IPv6 Potential Routing Table Size

Jason Schiller
schiller@uu.net
Sven Maduschke
sven.maduschke@verizonbusiness.com
IP Core Infrastructure Engineering
Verizon Business

Aggregation is Holy Grail

- IETF and ARIN recommendation is that aggregation is of the utmost importance for good IPv6 stewardship
- Must solve multi-homing, mobility, and provider independence without de-aggregation

Operator's Take on De-aggregation

- Lack of Provider Independent Addresses is preventing wide spread deployment and is leading to lack of IPv6 content
  - Even with stateless auto-config renumbering is difficult
  - Getting IP addresses from the up-stream ISP creates "provider lock-in"
  - ARIN members are pursuing ARIN policy 2005-1 and 2006-4
- Provider Independent (PI) space will add to the global routing table size
- PI space sets the precedent that de-aggregation is acceptable
  - De-aggregation may be used to solve other problems, multi-homing, mobility
  - De-aggregation of PI space will lead to de-aggregation of Provider Assigned (PA) space

Operator's Take on De-aggregation

- Shim6 is broken as a solution for large business customers
  - No transit AS TE
  - No inbound destination TE
  - Won't scale for content providers where end host (server) has 30,000 concurrent TCP sessions
  - Doesn't help for short lived traffic
  - Managed on the end host, and not in the network
    - End hosts managed by end users, not the owner of the network
    - Too many places to manage TE policy
- No good non-de-aggregation solution for multi-homing or Provider Independence
- Less then 1,000 IPv6 routes in the Internet routing table
- Less than 100 new IPv6 Internet routes a year
- 1,200 IPv6 Internet routes in two years will not be a problem
- Let's just de-aggregate

Long Term Commitment to IPv6 De-aggregation

If we decide to de-aggregate now, in the long term we commit to solving the routing table growth problem through hardware

- Are Service Provider Operators and their vendors looking at hardware capabilities and scaling functions at 5 or 10 years out?
- We have seen this problem already in IPv4
  - Do we want to repeat our mistakes?
  - Do we want to embark on a hardware / routing table scaling escalation?
- With a larger IPv6 address space the potential for growth is much higher

Impact of Routing Table Growth On Hardware

Extra routing state:

- Consumes routing memory (RIB)
- Consumes forwarding memory (FIB)
- Affects forwarding rate
  - (FIB lookup as a function of memory speed and size)
- Affects convergence
  - (SPF, RIB rewrite, RIB to FIB population)

Combating Routing Table Growth Long Term Through Hardware

- Commit to continuously scaling router memory size and speed to support very large RIB and FIB sizes
- Commit to continuously faster processors for SPF of larger tables
- Optimize FIB storage and SPF processes
- Hope hardware / software solution is available at least 5 years before wide spread adoption
- Use 5 years to depreciate and replace current hardware through normal refresh with new hardware capable of holding larger routing information
- Hope that newly deployed equipment will survive in the network for at least 5 years
- Hope that next generation of equipment will be ready in time, and will survive in the network for at least five years

IPv6 Address Size

- IPv4 has 2^32 IP addresses (4,294,967,296)
- IPv4 largest unicast Internet routable block /24 (16,777,184)
- Concerns about address exhaustion in some countries
- Use of Network Address Translation (NAT) to reduce consumption

- IPv6 has 2^128 IP addresses
- 64 bits reserved for host, 64 bits reserved for network
- IPv6 Unicast routable space 2000::/3 (2,305,843,009,213,693,952 /64s) (35,184,372,088,832 /48s)
- 137,439,215,616 times more IPv6 /64s than IPv4 /24s
- 2,097,152 times more IPv6 /48s than IPv4 /24s

Current IPv4 Route Classification

- Three basic types of IPv4 routes
  - Aggregates
  - De-aggregates from growth and assignment of a non-contiguous block
  - De-aggregates to perform traffic engineering

- Tony Bates CIDR report shows:

DatePrefixes	Prefixes	CIDR Agg
14-03-06	180,219	119,114

- Can assume that 61K intentional de-aggregates

Current IPv4/IPv6 Routing Table Size

- Assume that tomorrow everyone does dual stack

Current IPv4 Internet routing table	180 K routes
21K active ASes (1 IPv6 aggregate / AS)	+ 21 K routes
61K intentional IPv6 de-aggregates for traffic engineering (assuming IPv4 style TE)	+ 61 K routes
	262 K routes
Current tier 1 ISP internal routes	+50K to 150 K routes
	312K to 412 K routes
Internal IPv6 de-aggregates for customers (projected from number of customers)	+40K to 120 K routes
	352K to 532 K routes

Tier 1 ISPs require IP forwarding in hardware (6Mpps)

Easily exceed the current FIB limitations of some deployed routers

What This Interpolation Doesn't Account For

- A single AS that currently has multiple non-contiguous assignments that would still advertise the same number of prefixes to the Internet routing table if it had a single contiguous assignment
- All of the ASes that announce only a single /24 to the Internet routing table, but would announce more specifics if they were generally accepted (assume these customers get a /48 and up to /64 is generally accepted)
- All of the networks that hide behind multiple NAT addresses from multiple providers who change the NAT address for TE. With IPv6 and the removal of NAT, they may need a different TE mechanism.
- All of the new IPv6 only networks that may pop up: China, Cell phones, coffee makers, toasters, RFIDs, etc.

Projected IPv6 Routing Table Growth

- Let’s put aside the date when wide spread IPv6 adoption will occur
- Let’s assume that wide spread IPv6 adoption will occur at some point
- What is the projection of the of the current IPv4 growth
  - Internet routing table
  - International de-aggregates for TE in the Internet routing table
  - Number of Active ASes
- What is the IPv6 routing table size interpolated from the IPv4 growth projections assuming everyone is doing dual stack and IPv6 TE in the “traditional” IPv4 style?
- Add to this internal IPv4 de-aggregates and IPv6 internal de-aggregates
- Ask vendors and operators to plan to be at least five years ahead of the curve for the foreseeable future

verizon

Internet CIDR Information
Total Routes and Intentional de-aggregates

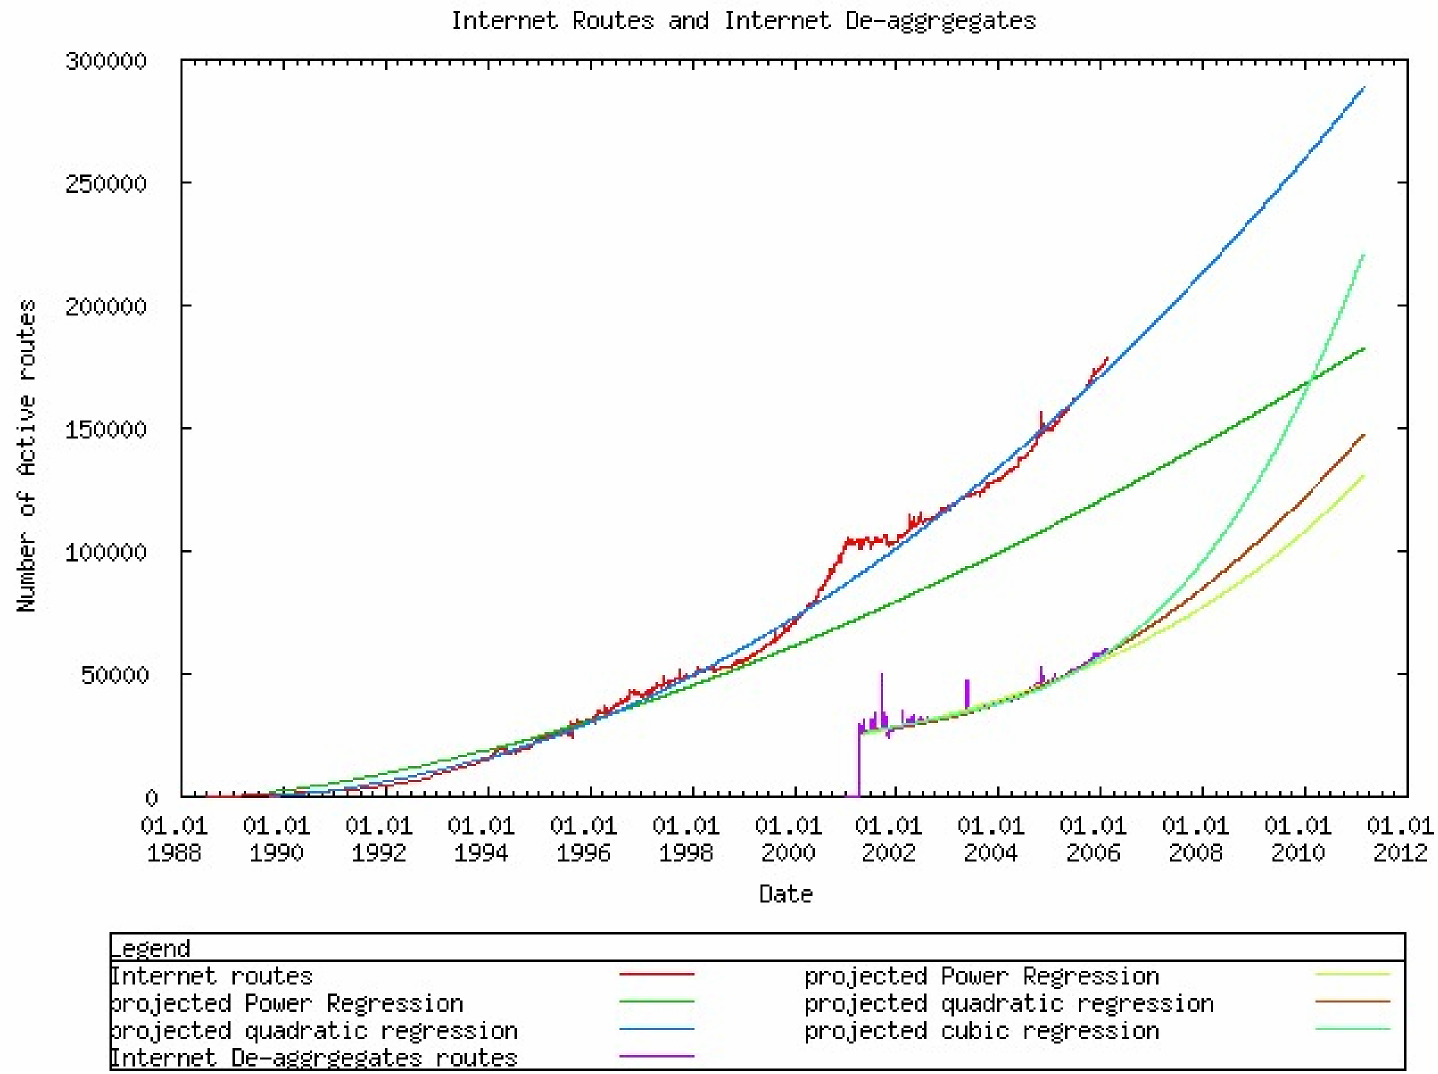

verizon

Internet CIDR Information
Active ASes

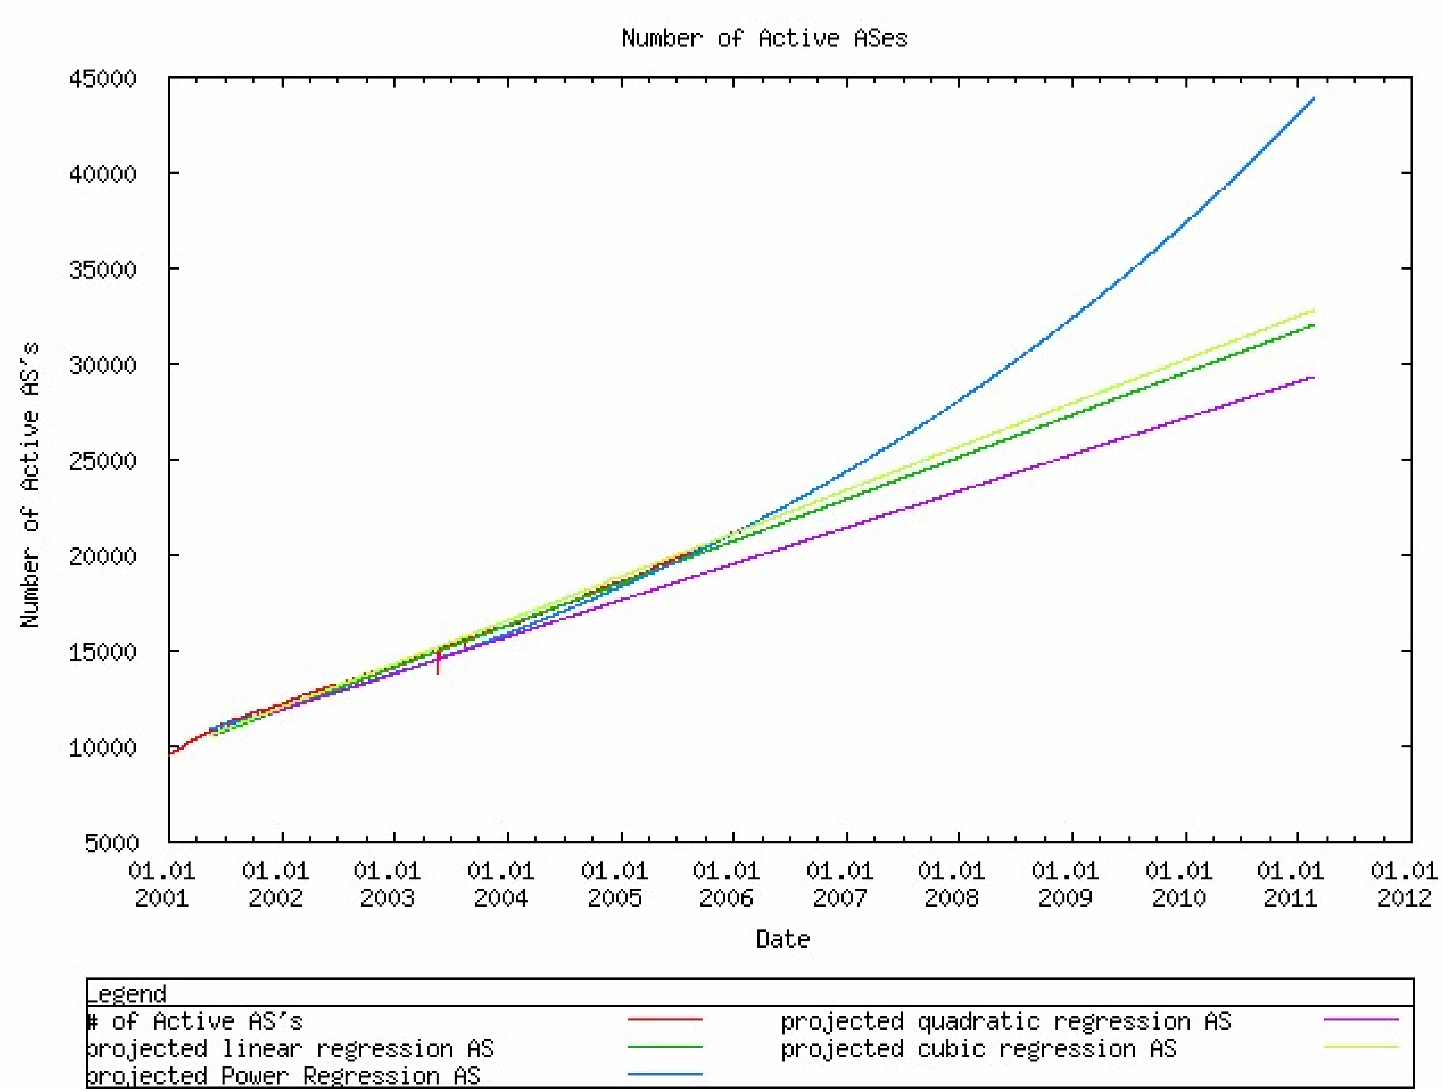

Future Projection of IPv6 Internet Growth (IPv4 Intentional De-aggregates + Active ASes)

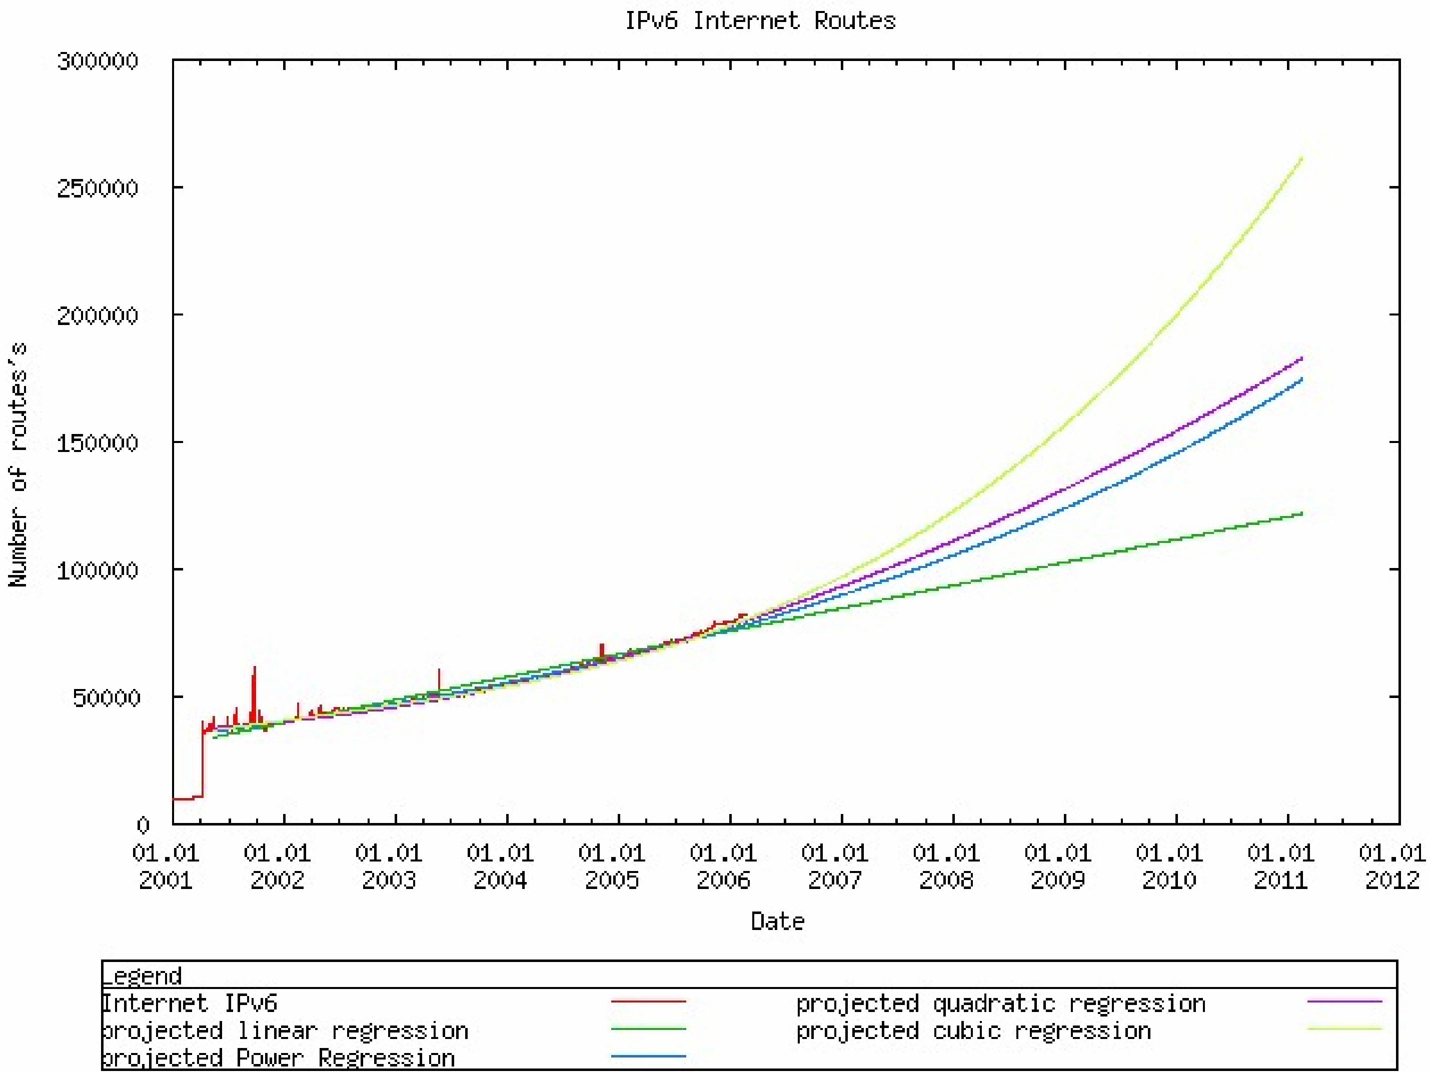

verizon

Future Projection of Combined IPv4 and IPv6 Internet Growth

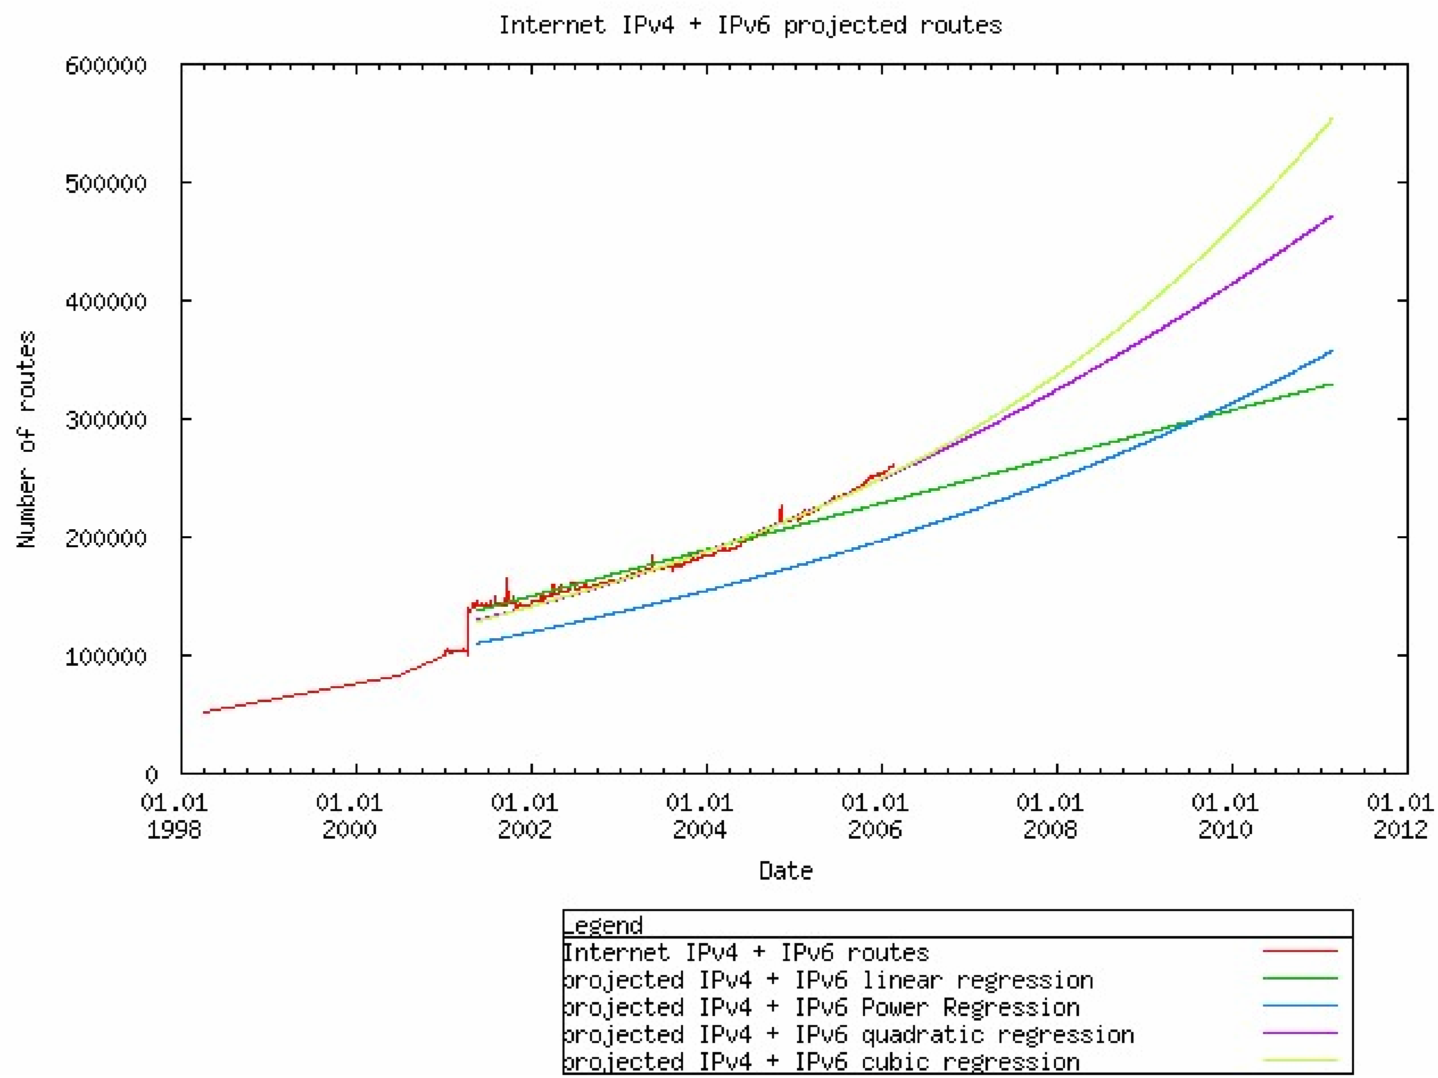

Tier 1 Service Provider
IPv4 Internal de-aggregates

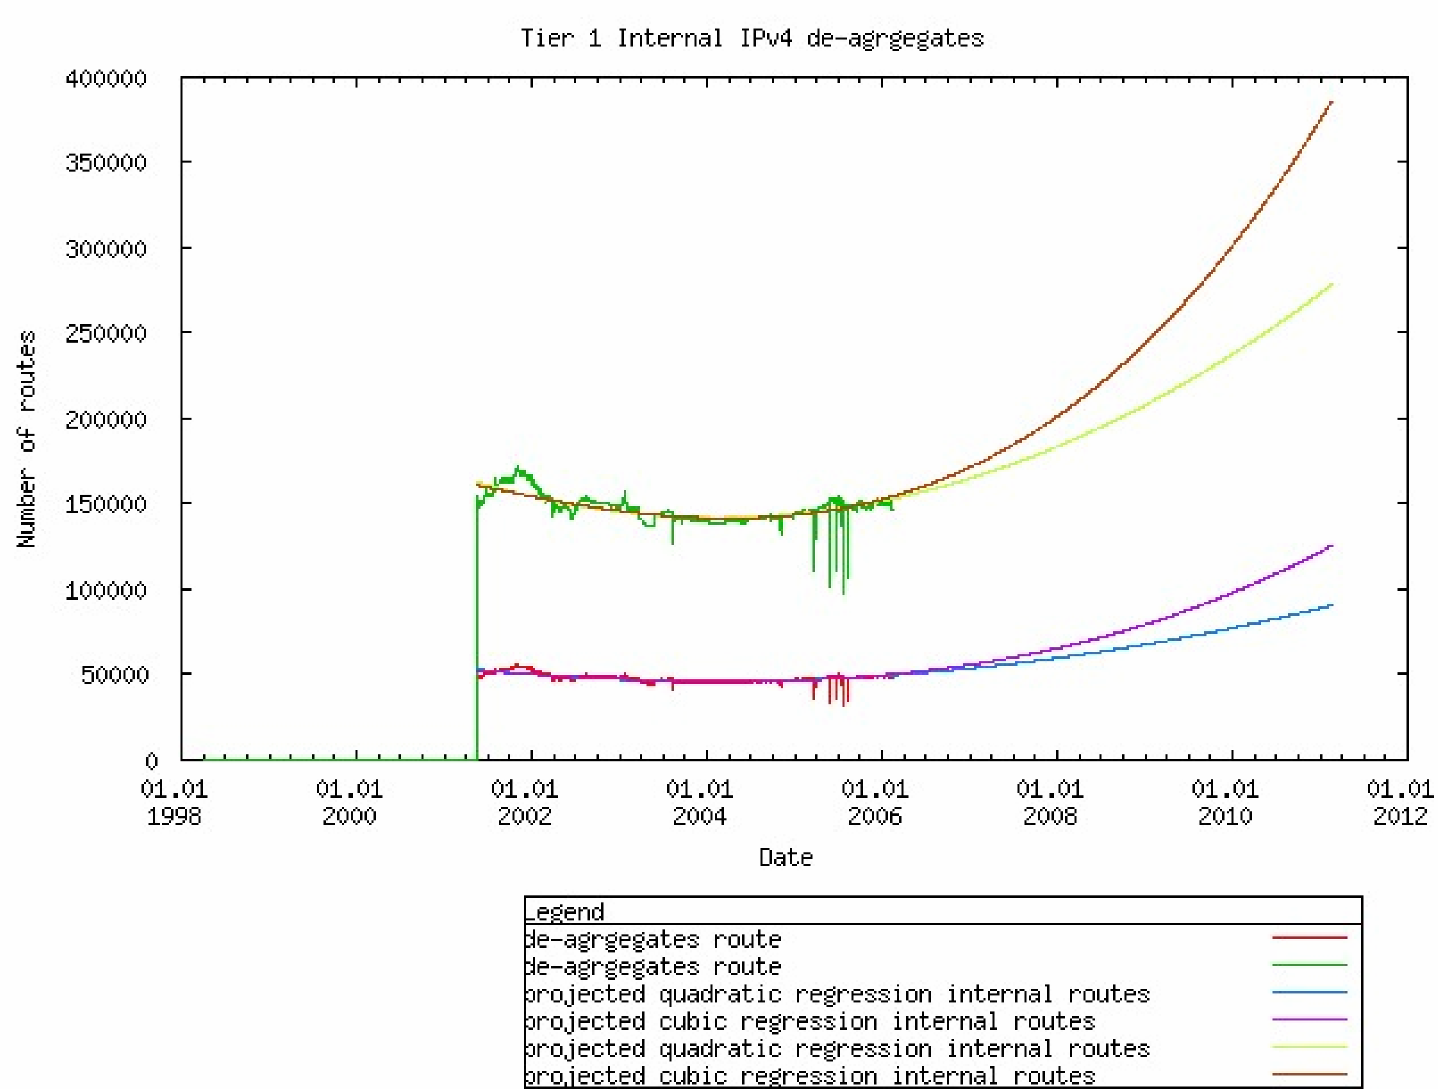

Future Projection Of Tier 1 Service Provider IPv4 and IPv6 Internal de-aggregates

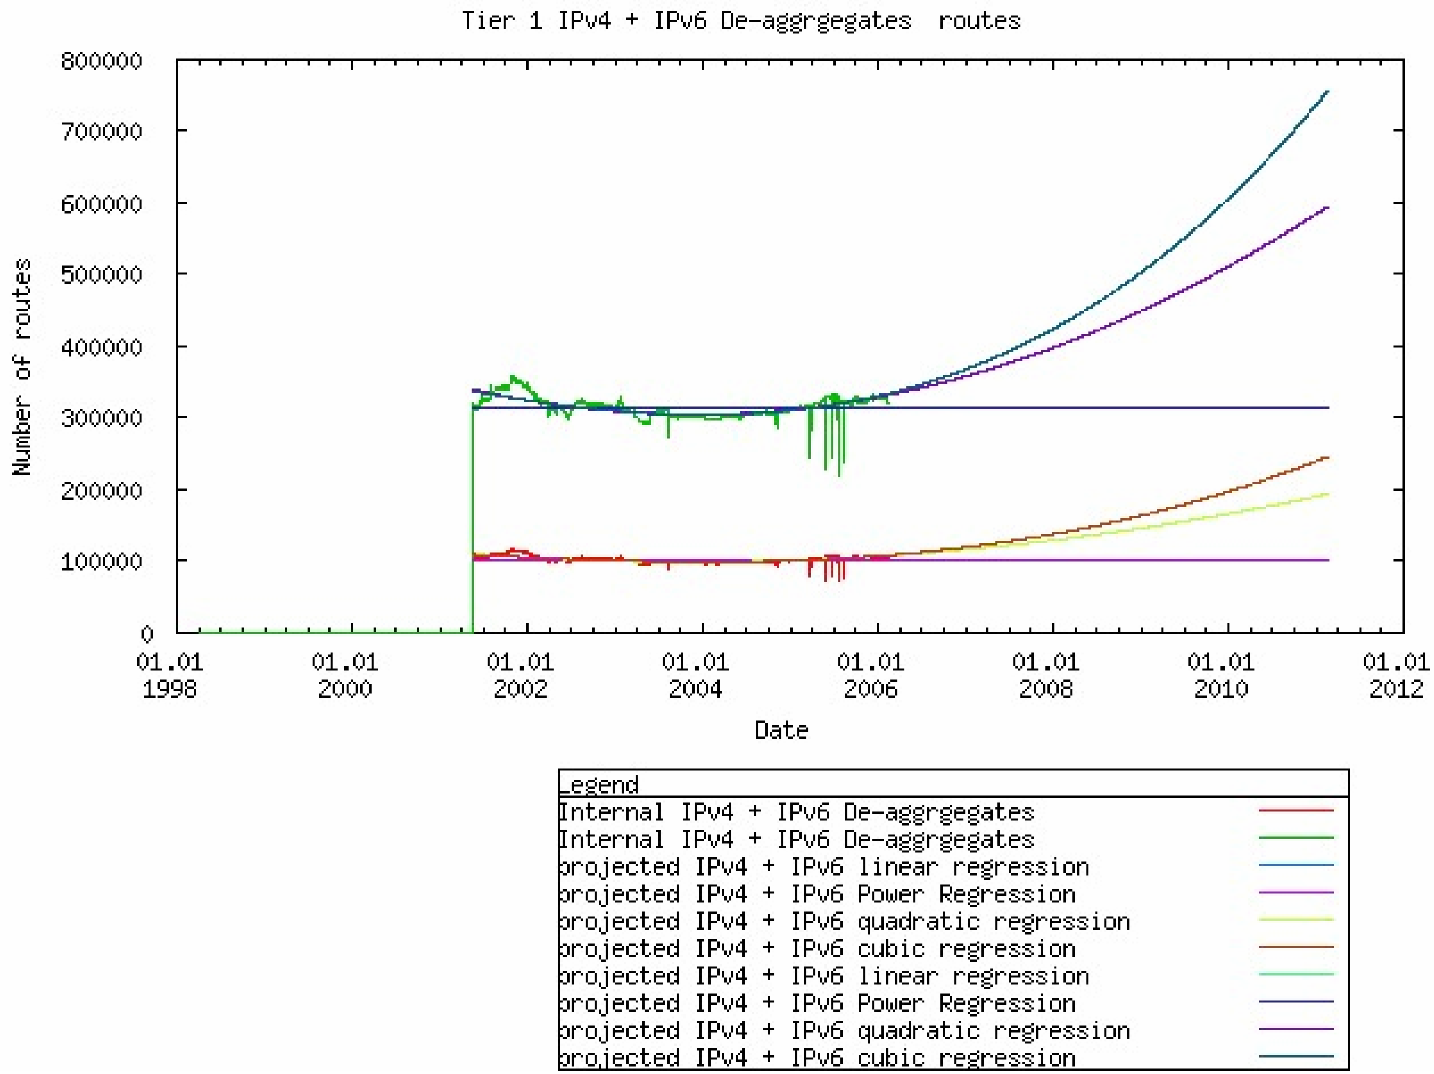

Future Projection Of Tier 1 Service Provider IPv4 and IPv6 Routing Table

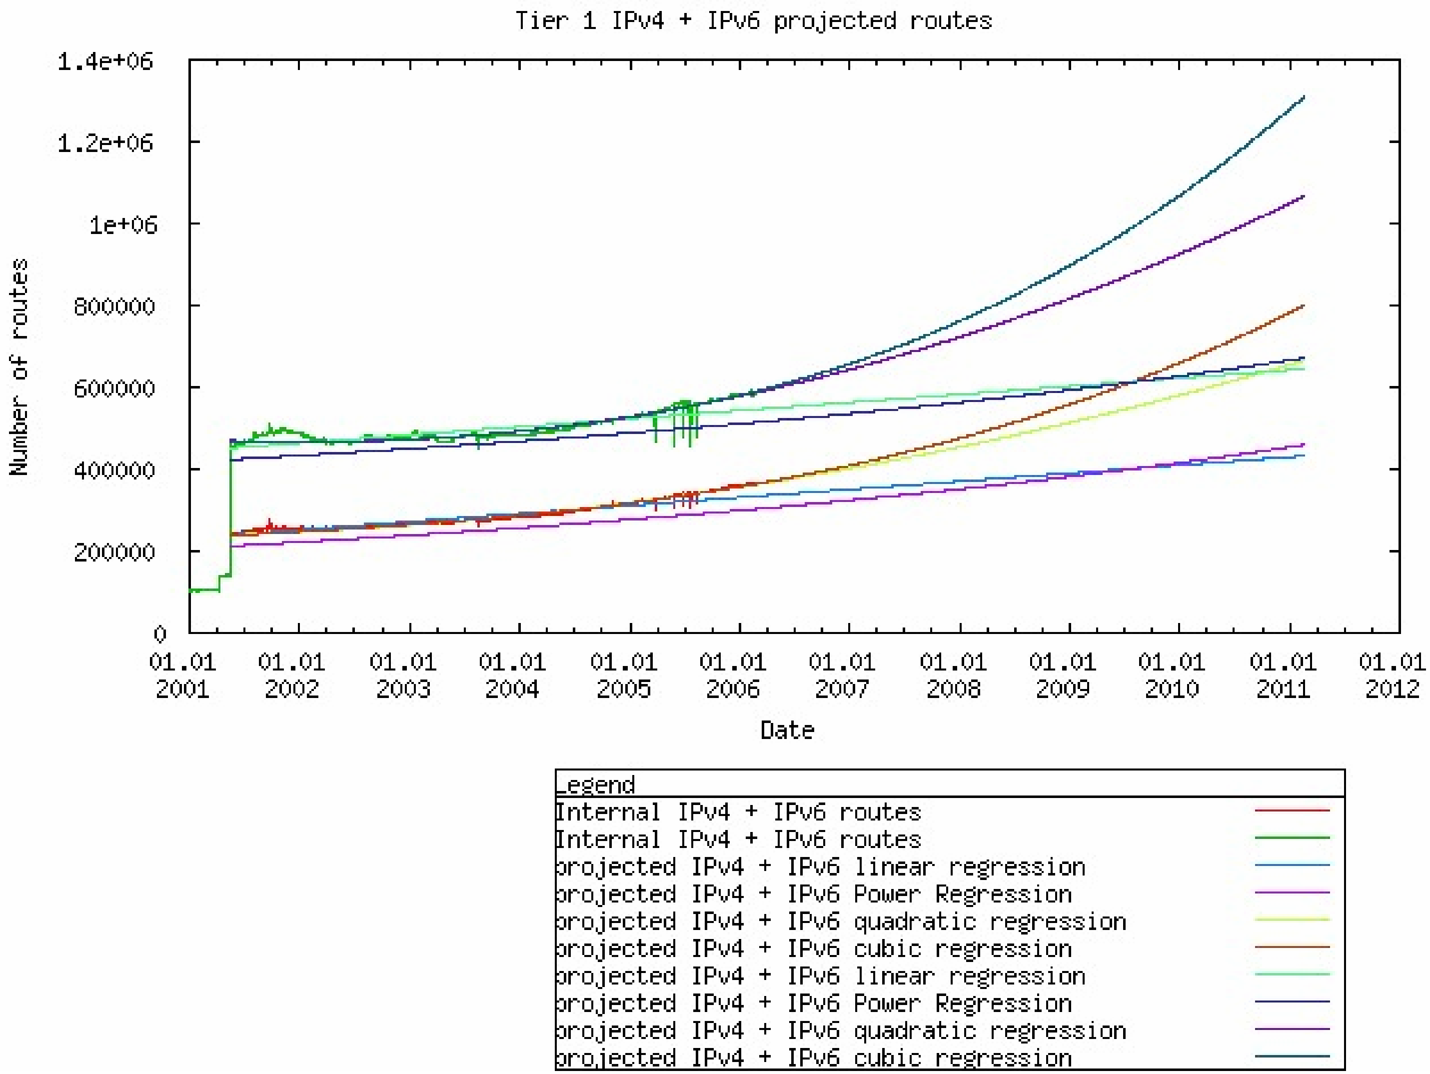

Conclusion

Route type	now	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	48,845	48,845	48,845	48,845
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

Conclusion

Current equipment purchases

- Assuming wide spread IPv6 adoption by 2011
- Assuming equipment purchased today should last in the network for 5 years
- All equipment purchased today should support 1M routes

Next generation equipment purchases

- Assuming wide spread IPv6 adoption by 2016
- Assuming equipment purchased in 2012 should last in the network for 5 years
- Vendors should be prepared to provide equipment that scales to 1.8M routes

Conclusion

- Can vendors plan to be at least five years ahead of the curve for the foreseeable future?
- How do operator certification and deployment plans lengthen the amount of time required to be ahead of the curve?
- Do we really want to embark on a routing table growth / hardware size escalation race for the foreseeable future? Will it be cost effective?
- Is it possible that routing table growth could be so rapid that operators will be required to start a new round of upgrades prior to finishing the current round?