Lightning Talk:
Help Me Fix IPv6!!

Brian Dickson

briand@ca.afilias.info

IPv6: a solution in need of a solution

- IPv6 has been long in development
- It includes lots of new ideas
- It has a few warts
- Some of those affect its ability to be used as an IPv4 replacement
- We can’t fault the IETF - they aren’t operators

IPv4 end-of-life

- IPv4 address assignments will be difficult after 2009
- IPv6 is a likely replacement
- Both will need to exist side-by-side
- Router slots are rare
- IPv6 needs to be a “wafer-thin mint” by comparison

IPv6 - bits, lots of bits

- IPv6 has 128 bits of address - this is well known
- However, the portion reserved for unicast is 2000::/3, or 125 bits
- That is still a lot of bits… isn’t it?

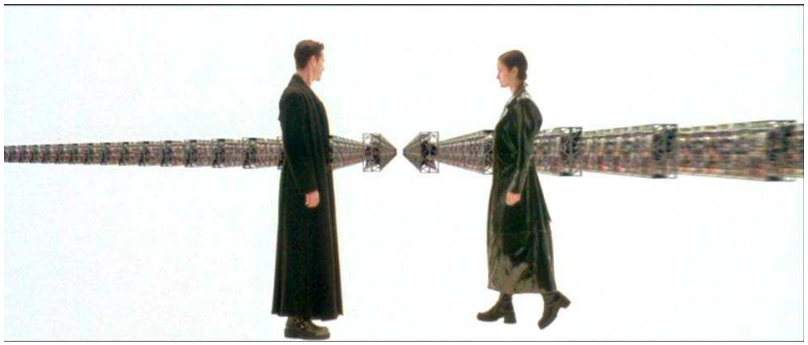

The problem: Interface Identifiers

- One “new” IPv6 feature is auto-configuration
- Hosts self-assign addresses without conflicts in IPv6 addresses
- EUI-64 was chosen for Interface Identifier
- One size fits all

What is EUI-64?

- An IEEE thing
- Bigger than EUI-48, aka MAC-48
- 64 bits - 24 bits of OUI (manufacturer), plus 40 bits of “serial number”
- Vs 48 bit MAC - 24 bits of OUI, 24 bits of s/n
- Used by… IEEE 1394, and not much else

Sidebar: Why EUI-64

- In 1997, Firewire (aka IEEE 1394) was the new hot thing
- Everything was going to have FW
- So everything needed an IPv6 address

What about 802.*??

- There’s a well-defined mapping, stuffing “FFFE” in the middle
- So, by using fixed size 64 bit Interface Identifiers, most hosts MUST have 16 wasted bits of FFFE

Can’t we ignore autoconf?

- No - the IETF mandates that a conforming IPv6 MUST support autoconf
- This means /64 is the smallest useable network size (for dumb users)

So what's the big deal?

- We are left with 61 bits of network for unicast
- IANA gives out /12's to RIRs
- RIRs give /32's to ISPs
- ISPs need to give out and/or reserve big chunks - /48 or /56
- Reserving 4 bits for growth leaves…

Aggregate with 12 bits?!?

- Yes - ISPs who get a /32 will likely reserve /44's for each customer
- 12 bits is very limiting for allocation and internal aggregation
- It effectively squeezes the ISP…

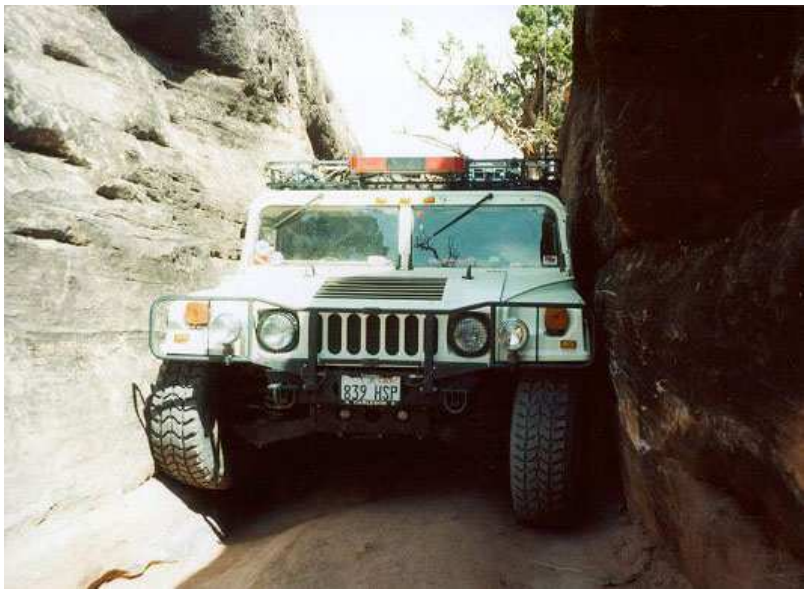

I still don’t get it…

- Router slots are expensive, and not many are left after IPv4 eats them up
- We need to minimize the prefix count from IPv6 for a long time
- Internal router slot is a big cost, too…

How is that different?

- In IPv4, we all got a bit of space, used it up, and got more, again and again
- We couldn’t allocate effectively - so we can’t aggregate well, ***internally*** (IGP)
- We don’t want either situation with IPv6
- Fragmentation is a pain

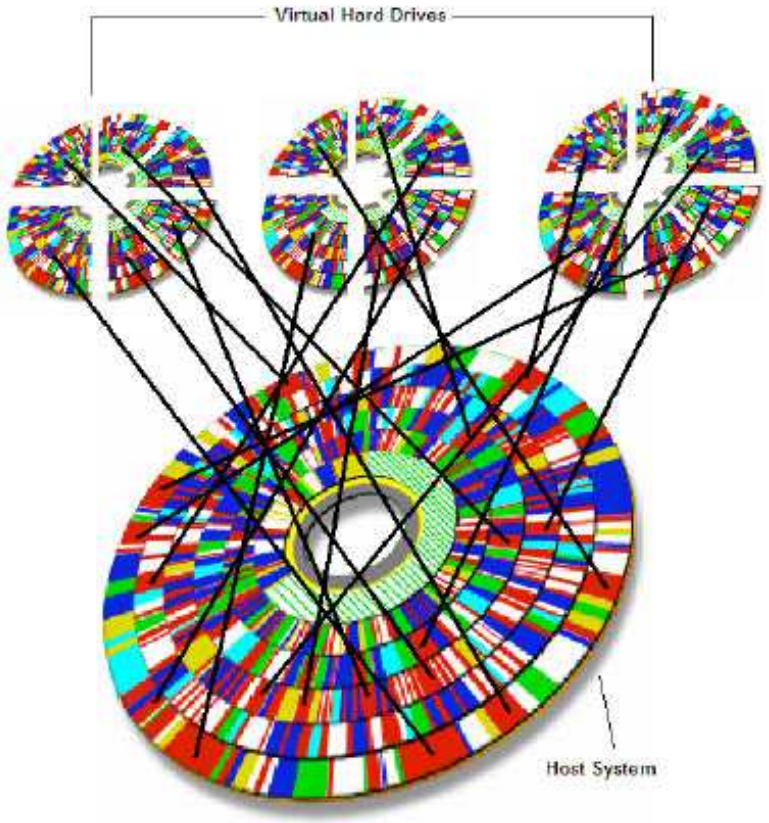

It’s about long-term resource *management*

- You need to make ONE prefix last a long, long time
- But you need to give customers prefixes that last a long, long time
- **Everybody** needs lots of bits for *hierarchical* allocations
- Otherwise, the internal use doesn’t scale and the IGP prefix count explodes
- 20 bits is about “right” for every level
- RIR -> ISP /40 -> customer /60 -> LAN /80
- Keeping everything nibble-aligned is good for DNS (ip6.arpa reverse mappings is nibble-based)

What does 20 bits do?

- It lets you do three levels of aggregation
- Each level is big enough
- No level is “too big”
- 5 bits = 32 blocks of ??
- 10 bits = 1024 (customer routes)
- 5/5/10 means edge routers have 1100 customer routes, everyone else only needs 64 routes!

Reminder: router slots = $$$$

- Even if you don’t have a problem…
- If anyone else need multiple prefixes, they eat slots in **your** routers
- It is in everyone’s interest to ensure everyone can aggregate effectively, and for a long time

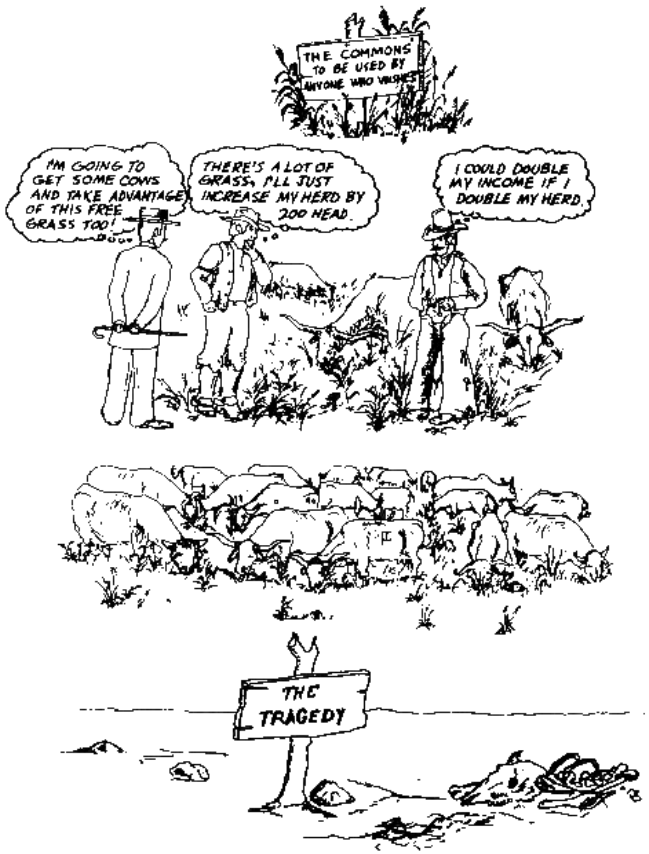

Isn’t IPv6 “done”?

- Yes…
- But, changes can still be made
- I have proposed allowing EUI-48 for Interface Identifiers
- E.g. for 802.*
- We need to “help” the IETF see the light
- It’s backward compatible

What next?

- Join the IPv6 wg mailing list
- Read the draft
- *READ* the draft
- Let me know where it needs fixing
- Tell the WG what you think
- Numbers count!!!

Then what?

- New or updated RFCs - a *very* long process
- Vendor updates
- Patching hosts (all of them!)
- Changing RIR policies
- Start allocating smaller blocks
- Some of these may be possible in parallel
- Need significant (overwhelming) demand from operators/ISPs

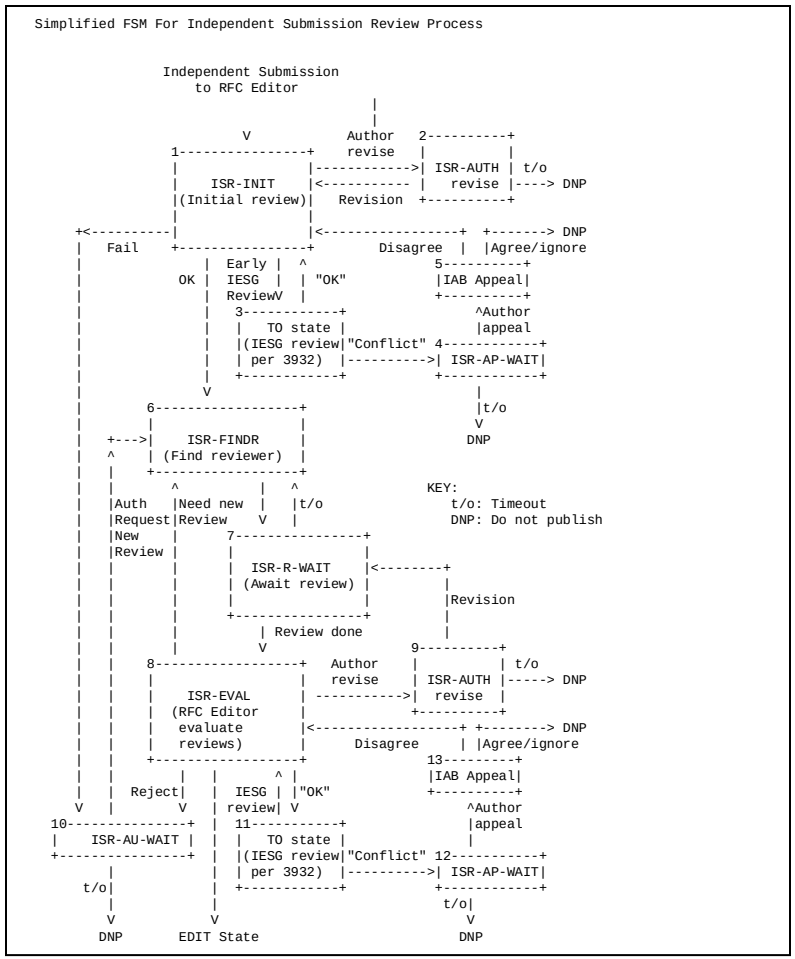

References

- https://datatracker.ietf.org/drafts/draft-dickson-v6man-new-autoconf/
- https://www1.ietf.org/mailman/listinfo/ipv6
- http://www.ietf.org/html.charters/6man-charter.html
- IETF 70 in Vancouver, BC, Dec 2-7 (please show up if you work for an ISP!)