IPv6 Keynote @SINOG4

ot@cisco.com

Per-Country IPv6 adoption

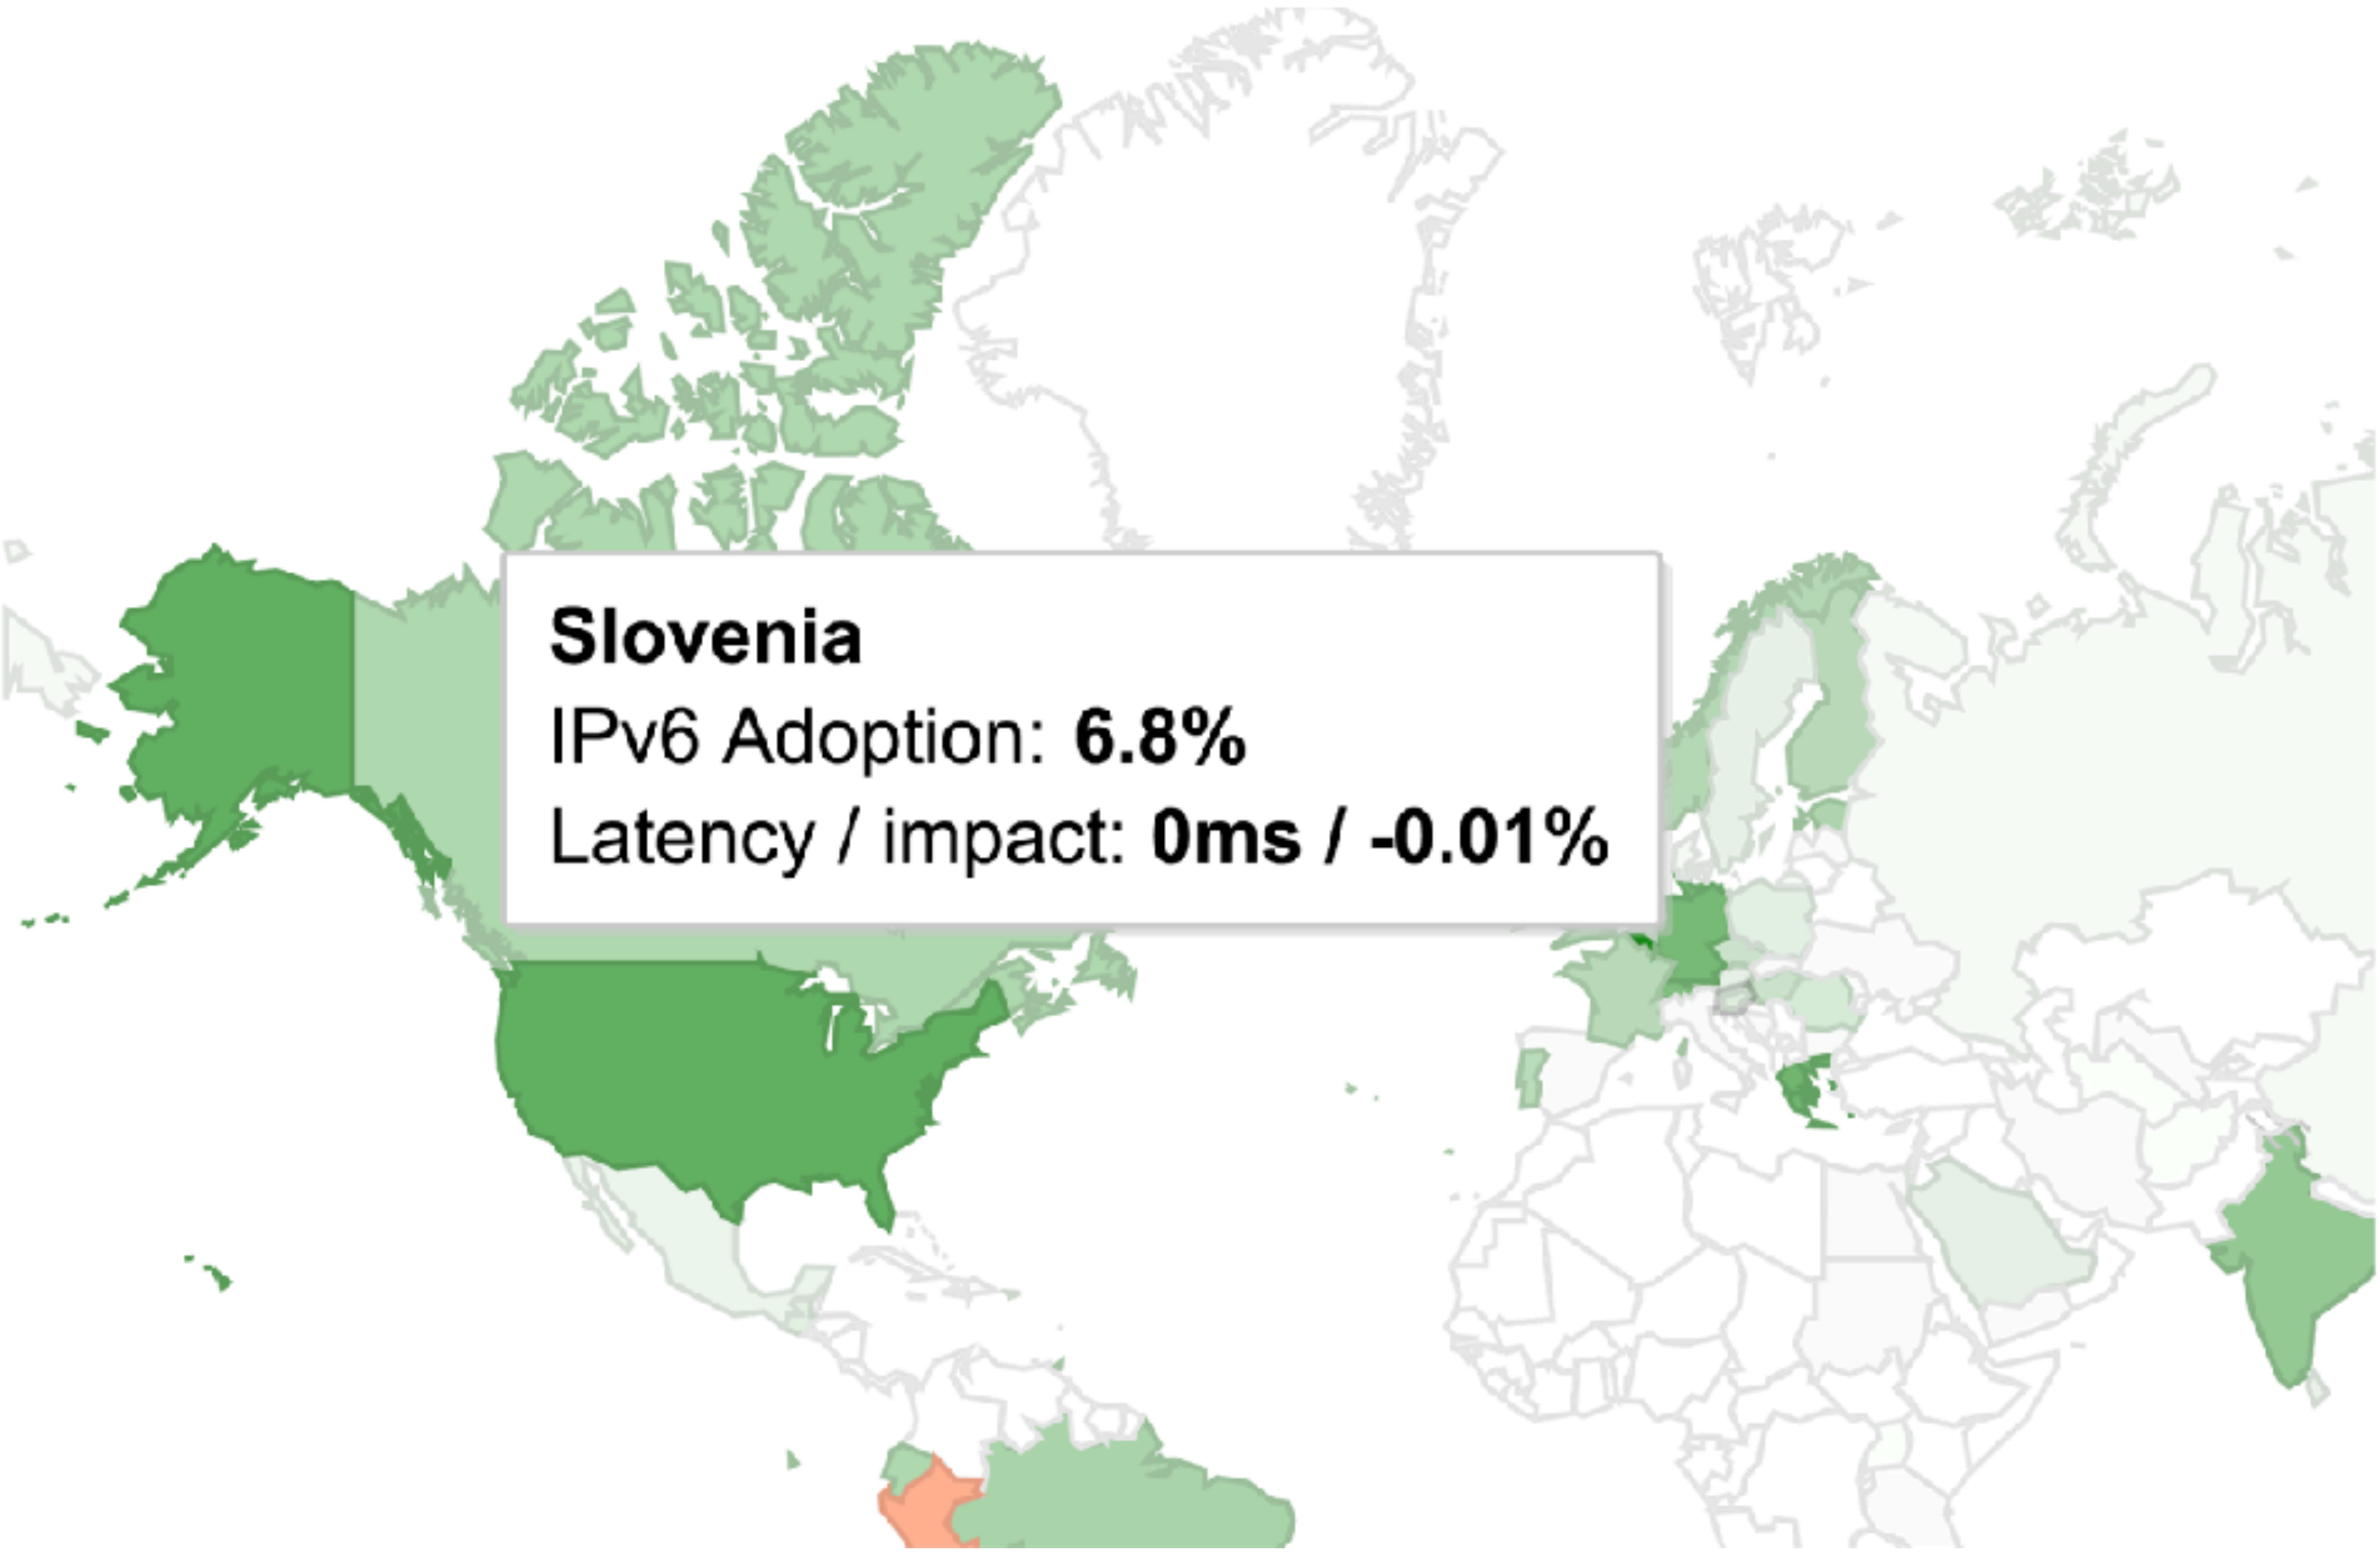

Use of IPv6 for Slovenia (SI)

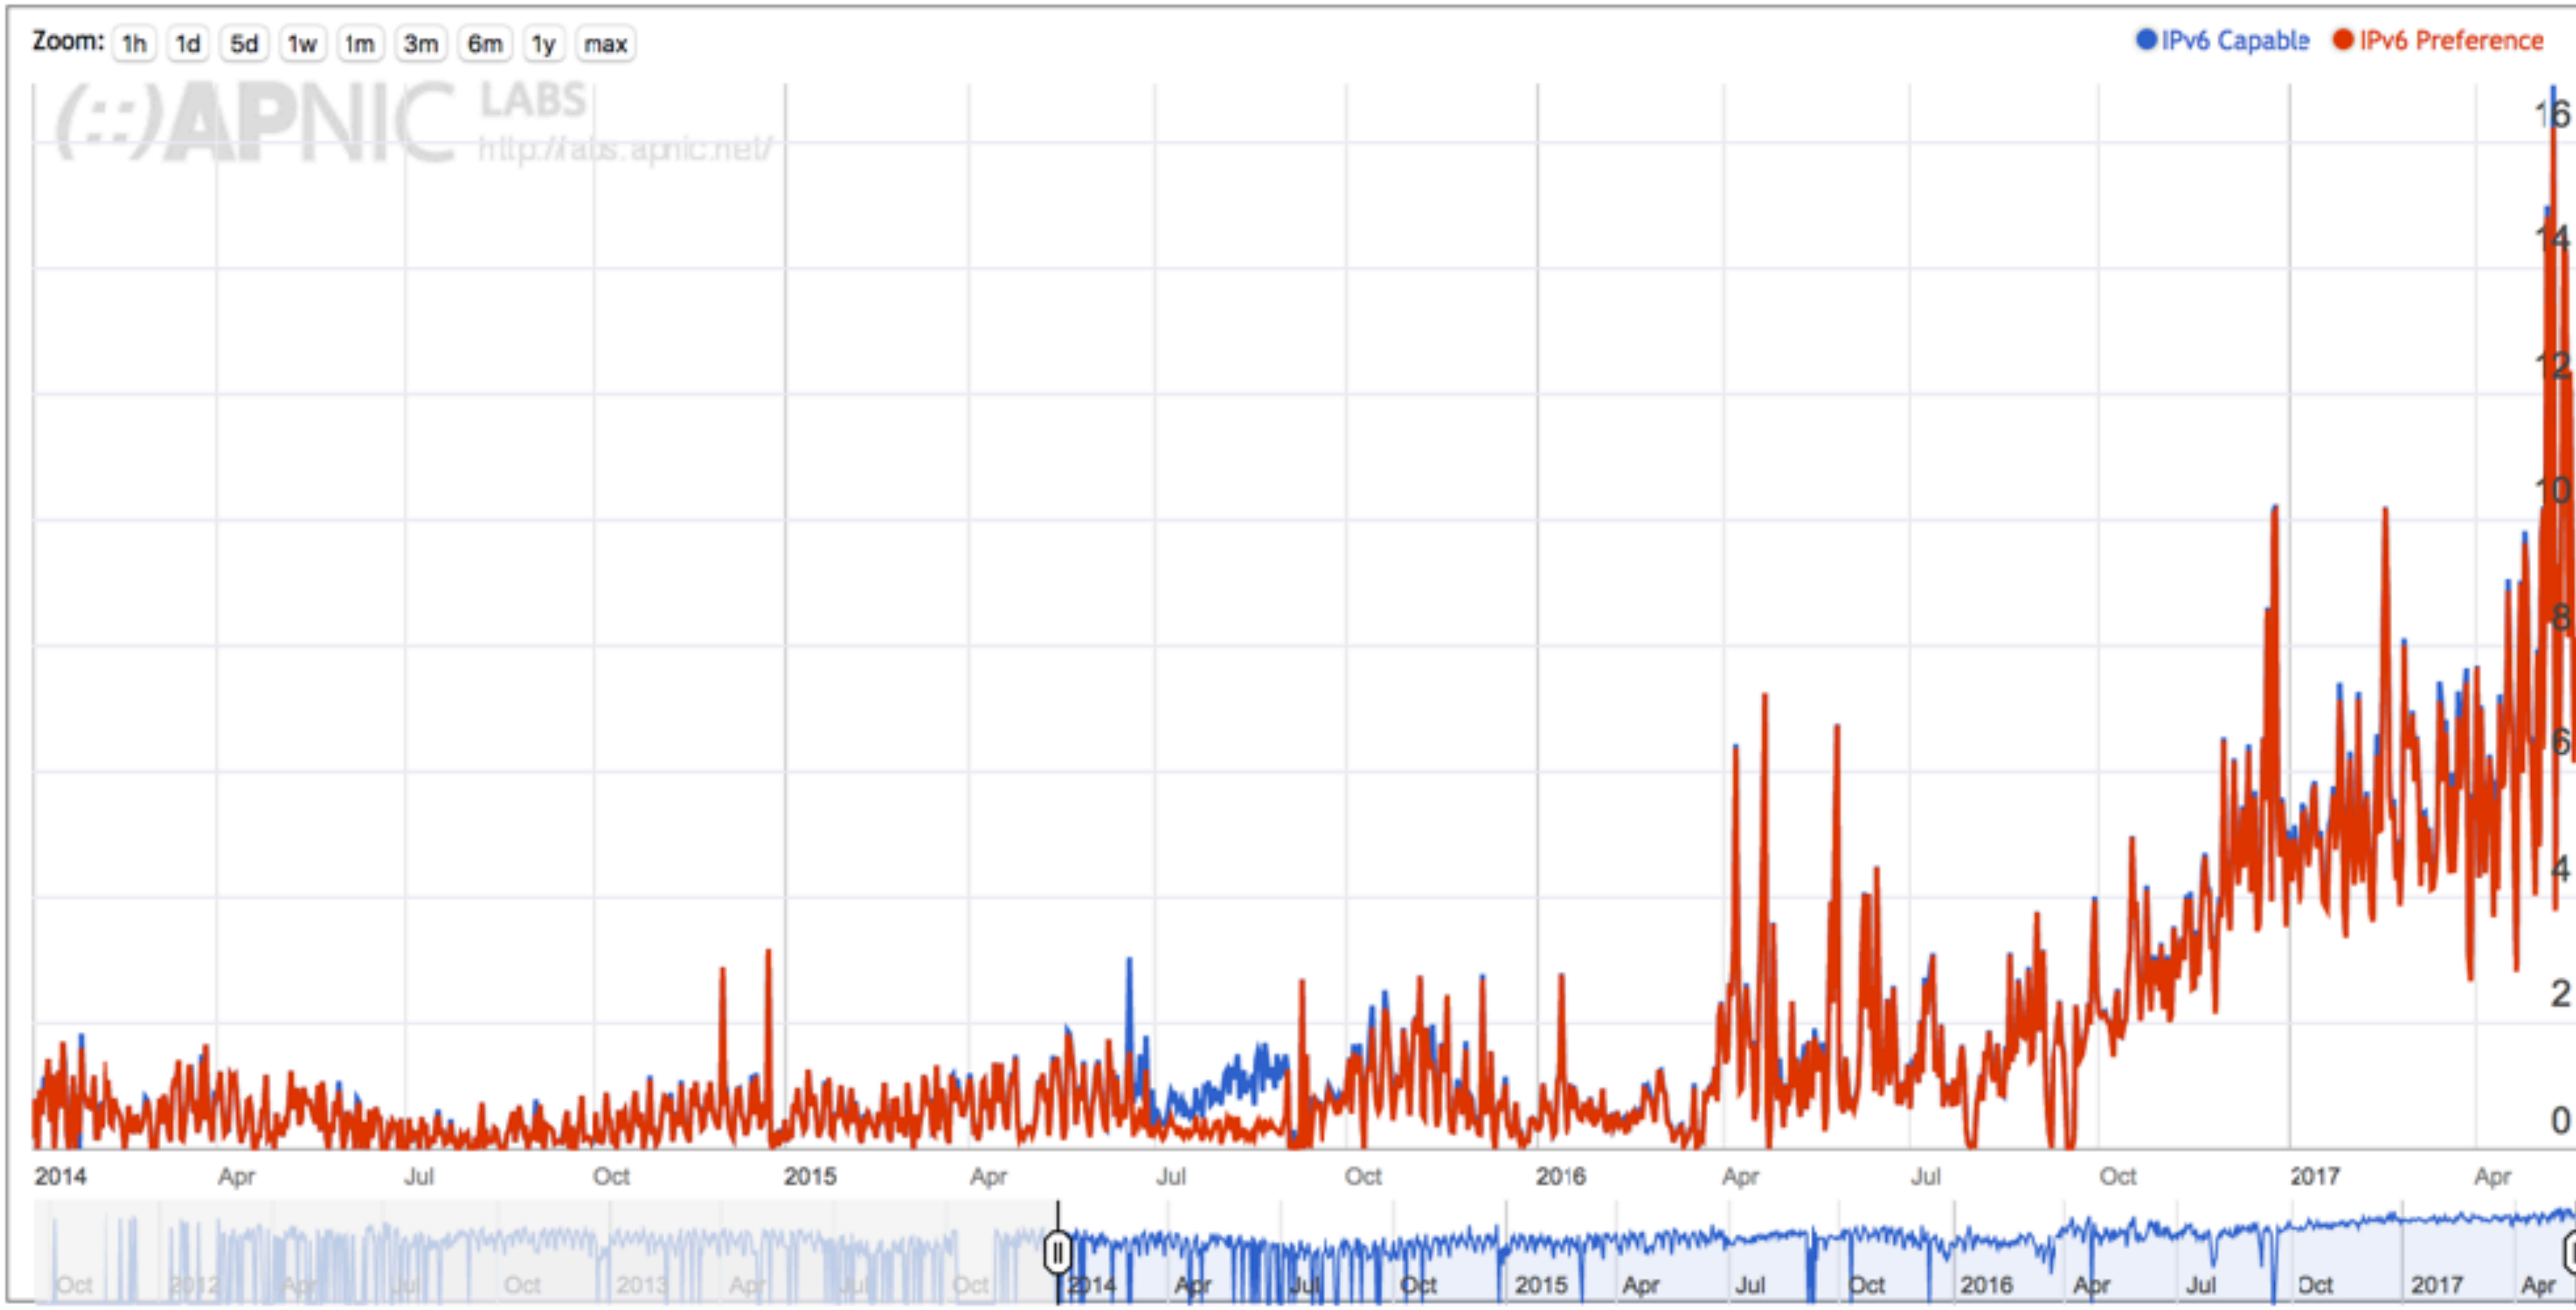

The IESG has approved the following document:
- 'Internet Protocol, Version 6
  (IPv6) Specification'
  (draft-ietf-6man-rfc2460bis-13.txt)
as Internet Standard

This document is the product of the IPv6 Maintenance Working Group.

The IESG contact persons are Suresh Krishnan and Terry Manderson.

A URL of this Internet Draft is:
https://datatracker.ietf.org/doc/draft-ietf-6man-rfc2460bis/

This talk

- Why is IPv6 what it is…
- The tussle
- The compromises
- Where we go from here

Why IPv6 is what it is…

- The Problem: We're running out of IPv4 addresses
- A balance of changing too *much* or changing too *little*
- Internet architecture goals and principles
- The tussle

“Why isn’t IPv6 backwards compatible with IPv4?”
“Because IPv4 offered no forward compatibility”

–Steve Deering

Lost features of the Internet

- transparency
- robustness through “fate sharing”
- dynamic routing
- unique addresses
- stable addresses
- connectionless service
- always-on service
- peer-to-peer communication model
- application independence

“IP should be as ubiquitous as electricity”

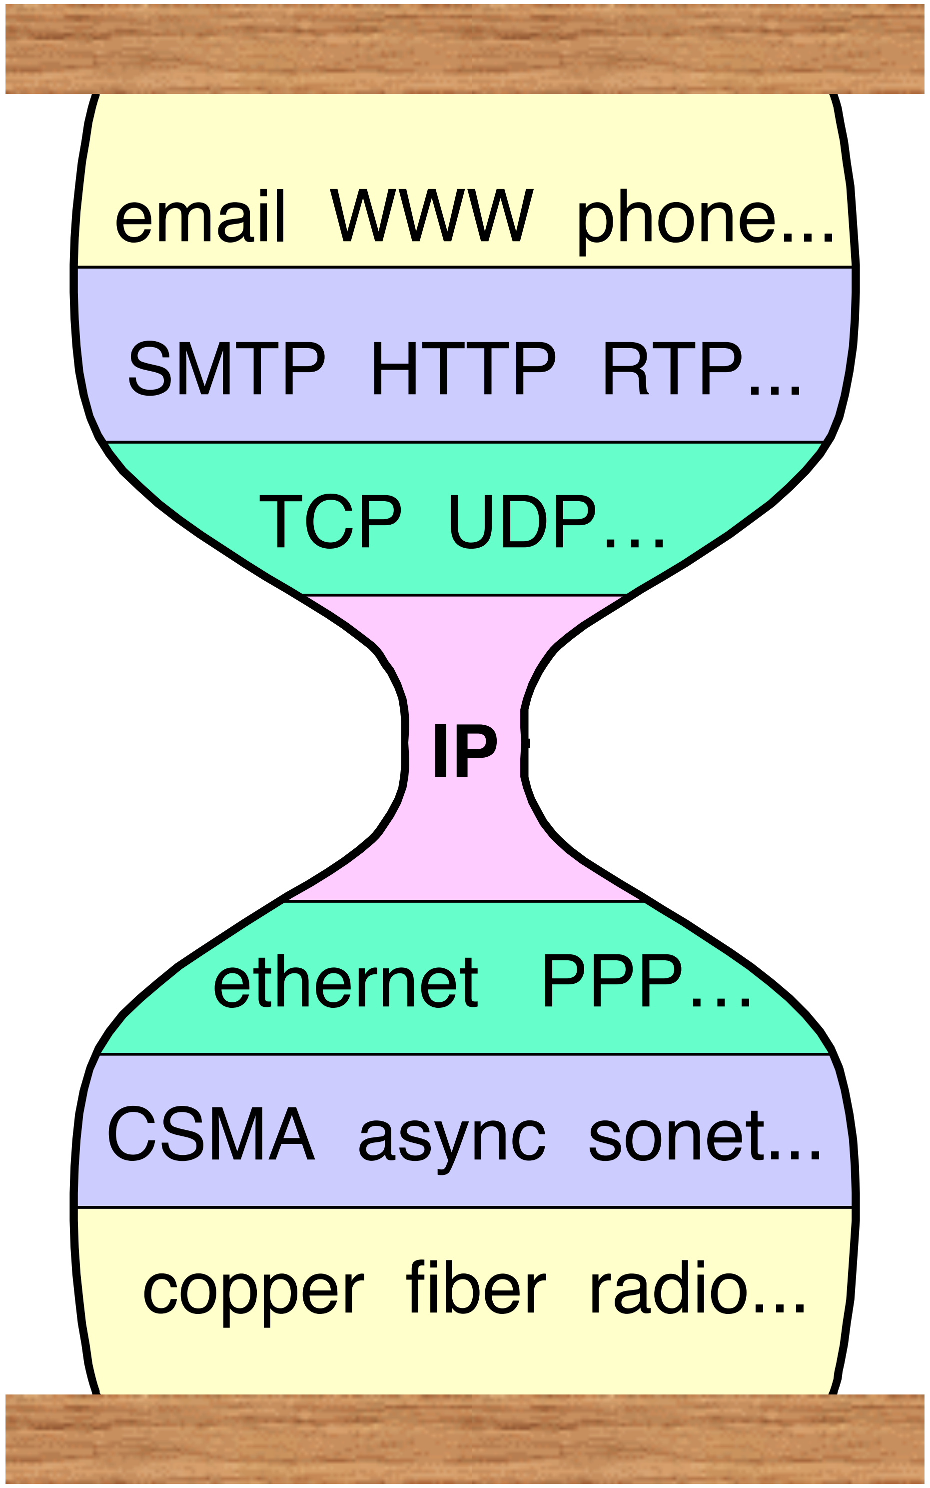

Steve Deering: Watching the waste of the protocol hourglass, IETF51

Putting on Weight

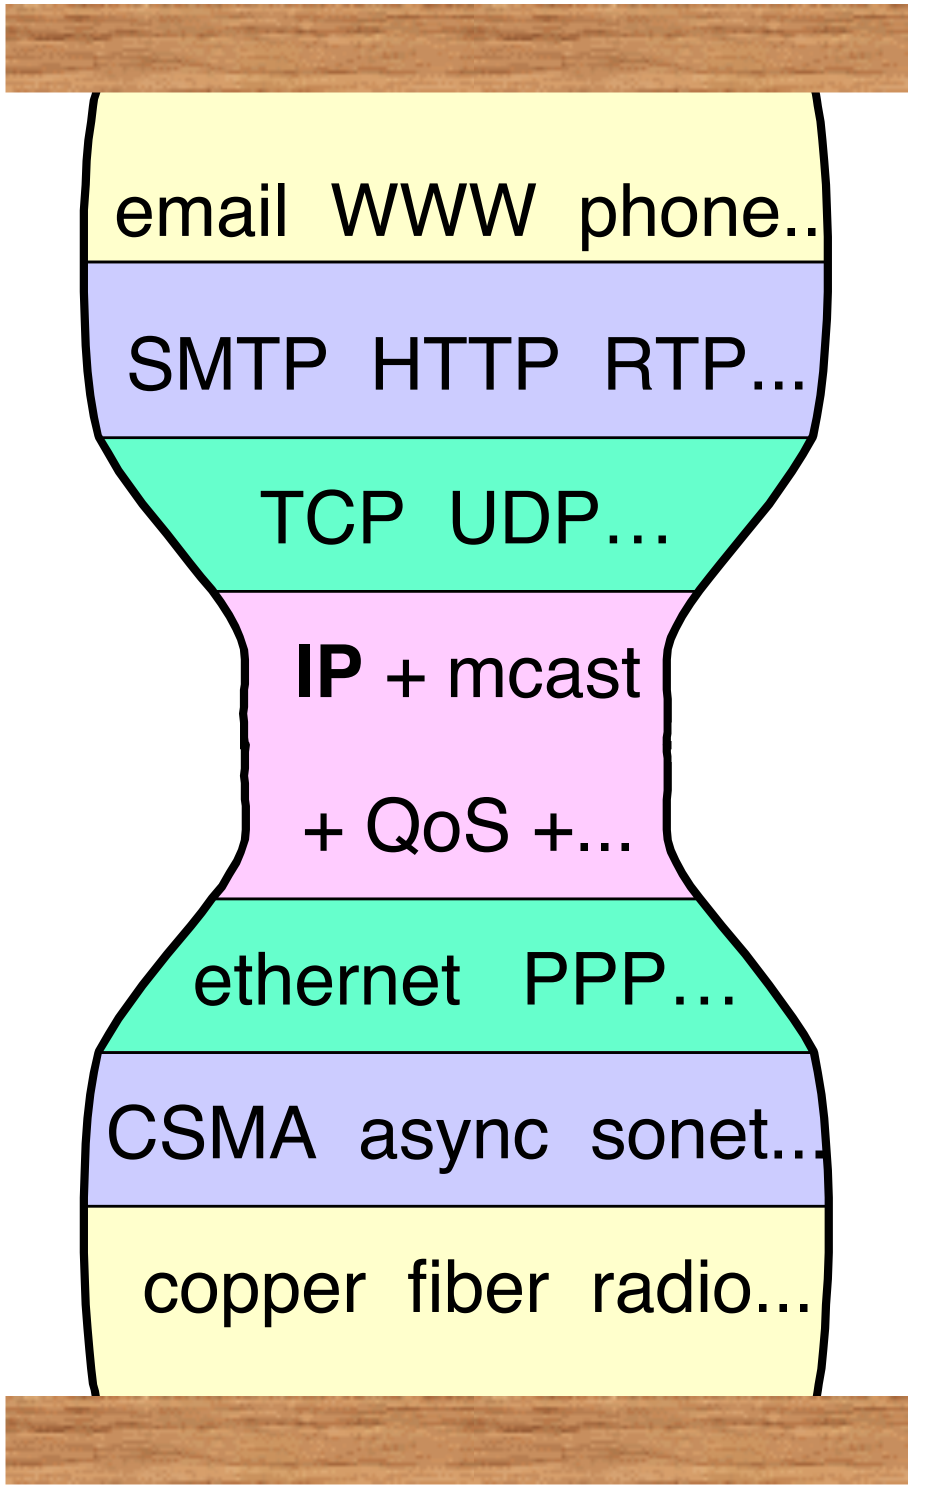

- requires more functionality from underlying networks

Mid-Life Crisis

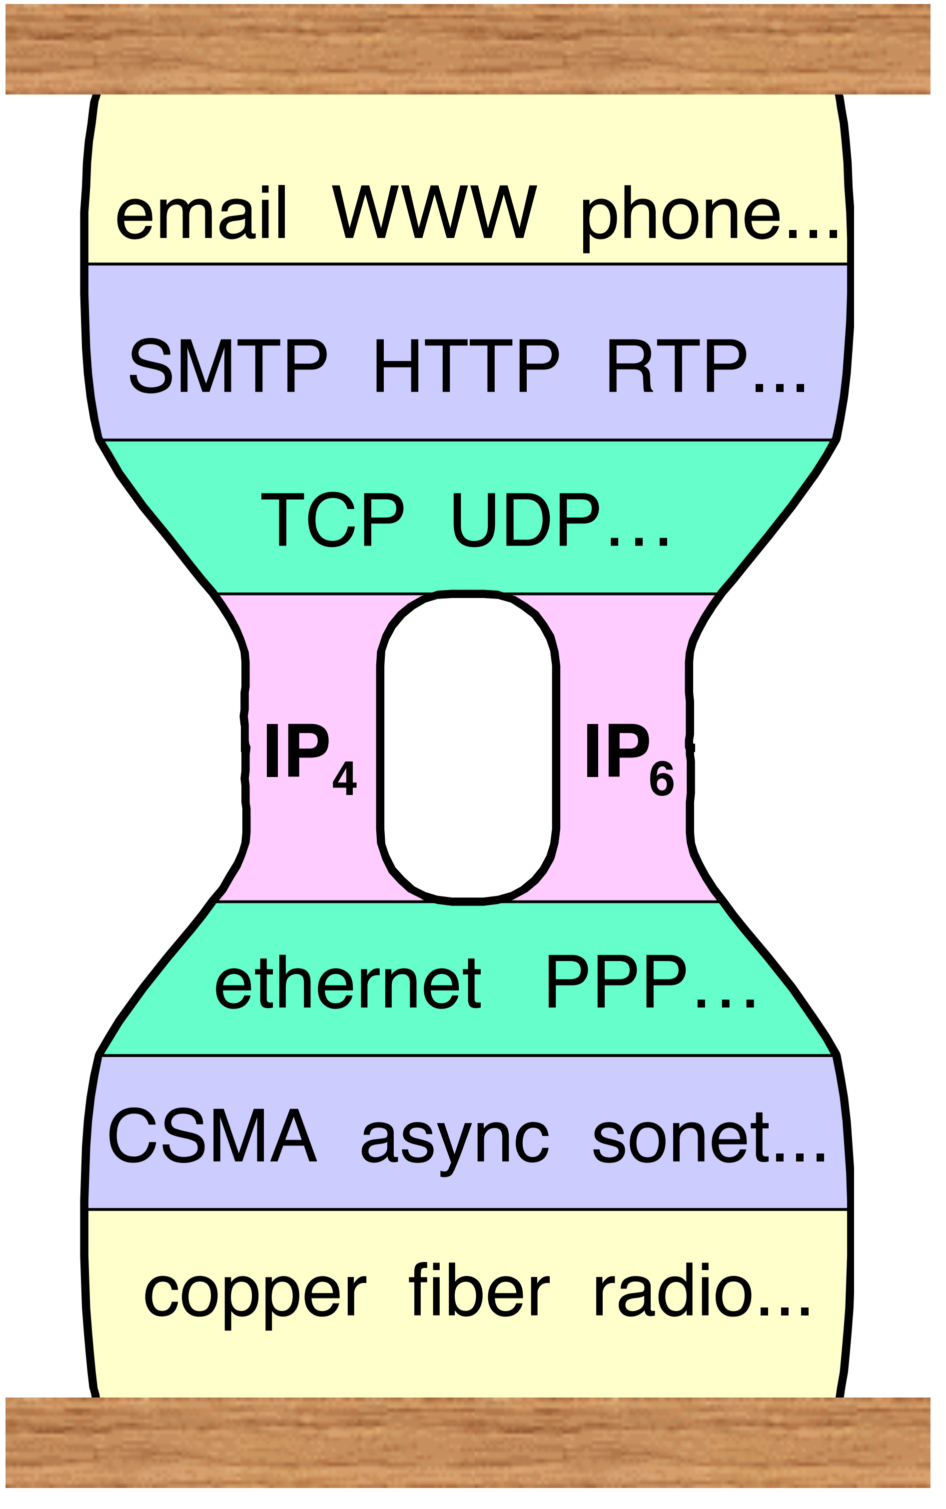

- doubles number of service interfaces
- requires changes above & below
- creates interoperability problems

Oops! An Accident

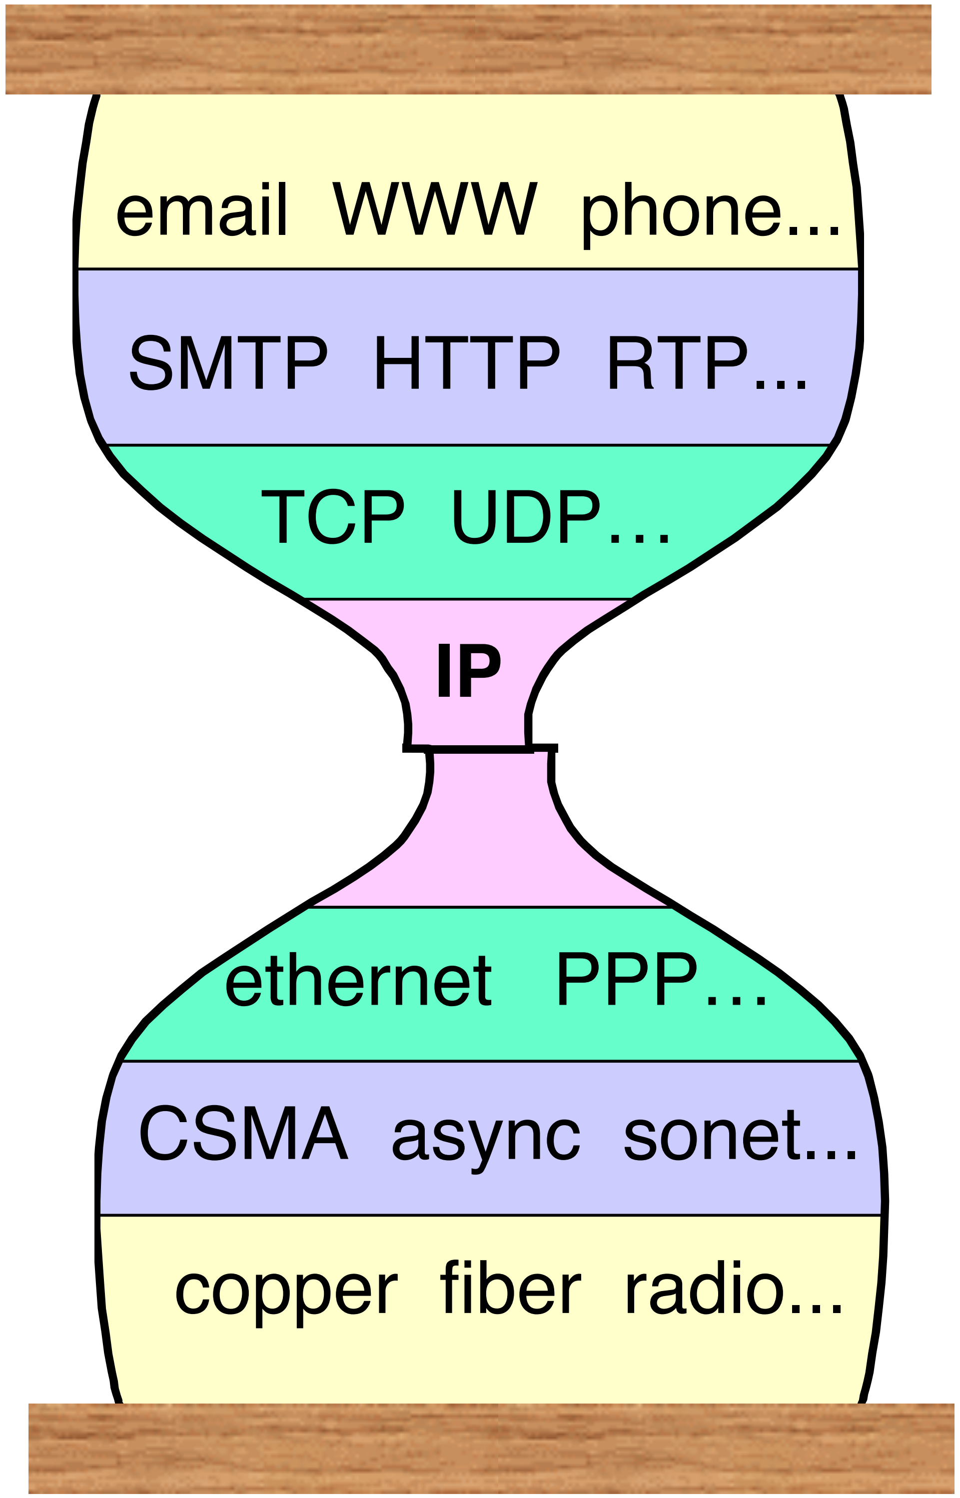

- NATs & ALGs used to glue the broken pieces
- lots of kinds of new glue being invented—ruins predictability
- some apps remain broken, since repairs are incomplete

More Fattening Temptations

- TCP “helpers”
- reliable multicast assists
- packet-intercepting caches
- “content-based routing”
- active networking

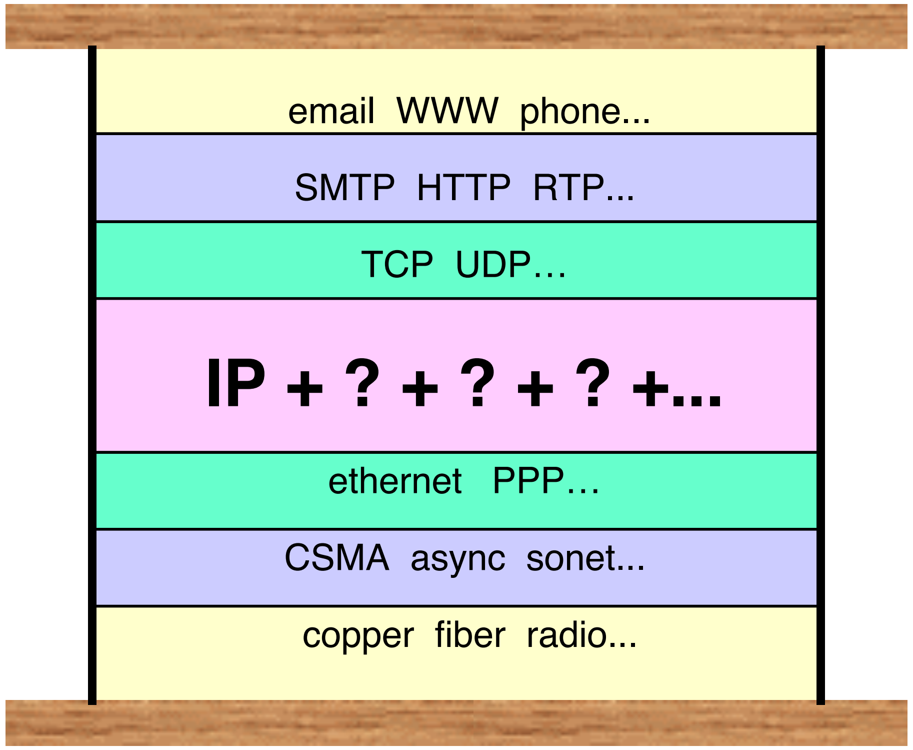

Threat-ened by Youths

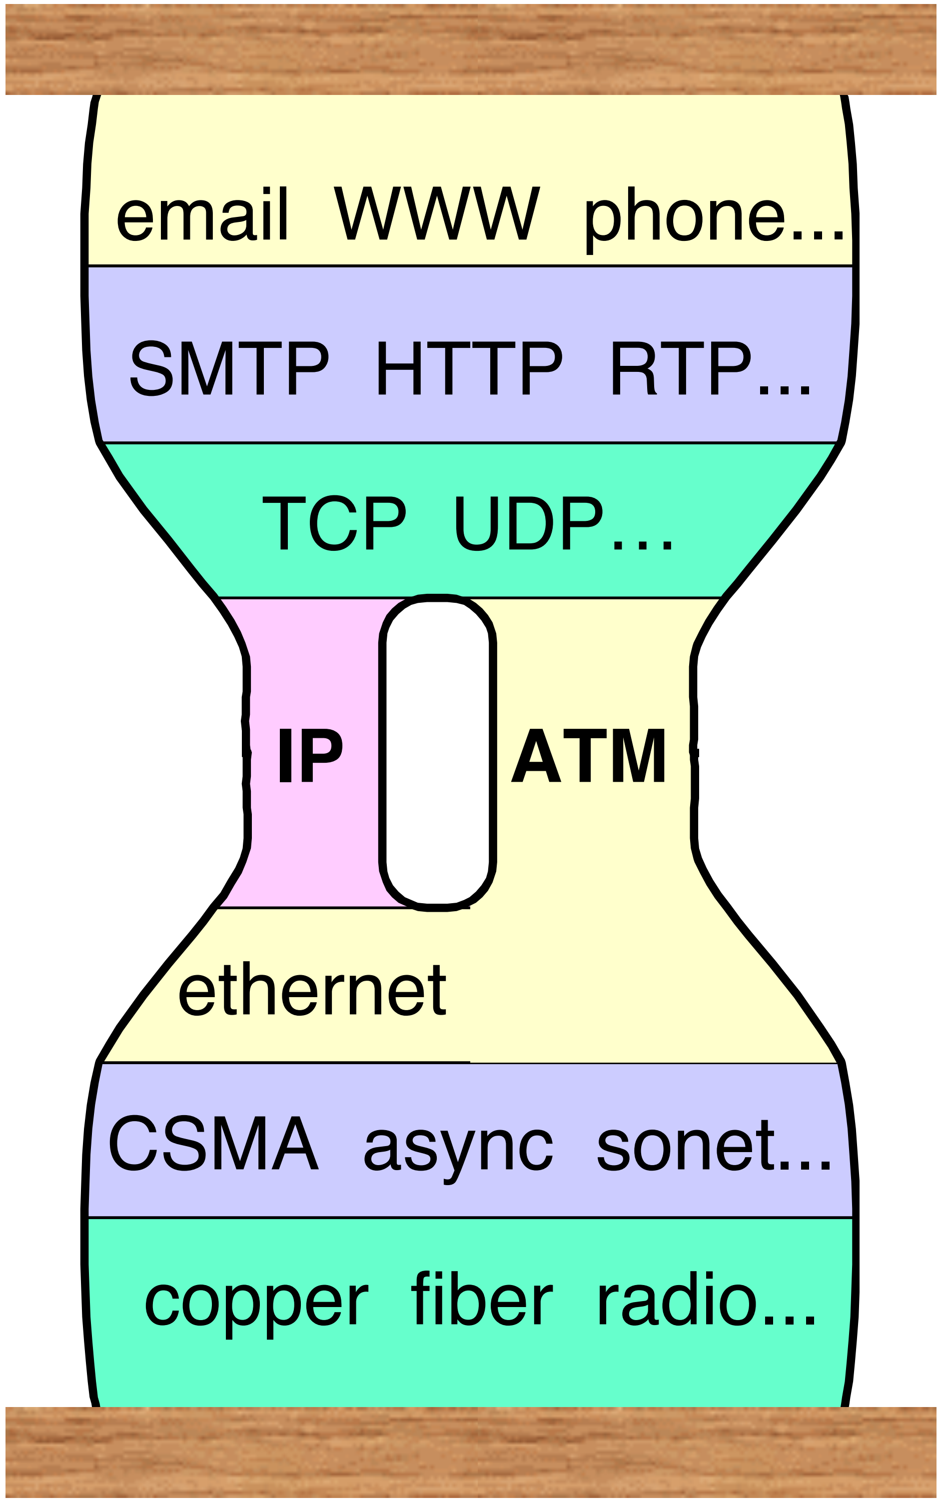

- danger : creeping dependencies on specific link-layers inhibit flexibility and evolution
- will never fully supplant IP, so end up with complicated hybrid & more address plans

But Still Supple

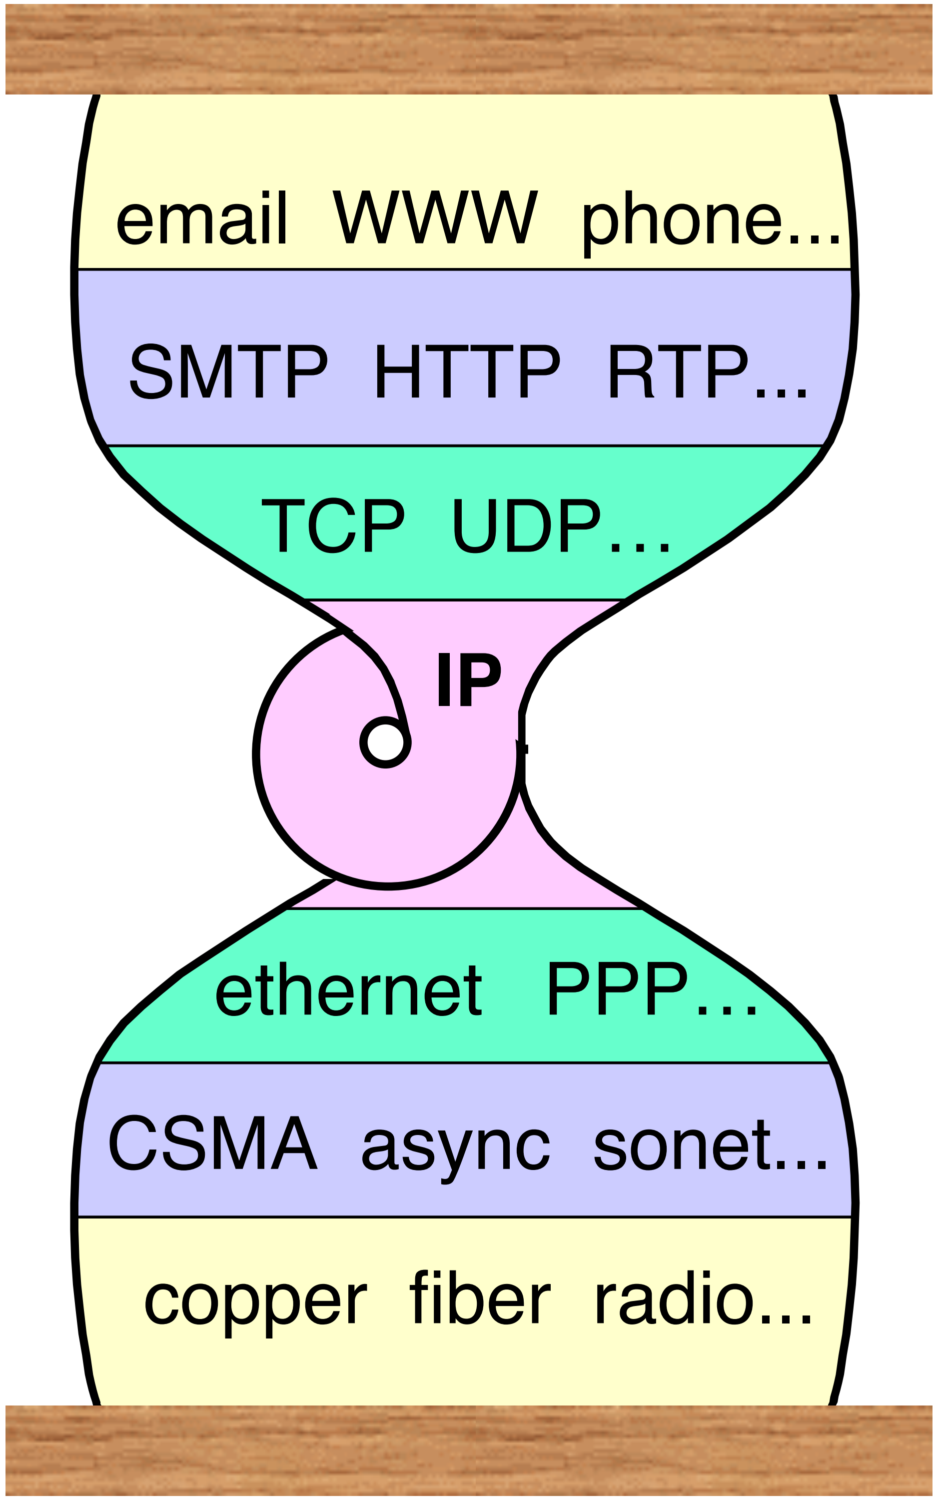

- IP-over-IP tunneling has become more and more common
- this is not so bad: retains benefits of hourglass model

A Fitness Goal

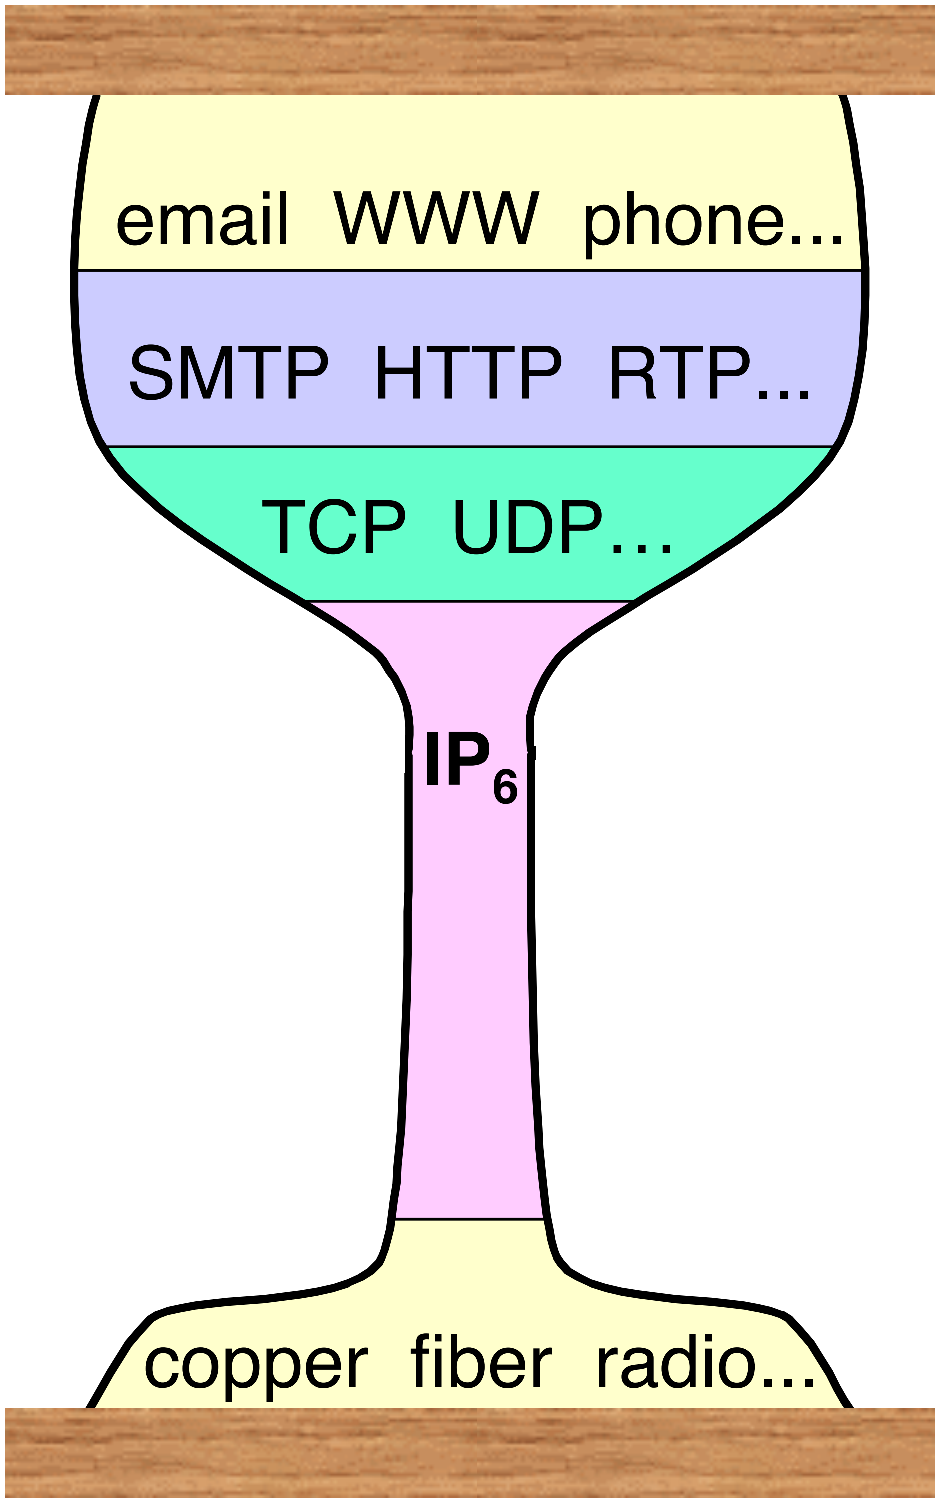

- perhaps we can trim down from an hourglass to a *wineglass*
- promising signs: IP-over-SONET, IP-over-WDM
- IPv6 to restore simplicity *and* functionality

IPv6 Solution

- No magic just 96 more bits - Simple evolution of IPv4 (SIP)
- 128 bit addressing
- Fixed size header (IPv4 has variable length). Optimized header (remove fragment information, checksum…)
- Replace IPv4 options with IPv6 Extension headers
- Generalise link-specific address resolution / host configuration into the network layer
  IPCP, ARP
- Limit changes to the network layer.
  No changes to transport protocols

Tussle in Cyberspace: Defining Tomorrow's Internet

David D. Clark
MIT Lab for Computer Science
ddc@lcs.mit.edu

John Wroclawski
MIT Lab for Computer Science
jtw@lcs.mit.edu

Karen R. Sollins
MIT Lab for Computer Science
sollins@lcs.mit.edu

Robert Braden
USC Information Sciences Institute
braden@isi.edu

- There are many players involved in the Internet with interests directly at odds with each other
- The technical architecture must accommodate societies tussle. While continuing to solve the traditional goals (i.e. solve problems)

The players

- Users
- Internet Service Providers
- Content and services providers
- Governments
- Intellectual property holders
- …

Dealing with the tussle

- Modularise the design around tussle boundaries
- Flexible design to permit different players to express their differences
- Tilting the playing field

Tussle spaces

- Economics
    - Consumers tussle with providers to the services they want at a low price
        - Lock in with IP addresses
        - Residential broadband access
- Trust
    - Open end to end communication in a low-trust environment
    - Users don’t trust the parties they want to talk with either
    - Nor trust the software they have to run
    - Content providers want to monetise information about the user, while the user wants privacy
- Openness
    - ISPs dislike and fear openness
    - Openness to innovation

Protocol Politics

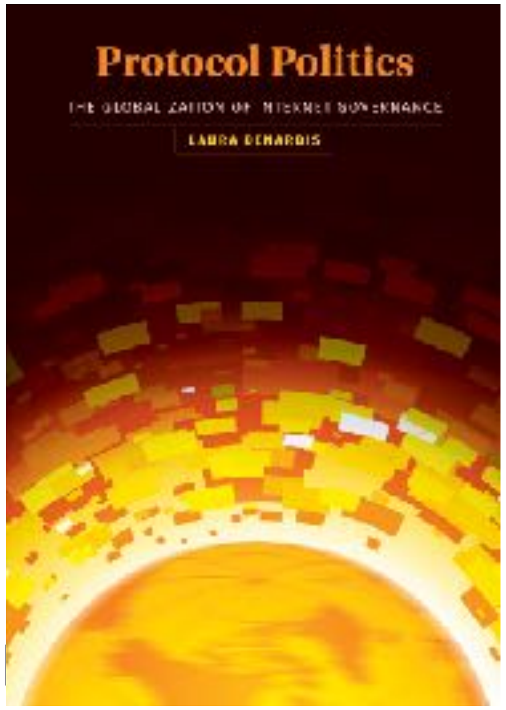

- Internet protocols are not value neutral
- Separation of policy and mechanism
  - Isolate parts of the system against the tussle
- End to end argument
  - State that a mechanism should not be placed in the network if it can be placed in the end node
- Cost and benefit must be aligned
- New protocols aren't deployed if they don't offer opportunity for competition
- **Keeping the net open and transparent for new applications is the most important goal**
- Peeking is irresistible

Compromises

- Addressing
- Extension headers
- Host configuration (DHCP, ND)
- Minimise changes to network layer

Tussle #1 - Addressing

- 64 bit addresses are clearly enough
- VLA vs 64-bit proponents => 128 bit addresses
  Variable length addresses decay to fixed length anyway
  Performance in routing lookup
- 8+8 proposal led to half of the bits to the hosts and half to the network. Resurfaced as ILNP
- Fifty eight ways of getting an IID

64 bit boundary (RFC7421)

```
|         n bits          |   m bits  |       128-n-m bits         |
+-------------------------+-----------+----------------------------+
| global routing prefix   | subnet ID |       interface ID         |
+-------------------------+-----------+----------------------------+
```

- 64 bit boundary ensures users would always get enough addresses. By numbering links providers can not give less than a /64
- Allow for 8+8, now ILNP
- Technically justifiable when IID was based on EUI-64… what about now?
- What about new proposals like instead of /64 to the link, /64 to the host? Or addressing of applications or addressing of chunks of data?
- Conundrum: Ensure implementations do the right thing, while at the same time…

Tussle #2:Host configuration

- SLAAC vs DHCP
  - Control in the network or Control by the host?
- ND RA or DHCP Default router configuration
- DNS recursive resolver configuration
- Duplicate all functions?

Tussle #3 Extension headers

```
+---------------+------------------------
|  IPv6 header  | TCP header + data
|               |
| Next Header = |
|      TCP      |
+---------------+------------------------

+---------------+----------------+------------------------
|  IPv6 header  | Routing header | TCP header + data
|               |                |
| Next Header = |  Next Header = |
|    Routing    |      TCP       |
+---------------+----------------+------------------------

+---------------+----------------+-----------------+-----------------
|  IPv6 header  | Routing header | Fragment header | fragment of TCP
|               |                |                 |  header + data
| Next Header = |  Next Header = |  Next Header =  |
|    Routing    |    Fragment    |       TCP       |
+---------------+----------------+-----------------+-----------------
```

In hindsight…

- Made “NAT” a part of the architecture. ILNP
- Removed fragmentation from the network layer
- Multi-access links are gone
- Not sure what to do with extension headers
- Required a session-layer / modified transport. Fundamental for multi-homing, mobility…
- Push the hard problems to transport
- Not expose IP addresses to the transport layer and above

Is the IPv6 transition a tussle?

IPv4 only => IPv6 over IPv4 => Dual stack =>
IPv4 over IPv6 => IPv6 only

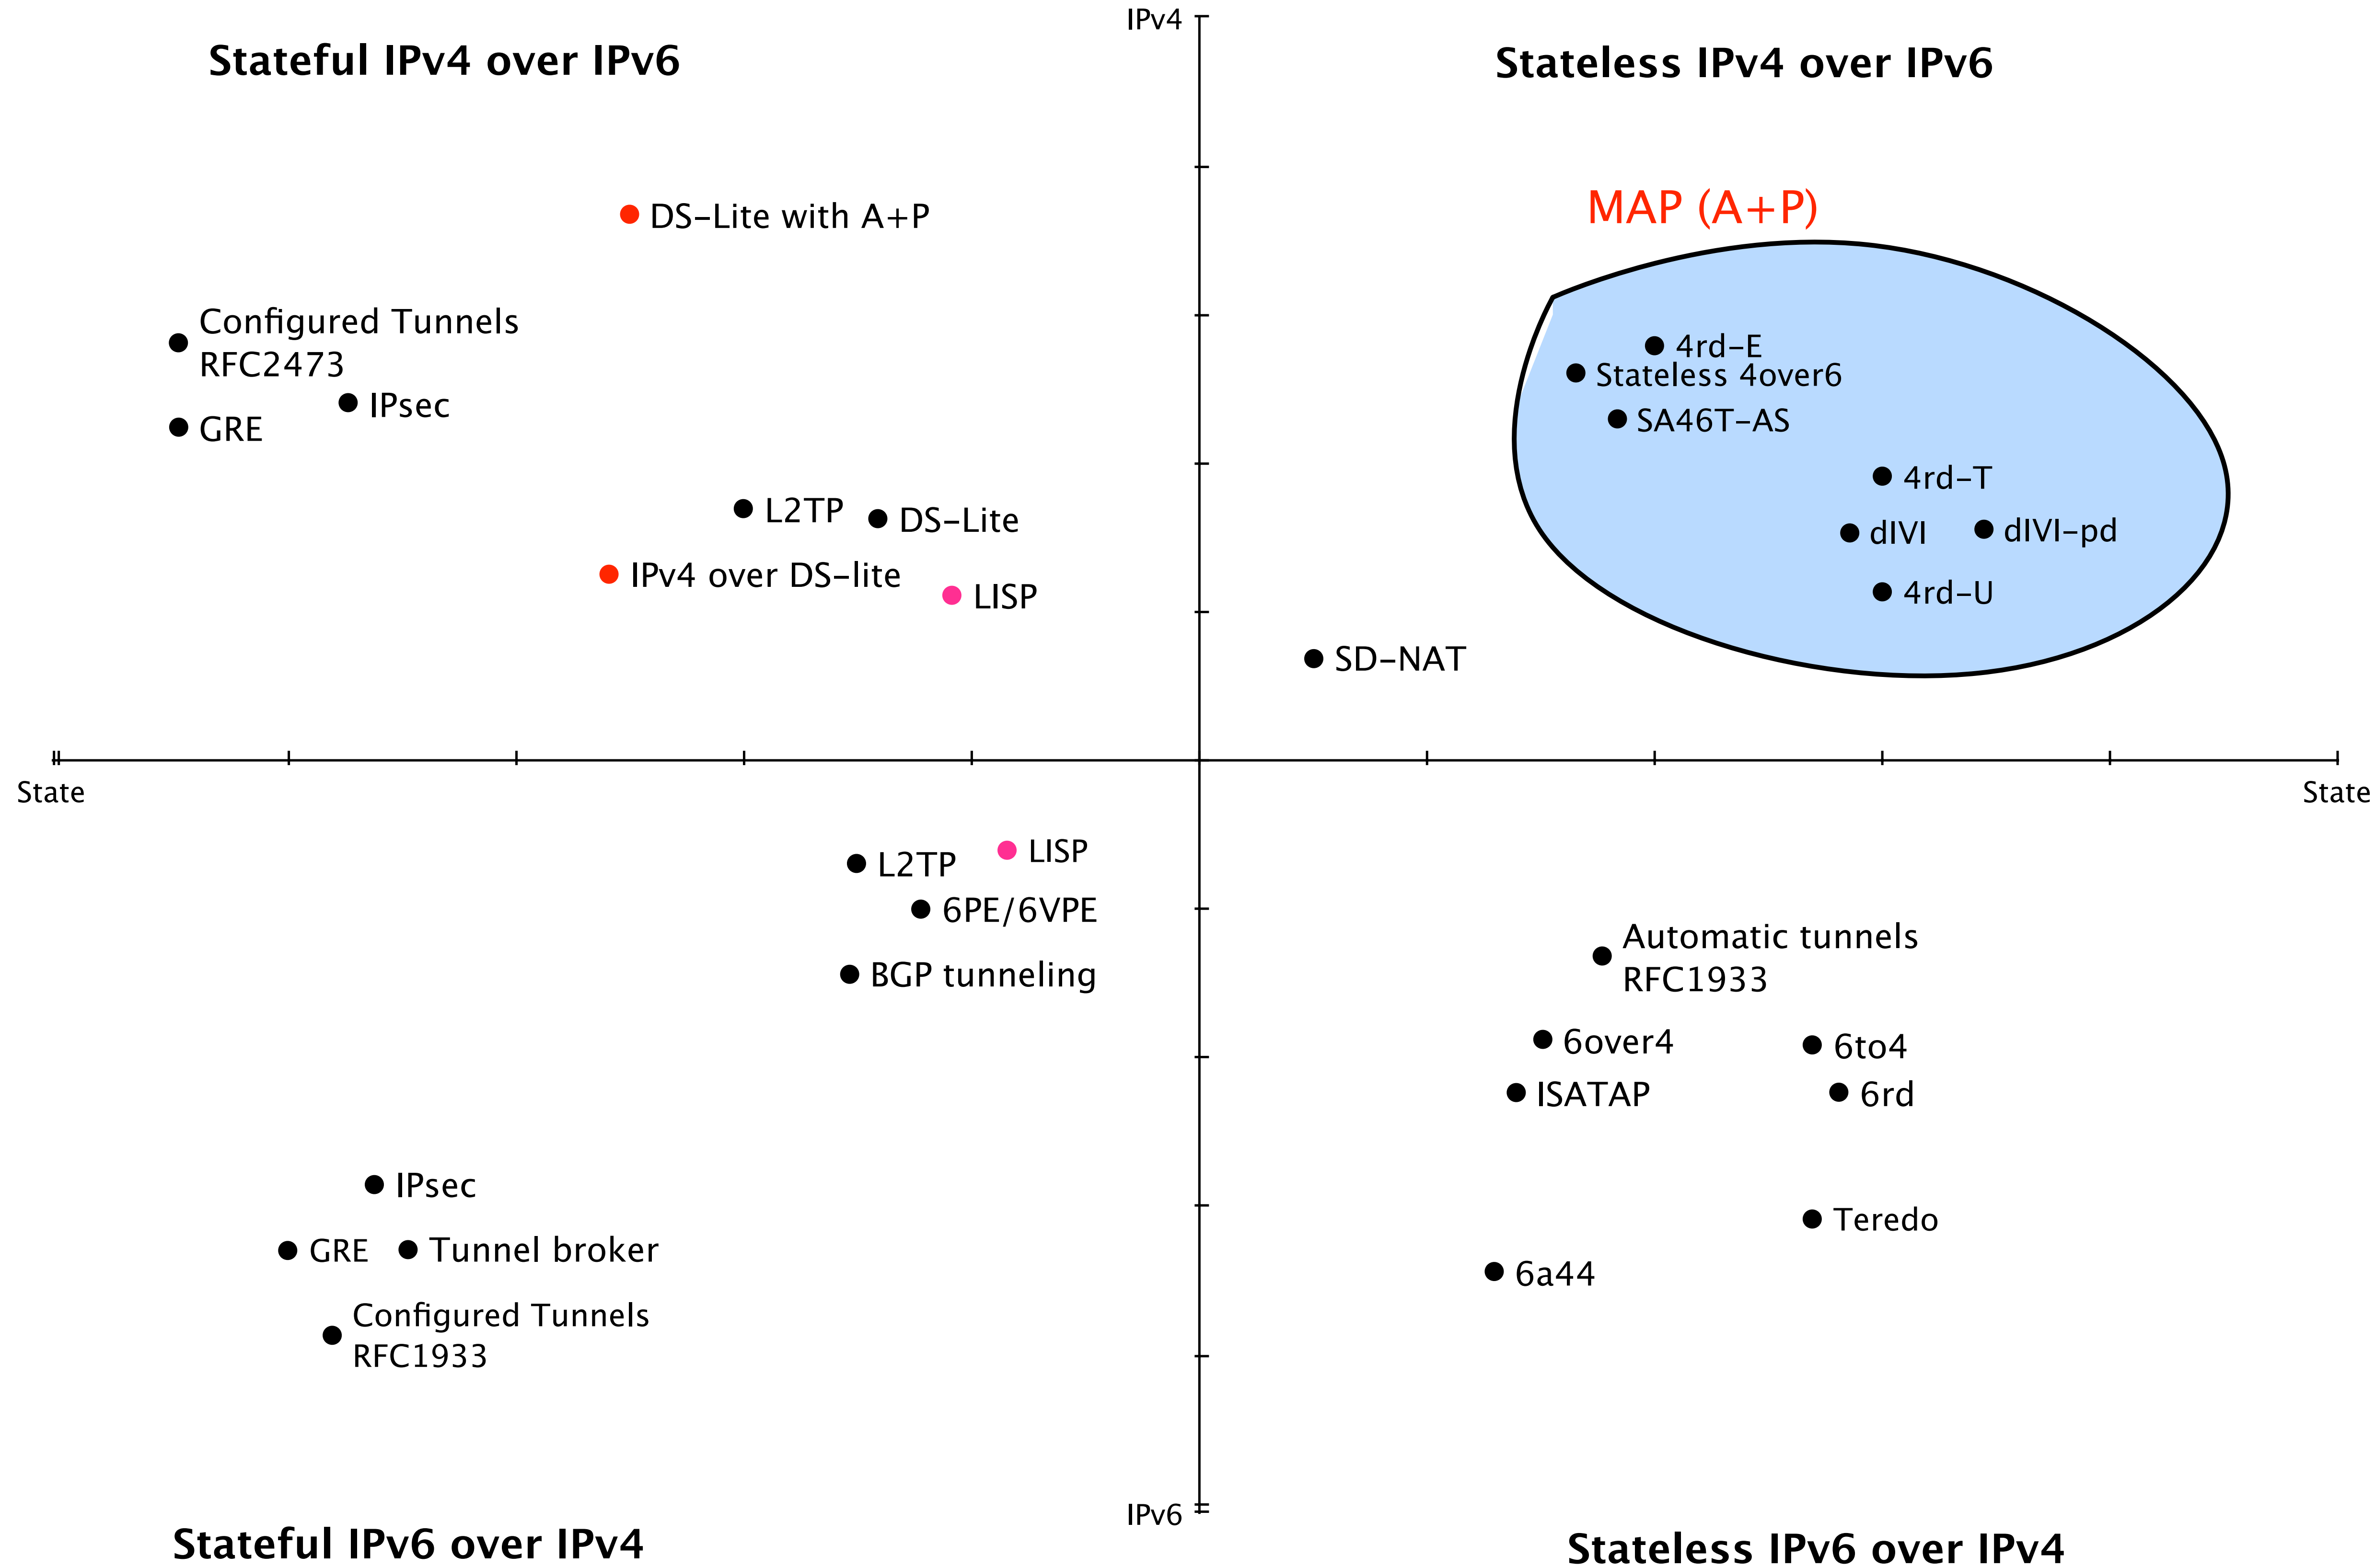

Stateful IPv4 over IPv6
Stateless IPv4 over IPv6
Stateful IPv6 over IPv4
Stateless IPv6 over IPv4
IPv4
IPv6
State
State
DS-Lite with A+P
Configured Tunnels
RFC2473
IPsec
GRE
L2TP
DS-Lite
IPv4 over DS-lite
LISP
MAP (A+P)
4rd-E
Stateless 4over6
SA46T-AS
4rd-T
dIVI
dIVI-pd
4rd-U
SD-NAT
L2TP
LISP
6PE/6VPE
BGP tunneling
Automatic tunnels
RFC1933
6over4
6to4
ISATAP
6rd
Teredo
6a44
IPsec
GRE
Tunnel broker
Configured Tunnels
RFC1933

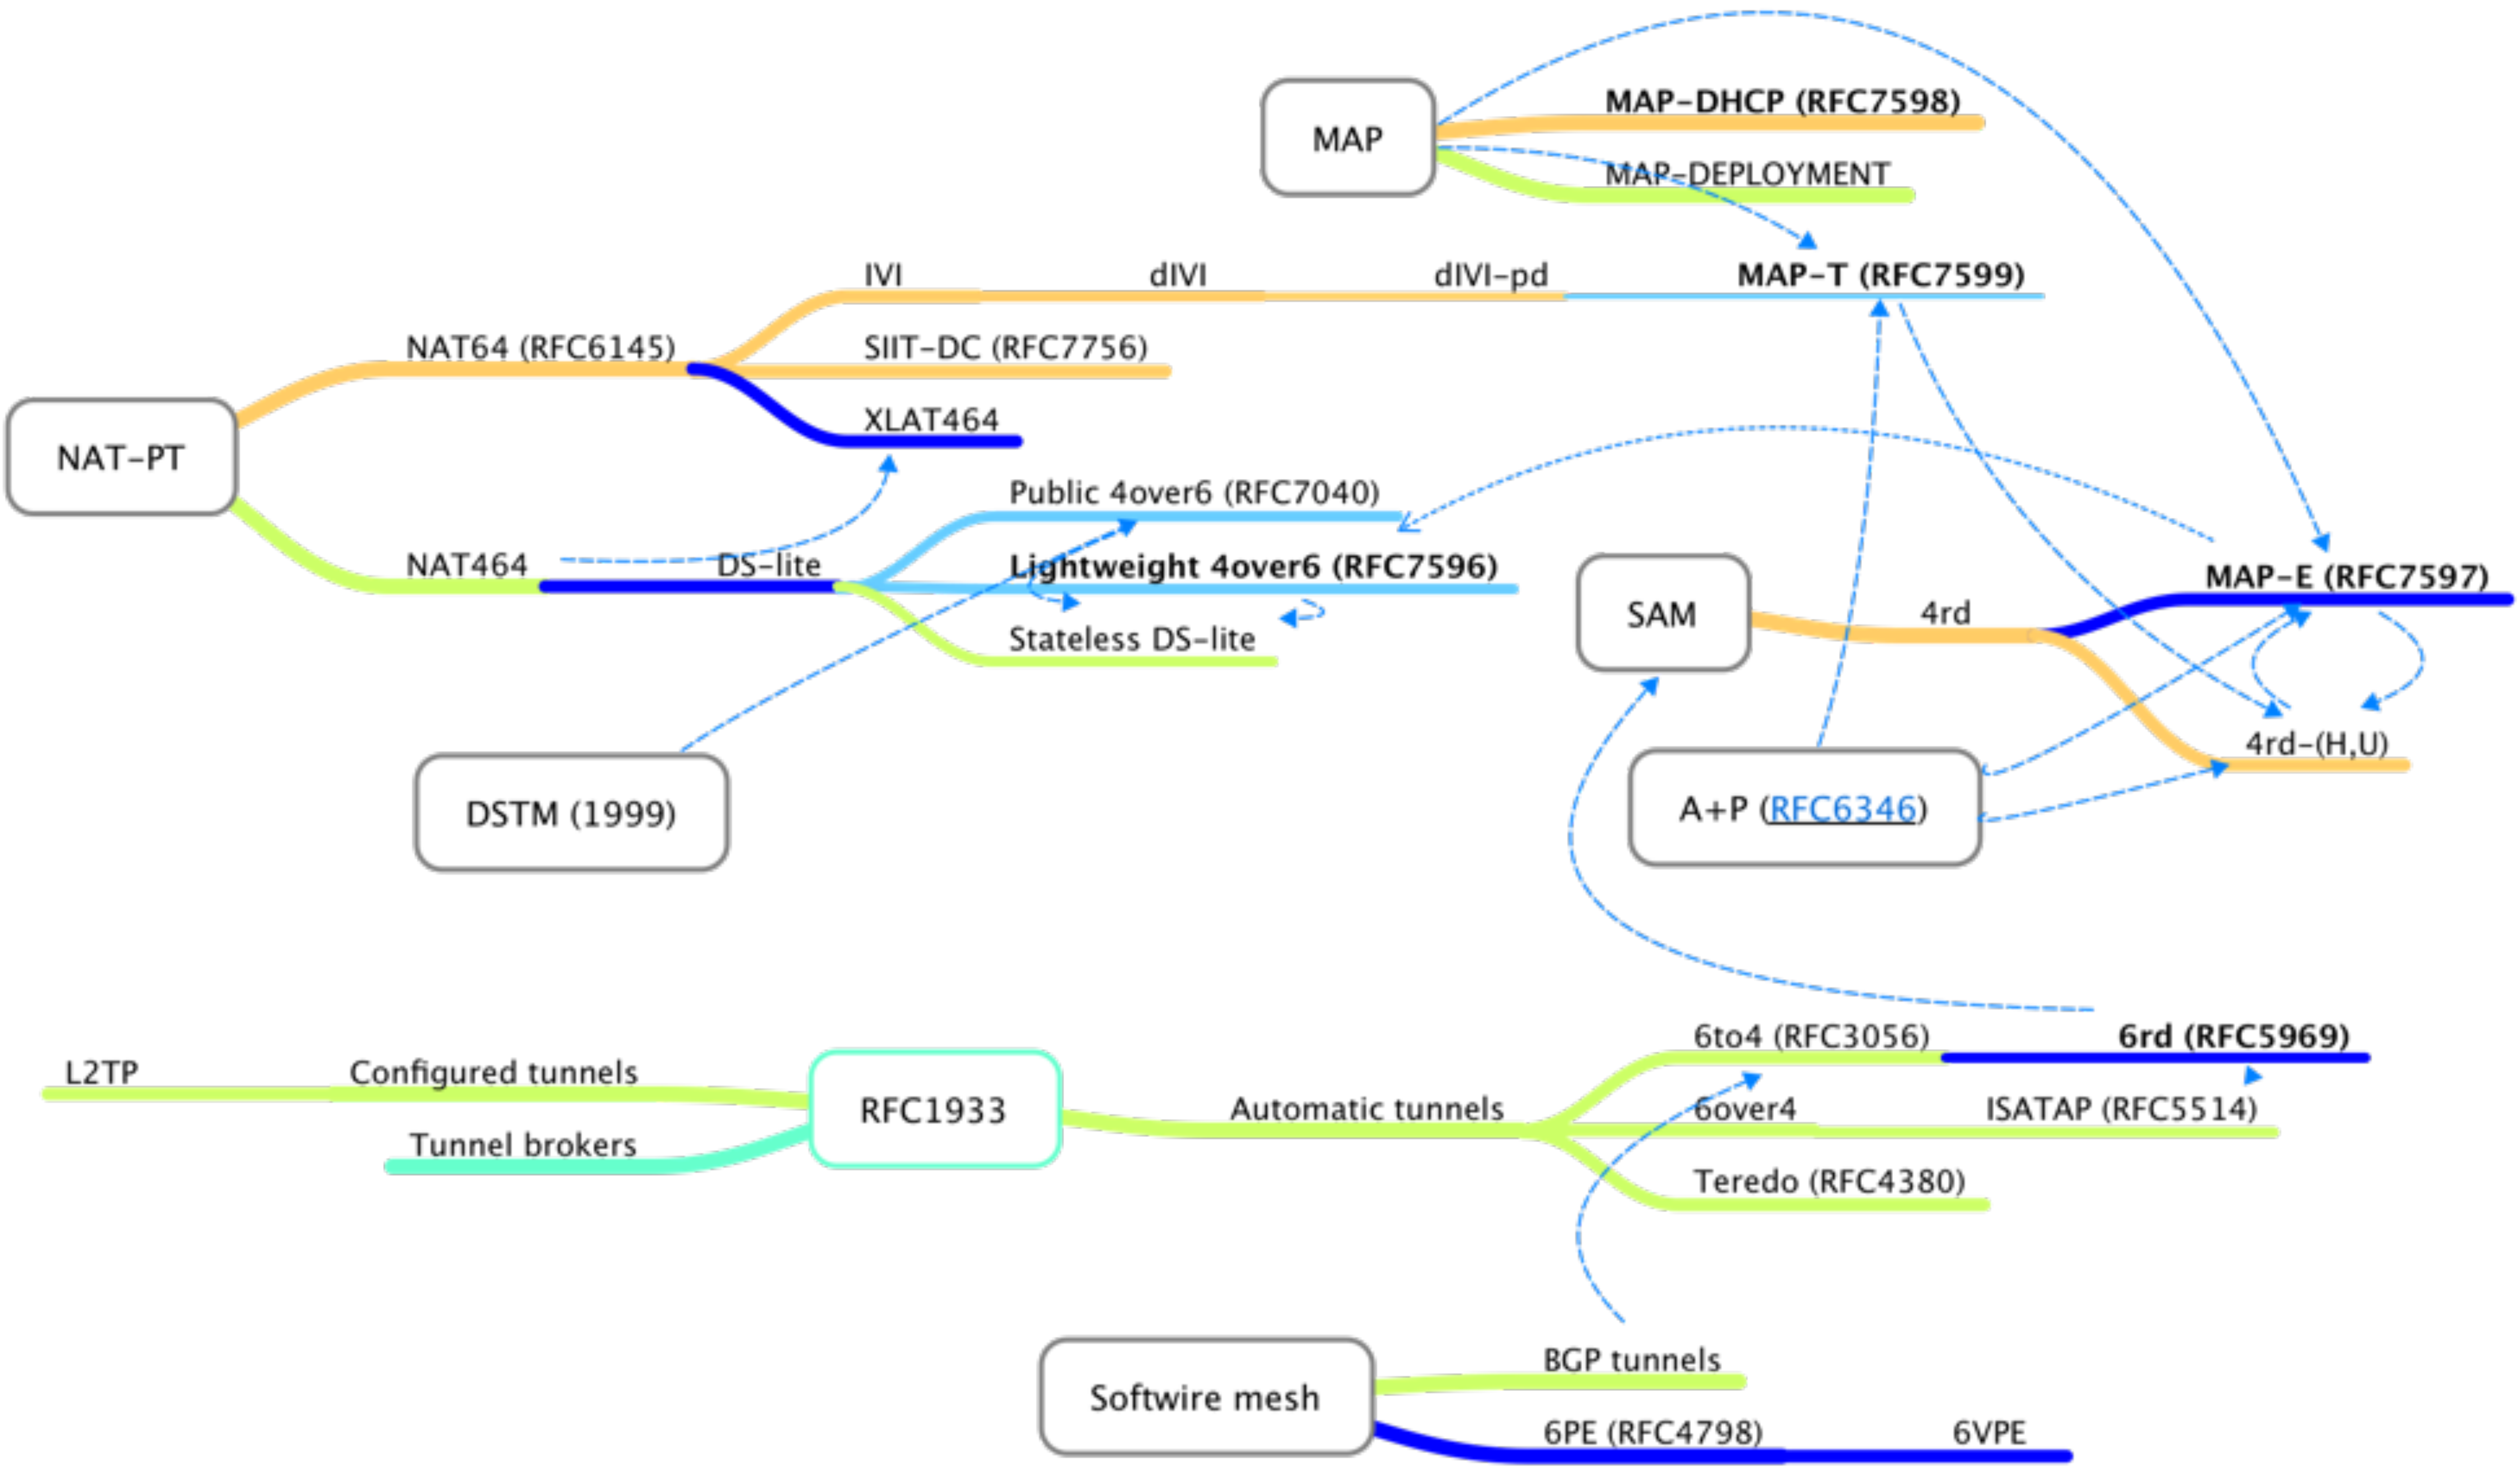
MAP
MAP-DHCP (RFC7598)
MAP-DEPLOYMENT
IVI
dIVI
dIVI-pd
MAP-T (RFC7599)
NAT64 (RFC6145)
SIIT-DC (RFC7756)
XLAT464
NAT-PT
Public 4over6 (RFC7040)
NAT464
DS-lite
Lightweight 4over6 (RFC7596)
Stateless DS-lite
SAM
4rd
MAP-E (RFC7597)
4rd-(H,U)
A+P (RFC6346)
DSTM (1999)
L2TP
Configured tunnels
Tunnel brokers
RFC1933
Automatic tunnels
6to4 (RFC3056)
6rd (RFC5969)
6over4
ISATAP (RFC5514)
Teredo (RFC4380)
Softwire mesh
BGP tunnels
6PE (RFC4798)
6VPE

IPv4 in the face of address exhaustion

- Sharing public IPv4 addresses among more and more users
- Network routing on IP addresses + UDP/TCP ports (A+P)
- What do you think will happen with packets without the L4 information? IP fragments…

$$32 + 16 > 128$$

Current Status

- IPv6 deployment growing healthily
- IPv6 becoming native transport and IPv4 as a service
- But no clear view on when IPv4 can be turned off. Perhaps IPv4 is just an “application” of the network forever, like any other “VPN”.
- IPv4 is evolving into A+P with new transport protocols on top.
- Lots of new development e.g. in open source land are IPv4 only

Future of networking

- Everything is becoming programmable
- Decomposition of the network functions
- Open source vs open standards
- The end to end Internet?

DFIU: Deploy IPv6