IPv6 deployment issues

Brian Carpenter
Department of Computer Science

THE UNIVERSITY OF AUCKLAND
NEW ZEALAND
Te Whare Wānanga o Tāmaki Makaurau

May 2008

Why we need IPv6

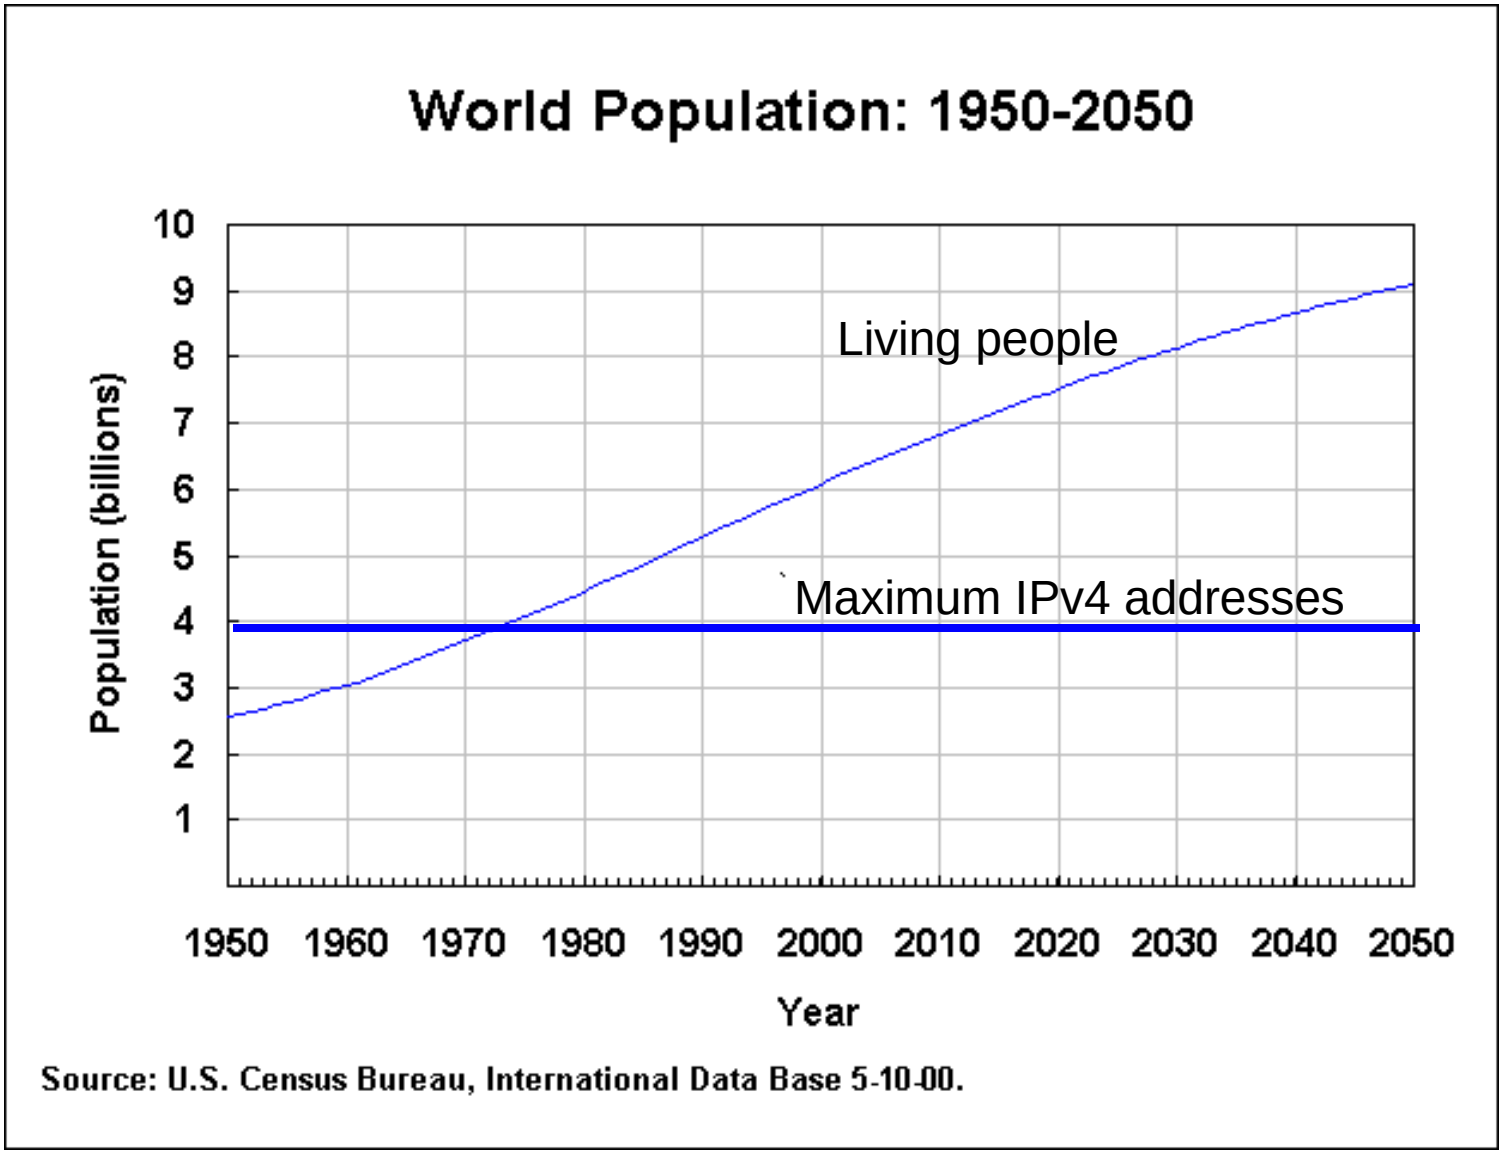

Obviously, having fewer addresses than people is silly

Slide stolen from
http://www.potaroo.net/presentations/2007-11-06-pita-ipv4.pdf

The IPv4 Consumption Model

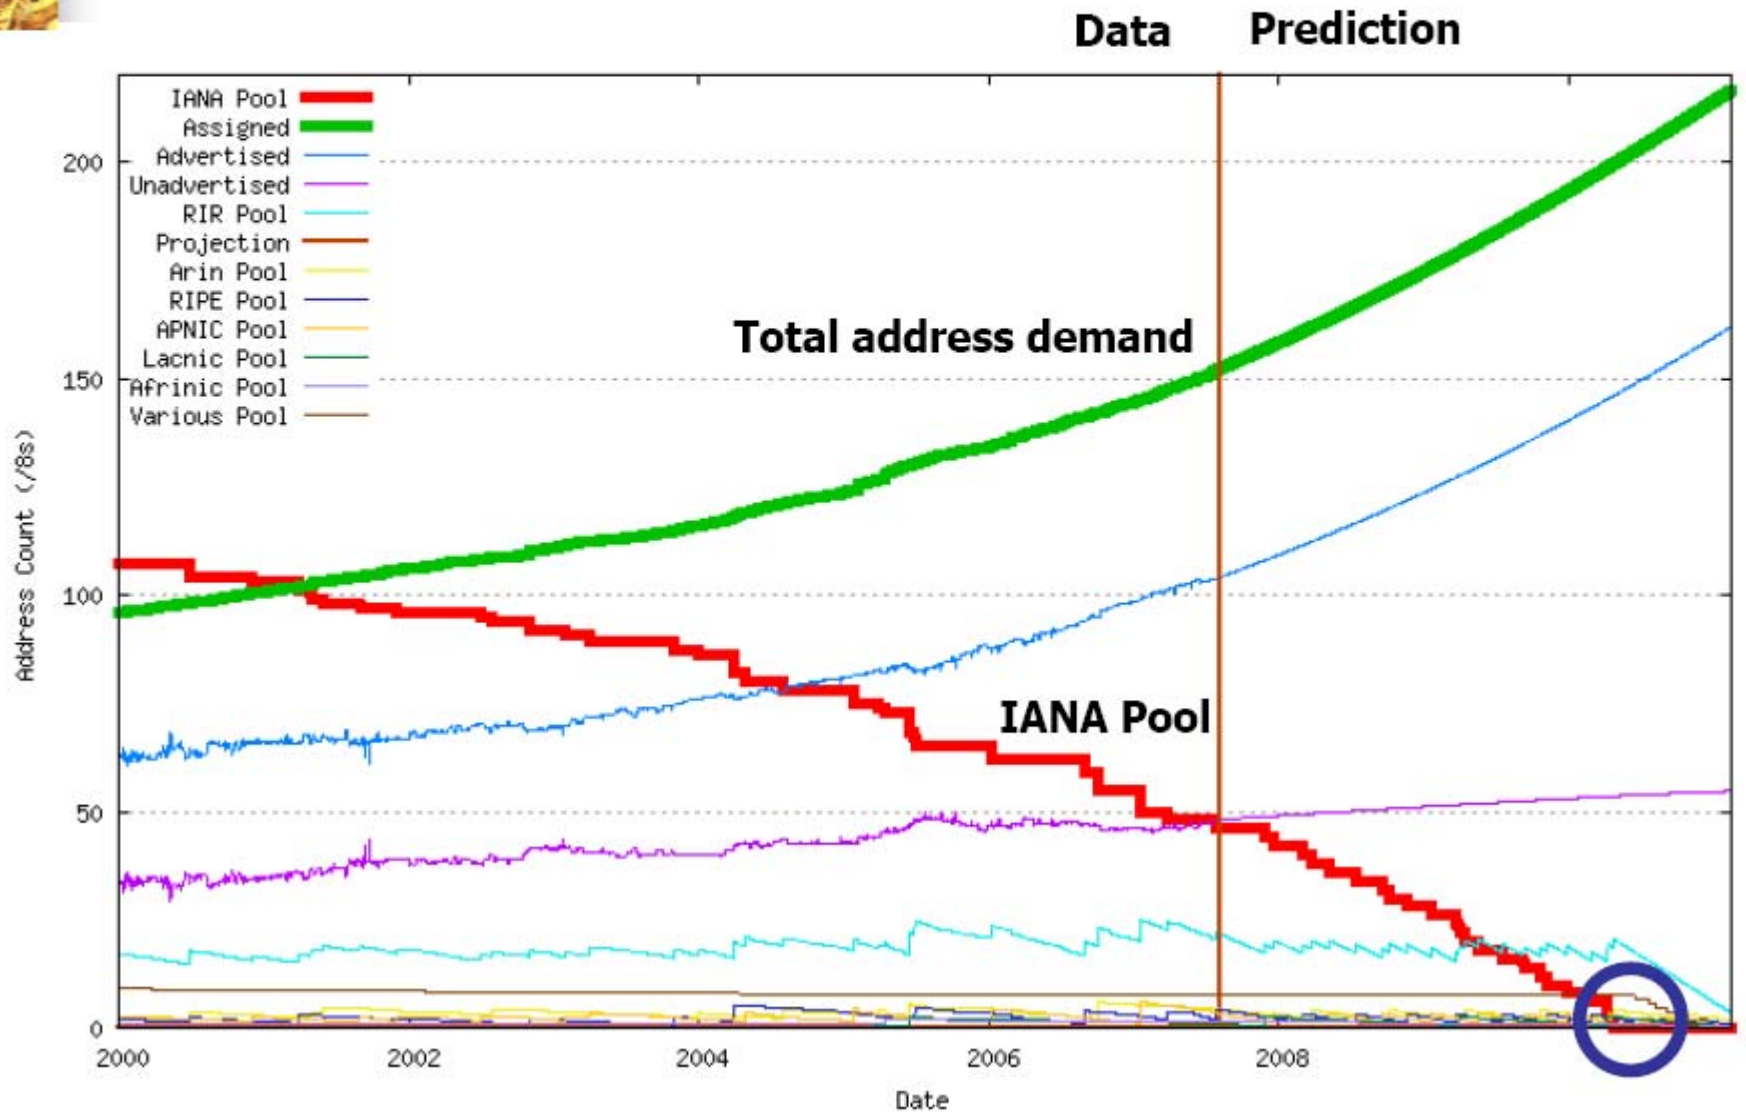

IPv4 and IPv6 coexistence

- The old and new versions will have to live together and work together for many years.
- IPv6 can be carried over IPv4 in tunnels
  - IPv6 packets encapsulated in IPv4 packets
- Servers and ISPs will become dual stack, able to support IPv4 and IPv6 clients simultaneously.
- Application proxies will be able to map IPv4 clients to IPv6 servers, or the opposite.
- Direct translation of v4 to v6 at packet level doesn't work well.

Coexistence mechanisms (simple version)

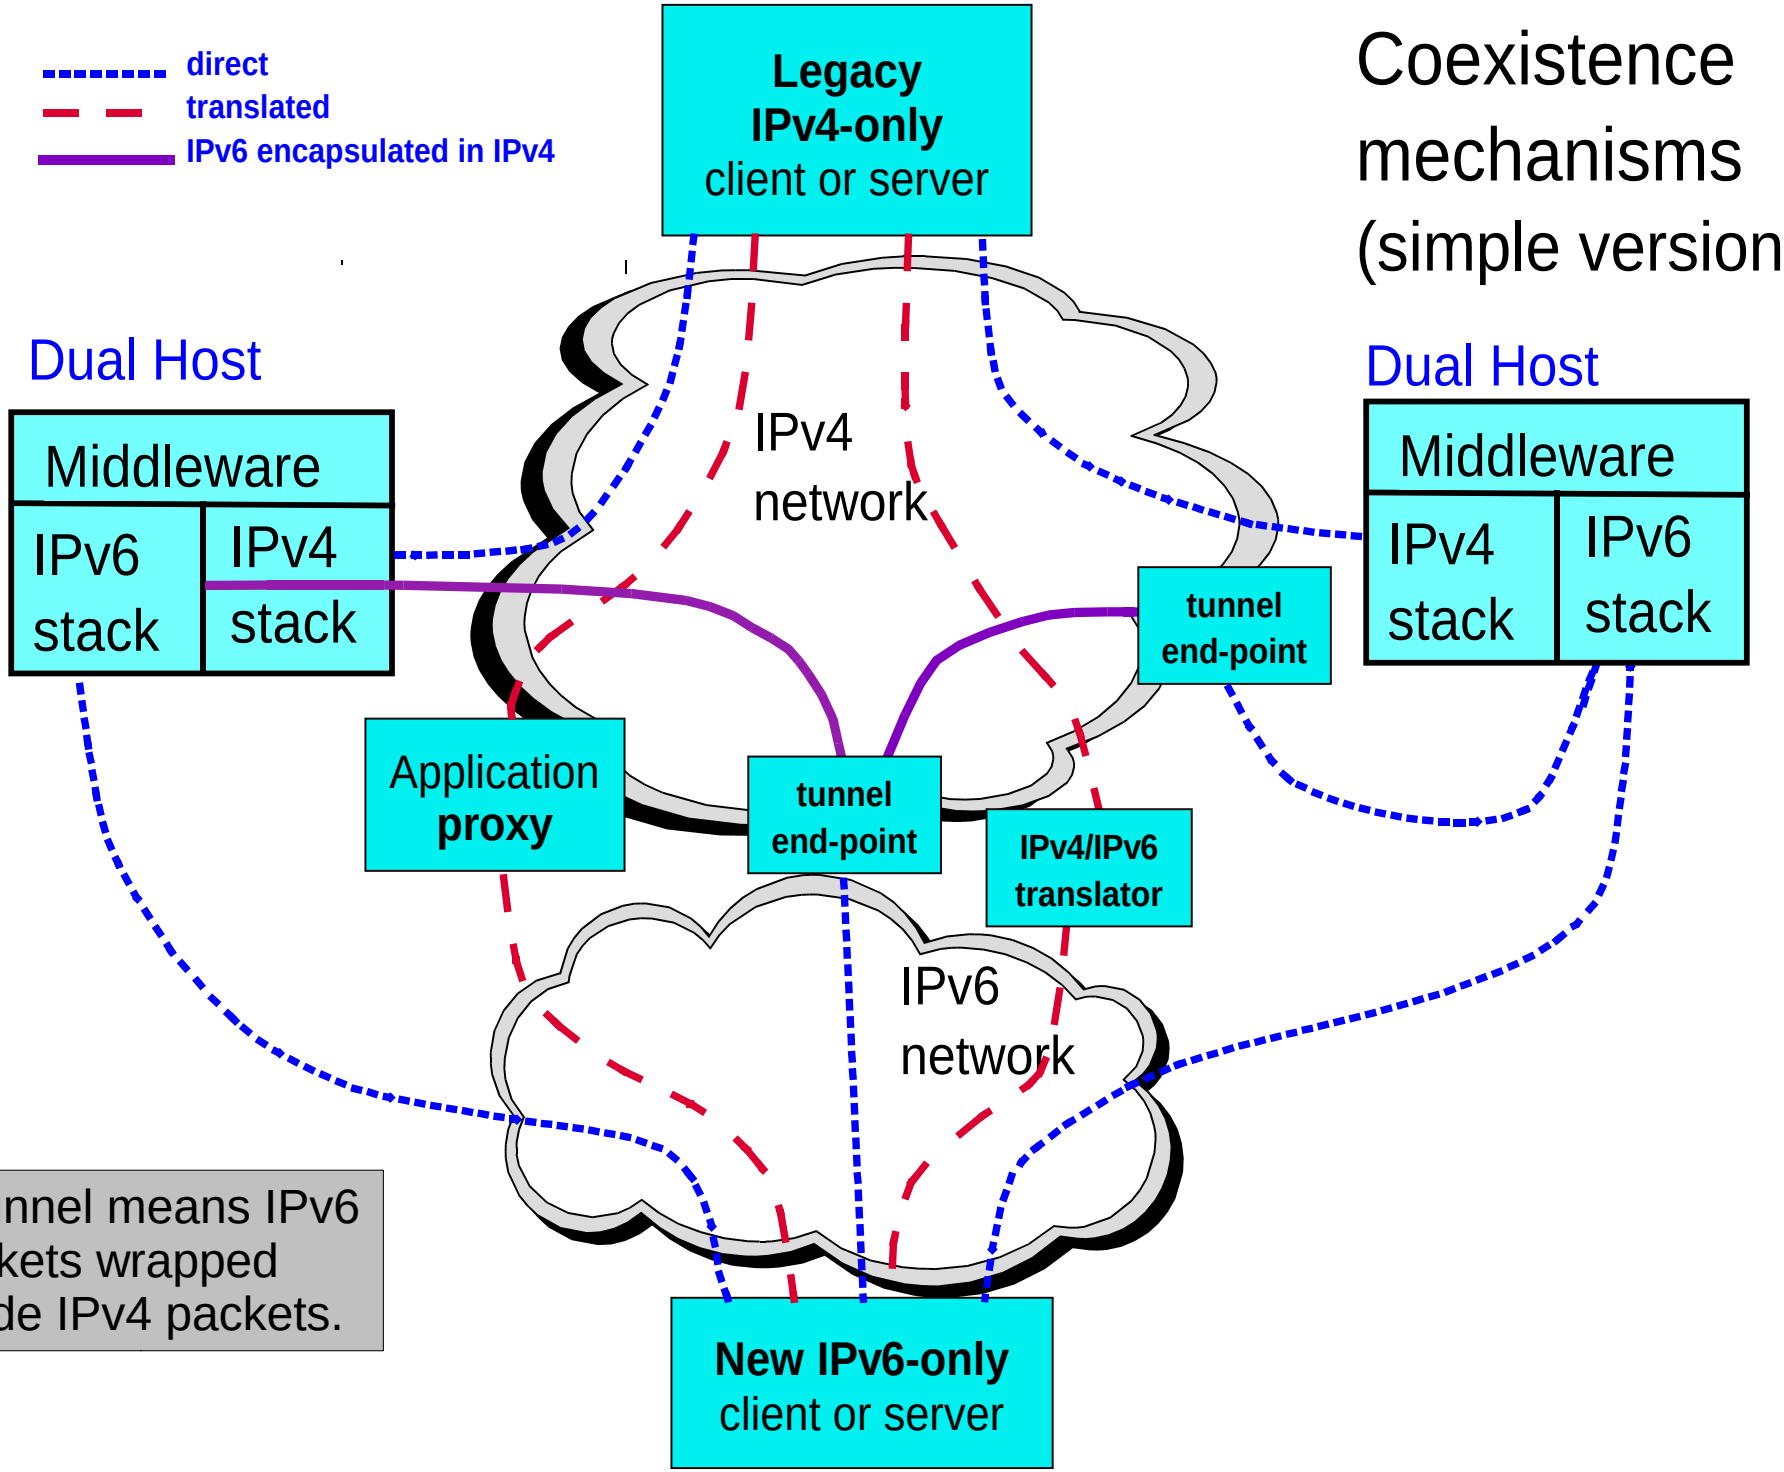

A tunnel means IPv6 packets wrapped inside IPv4 packets.

So if it's that simple, what's happening?

- About 1200 BGP4 entries (IPv4 has 280,000)
- No reliable traffic estimate, but indications are around 0.2% of IPv4 traffic (depending on where you look)

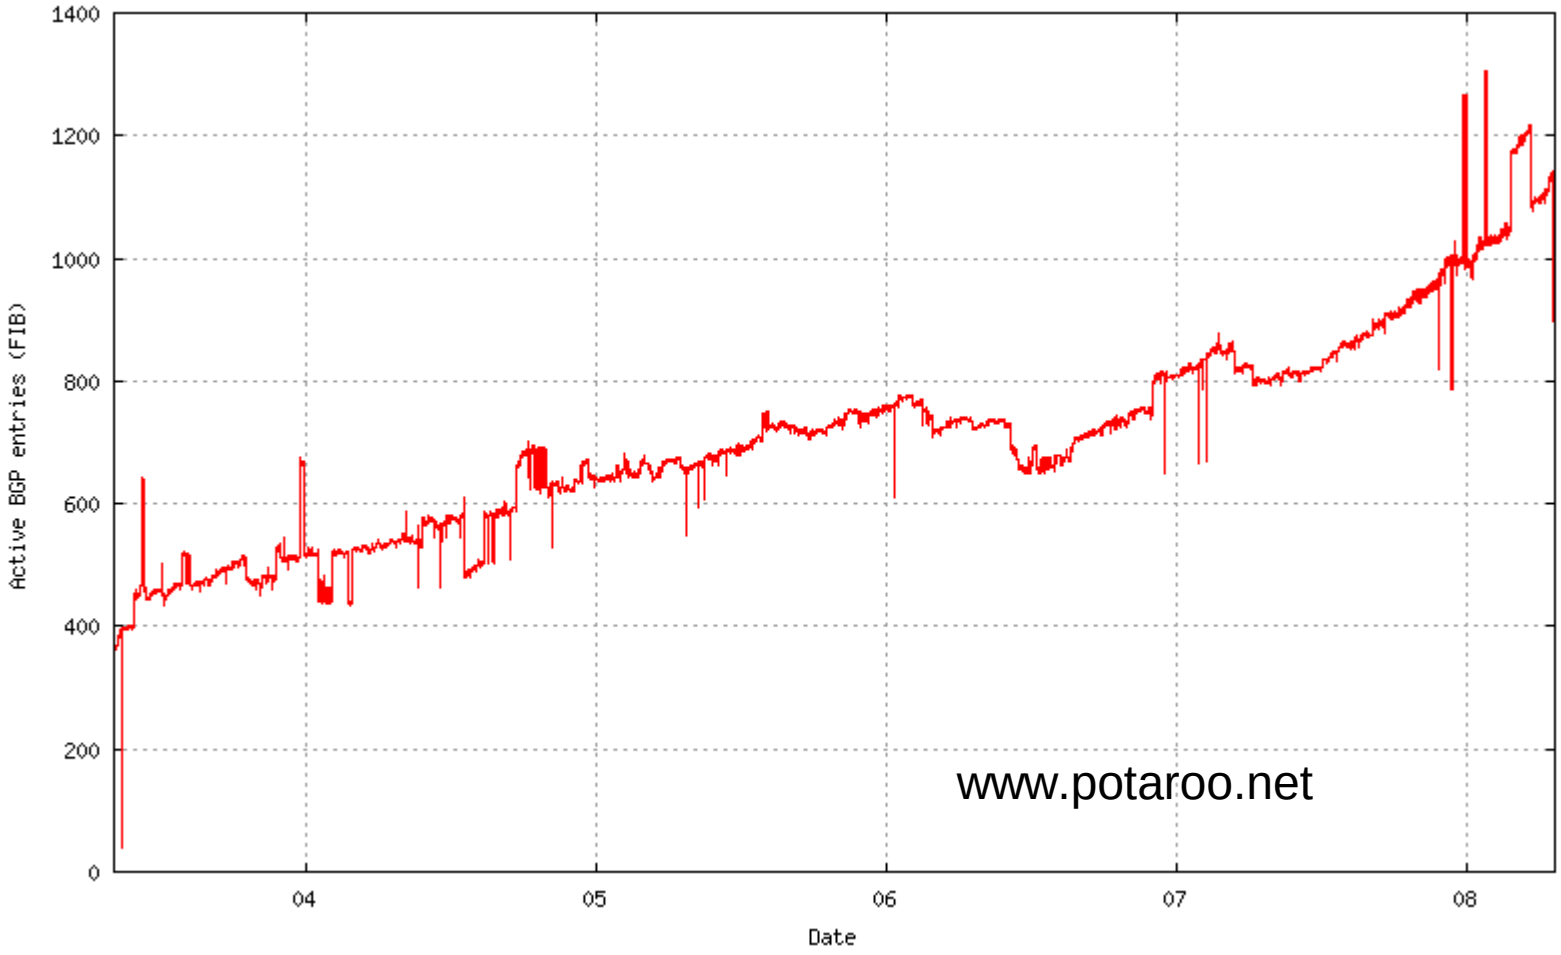

But, we have coexistence mechanisms coming out of our ears

- Dual stack (RFC 4213)
  - Socket API (RFC 3493)
  - DNS supports IPv4 and IPv6 (RFC 3596)
- IPv6 in IPv4 tunnels (RFC 4213)
- NAT-PT translation (RFC 2766)
  - IETF has deprecated this (RFC 4966)
- Tunnel Broker (RFC 3053)
- 6to4 implicit tunnels (RFC 3056)

… coexistence mechanisms coming out of our ears (2)

- Less favoured in IETF
  - Bump in the Stack (RFC 2767)
  - Bump in the API (RFC 3338)
  - SOCKS (RFC 3089)
  - Transport relay (RFC 3142)
  - 6over4 using IPv4 multicast (RFC 2529)
  - ISATAP (RFC 5214)
  - Teredo (RFC 4380)
- Still in draft (expired)
  - DSTM

So who's to blame?

- Not Microsoft (since XP SP1)?
- Not IBM?
- Not Google?
- Not the government?
- Not Cisco?
- Not the IETF?
- Not the ISPs?

➔ Then who?

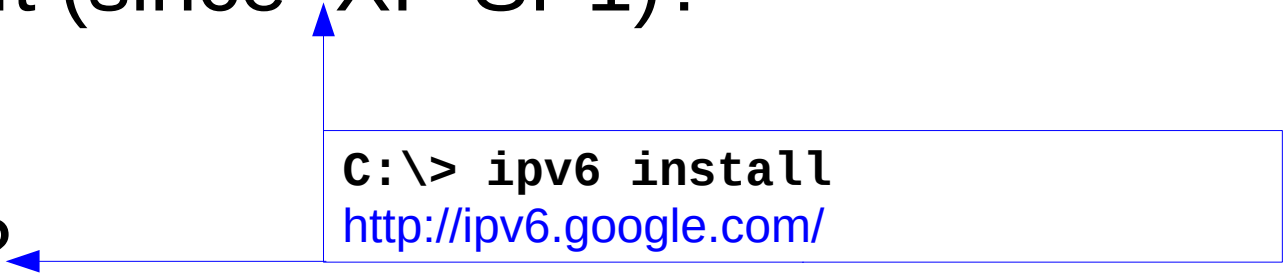

Missing bits and pieces

- Above all: compelling economic incentives
- Up to now, address sharing via NAT has appeared to be the low cost alternative to deploying IPv6
  - The Internet has come to tolerate the mess created by NAT, and has closed its eyes to the hidden cost
  - This will probably change by 2010, when the IPv4 address shortage will really inhibit business growth

Operator view of what's missing

- Connectivity things
  - How does a user at a v6-only site get to the [old] Internet, i.e. a v4-only site?
  - DNS registrars need to support delegation to IPv6 nameservers, and IPv6 glue records.
  - DOCSIS and 802.* must support IPv6 on media.

[This and the following 4 slides borrow heavily from Randy Bush]

Operator view of what's missing (2)

- Core ISP needs
  - Routers must support dual stack
  - Tools for Provisioning, Address Assignment, DHCPv6 and DNS Integration
  - Monitoring & Measurement over v6?
  - New line cards are often required!

Operator view of what's missing (3)

- Subscriber support
  - Authentication and session setup, e.g. PPPoE, IPoE, DHCP
  - Provisioning, back-end database, ...
  - "How to scale the routing/provisioning combo to deal with million of customers using stable prefix delegation?"

Operator view of what's missing (4)

- Consumer equipment
  - $50 DSL Modems do not support v6
  - $50 Firewalls do not support v6
  - Teredo does not really scale [and 6to4 cannot traverse a NAT]

Operator view of what's missing (5)

- Firewalls
  - Less than 1/3 had IPv6 Transport
  - 25% supported IPv6 Routing
    http://www.arin.net/meetings/minutes/ARIN_XX/PDF/thursday/Firewalls_Piscitello.pdf
- Enterprise applications
  - Open source and Java code not too big a problem, but proprietary applications present a very spotty picture.

Missing technical solutions

- The above are essentially product development and deployment issues - where we can hope that economic incentives will one day apply.
- We don't have a good solution for IPv6 multihoming
- We don't have a good solution for IPv4-IPv6 translation at the packet level

Multihoming

- Today's solution is for a multihomed site to have a provider-independent site prefix that is announced via multiple ISPs.
  - Because of the way route aggregation works in binary addressing, this simply doesn't scale as the number of multihomed sites increases.
  - Table size will go like *N* instead of *log(N)* or *sqrt(N)*.
- After years of concern, we only know two approaches
  1. Ignore the routing system; solve the problem end to end between hosts (using multiple addresses per host).
  2. Split addressing into two layers: a locator used for routing and traffic engineering, and an identifier used between the hosts.

Host-based multihoming: SHIM6

- Inserts shim code at the top of the IPv6 stack
  - remote host has several IPv6 addresses (one locator per ISP)
  - one of them is used as Upper Layer ID (i.e. the address used in socket calls, TCP checksums, IPsec, etc.)
  - the shim switches dynamically between the locators (i.e. the addresses used in the packet headers)
  - zero visibility at routing level; only host software is touched
  - host sites must operate one prefix per ISP
  - a bit more complicated than it sounds, due to reachability and security issues
- Takes traffic engineering out of the hands of ISPs
  - ISPs would like control of path selection, currently implemented by BGP4 policy

Routing-based multihoming: research

- Basic idea is not new: split apart the functions of an address*
  - identifier is used end-to-end (e.g. TCP checksum)
  - locator is used for routing site-to-site (and for traffic engineering)
- Not clear how to make this change successfully on a running Internet, even with only 0.2% of IPv6
  - cut the IPv6 address in two halves (64 bit locator and 64 bit identifier)?
  - encapsulate normal IP packets (with identifier-addresses) in tunnels (with locator-addresses)?
- Ongoing work in the IRTF Routing Research Group

*can arguably be traced as far back as a paper by Louis Pouzin in 1974

Packet-level translation

- NAT-PT (network address translation - protocol translation) was designed years ago (RFC 2766)
  - However, it suffers from all the problems of regular NAT *plus* some serious side-effects of DNS translation
  - The IETF has deprecated it (even though it works in some carefully managed scenarios) (RFC 4966)
- Two ways forward can be considered:
  1. An improved form of NAT-PT
     e.g. draft-van-beijnum-v6ops-mnat-pt
  2. On an IPv6-only network, use a dual stack and a tunnel to reach the IPv4 world
     e.g. draft-despres-v6ops-apbp

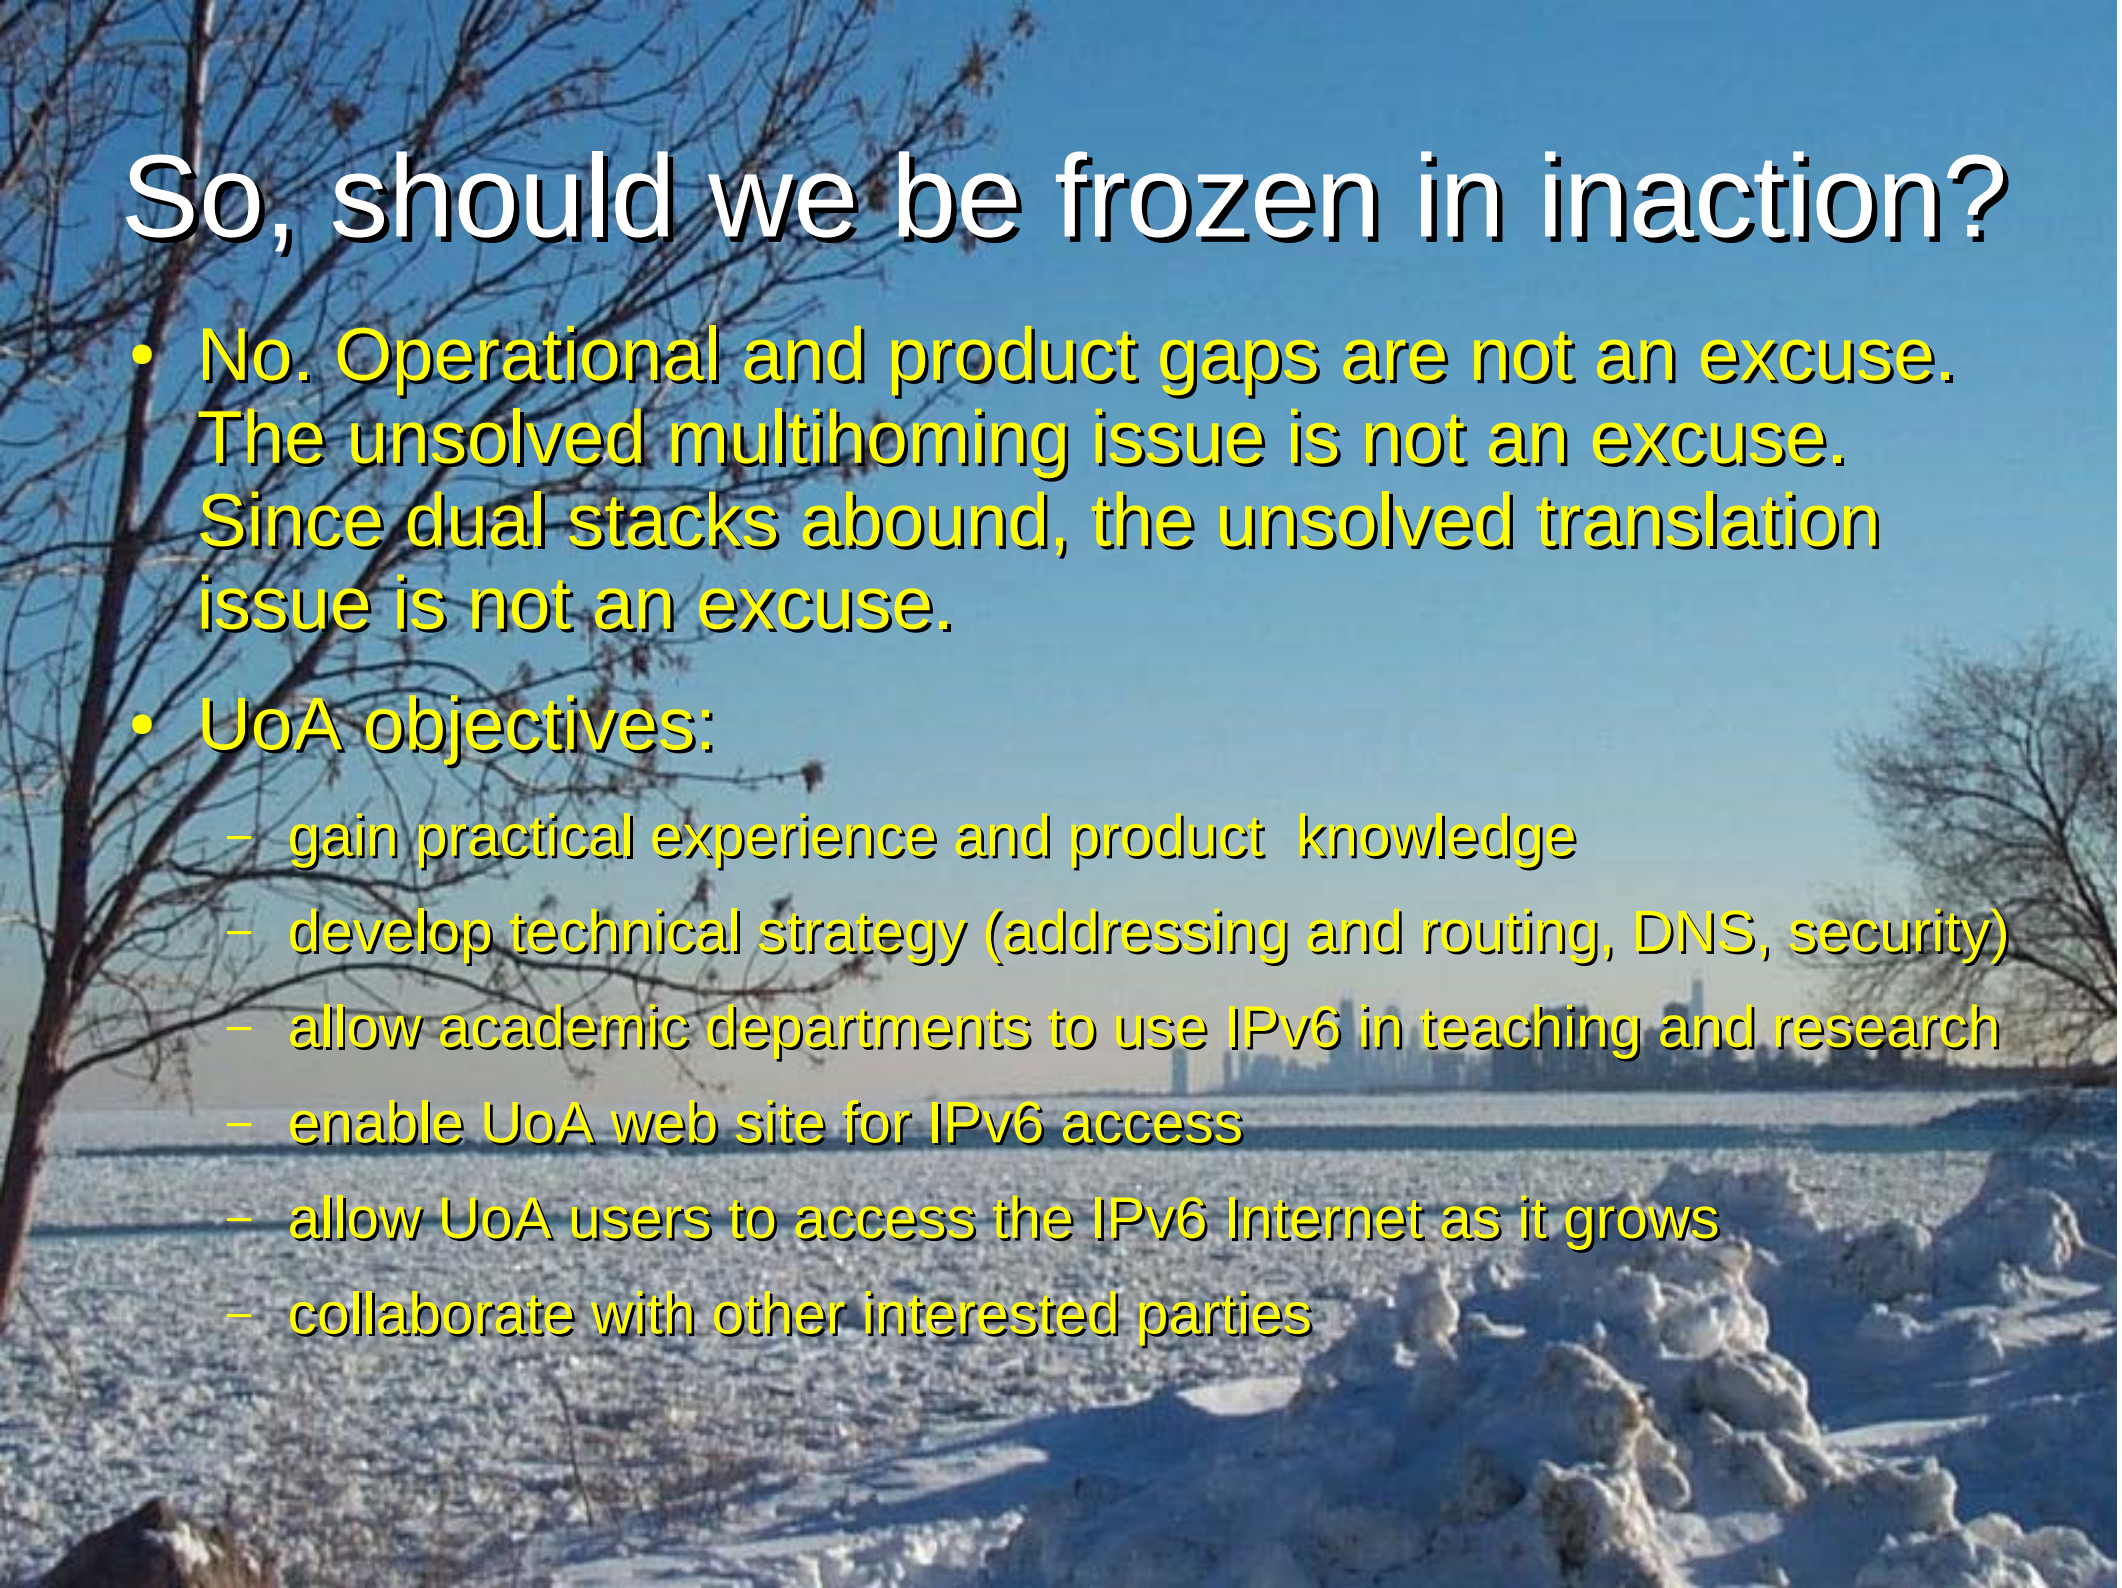

So, should we be frozen in inaction?

- No. Operational and product gaps are not an excuse. The unsolved multihoming issue is not an excuse. Since dual stacks abound, the unsolved translation issue is not an excuse.
- UoA objectives:
  - gain practical experience and product knowledge
  - develop technical strategy (addressing and routing, DNS, security)
  - allow academic departments to use IPv6 in teaching and research
  - enable UoA web site for IPv6 access
  - allow UoA users to access the IPv6 Internet as it grows
  - collaborate with other interested parties

UoA action plan (tentative)

- Educate IT staff.
- Verify transit connectivity and advertisement of address prefixes with ISPs
- Build list of products and applications, and check IPv6 support plan for each of them.
- Create IPv6 testbed subnet.
- Verify support of IPv6 firewalling.
- Enable DHCPv6, OSPFv3 and BGP4+ centrally.
- Enable AAAA records and dual stack access to DNS.
- Set up initial management and measurement for IPv6.
- Enable gateway for Teredo/6to4 access in collab with InternetNZ Tui project
- Route or tunnel testbed through to central facilities.
- Use testbed as a teaching network, or create a clone as teaching network.
- If applicable, migrate testbed from tunnelled to native IPv6 (i.e. OSPFv3 to testbed).
- Create test IPv6 UoA web server - progressively duplicate the real site there.
- Dual stack the real web site.

Sources

- http://www.potaroo.net/ispcol/2008-04/ipv6.html (lots of good stuff at potaroo.net)
- http://cenic08.cenic.org/program/slides/080312.cenic-v6-op-reality.pdf
- http://conference.nznog.org/presentations/20080124_04-6to4-teredo-tui_nathan-ward.pdf
- http://www.civil-tongue.net/6and4/
- http://www.irtf.org/charter?gtype=rg&group=rrg
- http://www.isoc.org/educpillar/resources/ipv6_faq.shtml
- http://penrose.uk6x.com/