Scaling issues with ipv6 routing+multihoming

Vince Fuller, Cisco Systems

Why doesn't ipv6 (or IPv4) routing scale?

- It's all about the schizophrenic nature of addresses
  - they need to be "locators" for routing information
  - but also serve as "endpoint id's" for the transport layer
- For routing to scale, locators need to be assigned according to topology and change as topology changes ("*Addressing can follow topology or topology can follow addressing; choose one" – Y. Rekhter)*
- But as identifiers, assignment is along organizational hierarchy and stability is needed – users and applications don't want renumbering when network attachment points change
- A single numbering space cannot serve both of these needs in a scalable way (see "further reading" section for a more in depth discussion of this)
- The really scary thing is that the scaling problem won't become obvious until (and if) ipv6 becomes widely-deployed

What if we do nothing? Assume & project

- ipv6 widely deployed in parallel with IPv4
  - Need to carry global state for both indefinitely
- Multihoming trends continue unchanged (valid?)
- ipv6 does IPv4-like mulithoming/traffic engineering
  - “PI” prefixes, no significant uptake of shim6
- Infer ipv6 table size from existing IPv4 deployment
  - One ipv6 prefix per ASN
  - One ipv6 more-specific per observed IPv4 more-specific
- Project historic growth trends forward
- Caveat: lots of scenarios for additional growth

Geoff's BGP report projections

- How bad are the growth trends? Geoff's BGP reports show:
  - Prefixes: 130K to 170K in 2005 (200K as of 10/2006)
    - projected increase to ~370K within 5 years
    - global routes only – each SP has additional internal routes
  - Churn: 0.7M/0.4M updates/withdrawals per day
    - projected increase to 2.8M/1.6M within 5 years
  - CPU use: 30% at 1.5Ghz (average) today
    - projected increase to 120% within 5 years
- These are guesses based on a limited view of the routing system and on low-confidence projections (cloudy crystal ball); the truth could be worse, especially for peak demands
- No attempt to consider higher overhead (i.e. SBGP/SoBGP)
- These kinda look exponential or quadratic; this is bad… (Tony will say more about why fixing it isn't just about buying more cheap memory)

Estimated IPv4+ipv6 Routing Table (Jason, 6/06)

Assume that tomorrow everyone does dual stack...	
Current IPv4 Internet routing table:	**180K routes**
New ipv6 routes (based on 1 prefix per AS):	**+ 21K routes**
Intentional de-aggregates for IPv4-style TE:	**+ 61K routes**
Internal IPv4 customer de-aggregates	**+ 50K to 150K routes**
Internal ipv6 customer de-aggregates	**+ 40K to 120K routes**
(projected from number IPv4 of customers)	
Total size of tier-1 ISP routing table	**352K to 532K routes**

These numbers exceed the FIB limits of a lot of currently-deployed equipment

Future Projection of Combined IPv4 and ipv6 Internet Growth

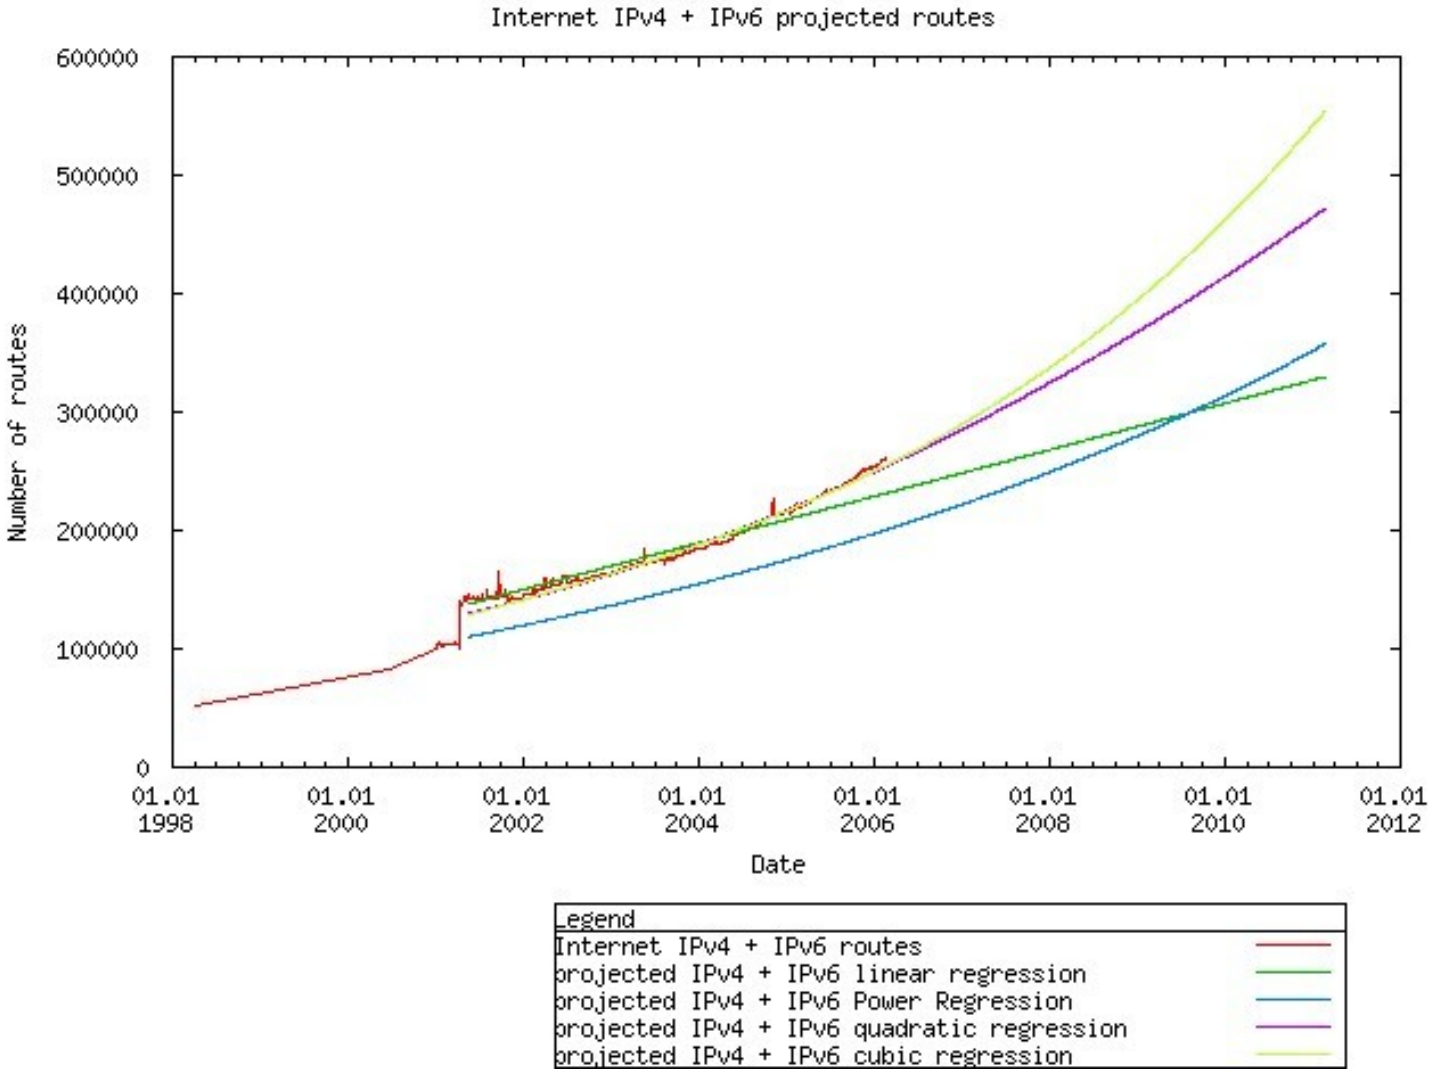

Summary of scary numbers

Route type	2006.06	5 years	7 years	10 Years	14 years
IPv4 Internet routes	180,219	285,064	338,567	427,300	492,269
IPv4 CIDR Aggregates	119,114				
IPv4 intentional de-aggregates	61,105	144,253	195,176	288,554	362,304
Active Ases	21,646	31,752	36,161	42,766	47,176
Projected ipv6 Internet routes	82,751	179,481	237,195	341,852	423,871
Total IPv4/ipv6 Internet routes	262,970	464,545	575,762	769,152	916,140
Internal IPv4 low number	48,845	88,853	117,296	173,422	219,916
Internal IPv4 high number	150,109	273,061	360,471	532,955	675,840
Projected internal ipv6 (low)	39,076	101,390	131,532	190,245	238,494
Projected internal ipv6 (high)	120,087	311,588	404,221	584,655	732,933
Total IPv4/ipv6 routes (low)	350,891	654,788	824,590	1,132,819	1,374,550
Total IPv4/ipv6 routes (high)	533,166	1,049,194	1,340,453	1,886,762	2,324,913

Recommended Reading

"The Long and Winding ROAD", a brief history of Internet routing and address evolution, http://rms46.vlsm.org/1/42.html

"Endpoints and Endpoint names: A Proposed Enhancement to the Internet Architecture", J. Noel Chiappa, 1999, http://users.exis.net/~jnc/tech/endpoints.txt

"On the Naming and Binding of Network Destinations", J. Saltzer, August, 1993, published as RFC1498, http://www.ietf.org/rfc/rfc1498.txt?number=1498

"The NIMROD Routing Architecture", I. Castineyra, N. Chiappa, M. Steenstrup. February 2006, published as RFC1992, http://www.ietf.org/rfc/rfc1992.txt?number=1992

"2005 – A BGP Year in Review", G. Huston, APRICOT 2006, http://www.apnic.net/meetings/21/docs/sigs/routing/routing-pr