IPv4 Unicast Extensions

“Legacy IPv4 will coexist with IPv6 indefinitely.”

- The Hidden Standards War: Economic Factors Affecting IPv6 Deployment

John Gilmore, Dave Täht
Paul Wouters
Linux Plumbers, Sept 10th, 2019

IPv4 Unicast Extensions

“Even if they have deployed IPv6, growing networks must continue to acquire scarce, increasingly expensive IPv4 addresses to interconnect with the rest of the Internet.”

- The Hidden Standards War: Economic Factors Affecting IPv6 Deployment

Unicast Addressing Won

- Globally routed Unicast is the success story of the Internet
- Large % of the traffic is globally routed unicast (some translated from behind NAT)
- Global Unicast addresses are the ones we're running out of
- All other kinds of IPv4 addresses are tiny niches

- Current IPv4 address allocation doesn't reflect that.

IPv4 Addresses now COST

- An IPv4 address market exists
- Current costs are ~US$20 per addr, and rising
- Why do we care?
  - Innovators need addresses
  - Big incumbents are buying them wholesale
    - So they won't ever run out
    - And so startups won't be able to afford to compete
  - A barrier to competition from startups
- It's all fine if we want Internet innovation to stop
- Or be run by monopolists...

Reducing the Cost

- Land speculators say “Buy land, they ain’t making more of it” -
  - But we CAN make more IPv4 addresses
  - It’s not hard.
  - It’s just a few patches...
  - A spec change...
  - And 5-7 years to deploy

Who Are We?

- Tech geeks who do protocols & policy sometimes
- We noticed IPv4 addresses are getting expensive and scarce
- Investigating what it would take to make more

- This is a **moonshot** talk.

- This is not a Linux issue. Or a *BSD issue. Or a Windows issue. It’s a protocol issue with both technical and political aspects.

Who Are We?

- John Gilmore
  - BOOTP/DHCP, DNSSEC, IPSEC, Cygnus Support
  - Co-founder of the Electronic Frontier Foundation (EFF)
- Dave Täht
  - IETF AQM and BABEL WGs, co-author, FQ_CoDel (RFC8290)
  - Director of the Make-Wifi-Fast, CeroWrt, and sch_cake projects
- Paul Wouters
  - IPSEC, DNS/DNSSEC/, LibreSwan, Red Hat.

We believe in RUNNING CODE... and rough consensus.

The Internet Isn’t Finished

- It’s an experiment. It’s a success disaster. The net arose from the contributions of many people. There were no singular geniuses who gave birth to the internet; rather the internet is a collage of many minds.
- There was no grand plan. The internet could have evolved into something quite different than what it is today. Or it could have not evolved at all.
- There were (and are) no internet deities.
- The internet could have gone other ways (or not at all.)
- The internet is more than the World Wide Web.
- Government initiatives can produce great things.
- “The internet is not finished. There is much yet to be created” - Karl Auerbach

Some IPv4 Address history

- Class A, B, C addrs. Now known as /8, /16, /24.
- 0/8 was “find my network number” in 1984, but didn’t work on LANs. Oops. Retired in 1989, RFC 1122. Replaced by BOOTP, DHCP
- 127/8 - Loopback got its own Class A network number
- 224/4 and 240/4 reserved in 1984 for future experiments. No experiment ever took place in 240.
- Class A/B/C didn’t fit real networks. CIDR replaced them. Took years to deploy. Required changing every Internet node.
- 224/4 used in 1988 for multicast, but it never scaled like unicast

Make New IPv4 Addrs How?

- A small specification change
- Small patches to kernels, userspaces, configs, routers
- A set of testbeds – local, then global
- Iterate the above until it all works

- Only then tackle politics of how to allocate them
- Make “running code” to enable later “rough consensus”
- “Consensus first” screwed it up 10 years ago. Running code first.

Reserved for Future Use?

- The Future is Now.
- 240/4 as Global Unicast
  - Has worked in Linux, MacOS, Android since ~2008
  - Last nit fix for linux landed in December
  - Patches now available for *BSD
- 0/8 as Global Unicast
  - Never used except 0.0.0.0.

Underutilized Addresses?

- While updating every node, extend these too:
- 127/8 - Loopback
  - Only tiny numbers of /24s seen in use
  - The other 16 million addresses are unused
  - Let’s make them unicast
- 224/4 - Multicast
  - Currently has 268 million addresses
  - 128m never ever allocated, never used. Make them unicast!
  - Reclaiming more is probably feasible.

Patch: 0.0.0.0/8 for Linux

```
 static inline bool ipv4_is_zeronet(__be32 addr)
 {
-    return (addr & htonl(0xff000000)) == htonl(0x00000000);
+    return (addr == 0);
 }
```

- Arguably the function should be renamed!

16,777,215 new IPv4 addresses!

Extend Every Subnet Too

- Zeroth address in subnet: fully usable as ordinary host
  - Was reserved in 1980s due to 4.2BSD using it for Broadcast (Oops)
  - Which made broadcast storms when talking to standard nodes.
  - 4.2BSD long gone; let users put nodes there!
- Final address in subnet: fully usable in Point-to-Point nets
  - In LANs, still reserved for Broadcast
  - Fully usable as ordinary host in non-LAN subnets.
- This extends each /29 from 6 to 7 usable addrs!
- And makes P2P interfaces only consume a /31. (RFC3021)

Next Steps

- Keep landing patches, testing, doing interop
- Document and fix all the problems
- Once we can prove everything is working…

- Then deal with IETF, IANA , ICANN, RIRs, etc.

AWS route table

Route Table: rtb-00b1b854d386ea8fb

Summary | **Routes** | Subnet Associations | Route Propagation | Tags

Edit routes

View IPv4 routes only

Destination	Target	Status
172.28.0.0/16	local	active
0.0.0.0/0	igw-[redacted]	active
0.1.1.0/24	vgw-[redacted]	active
127.1.0.0/24	vgw-[redacted]	active
240.0.0.0/4	vgw-[redacted]	active

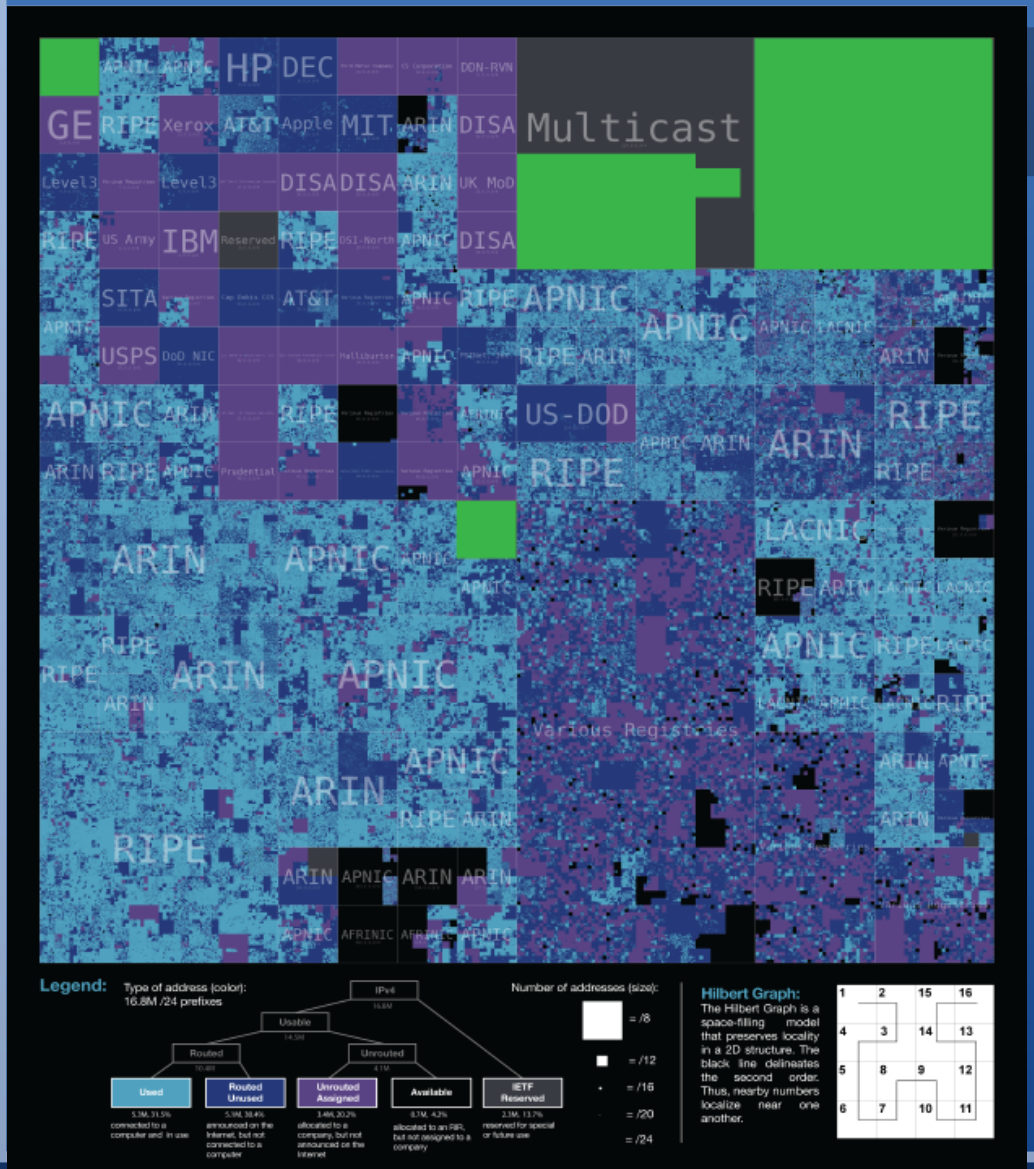

Any Questions?