It's The End Of The World As We Know It

(aka "The New Internet Architecture")

NANOG 45
Jan 25-28, 2009
Santo Domingo, DR

David Meyer
dmm@1-4-5.net
(Speaking for myself)
http://www.1-4-5.net/~dmm/talks/NANOG45/iteotwawki-nanog45.pdf

Agenda

- What's the Problem?
- What Is The New Architecture?
  - And what world is ending?
- What My (Very Cloudy) Crystal Ball Tells Me
- So What's Next?
- Q&A

Before We Dive Into All Of This…

Notwithstanding reports to the contrary, the sky hasn't fallen (yet)
But if Form Enables Function…

What's the Problem?

- Data Plane under attack due to panic based on IPv4 run out
  - O(30) /8s left in the IANA Free Pool
  - Dual-stack transition to IPv6 abandoned
  - All varieties of NAT (and beyond) being proposed
- Control Plane under duress (crumbling?) due to various operational practices and economic concerns
  - And we haven't even seen widespread adoption of IPv6
  - That combined with the RIRs "PI-for-all" IPv6 allocation strategy means more rate*state in store for the control plane
  - And deaggregation is on the rise (for various reasons)
    - And just wait until the gray/black address trading market turns white
    - ARIN Policy Proposal 2008-6, similar proposal at RIPE
    - And what does this say about the future of consensus-based self-governance?
  - I'll just note here that the IPv6 designers never successfully dealt with control plane issues

Ok, But What's The Concern?

- Assertion: The lack of a reasoned approach to both the IPv4 run-out problem (data plane) and the growth of routing state (control plane) are life-threatening to the (end-to-end) Internet we all know and love

- I want to focus on the data plane (because that's the panic de jour), but let's overview the control plane issue(s) for a minute….

Internet Control Plane
What's the Issue?

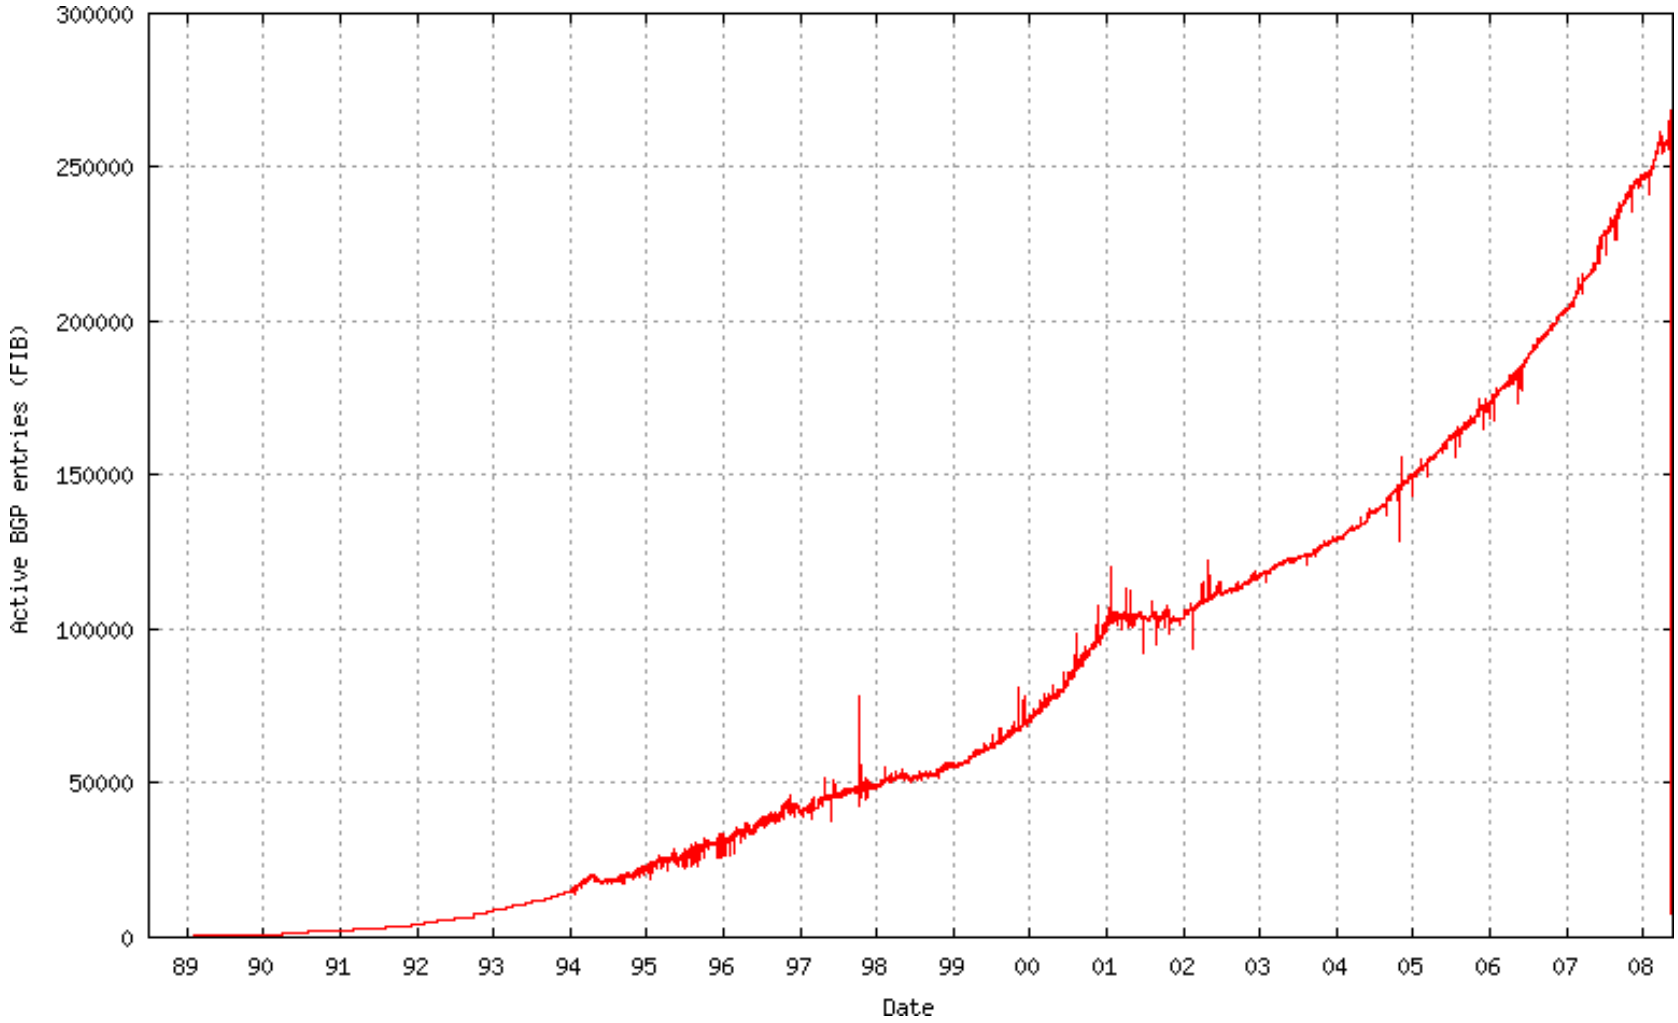

Is Locator/ID Split the Solution?

Changing the semantics of the IP address

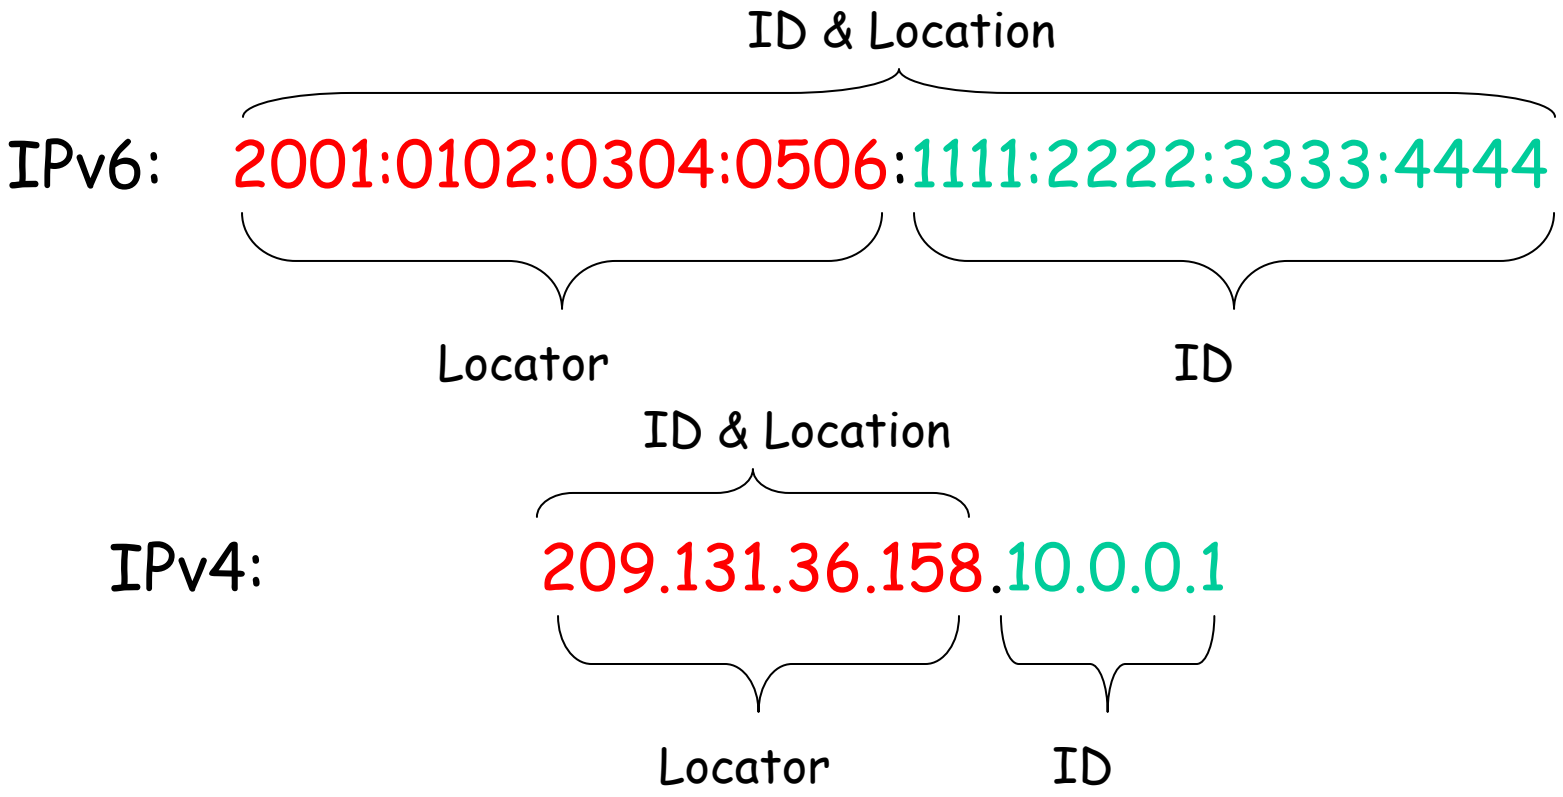

Scaling the Control Plane

- Lots of solutions based on the Loc/ID split idea
  - See Noel's page for some of the definitive work on Loc/ID
    - http://ana.lcs.mit.edu/~jnc/tech/endpoints.txt
  - Basically, you have one (blunt) instrument to scale the control plane: topological aggregation
- 8+8/GSE, Six/One Router, IvIP,LISP,…
  - Check out the RRG mailing list for more
- None has seen serious implementation other than LISP, and none has seen serious production deployment
- See http://www.lisp4.net for some information on the LISP protocol and deployment status
- So lets get on to the data plane

Scaling Internet Data Plane: What Was The Plan?

- IPv6
  - But, how to get there from here?
- Well, we had Dual Stack, and we had…
- Dual Stack, and we had…
- Dual Stack
- Dual stack turns out to be an inherently flawed approach
  - While I can signal that the correspondent host is IPv6 capable with a *AAAA* record
  - This tells me nothing about the capability of the data path
  - Consider Vista's behavior

Dual-Stack IPv6 Uptake Model*

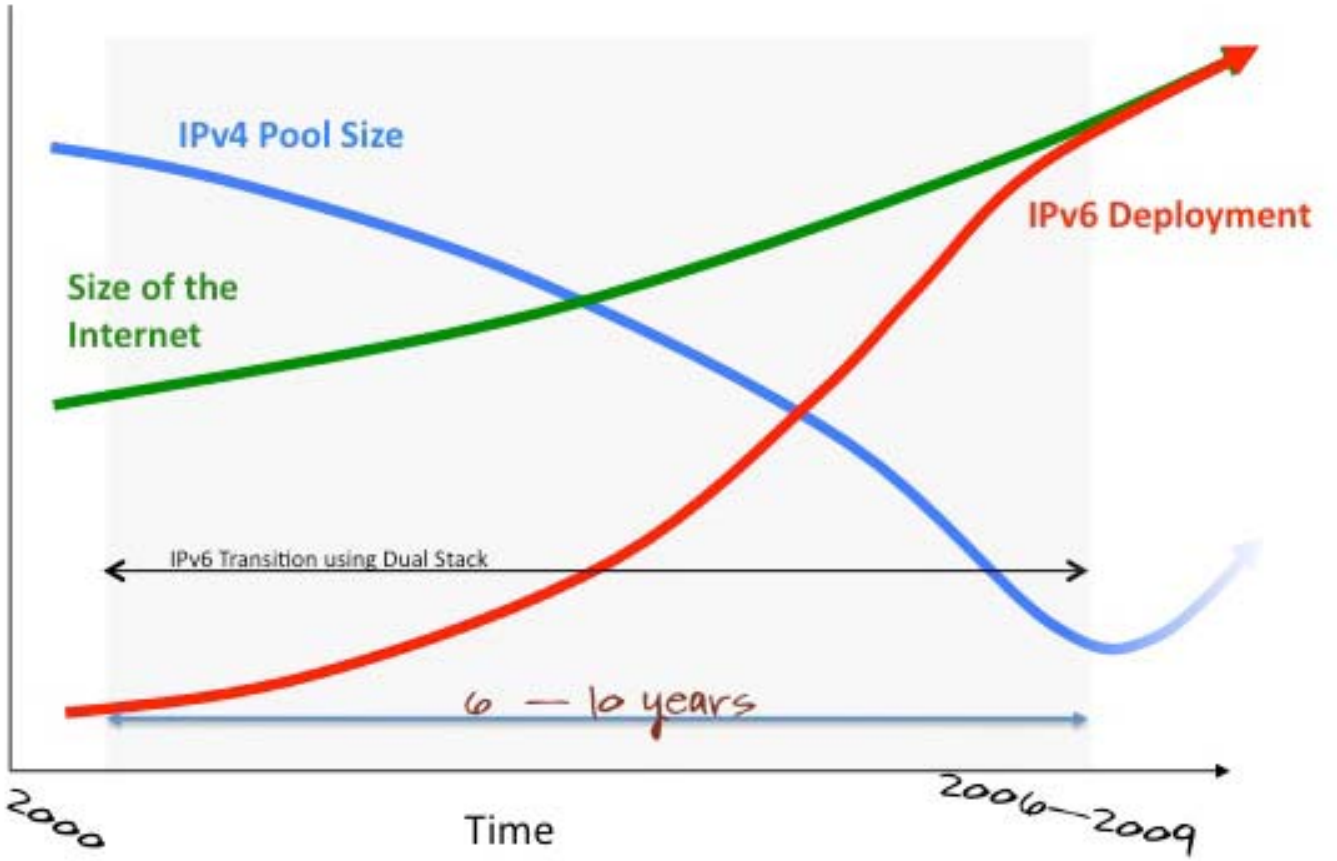

* Courtesy Geoff Huston

What Really Happened*

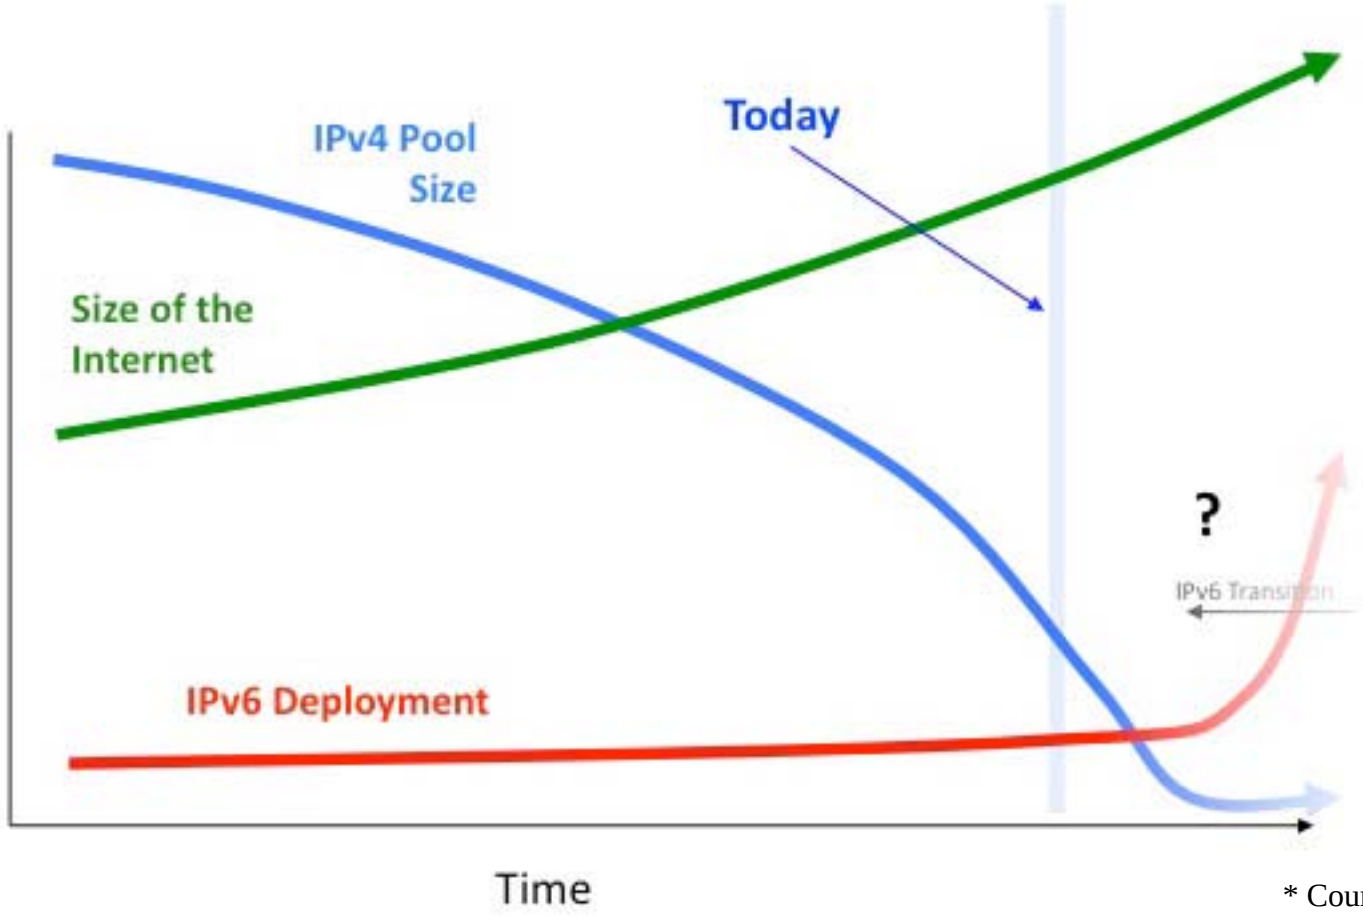

* Courtesy Geoff Huston

So What Is The "New Internet Data Plane"

- Dual stack as a transition mechanism has been effectively abandoned
  - draft-arkko-townsley-coexistence-00.txt
  - Couple that with "disappointing" IPv6 uptake
- So what do we see emerging?
  - Carrier Grade NAT-PT (really big double/triple NAT + PT)
  - A+P (Steal some bits from the port)
  - Dual-stack Lite (CGN + tunneling)
  - IPv6 edge with IPv4 core (664)
    - Does this work?
    - ipv6.google.com or "Google over IPv6"
  - IPv4 edge with IPv6 core (446)
  - NAT66
  - …
- 7 meetings of the BEHAVE WG at the last IETF

Crystal Balls and the Like

- Carrier Grade NAT will be deployed
  - Note "synergy" here
- Dual-Stack Lite will be standardized
  - Recently added to the softwire WG charter as a work item
- A+P will be picked up by those who "dislike" CGN
  - Some indication that A+P might have been abandoned
- Potential for IPv6 to be confined to the edge
  - Would this be fatal for IPv6?
  - i.e., can IPv6 survive as a purely edge technology?
  - Does an IPv6 edge with an IPv4 core solve any problem?
    - Still have the "who can I talk to with IPv6" problem
  - Second system syndrome?
- So what are we left with?

Crystal Balls and the Like

- Ok, Carrier Grade NAT (et al) will be deployed
  - Get used to it
  - And what is the effect of CGN deployment on the viability of IPv6?
  - CGN and IPv6 deployments are not independent

- But what does it mean?
  - Escalating cost
    - At the very least the carrier now bears the support/customization costs on a per-customer basis
  - Escalating application complexity and fragility
  - Reduced flexibility
  - Increased risks of failure
  - → End-to-end fragile, if at all

Crystal Balls and the Like

- And as Van pointed out NANOG/ARIN LA meeting
  - NAT couples inside address/port to outside address/port in a many:1 relationship
  - It is just this kind of coupling that causes complexity (and hence cost, fragility)
    - http://www.1-4-5.net/~dmm/talks/NANOG26/complexity_panel
    - Or RFC 3439 (Some Internet Architectural Guidelines and Philosophy)
- This will induce an structure on the SP industry
  - In the same way that super-linear growth of control plane rate*state does
- And where can we look for solutions?
  - Think the IETF or the RRG is the right place?

If a Picture is Worth 2^10 Words…

So What's Next?

- For SPs, the Internet is about to become a lot more expensive to deploy and operate
- For Users, the Internet is about to become a lot more expensive and a lot less reliable
  - And a lot more balkanized
- So now the question isn't IPv4 vs. IPv6
  - Or what the IPv6 value proposition is, or what is Metcalfe's Law for IPv6, or …

What Is The Question Now?

The question now is how can we transition from the heavily NAT-PT'ed world we are faced with to something that more closely resembles the end-to-end Internet we all know and love

More Questions Than Answers

- What's needed?
- Serious research into what we can do/deploy effectively in the near-to-medium term to combat the effects of CGN (in a scalable manner, of course)
- Serious research into what kind of Internet-scale data and control planes can be designed and importantly, deployed
- Coordinated/cooperative effort across a wide variety of disciplines

Q&A

Thanks!