ELEC / COMP 177 – Fall 2015

Computer Networking
→ Future of the Internet

Final Exam

- **Thursday, December 10th – 8am-11am**
- Same format as midterm
  - Open notes, open computer, open internet
  - 1 programming problem using Python
- Time limited – 3 hours max
- **Bring your Linux laptop/USB key!**

- Topic: Text-only "Wireshark" tool

Future of the Internet

IPv6

IP Versions

ID	Description
0-3	Unused: Development versions of IP
4	Current network-layer protocol
5	Unused: Experimental stream protocol – ST
6	New network-layer protocol (1996)
7-9	Unused: Experimental protocols – TP/IX, PIP, TUBA
10-15	Not allocated

Why Replace IPv4?

- Performance?
  - Quality of Service?
- Security?
  - Authenticating users?
  - Tracking criminals / spammers?
  - Denial of service?
- New applications?
  - Mobile devices?
- Routing is scaling out of control?

- IPv6 does not attempt to solve all problems with computer networks
  - Actually, it only really solves one of them!

Why Replace IPv4?

- The problem
  - IPv4 has ~4.3 billion addresses
  - World has ~7.3 billion people!
    - How many internet-capable devices per person?
- IP address exhaustion
  - Internet will not "collapse", but new devices / networks will not be able to join
- When? Now! *(back in 2011, actually)*
  - Consuming about one /8 block (16.78 million addresses) per month

IP Address Distribution

- Global pool: *Internet Assigned Numbers Authority* (IANA)
- Regional Internet Registry (RIR) <- "Local distributor"
  - ARIN, LACNIC, AfriNIC, RIPE, APNIC

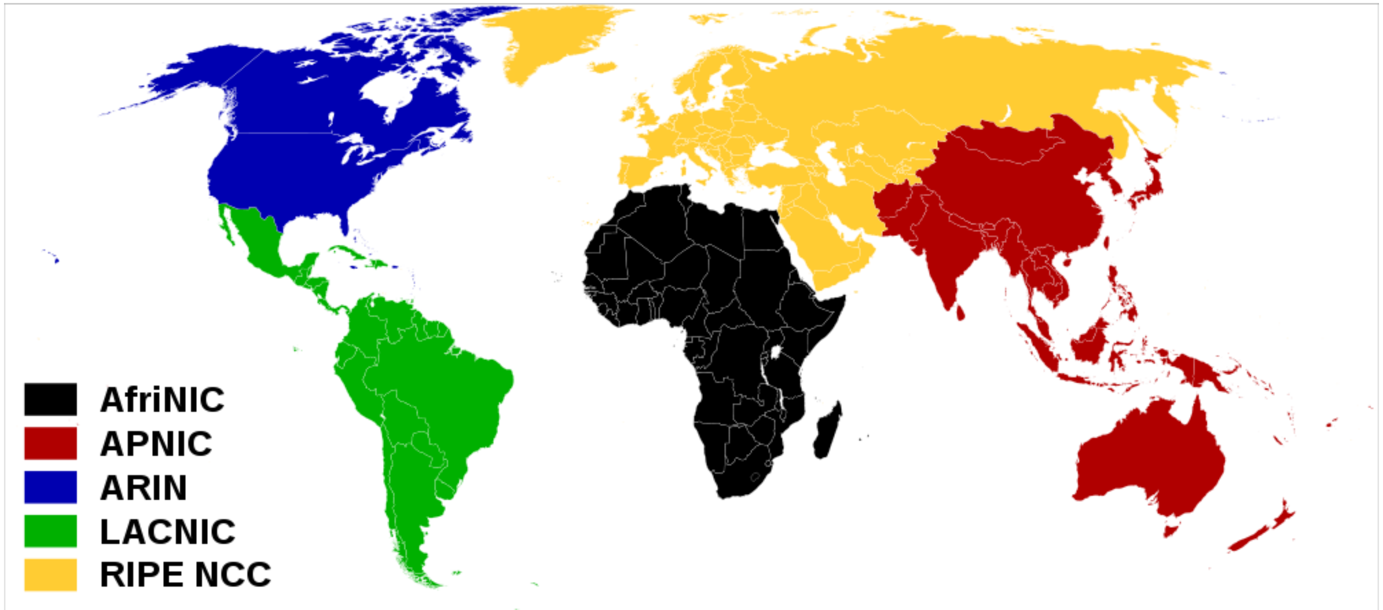

End of IPv4

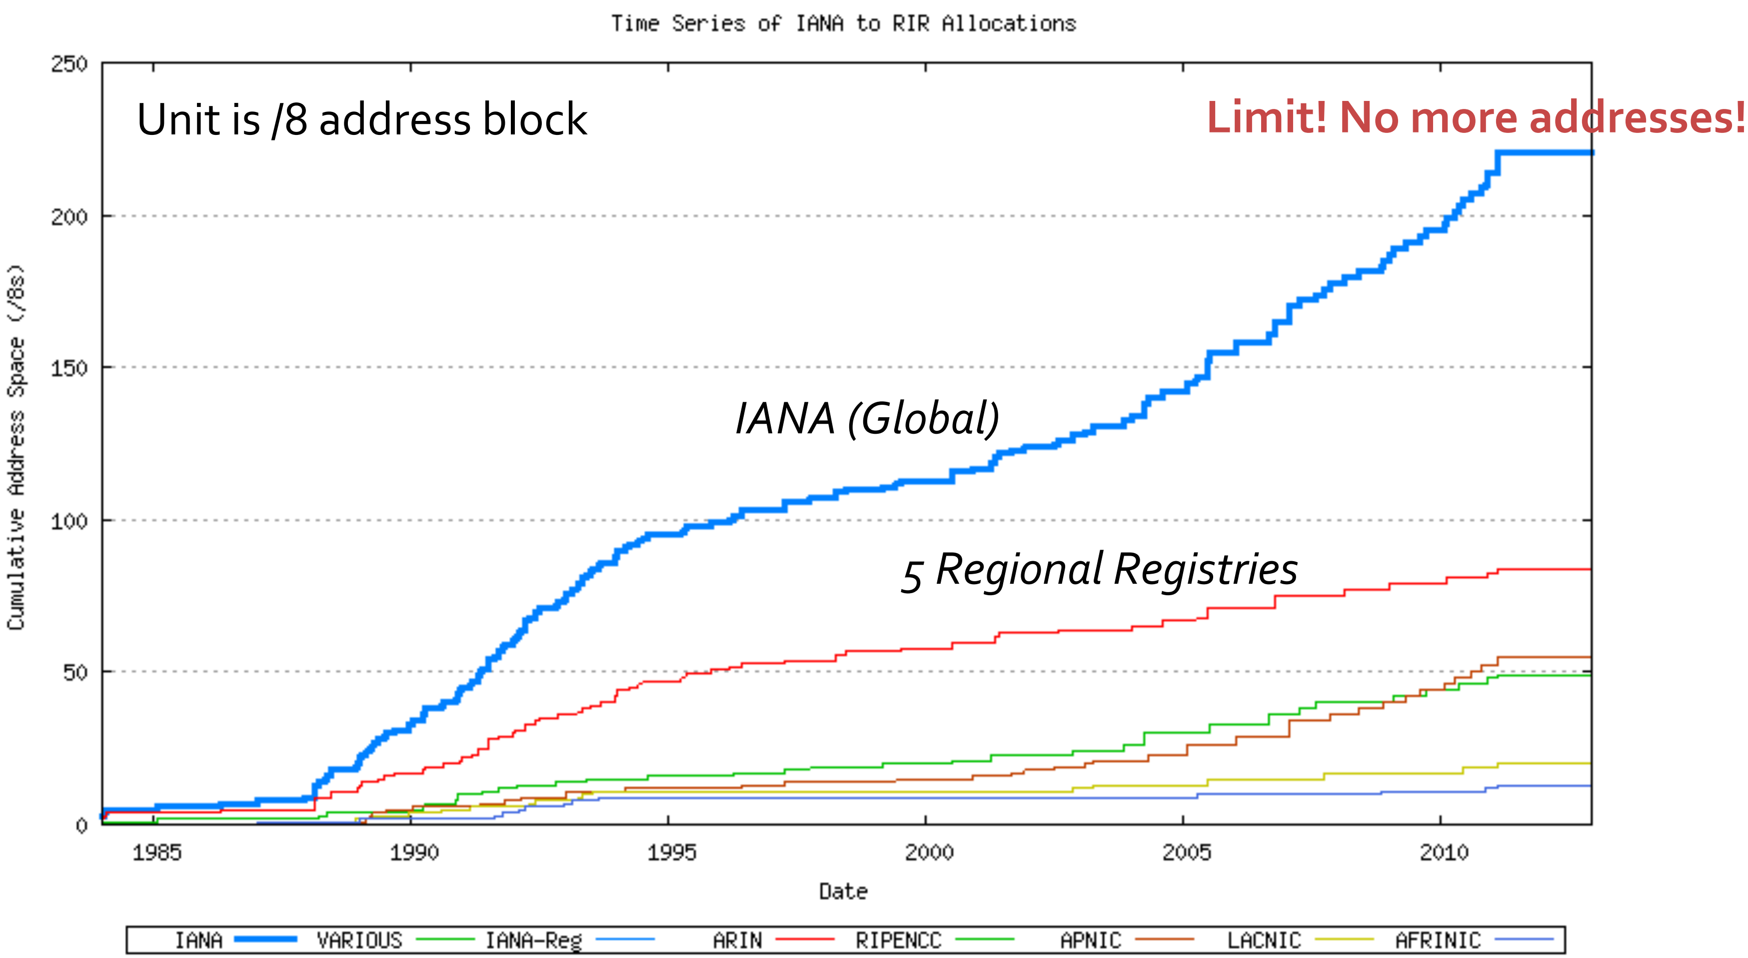

http://www.potaroo.net/tools/ipv4/index.html

End of IPv4

- February 3rd, 2011
  - Internet Assigned Numbers Authority (IANA) distributed final five /8 IPv4 blocks
  - One to each Regional Internet Registry (RIR)
- Registries now only allocate /22 blocks
  - *1024 addresses = tiny!*
- When /22 blocks are exhausted, game over
- **When will the regional internet registries run out of addresses?**

End of IPv4

- Exhaustion Dates
  - APNIC (Asia) – April 14th, 2011
  - RIPE NCC (Europe) – Sept 14th, 2012
  - LACNIC (South/Latin America) – June 10th, 2014
  - **ARIN (North America) – Sept 24th, 2015**
  - AfriNIC (Africa)
    - Slow rate of consumption
    - Several years remaining of supply

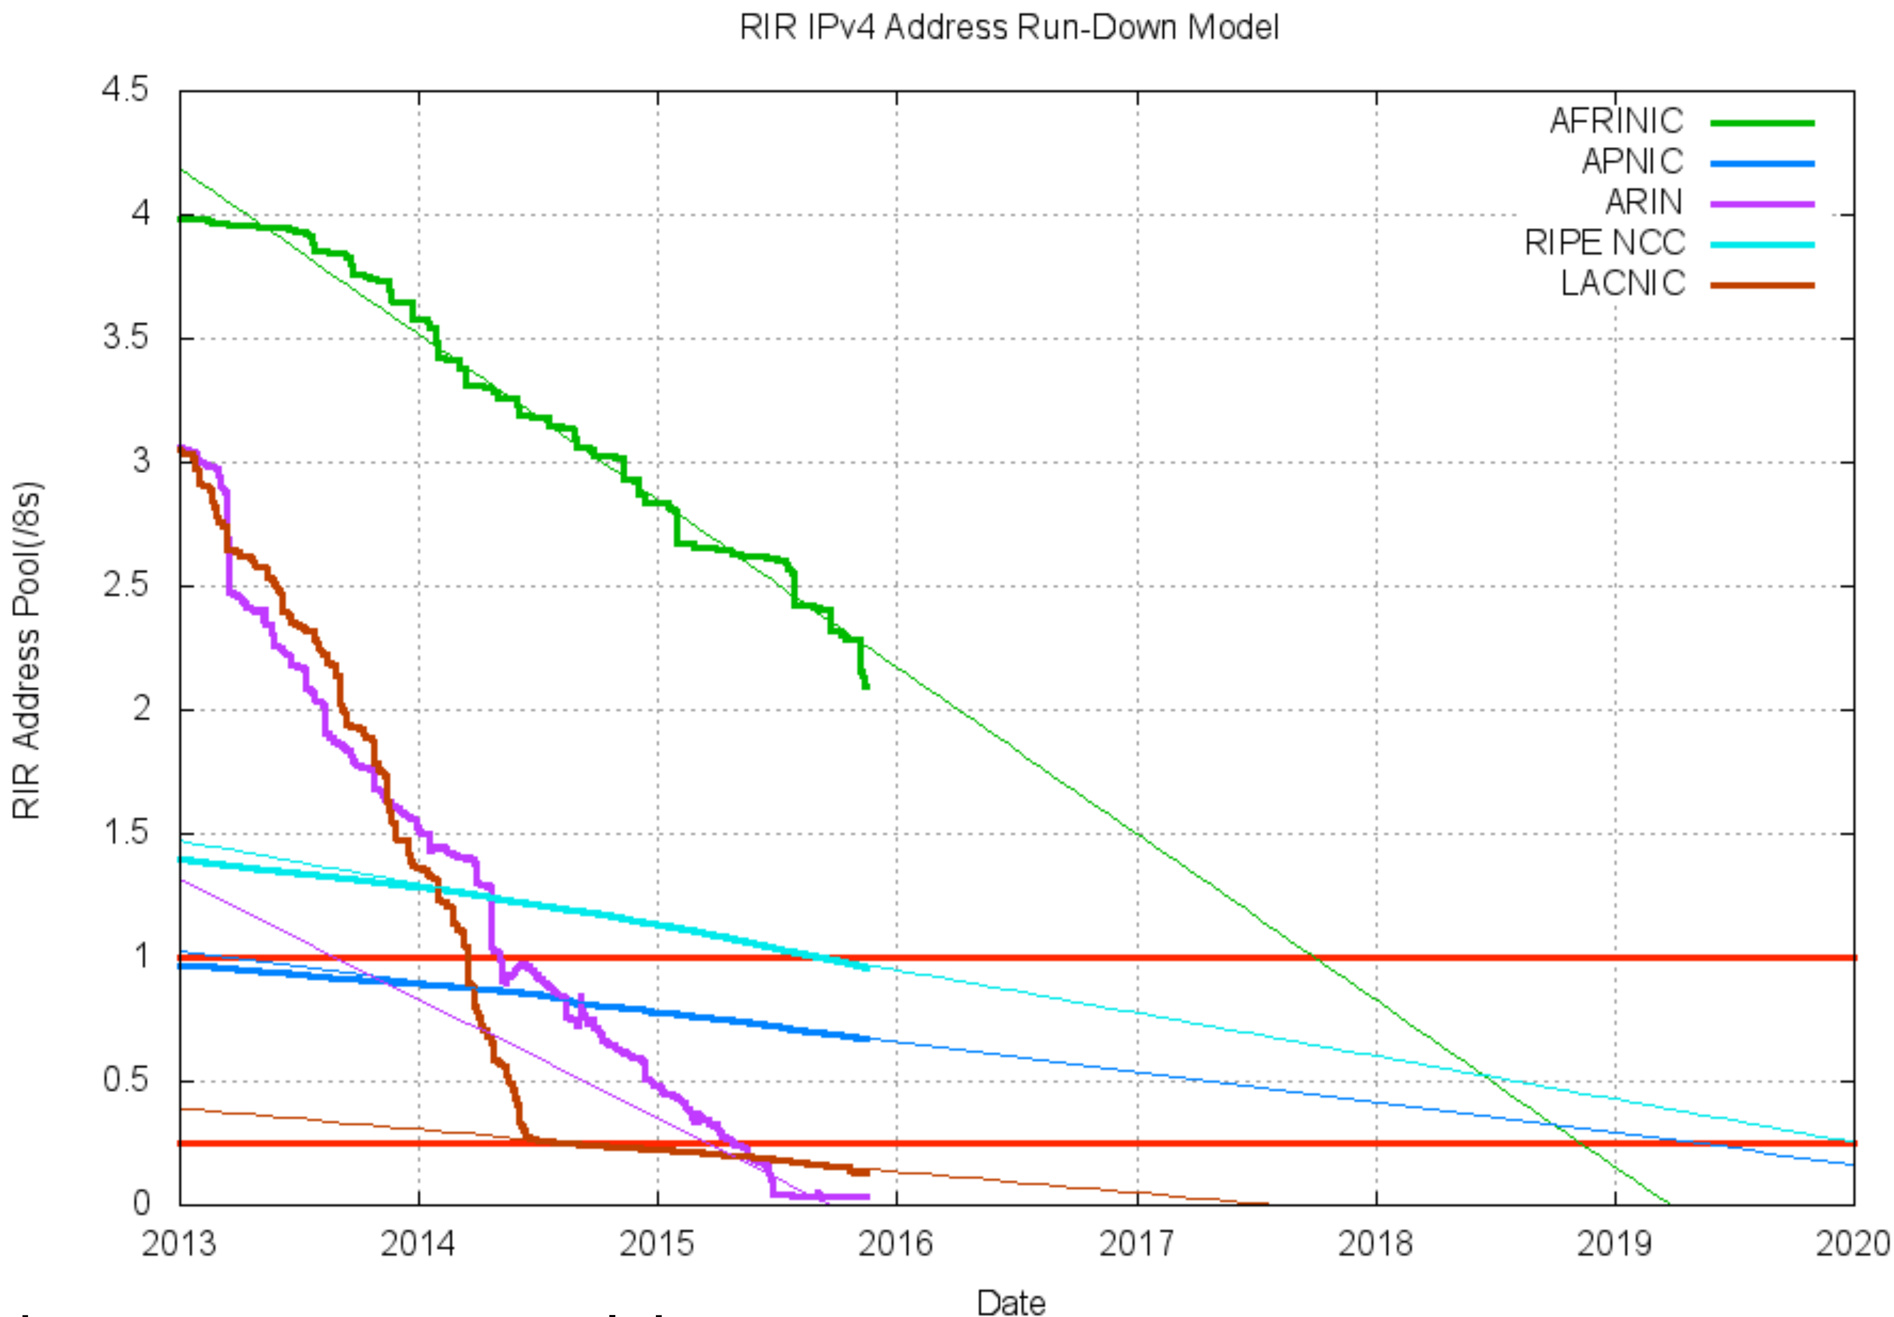

http://www.potaroo.net/tools/ipv4/

IPv4 Address Space

- Unavailable Addresses
  - 10.x – Private Addresses
    - Along with 192.168.x and 172.16.x to 172.31.x
  - 127.x – Local Loopback Addresses
    - Why an entire /8?
  - 224.x to 239.x — Multicast groups
  - 240.x to 254.x — Reserved for "future use"
    - Waste of address space
    - Impossible to re-use today because most IP software flags these addresses as invalid
  - 91 entities with entire class A's (Govt, IBM, GE, HP, MIT, …)
- Current Allocation
  - http://www.iana.org/assignments/ipv4-address-space

Used –vs– Unused Addresses

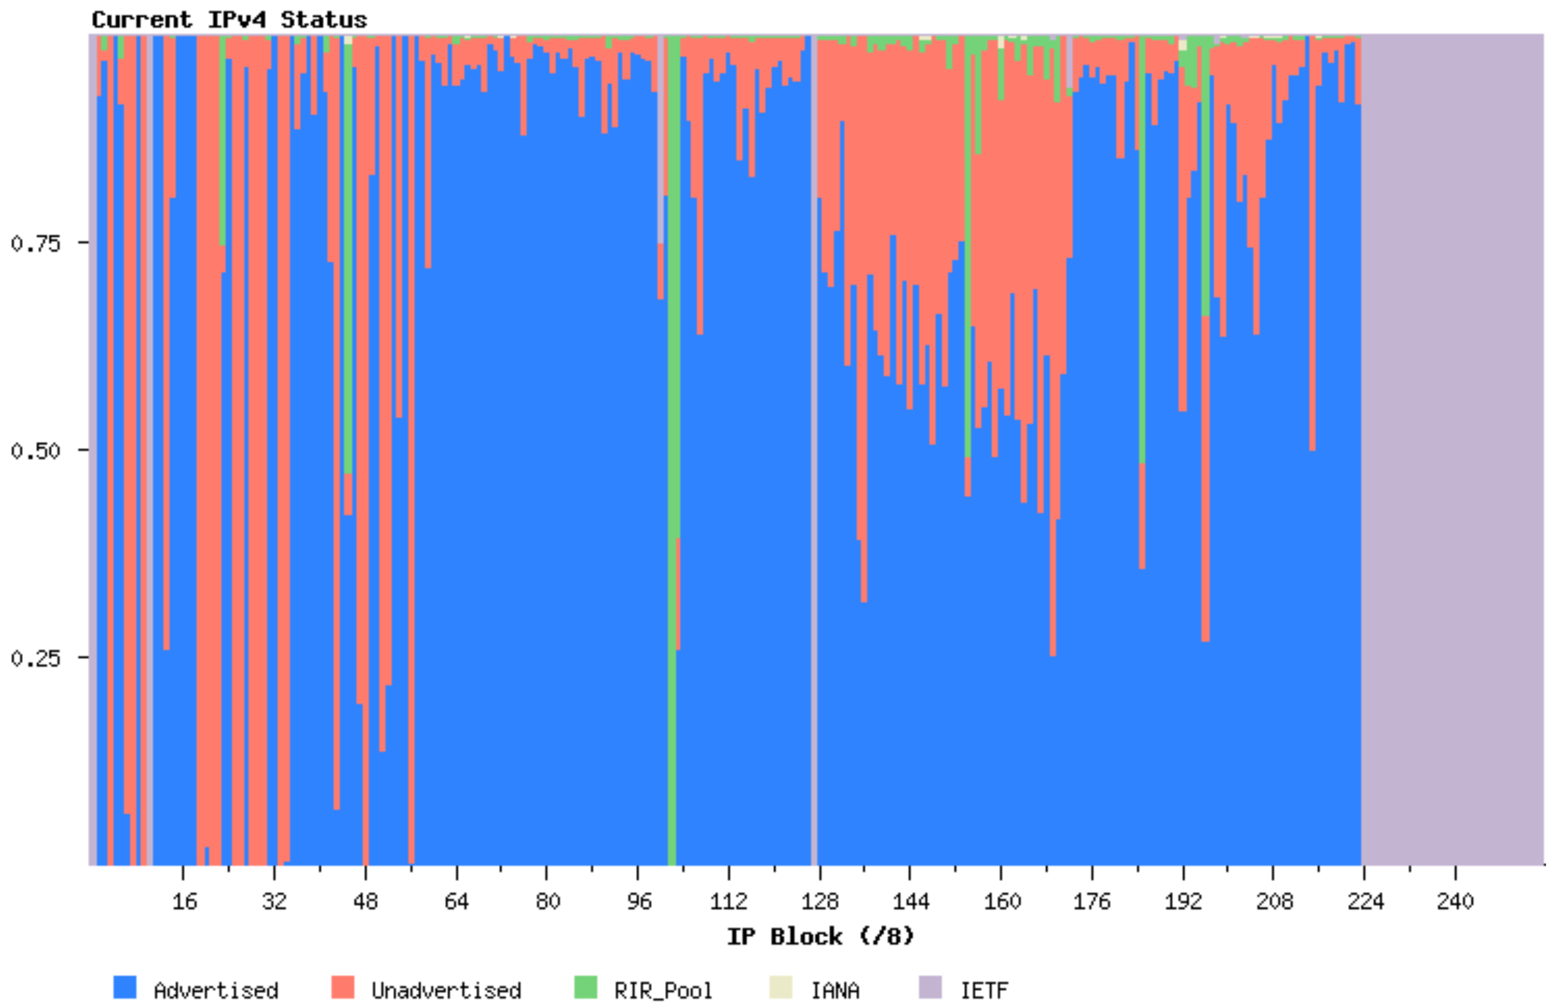

http://www.potaroo.net/tools/ipv4/

IPv4 vs IPv6 - Similarities

- Datagram
  - Each packet is individually routed
  - Packets may be fragmented or duplicated
- Connectionless
  - No guarantee of delivery in sequence
- Unreliable
  - No guarantee of delivery
  - No guarantee of integrity of data
- Best effort
  - Only drop packets when necessary
  - No time guarantee for delivery

IPv4 vs IPv6 - Differences

- Address Length
  - IPv4 – 32 bits (2^{32} = ~4 billion)
  - IPv6 – 128 bits (2^{128}= ~340 trillion, trillion, trillion)
- Security – ~~IPSec support required in IPv6~~
  - *Dec 2011: IPSec support "recommended" in IPv6*
  - IPSec encrypts each IP packet independently
- Reliability – No Header Checksum in IPv6
  - Easier for routers – No need to update checksum after decrementing TTL
  - Relies on link-level error checking
- Quality of Service
  - Label data flows for special priority levels at routers
- Simplified Header Format
  - Infrequently used fields are optional

IPv6 Address Notation

- 128 bits – 8 groups of 4 hex digits
  - `2001:0db8:85a3:08d3:1319:8a2e:0370:7334`
- User friendly! Easy to remember!
- “Helpful” Shortcuts:
  - Omit leading zeros in a group (`0005:0db8:`… is equivalent to `5:db8:`…)
  - Collapse groups of all-zeros with :: (`2001:0000:0000:0000:0000:8a2e:0370:7334` is equivalent to `2001::8a2e:0370:7334`)

IPv6 – Routing

- **How can having bigger IP addresses (128 bits) make routing easier?**
  - Larger address space allows more intelligent network organization
    - Addresses match physical network organization
    - Collapse routing table entries
  - Basic idea
    - Use upper 64 bits for routing
    - Use lower 64 bits for interface ID
      (clients pick this randomly or based on MAC address)

Routing

- **Besides the address layout, how does IPv6 make routing easier?**
  - No checksum calculation
  - No fragmentation
  - Infrequently used headers are optional
- **How does IPv6 make routing harder?**
  - Forwarding table entries 2x-4x larger
  - Need to route both IPv4 and IPv6 for the foreseeable future

Deployment

- Why should I deploy IPv6 today?
  - My customers can reach anywhere on the Internet today
  - “Famous services” (e.g. Google) will always be reachable
  - Only new applications / users will suffer
- How do I deploy IPv6?
  - Flip a switch across the internet?
- Legacy routers may not be upgradeable
  - Hardware implementations cannot be changed
  - Software workarounds offer low performance
- Islands of IPv6 in the sea of IPv4
  - Dual network stacks support both IPv4 and IPv6
  - Tunnel IPv6 across IPv4 networks
- Need to upgrade other systems
  - DHCP
  - DNS (11 of 13 root nameservers as of Nov 2015)
  - Firewalls, traffic shapers, etc.

IPv6 – Original Plan

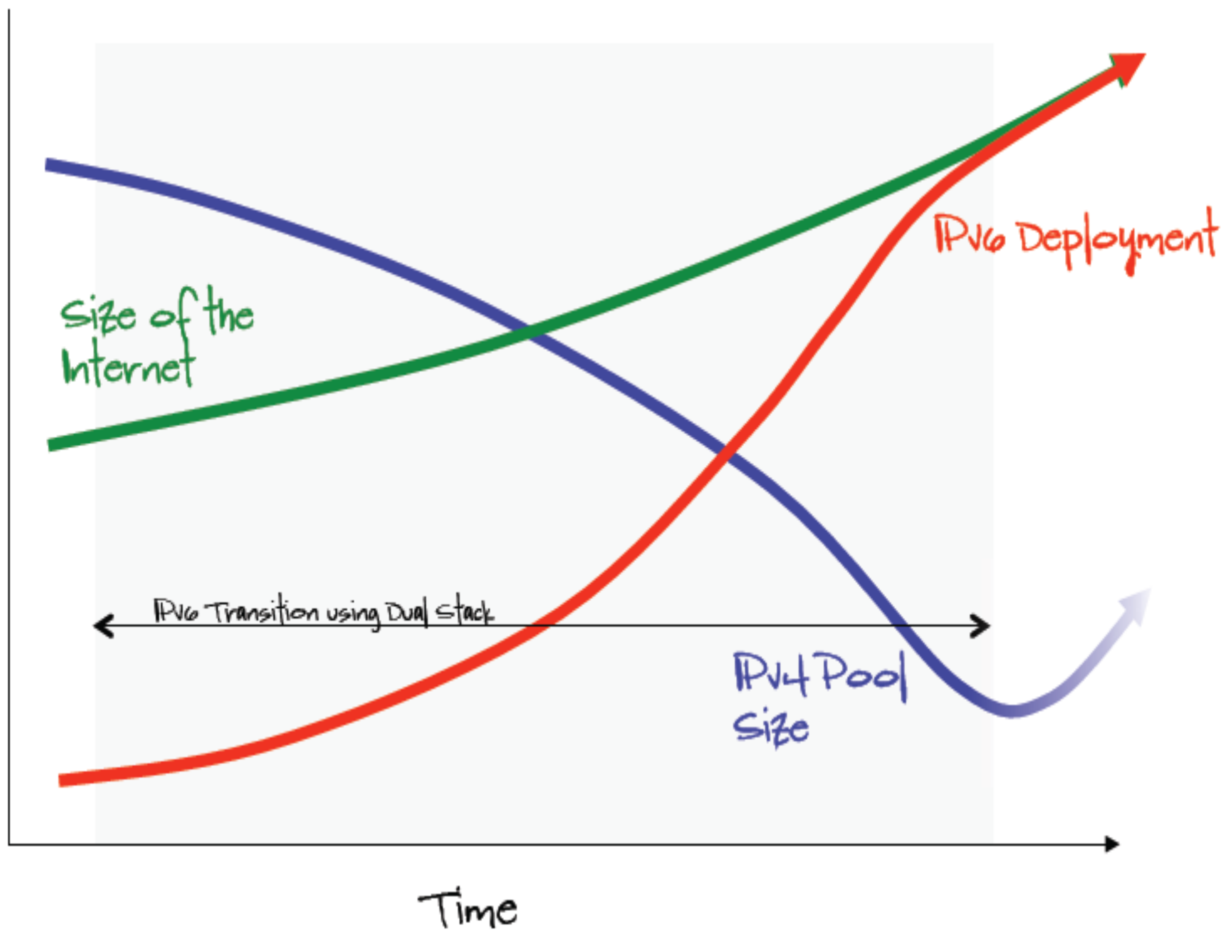

Figure from http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf

IPv6

Current Status

Actual book cover!

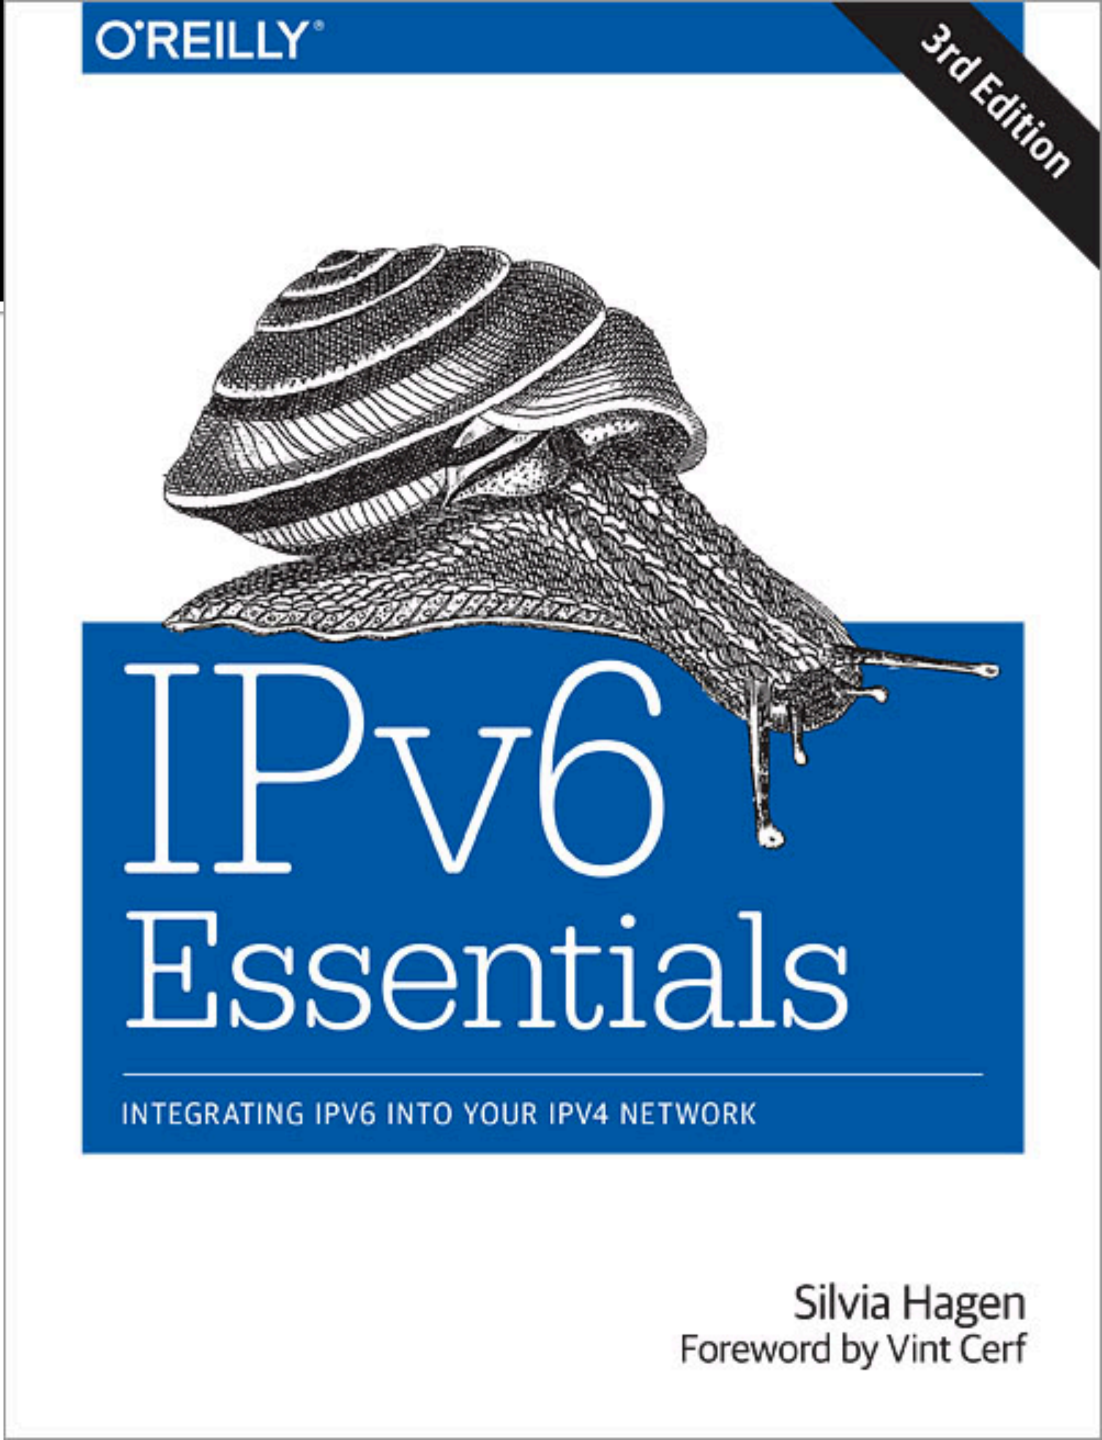

IPv6 – The New "Plan" (?)

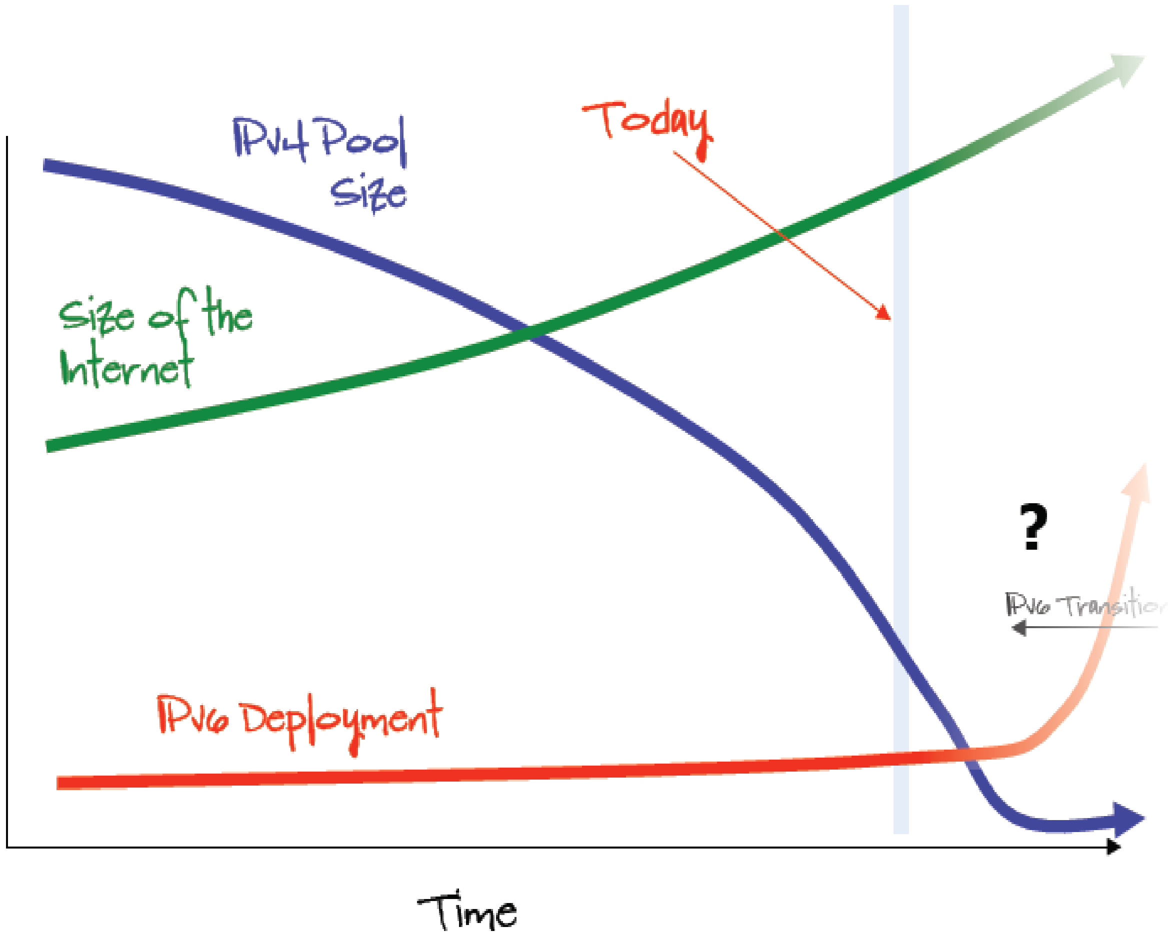

Figure from http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf

Google IPv6 Usage

IPv6 Adoption

We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.

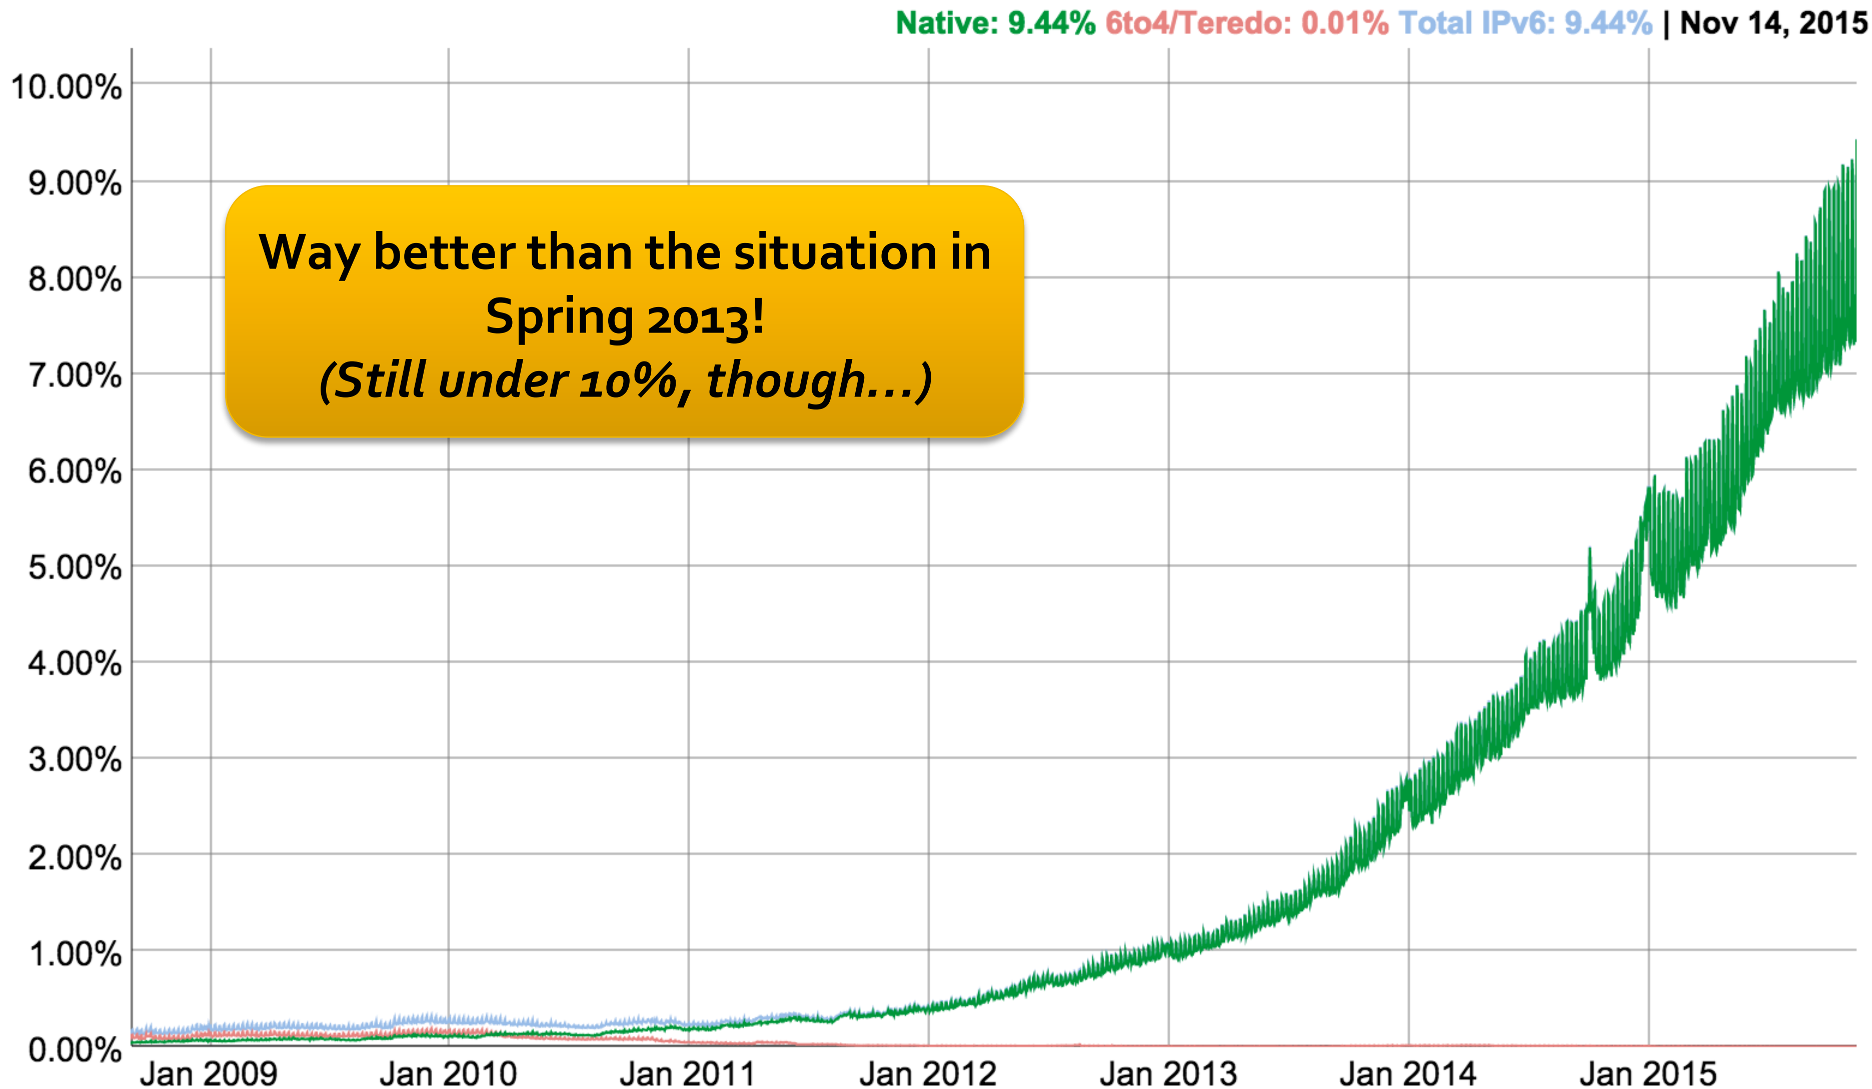

http://www.google.com/intl/en/ipv6/statistics.html

IPv6 – Failure is an Option

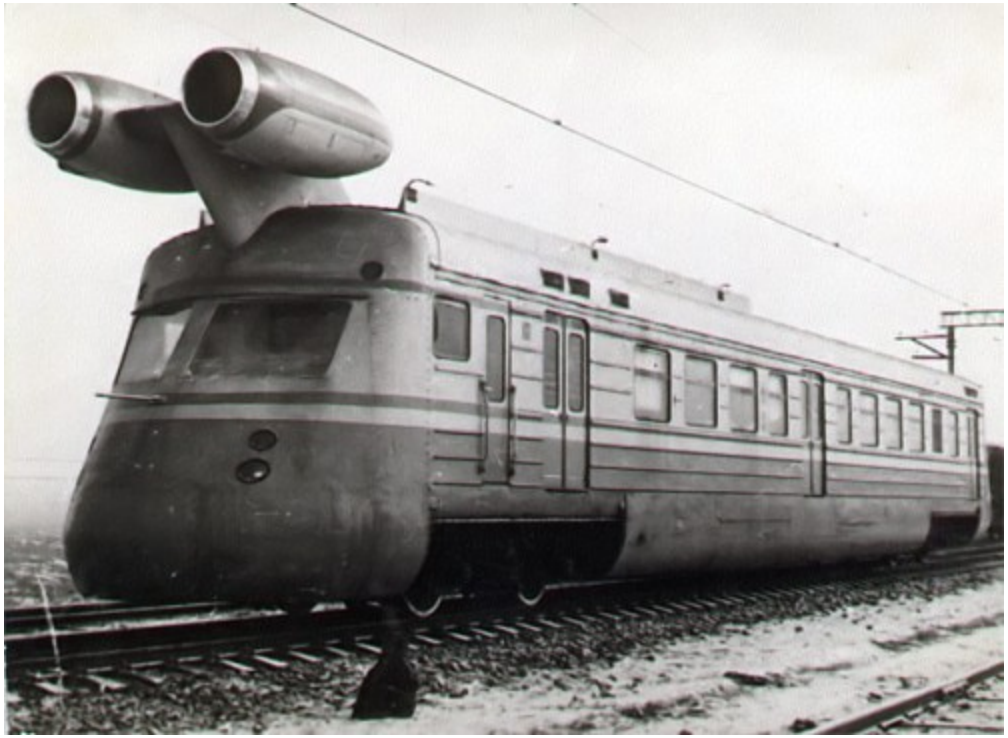

Is this IPv6?

IPv6 – Failure is an Option

- What happens if IPv6 "fails"?
  - Failure is defined as anything less than a complete migration from IPv4 to IPv6
  - Do we stop allowing new hosts to connect to the internet?
- What about using NAT? (address translation)
  - Observation: Only 5-20% of assigned IPs are actually used by hosts.
  - Solution: Use lots of NAT to reclaim unused addressed
- What happens if this works, and we build "carrier-grade" NAT everywhere?
  - No more end-to-end connectivity?
  - Need coordination with ISP to deploy new services?
  - New opportunities for ISPs to filter traffic and charge for services?

http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf

IPv6 – Failure is an Option

- If an organization deploys NAT extensively, how can you get them to give up the reclaimed addresses?
- IP Address Marketplace
  - Can we create a marketplace? (Currently "forbidden" to sell IP blocks)
  - Imagine: *"For Sale: One Lightly-Used IP Block (only used by grandma to check email on Sunday)"*
  - Same problems as buying a used car:
    - Does the person selling the IP block actually "own" it?
    - What is the condition of the IP block? (If used for spam or illicit activity, IP block may be in blacklists worldwide)

http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf