddm-rdb         446/udp     # DDM-RDB
ddm-rdb         446/tcp     # DDM-RDB
ddm-dfm         447/udp     # DDM-RFM
ddm-dfm         447/tcp     # DDM-RFM
#                           Jan David Fisher <jdfisher@VNET.IBM.COM>
ddm-ssl         448/udp     # DDM-SSL
ddm-ssl         448/tcp     # DDM-SSL
#                           Steve Ritland <srr@vnet.ibm.com>
as-servermap    449/udp     # AS Server Mapper
as-servermap    449/tcp     # AS Server Mapper
#                           Barbara Foss <BGFOSS@rchvmv.vnet.ibm.com>
tserver         450/udp     # Computer Supported Telecomunication Applications
tserver         450/tcp     # Computer Supported Telecomunication Applications
#                           Harvey S. Schultz <harvey@acm.org>
sfs-smp-net     451/udp     # Cray Network Semaphore server
sfs-smp-net     451/tcp     # Cray Network Semaphore server
sfs-config      452/udp     # Cray SFS config server
sfs-config      452/tcp     # Cray SFS config server
#                           Walter Poxon <wdp@ironwood.cray.com>
creativeserver  453/udp     # CreativeServer
creativeserver  453/tcp     # CreativeServer
contentserver   454/udp     # ContentServer
contentserver   454/tcp     # ContentServer
creativepartnr  455/udp     # CreativePartnr
creativepartnr  455/tcp     # CreativePartnr
```

HTTPs over TCP is the last remaining bastion of end-to-end coherency in today's Natted net!

What does this imply about open networking?

Openness is difficult to sustain in a restricted environment dominated by scarcity – the space for innovation is limited and open spaces are consumed by incumbents

So we need to chose carefully!

We need to think about how to build a post-PC world where content, computation, storage and communications are sustainable abundant and openly available commodities.

This objective is basically incompatible with the current momentum of the Internet

And its not yet clear which path the internet will take!

And its not yet clear which path ~~the internet~~ will take!

market forces

Thank You!