“apocalypse then”: ipv4 address space depletion

cooperative association for internet data analysis

26 october 2005

problem statement

environmental problem: running out of addresses!

solution started in 90s, but uptake slower than "expected"

important connection:
the research community has same problem.

underlying question:
how do we innovate architecturally?
deeper question: how do we leave the Internet better than we found it?

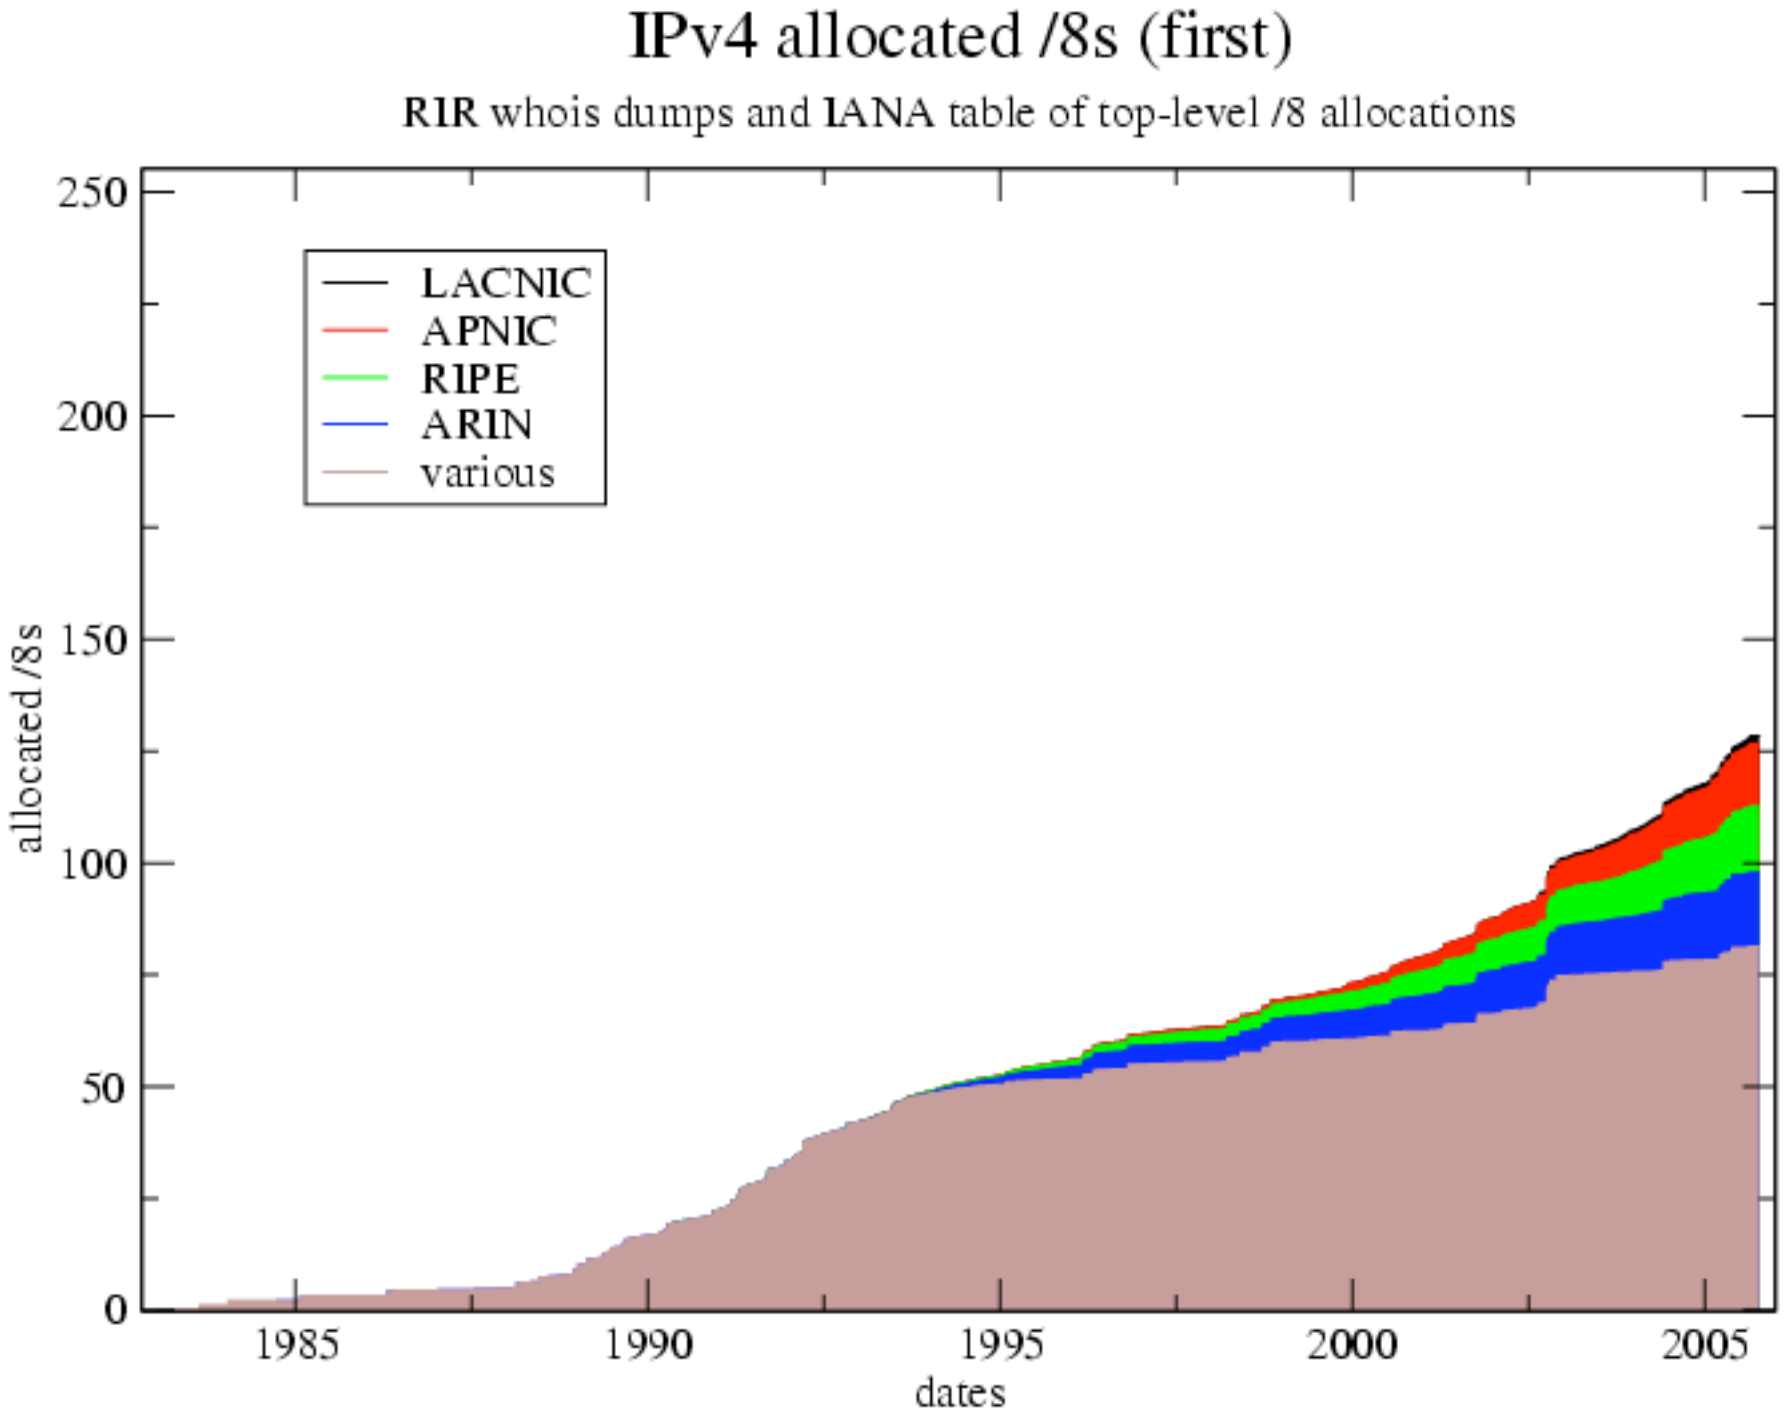
IPv4 allocated /8s (first)
R1R whois dumps and IANA table of top-level /8 allocations
LACNIC
APNIC
R1PE
ARIN
various
allocated /8s
250
200
150
100
50
0
1985
1990
1995
2000
2005
dates

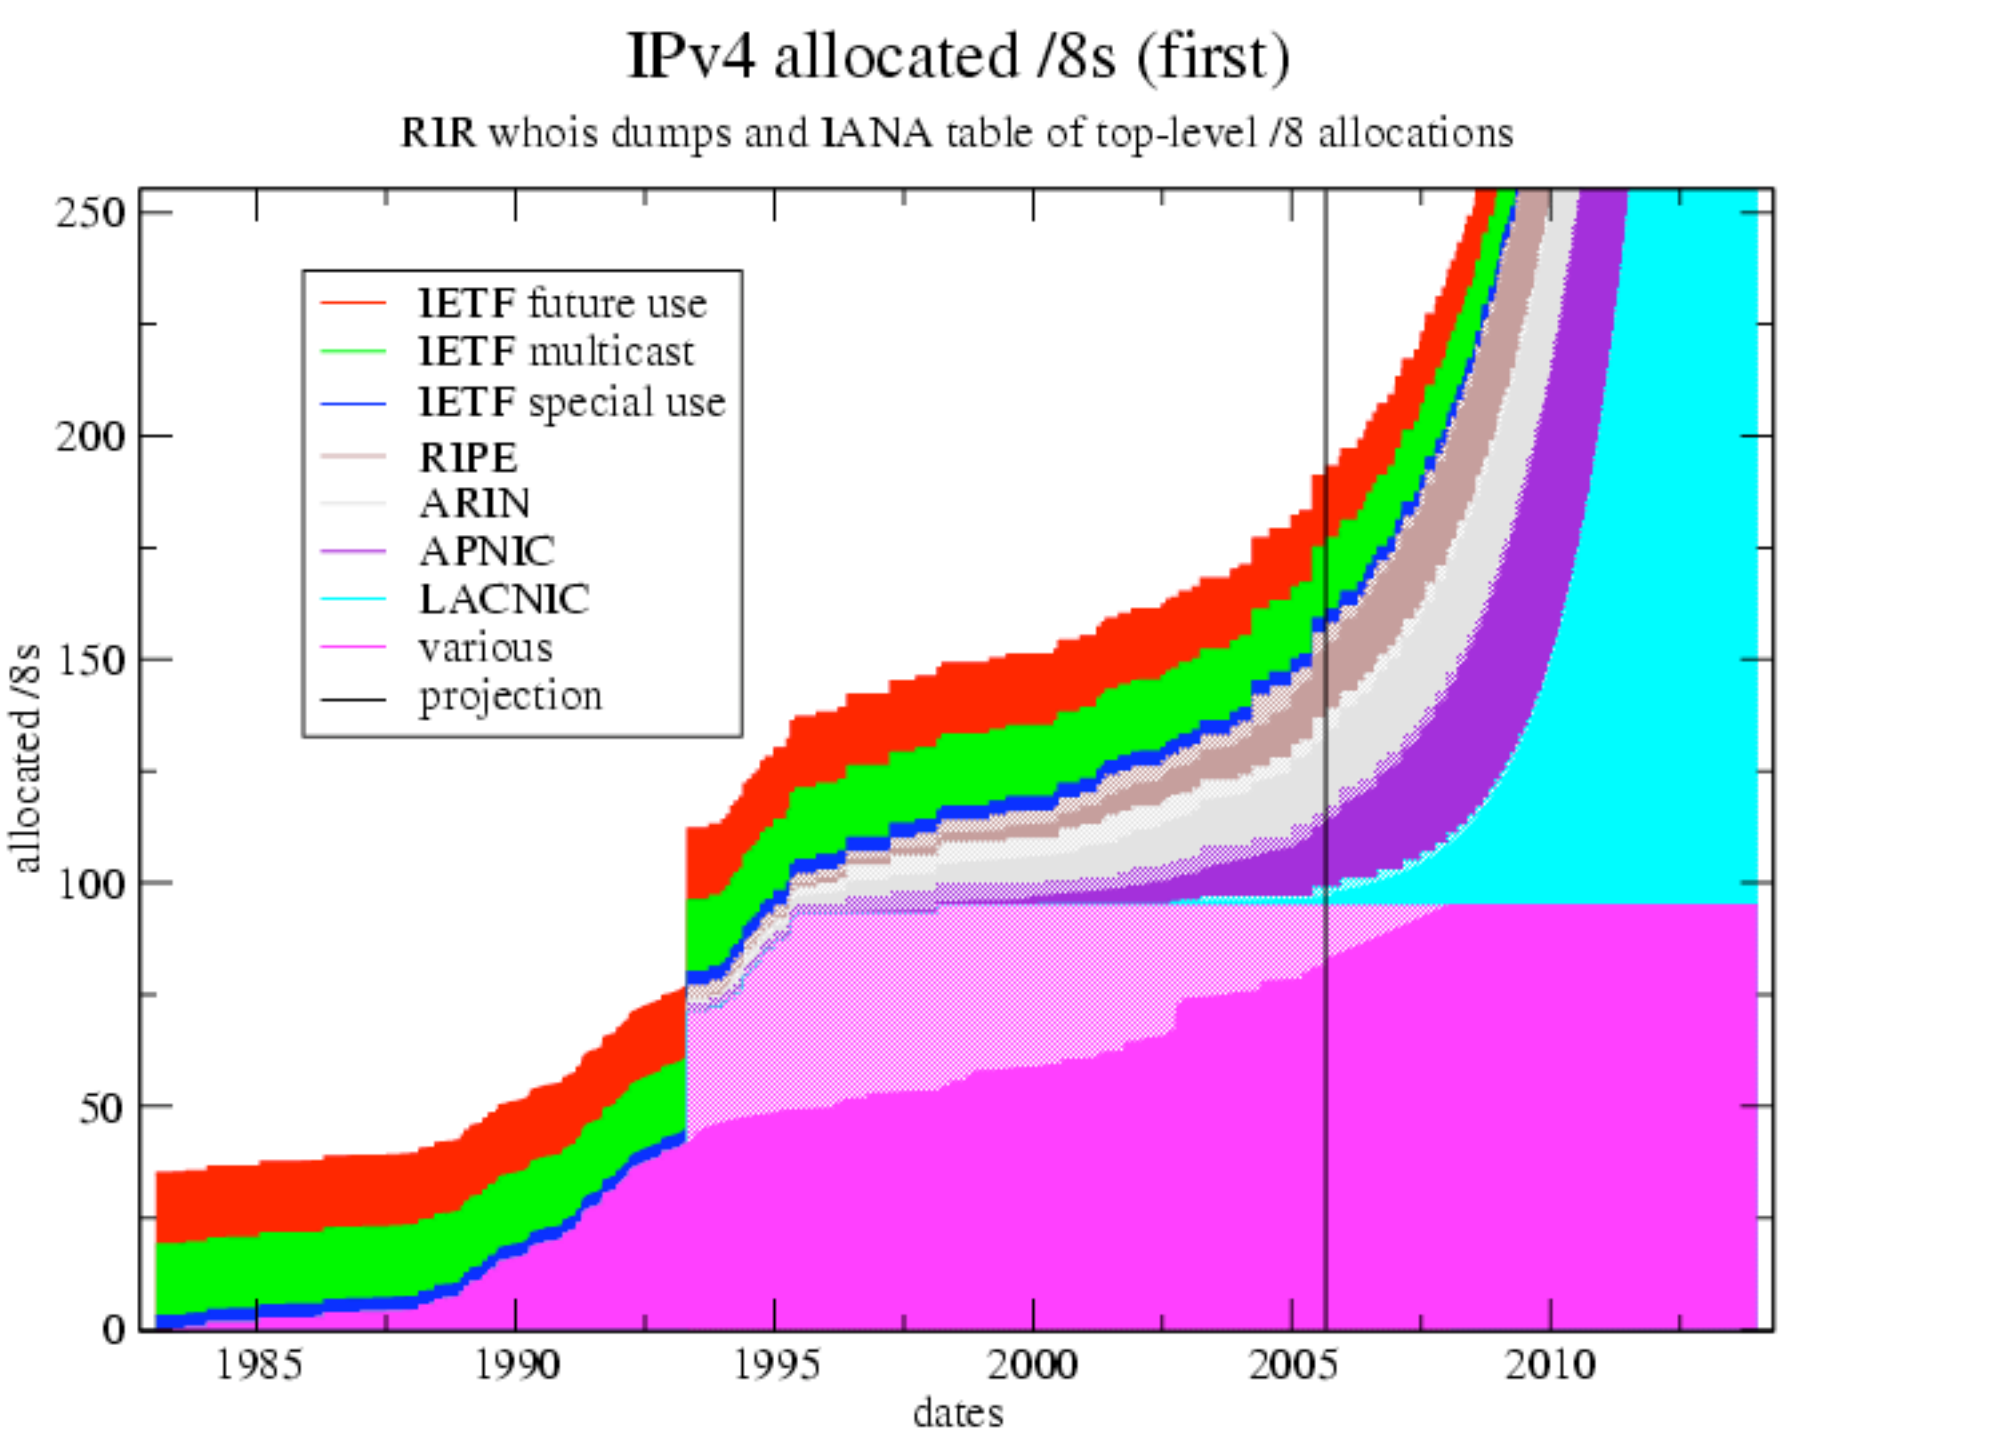
IPv4 allocated /8s (first)
RIR whois dumps and IANA table of top-level /8 allocations
IETF future use
IETF multicast
IETF special use
RIPE
ARIN
APNIC
LACNIC
various
projection
allocated /8s
250
200
150
100
50
0
1985
1990
1995
2000
2005
2010
dates

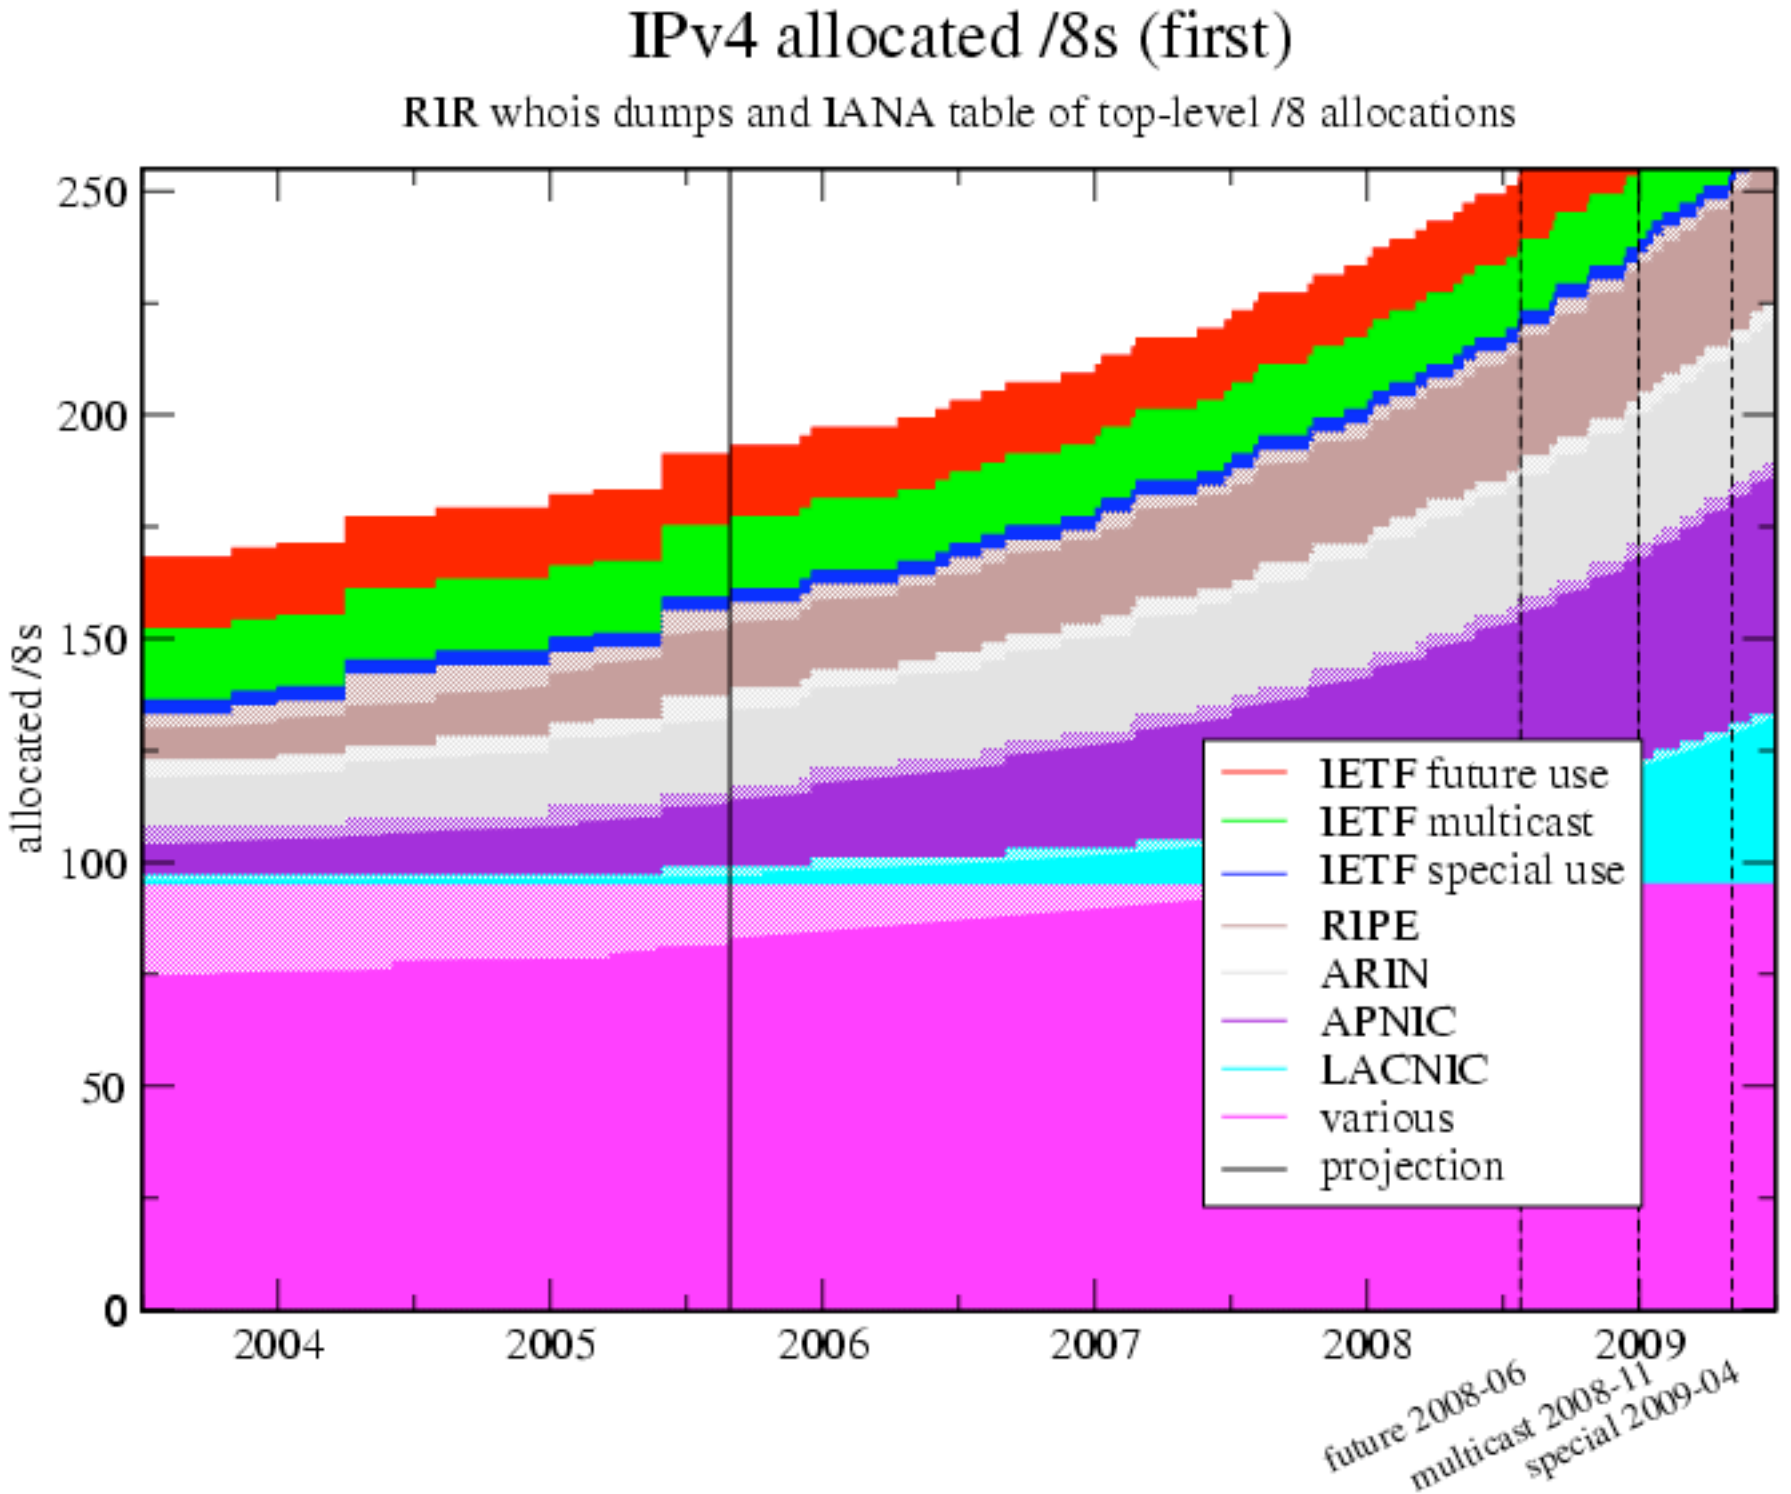
IPv4 allocated /8s (first)
RIR whois dumps and IANA table of top-level /8 allocations
allocated /8s
250
200
150
100
50
0
2004
2005
2006
2007
2008
2009
IETF future use
IETF multicast
IETF special use
RIPE
ARIN
APNIC
LACNIC
various
projection
future 2008-06
multicast 2008-11
special 2009-04

Breakdown by Num Allocations per Organization of ARIN IPv4 Space

ARIN whois data (20050831); excluding DoDNIC, JPNIC, and pre-RIR /8 allocations; stacked plot

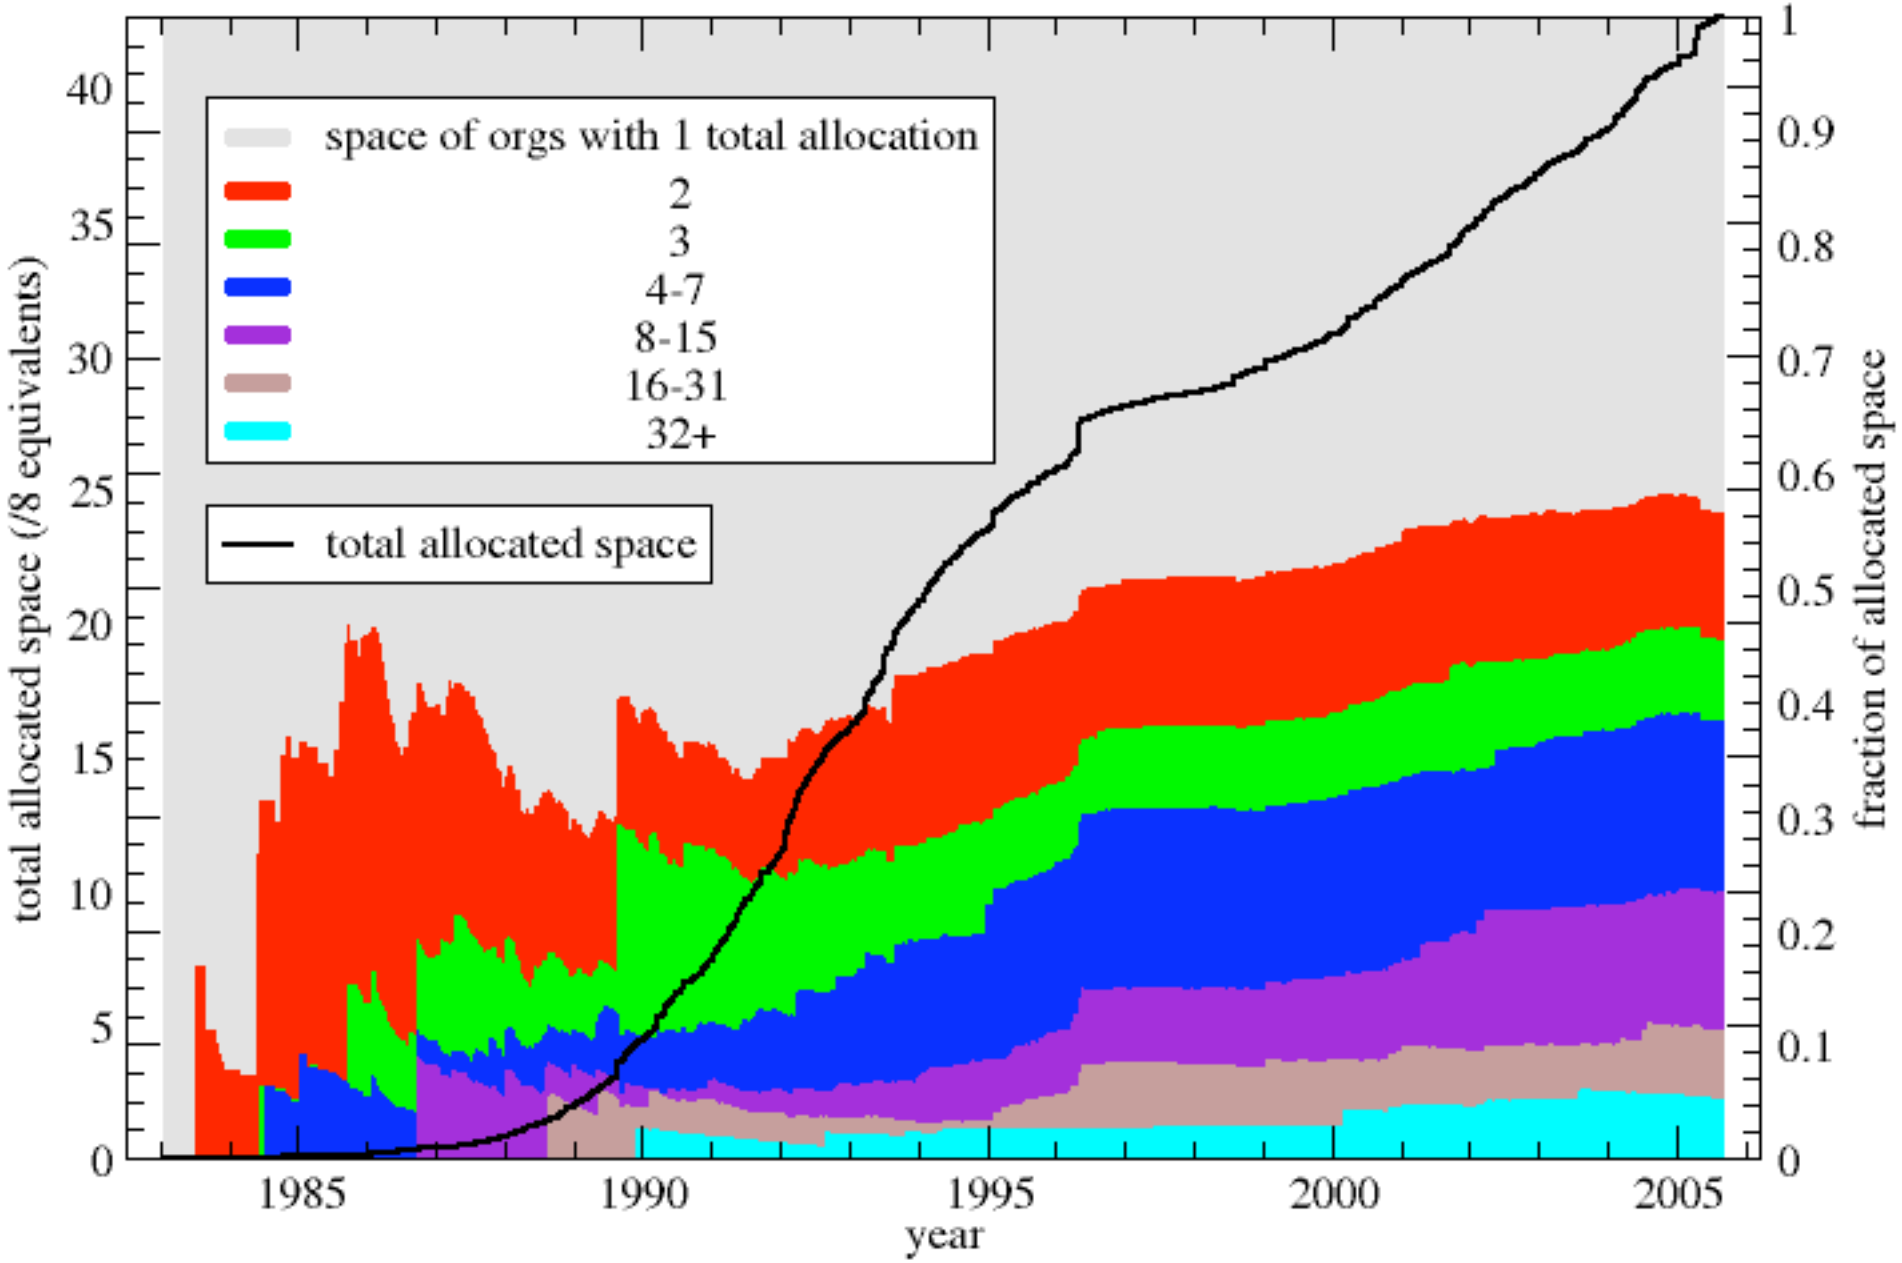

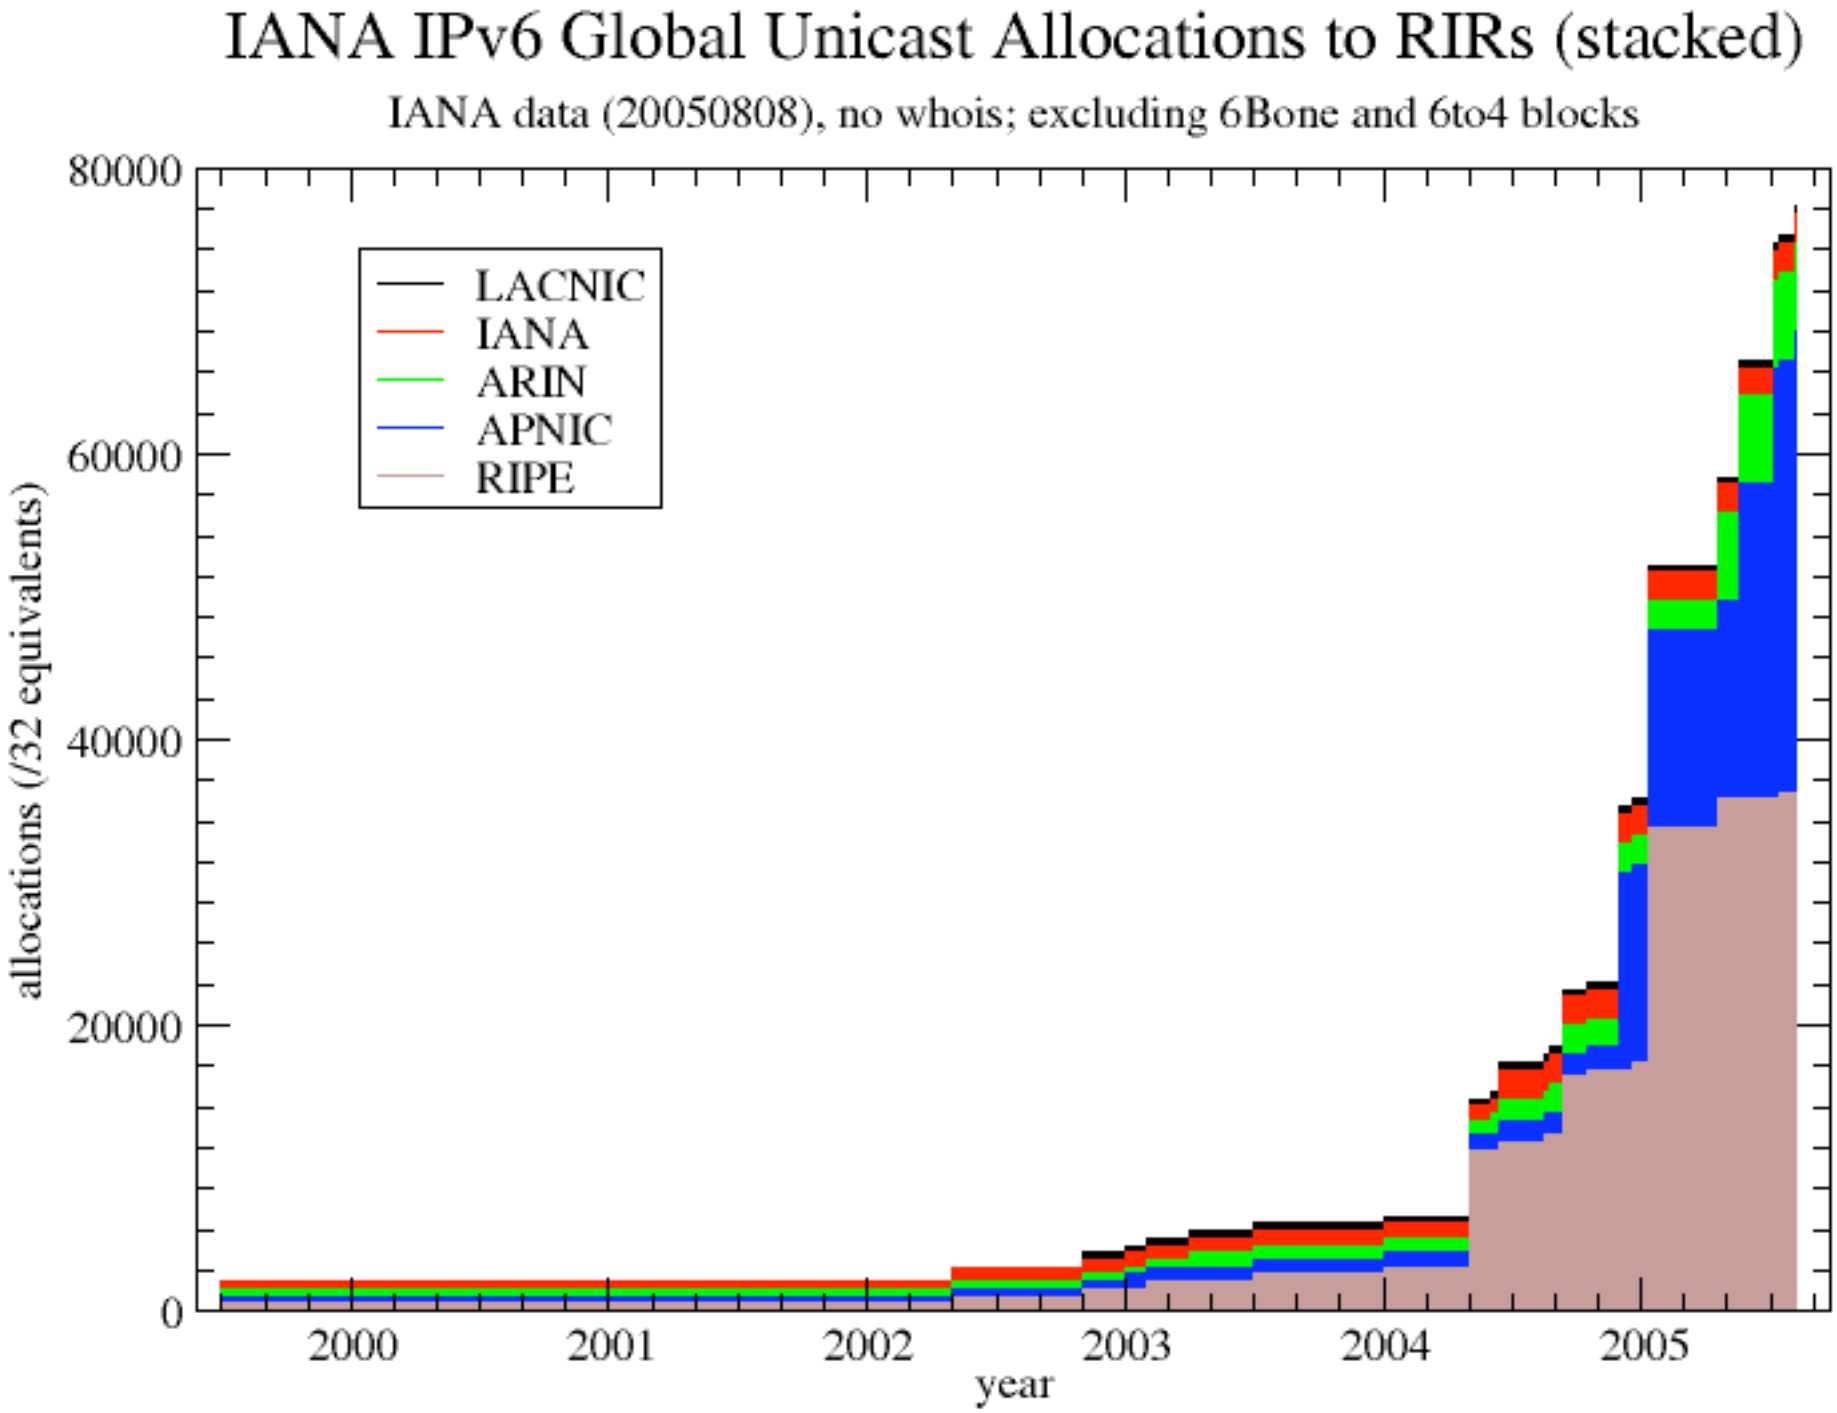
IANA IPv6 Global Unicast Allocations to RIRs (stacked)
IANA data (20050808), no whois; excluding 6Bone and 6to4 blocks
LACNIC
IANA
ARIN
APNIC
RIPE
allocations (/32 equivalents)
80000
60000
40000
20000
0
2000
2001
2002
2003
2004
2005
year

the news gets worse.

this isn't even the data we need
to inform a discussion about IPv6 uptake.

what [predicting] the future needs

the numbers that will drive our future have different units

innovation requires capital.

INNOVATOR	EPS ($)	MKT CAP ($B)
MCIW	-11.22	6.5
SPRNT/NXTL	-0.31	34
VERIO/NTT	1.98	71.6
LEVEL3	-0.74	1.9
SBC/T	1.41	78
QWEST	-0.45	7.7
COGENT	-7.42	0.2
GLBC	-13.84	0.3
SAVVIS	-0.90	0.12
ABOVENET	n/a	n/a
WILTEL	n/a	n/a
TELEGLOBE	-0.74	0.2
C&W	0.70	4.7B
TWTELCOM	-1.12	1.0
(TWARNER)	0.48	82
XO	-2.18	0.4

source: finance.yahoo.com, 25 oct 2005

where is the capital for innovation?

ironically, it's from where the innovation is happening.

we're trying to make the future safe for innovation, but we have to innovate to do so.

INNOVATOR	EPS ($)	MKT CAP ($B)
CISCO	0.87	108
GOOGLE	3.41	97
AMAZON	1.25	19
YAHOO	1.07	49
EBAY	0.73	51
JUNIPER	0.53	13
APPLE	1.56	47.
INTEL	1.33	141
VERISIGN	0.93	6.15
DELL	1.27	76.3
MICROSOFT	1.12	269B

source: finance.yahoo.com, 25 oct 2005

capital distribution problem

(the ones who need to innovate in the core don't have capital)

INNOVATO	EPS ($)	MKT CAP ($B)
MCIW	-11.22	6.5
SPRNT/NXTL	-0.31	34
VERIO/NTT	1.98	71.6
LEVEL3	-0.74	1.9
SBC/T	1.41	78
QWEST	-0.45	7.7
COGENT	-7.42	0.2
GLBC	-13.84	0.3
SAVVIS	-0.90	0.12
ABOVENET	n/a	n/a
WILTEL	n/a	n/a
TELEGLOBE	-0.74	0.2
C&W	0.70	4.7B
TWTELCOM	-1.12	1.0
(TWARNER)	0.48	82
XO	-2.18	0.4

INNOVATOR	EPS ($)	MKT CAP ($B)
CISCO	0.87	108
GOOGLE	3.41	97
AMAZON	1.25	19
YAHOO	1.07	49
EBAY	0.73	51
JUNIPER	0.53	13
APPLE	1.56	47.
INTEL	1.33	141
VERISIGN	0.93	6.15
DELL	1.27	76.3
MICROSOFT	1.12	269B

source: finance.yahoo.com, 25 oct 2005

how are we at innovating anyway?

failures:	successes:
atm multicast qos rsvp diffserv s*bgp dnssec	email http browser blogs bittorrent search engines voip
what's in common: competing in middle	what's in common: cooperating at edge

IPv6 (in US) blocked on two things

(1) capital

applications, hardware, stable router OS, ops training, users

(2) incentive

"it would have to stop a fatal threat to my business" (cidr,y2k)
"it would have to support things i want to do now but can't"
"it would have to be free"
"it would have to be fun"

(source: nanog voices)

not only did we forget to include (a source of capital for) a 'bell labs',
we killed the last one the communications industry had. oops.

choices

- confront how to structure a market in IPv4 addresses (we gave up once)
- confront how to govern a transition to IPv6
- govern reclamation/use of IPv4 addresses
- ask IETF to ‘go back and try again’
- ask research community to save us

who pays for any of those? what do they cost?

reclamation 'potential'

(ISI.edu pung entire IPv4 address space in 2003)

2003-06-01
- 20.2% of prefixes did not respond to probe
- 3.2 /8 equivalents (4.4% of total routed) could be returned

2003-10-08
- 18.9% of prefixes did not respond to probe
- 2.9 /8 equivalents (3.9% of total routed) could be returned

2003-06-01
- 4,764,826 /24 equivalents in the routing table
- 74.5% of /24s did not respond to probe
- 54.2 /8 equivalents (74.5% of total routed) could be returned

2003-10-08
- 4,872,851 /24 equivalents in the routing table
- 75.1% of /24s did not respond to probe
- 55.8 /8 equivalents (75.1% of total routed) could be returned

sources: isi.edu via predict.org (ping); routeviews (BGP)

historical justification notwithstanding,
policymakers may conclude we haven't nailed stewardship.

historical context

1966: Larry Roberts, “Towards a Cooperative Network of Time-Shared Computers” (first ARPANET plan)
(*pssst*: *you are still using the same stuff*)

1969: ARPANET commissioned by DoD for research

1977: Kleinrock’s paper “Hierarchical Routing for large networks; performance evaluation and optimization”
(pssst: you are still using the same stuff)

1980: ARPANET grinds to complete halt due to (statusmsg) virus

1986: NSFNET backbone, 56Kbps. NSF-funded regionals.
IETF, IRTF. MX records (NAT for mail)

1991: CIX, NSFNET upgrades to T3, allows .com. web. PGP.

1995: under pressure from USG, NSF transitions backbone to competitive market. no consideration of economics or security.

“according to best available data”

what do we do in provisioning crises?
and how successful have we been?

- backbone provisioning: free market
- DNS: free market
- address space: _____?

how have free markets been to infrastructure?

- backbone provisioning: profits toward (& below) zero, consolidation toward monopoly, no security, no innovation.

"The design was immediately and grievously flawed for not only was there no plan for the privatization and no criteria by which to measure its success or failure. Furthermore it held unacknowledged economic implications for what was being privatized was the only part of the network that had no customers." --Gordon Cook, 1992

- DNS: profits toward zero, consolidation, no security, unfinished internationalization. some innovation ("sitefinder")

"Turning hegemony into democracy by peaceful means has been done only a few time in human history, and the outlook for this time isn't good. --Paul Vixie, 2005, fm.vix.com

- address space: _____?

learning from mistakes

- most important thing we've learn so far: society has decided IP is like water.
- strong implications for an industry wanting to sell wine (need complex aromas). but that's what the data shows.
- when you want to move water, you care about 4 things: safe, scalable, sustainable, stewardship.

how have we done?

- how safe is the Internet? data doesn't look good.
- how scalable is the Internet? data doesn't look good.
- how sustainable is the Internet? data doesn't look good.
- how did we do on stewardship? data doesn't look good.

there is good news

- we made something so great, everyone wants it.
- in fact many of us want it more than once! (um..)
- the current industry is a historical artifact of technical and (science & regulatory) policy ‘innovations’ in the 60s, 70s, 80s, 90s, and 00s
- in the meantime, it became global critical infrastructure. oops.

who pays for critical infrastructure?

- how is the Internet different from other critical infrastructure? (hint: who protects public interest?)
- what will give? how long will it take?
- according to history: capital will be allocated to architectural innovation

(rural electrification: 48 years, finally using cooperatively owned companies and federal funding act)

what worries me most

- that we won't learn from our own mistakes
- that we will design another architecture with no actual plan for economic sustainability (much less incenting further innovation in a competitive market!)
- that other forces will "code" innovation into the architecture
- that "denial is a special case of optimism"

reality

this isn't a data analysis problem.
it's a scenario planning problem.

scenario planning

Peter Schwartz, The Art of the Long View

- of course we should try to change the future: it's the only thing we can change.
- goal: not an accurate picture of tomorrow, but rather making better decisions about the future
- what are long term forces and how do they interact? how are they measured? who is measuring? what tradeoffs do they imply?
- are we leaving a better future than we found? or at least as good?
- are we net gaining or net losing freedom?

scenario planning

example, oil crisis 1973

japan:	us:
rebuilt everything, made themselves the most energy efficient industrial economy on the planet in less than 3 years.	political leaders fought to preserve spirit of 'manifest destiny'. treat crisis as temporary setback in a land of opportunity.
tradeoff: severe recession as they rebuilt capital infrastructure	tradeoff: need for hegemonic influence over oil-producing regions

source: Peter Schwartz's Art of the Long View

network scenario planning

bring economists, technologists, policymakers, researchers together to talk about the long view of progress on

- ubiquitous, low-cost, open, secure infrastructure
- modern emergency communications services
- diagnosis and configuration technologies
- economic analysis and regulatory policy

how caida is trying to help

1. scenario planning workshop -- early 2005, present results at next ARIN meeting.
   (email kc if interested)
2. wiki -- dec 2005
3. internet id consumption mailing list, now -- iic-wg@caida.org
4. open to your suggestions

a pessimist complains about the wind;
an optimist expects the wind;
a realist adjust the sails. --william ward