renesys

The IPv6 Routing Table in 2010

James Cowie
Clint Hepner

NANOG49
14 June 2010

Overview

- Compare relationships found in today's IPv4 and IPv6 tables
  - Which IPv4 ASNs have set up camp in IPv6?
  - How does their interconnection compare to IPv4?
  - What kind of IPv6 prefixes are they advertising?
  - How's the general reachability at each length?
  - What's the "/24 of the IPv6 table" going to be?
  - Unsettling implications for likely rates of IPv6 growth in 2010-2011, …
    - … even as the world is faced with IPv4 free space exhaustion

What Took You So Long?

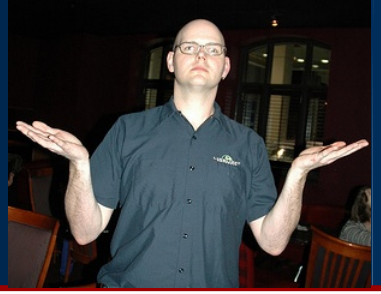

- We've been watching and waiting for ten years
- Only in the last two years have we heard rumbles of interest from customers
- Only in the last 6 months have we had enough offers of peering to start supporting analysis
- We're integrating 80+ live full-table feeds now
- We still don't have enough data to draw the kind of conclusions we'd like to about the IPv6 ecosystem .. That's worrisome.

Let's give it a shot anyway.

- Think about the implications of the fact that there are two tables (e.g., "two internets").
- To avoid disruptions, "most" of the business relationships that make up the Internet ecosystem will need to port, or die.
- The existing Internet grew organically, and laughs at global engineering solutions.
- Can we replicate the existing Internet, piece by piece, in IPv6? How far along are we?

Let's Discuss Reasonable Expectations

The IPv6 table **will** contain fewer prefixes

..but presumably all the same ASNs and relationships among ASNs, .. Right?

- Do we expect IPv6 to converge to substantially fewer ASN participants? Fewer multihomers?
  - This would be a significant reduction in the choices available to enterprises at the edge.
- Economic balance of power between Internet's edge and wholesale transit core unlikely to tip back, good engineering intentions notwithstanding.

Autonomous Systems

- IPv4: 34,500+ autonomous systems in use
  - Believed to represent at least 30,000 organizations
  - 14,000 (42%) are single-homed
  - 14,500 (43%) are actively dual-homed
  - 5,200 (15%) are tri-homed or more
  - Every one of these relationships may need porting.
- IPv6: just over 2,200 ASNs in use (~1:16)
  - 1,100 (**50%**) are single-homed
  - Paucity of data to study detailed preferences yet

Autonomous System Relationships

- 107,000+ edges in the IPv4 ASN graph
  - 63,000+ are transit (directed provider-customer)
  - 44,000+ more are probably peering
  - Each represents a negotiated/contractual relationship that needs to find its way onto the new IPv6 Internet
- Only 8,500 edges in the IPv6 ASN graph …
  - ~1:13 compared to the IPv4 graph, but …
  - ~2,200 of these are existing transit edges in IPv4
  - ~3,800 are existing peering edges in IPv4
  - **~2,400 (29%) are novel edges (not in IPv4)**

Implications

- **29%** of the relationships in IPv6 table are new (2,400 not seen in the existing IPv4 Internet)?!
- If real, this would imply a serious discontinuity in transit preferences between IPv4 and IPv6
- Are people really unable to get IPv6 satisfaction from their current provider mixture?
- Or are these edges just experiments?
- Let's see what kind of ASNs added adjacencies.

IPv6 Driving New Relationships Among ASNs

- 56% of all ASNs in the IPv6 table had to form at least **one new relationship** to get there.
- 1,277 ASNs forming 2,400 new relationships
  - Hurricane Electric 6939 has 400+ new adjacencies
  - Tinet 3257 has 130+ new adjacencies
  - 50+ other ASNs have at least 10 new adjacencies
  - 500+ have 2-10 new adjacencies
  - 700+ have exactly one new adjacency
- Classic success-breeds-success growth pattern

Who added these adjacencies?

- ~90 (7%) of ASNs with new adjacencies were not previously seen in IPv4 table at all
  - Typically ASNs created or used for IPv6 alone
  - Or mistakes ☺
- 220 (17%) single-homed under IPv4
- 350 (28%) dual-homed under IPv4
- **580 (46%) triple-homed or better** under IPv4
- In other words, highly-connected sophisticates. Not your average IPv4 ASN.

Where's the IPv4 edge in this transition?

- Of the ~34,500 IPv4 ASNs, about 29,000 are "edge" (no ASN customers) and about half of those are multihomed.
- But of the 2,200 ASNs originating IPv6 prefixes:
  - ~1,200 (54%) are IPv4 **provider** ASNs (16% of IPv4)
  - ~950 (46%) are IPv4 **edge** ASNs (84% of ipv4)
  - This is badly inverted **-- the edge is slow to join**
  - Are we missing some incentive for non-provider enterprises to join the IPv6 party?

The IPv4 “Solar System”

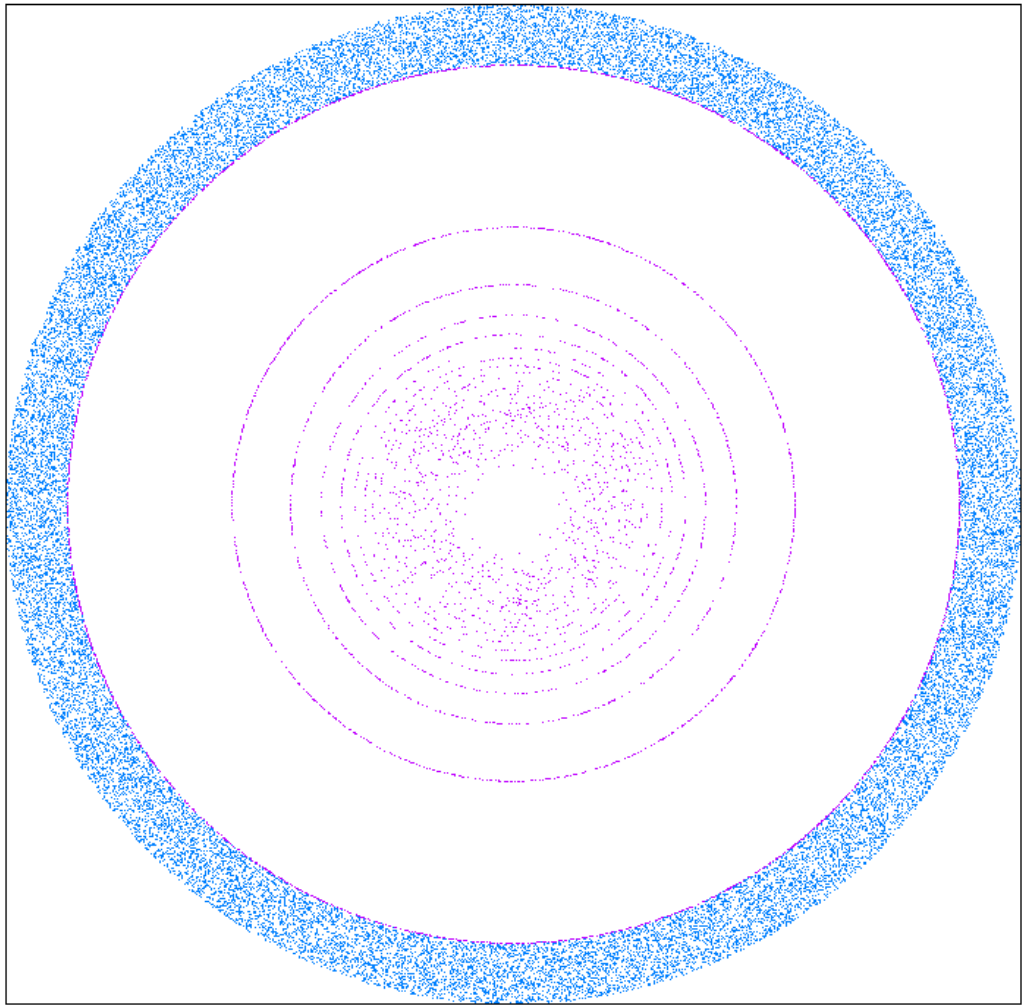

- Purple dots (providers) have downstreams, blue dots do not.
- Innermost providers have thousands of downstream ASNs.
- You’re looking at nearly 40,000 autonomous systems, most in the “Oort Cloud” of enterprise space

The IPv6 “Solar System”

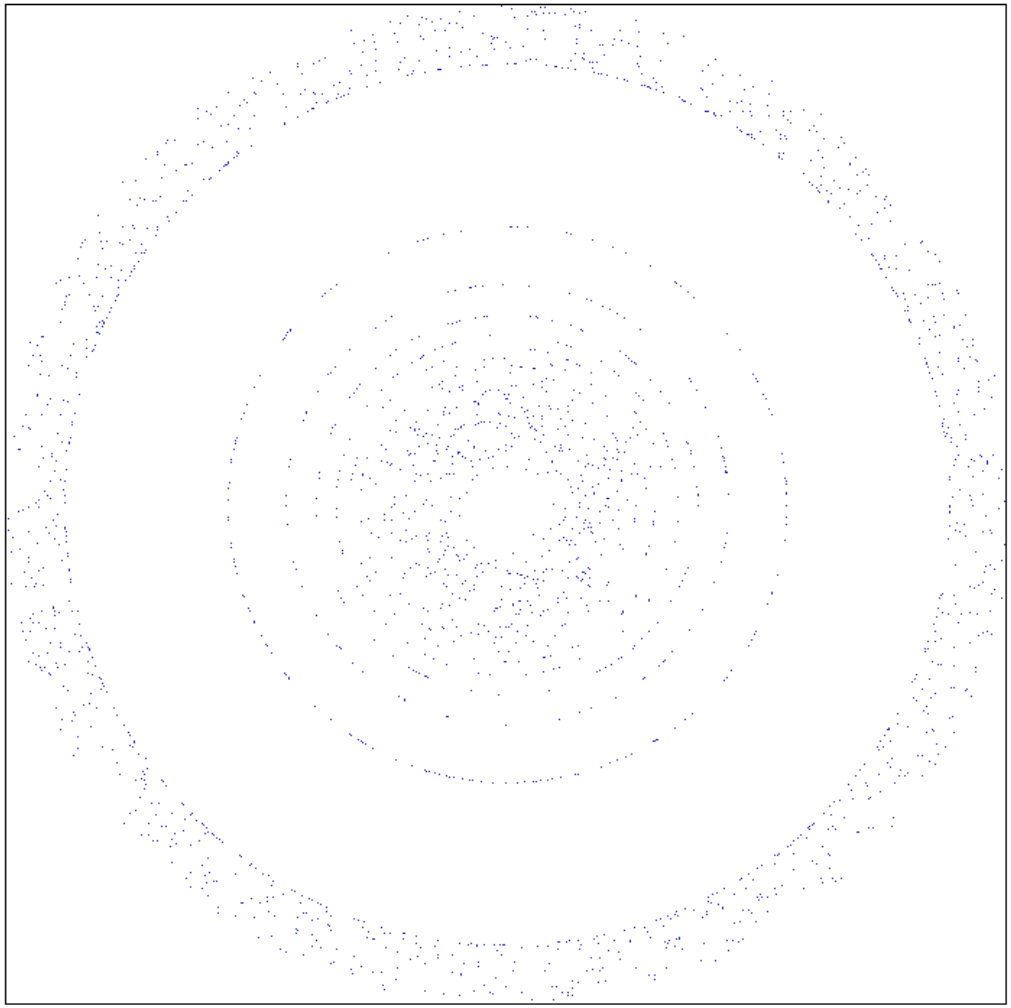

- ASNs have the same placement here as in the IPv4 solar system
- Note good density in the core, sparse coverage at the edge
- Only about 1 in 16 has made the leap to IPv6 space
- All of the easily converted, have been.

Combined View

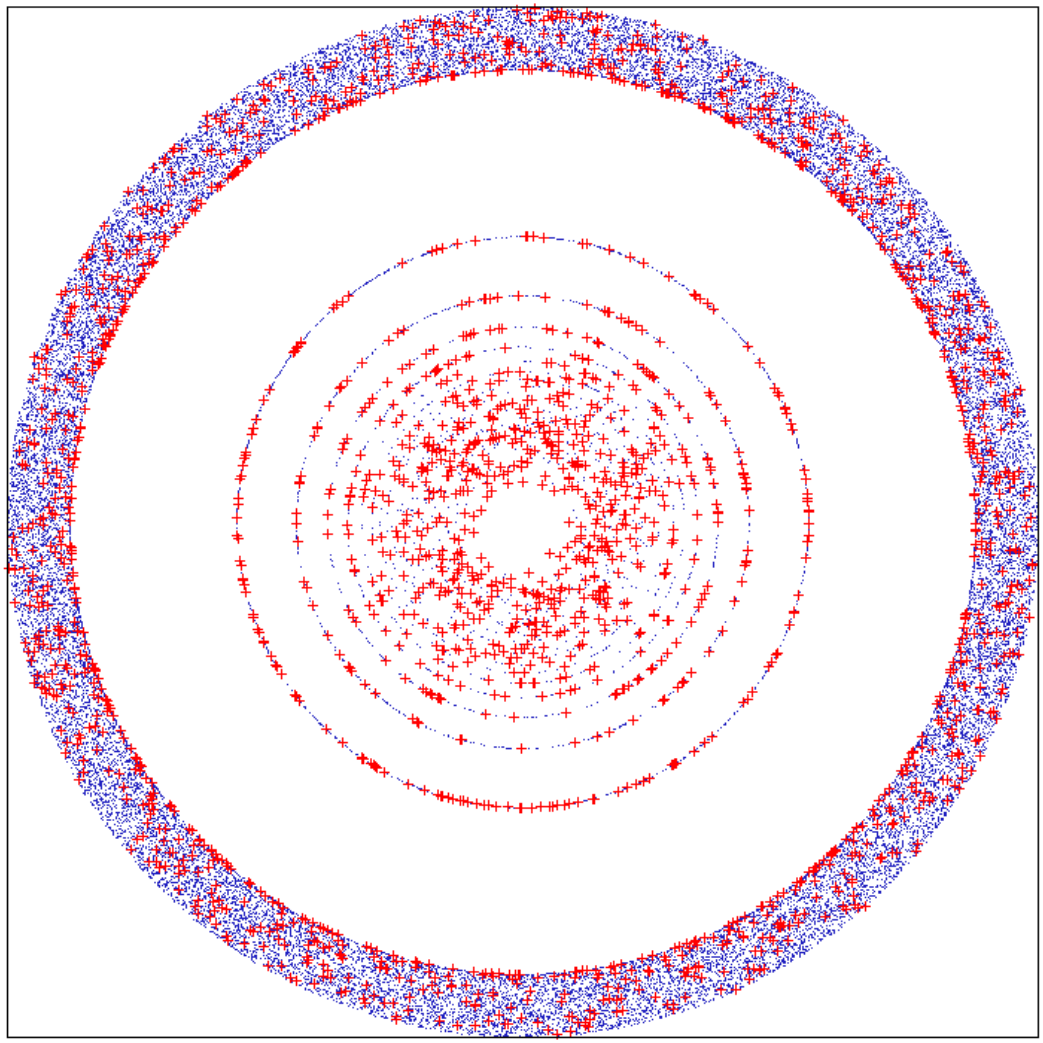

- **32,000+** teeny tiny blue dots: IPv4 ASNs who have made no move to join the IPv6 alternative Internet
- **2,200+** much larger red crosses, magnified for visibility: ASNs visible in at least one IPv6 route from at least one peer

IPv6 ASNs: Participation by Continent

	Top20	Top100	Top1000	Overall		Total IPv6	Total IPv4
North America	90%	45%	27%	**6%**		867	14996
South America	70%	42%	14%	**12%**		142	1144
Europe	90%	70%	42%	**9%**		1260	14638
Asia	95%	58%	26%	**11%**		489	4477
Africa	85%	49%	12%	**21%**		117	545
Australia	85%	57%	21%	**18%**		213	1175
Earth	**100%**	**70%**	**41%**	**7%**		2280	34534

A quick look at IPv6 Routes

- About **3,500** prefixes total (counting generously*).
  - 45 shorter-than-/32s (1%)
  - **1,825 /32 (51%)**
  - 315 /33-/47 (9%)
  - **910 /48 (25%)**
  - 40 /49-63 (2%)
  - 255 /64 (7%)
  - 190 longer-than-/65 (5%)

** At 90%+ visibility, full table only ~2,500 routes*

The Critical Question: Visibility

- In the IPv4 routing table, rough consensus allows global visibility between /8 and /24.
- You'll often see things </8 and >/24 in "full tables" .. But they are not globally propagated.
- What are the equivalent bounds for IPv6 shaping up to be?
- Practical implications for edge systems and others who expect olde-school multihoming and portable addressing to keep working
- Also those who trade in routing table size futures

Global Reachability

- On the Internet, it can be rather hard to say when a prefix is Globally Reachable
- We use various common-sense rules of thumb involving number of peers who will offer the route as part of a “full table” (again, whatever that means).
- Try this relatively generous definition on for size:
  - *A globally reachable prefix has routes known to at least 90% of the surveyed peers.*

Global Reachability in IPv4

- After removing redundant more-specific routes..
  - 86% of all IPv4 /24s are seen by 90%+ of peers
  - 81% of all IPv4 /22s are seen by 90%+ of peers
  - 87% of all IPv4 /18s are seen by 90%+ of peers
  - **96%** of all IPv4 /16s are seen by 90%+ of peers
- This is a byproduct of our culture: rough consensus among operators about what should be Generally Visible and readvertised.
- *This is what we want to see in IPv6 as well…*

Global Reachability in IPv6

- If we condition on full tables (peers with at least 1500 routes of the 3500 known)…
- 92% of /32s are seen by 90%+ of peers.
- 73% of /40s are seen by 90%+ of peers.
- 67% of /48s are seen by 90%+ of peers.
- 71% of all prefixes are seen by 90%+ of peers.
- **The edge enterprises are even more visible by their absence here. Presumably the majority of IPv4 non-provider ASNs will manifest as a portable /48.**
- **Where are all the portable /48s?**

Good Reachability in PI space ….

- 408 PI allocations visible by 50%+ of peers
- **94%** aggregate reachability across all lengths
- If only there were more of them ….

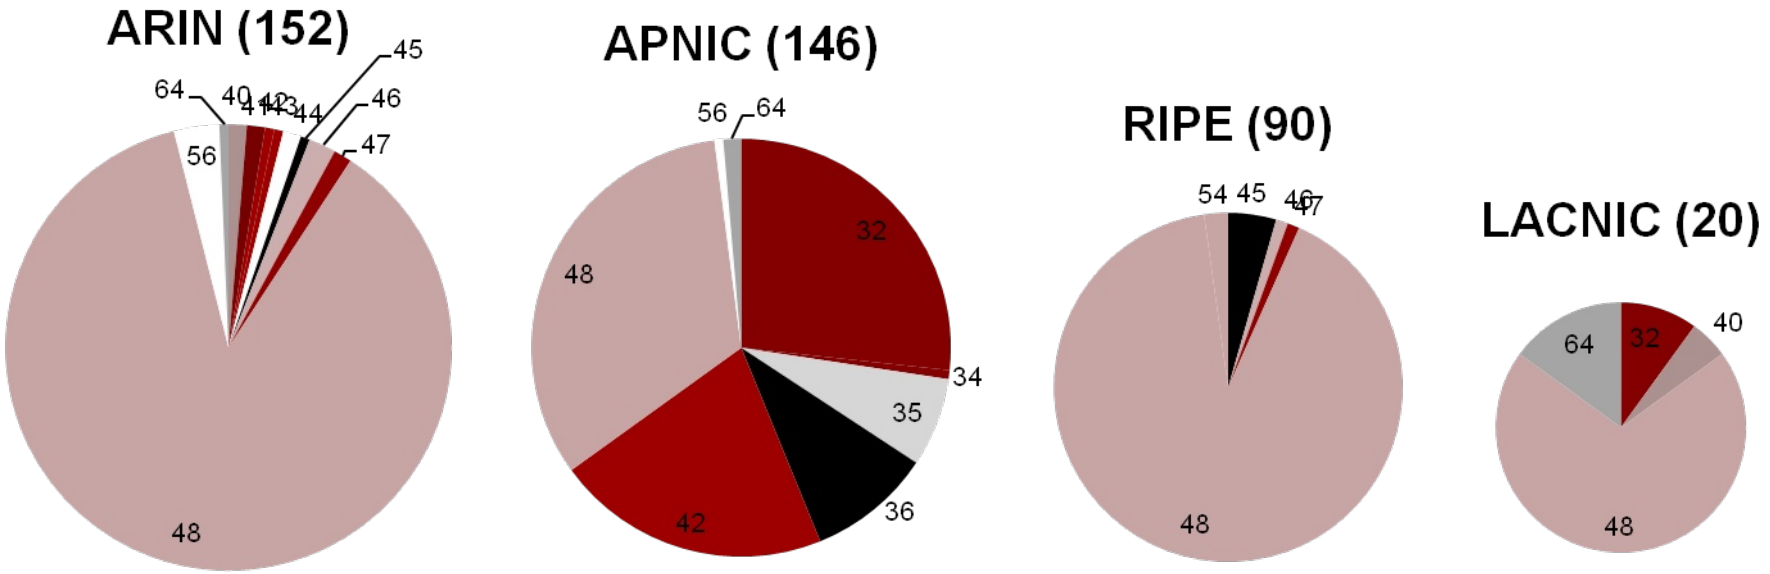

Summary: Go Recruit The Edge

- Out of the 107,000 peer-peer and provider-customer relationships that make up the global IPv4 ecosystem, only about **6%** (appx. 8500) have materialized in the IPv6 table so far.
- These edges are biased towards replication of the IPv4 core – highly multihomed providers adding IPv6
- Non-provider enterprises are staying away in droves (939 of 28,944, just **3%**, even originate a single IPv6 prefix). Multihomed enterprises: **4%**.

The core is leading and the edge is not following.

Summary: Changes in the Core

- Fully **29%** of the relationships in the IPv6 table are new adjacencies, perhaps implying that providers are failing to meet existing customers' demand for IPv6 services.
- **Very similar to NANOG meeting attendance statistics!**
- We appear to still be in the early days of a migration of IPv4's rich ecosystem of contractual relationships.
- PEER WITH US and we'll keep watching this evolution

Thank you!

http://www.renesys.com/tech/peering.shtml

IPv6 Peering Inquiries: peering@renesys.com