Implications of Global IPv4/v6 Routing Table Growth

10/01/2007

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

Acknowledgements

- Noel Chiappa for his extensive writings over the years on ID/Locator split
- Mike O'Dell for developing GSE/8+8
- Geoff Huston for his ongoing global routing system analysis work (CIDR report, BGP report, etc.)
- Sven Maduschke for the growth projection section
- Tony Li for the information on hardware scaling
- Vince Fuller for drawing together information on Locator/ID split
- Vince Fuller for a brief history of IPv6
- Marshall Eubanks for finding and projecting the number of businesses (potential multi-homers) in the U.S. and the world

verizon

Aggregation is Holy Grail

- RFC-1887 (12/1995)
  - –RIRs to assign large amounts of IPv6 space to ISPs and transit providers (network providers)
  - –All down stream end-sites (network service subscribers) to use a more specific route of a provider aggregate
- 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

verizon

Operational Reality And De-aggregation

- There are various factors contributing to de-aggregation and growth of the routing table
  - Provider Independence
  - Multi-homing
  - Traffic Engineering
  - Acquisitions and mergers
  - Dual stack
  - Internal routes
  - Fragmentation from IPv4 exhaustion
  - IPv6 on everything

verizon

Operator's Take on Provider Independence

- Nearly all RIR communities have demanded and passed IPv6 PI policies
  - Lack of IPv6 is preventing wide spread IPv6 adoption
  - Even with stateless auto-config renumbering is difficult
  - Getting IP addresses from the up-stream ISP creates "provider lock-in"
  - End sites want to multi-home
- 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

verizon

Multi-homing

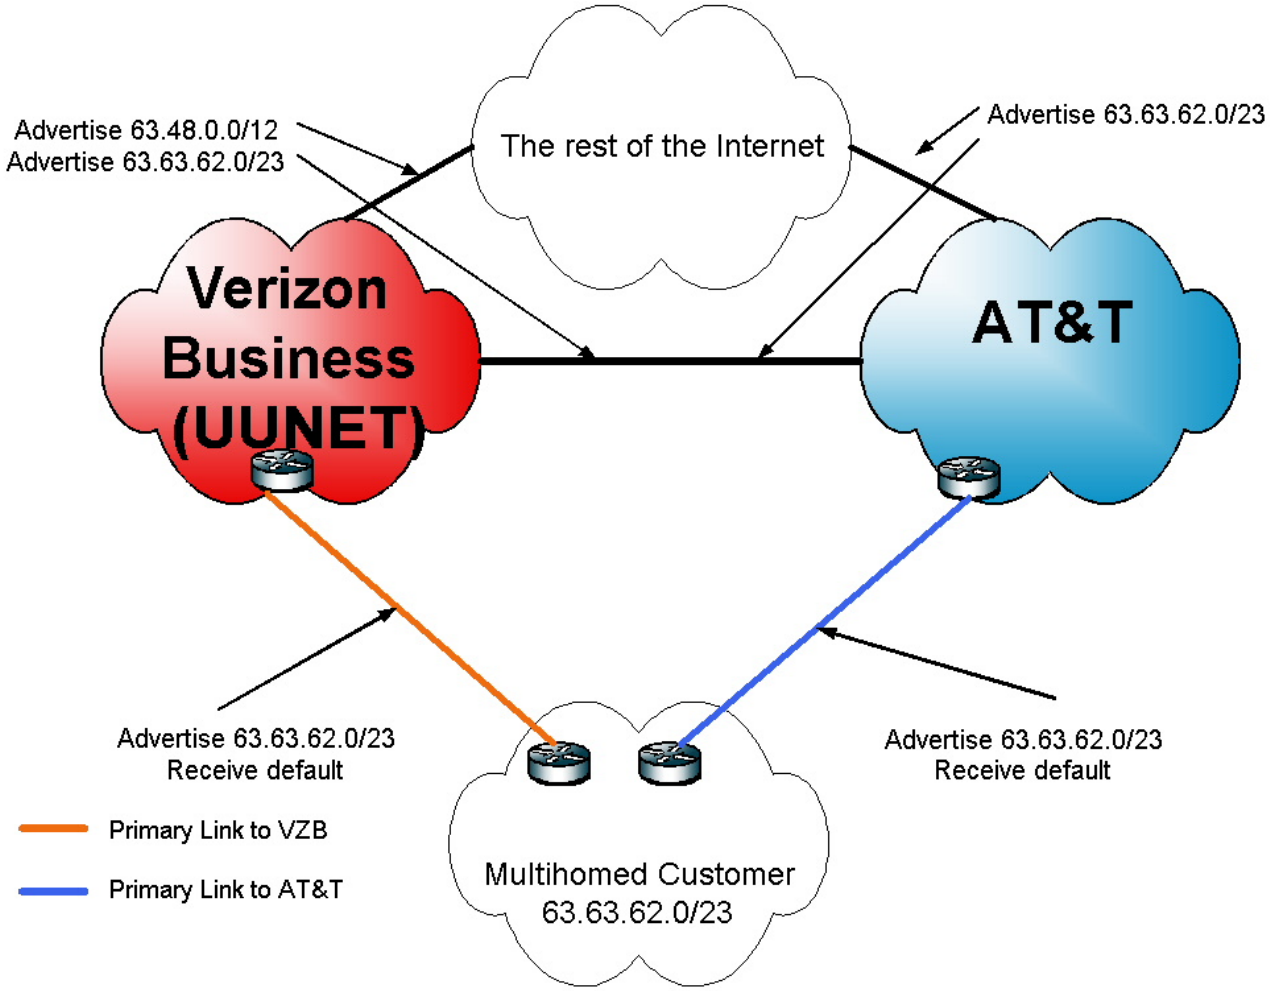

Multi-homing Traffic Engineering

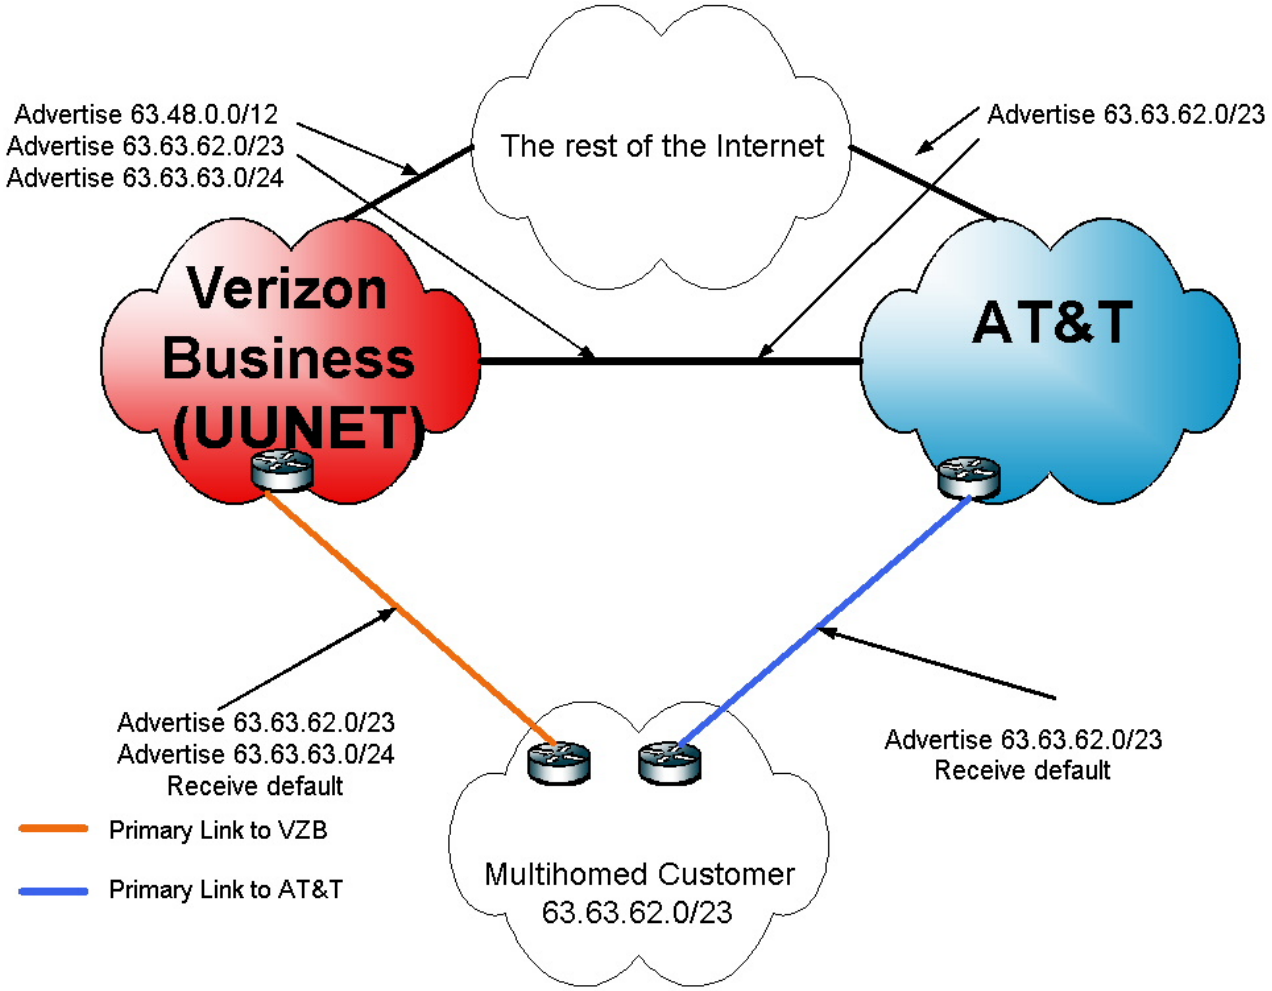

Operator's Conclusion on De-aggregation

- No good non-de-aggregation solution for multi-homing, TE 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

verizon

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)

verizon

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 BGP path selection of larger tables
- Optimize FIB storage and BGP path selection processes
- Hope hardware / software solution is available at least 5 years before wide spread IPv6 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
- … Rinse and repeat

IPv6 Goal: More IP Addresses!

- 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 (340,282,366,920,938,000,000,000,000,000,000,000,000 [340 undecillion])
- 64 bits reserved for host, 64 bits reserved for network
- IPv6 Unicast routable space 2000::/3 [1/8 of all IPv6 addresses] (2,305,843,009,213,693,952 /64s or 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

"Mission Accomplished"

verizon

Projected IPv6 Routing Table Growth

- Let's put aside the date when wide spread IPv6 adoption will occur
- How big would the routing table be if wide spread IPv6 adoption occurred now?
- 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

Assumptions:
Projected IPv6 Routing Table Growth

- Assume current IPv4 multi-homing trends occur in IPv6
- Assume dual stack IPv4 and IPv6 for the foreseeable future
- Assume current address consumption continues unchanged
  - Future growth is unbounded and could be far worse
    - Fragmentation due to IPv4 exhaustion
    - IPv6 on everything
- Ask vendors and operators to plan to be at least five years ahead of the curve for the foreseeable future

verizon

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
10-17-07	239,965	153,505

- Can assume that 86.4K intentional de-aggregates

verizon

Estimated IPv4+ipv6 Routing Table (10/07)

Assume that tomorrow everyone does dual stack...

Current IPv4 Internet routing table:	240K routes
New IPv6 routes (based on 1 prefix per AS):	+ 26K routes
Intentional de-aggregates for IPv4-style TE:	+ 86K routes
Internal IPv4 customer de-aggregates	+ 50K to 150K routes
Internal ipv6 customer de-aggregates (projected from number IPv4 of customers)	+ 40K to 120K routes
Total size of tier-1 ISP routing table	443K to 623K routes

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

verizon

What This Extrapolation 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 extra addresses from IPv4 fragmentation due to IPv4 exhaustion
- All of the new IPv6 only networks that may pop up: China, Cell phones, coffee makers, toasters, RFIDs, etc.
- VPN routes

verizon

Trend: Internet CIDR Information
Total Routes and Intentional de-aggregates

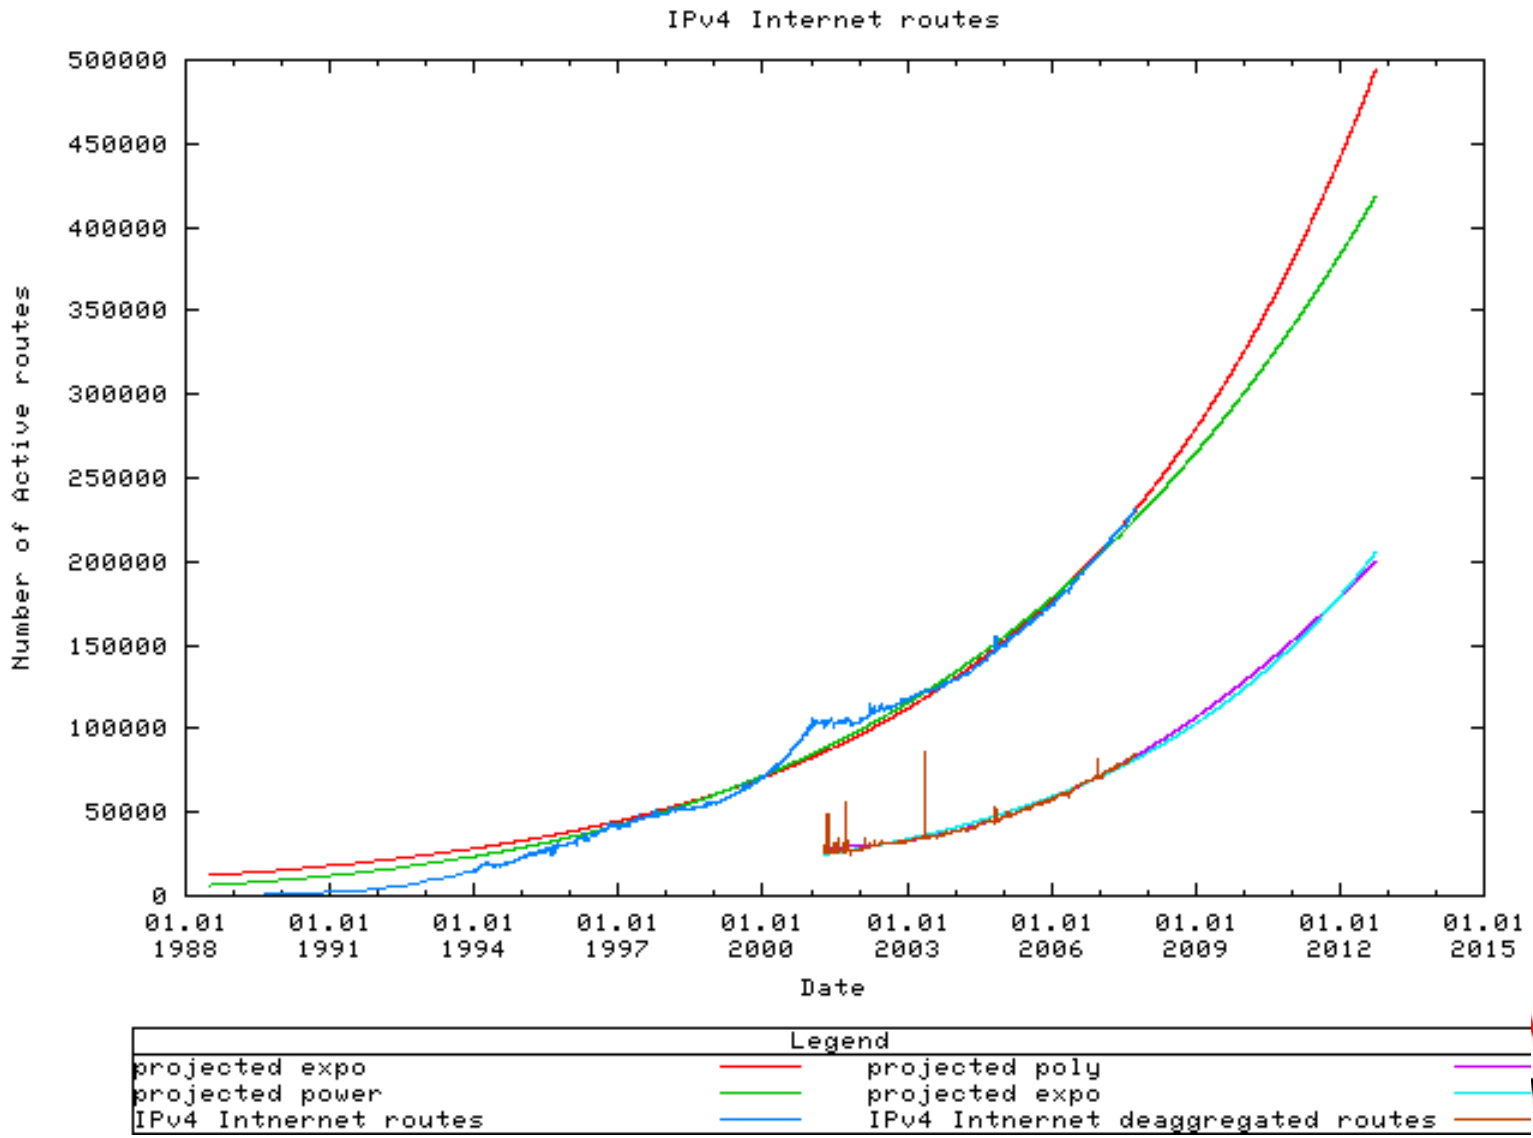

Trend: Internet CIDR Information
Active ASes

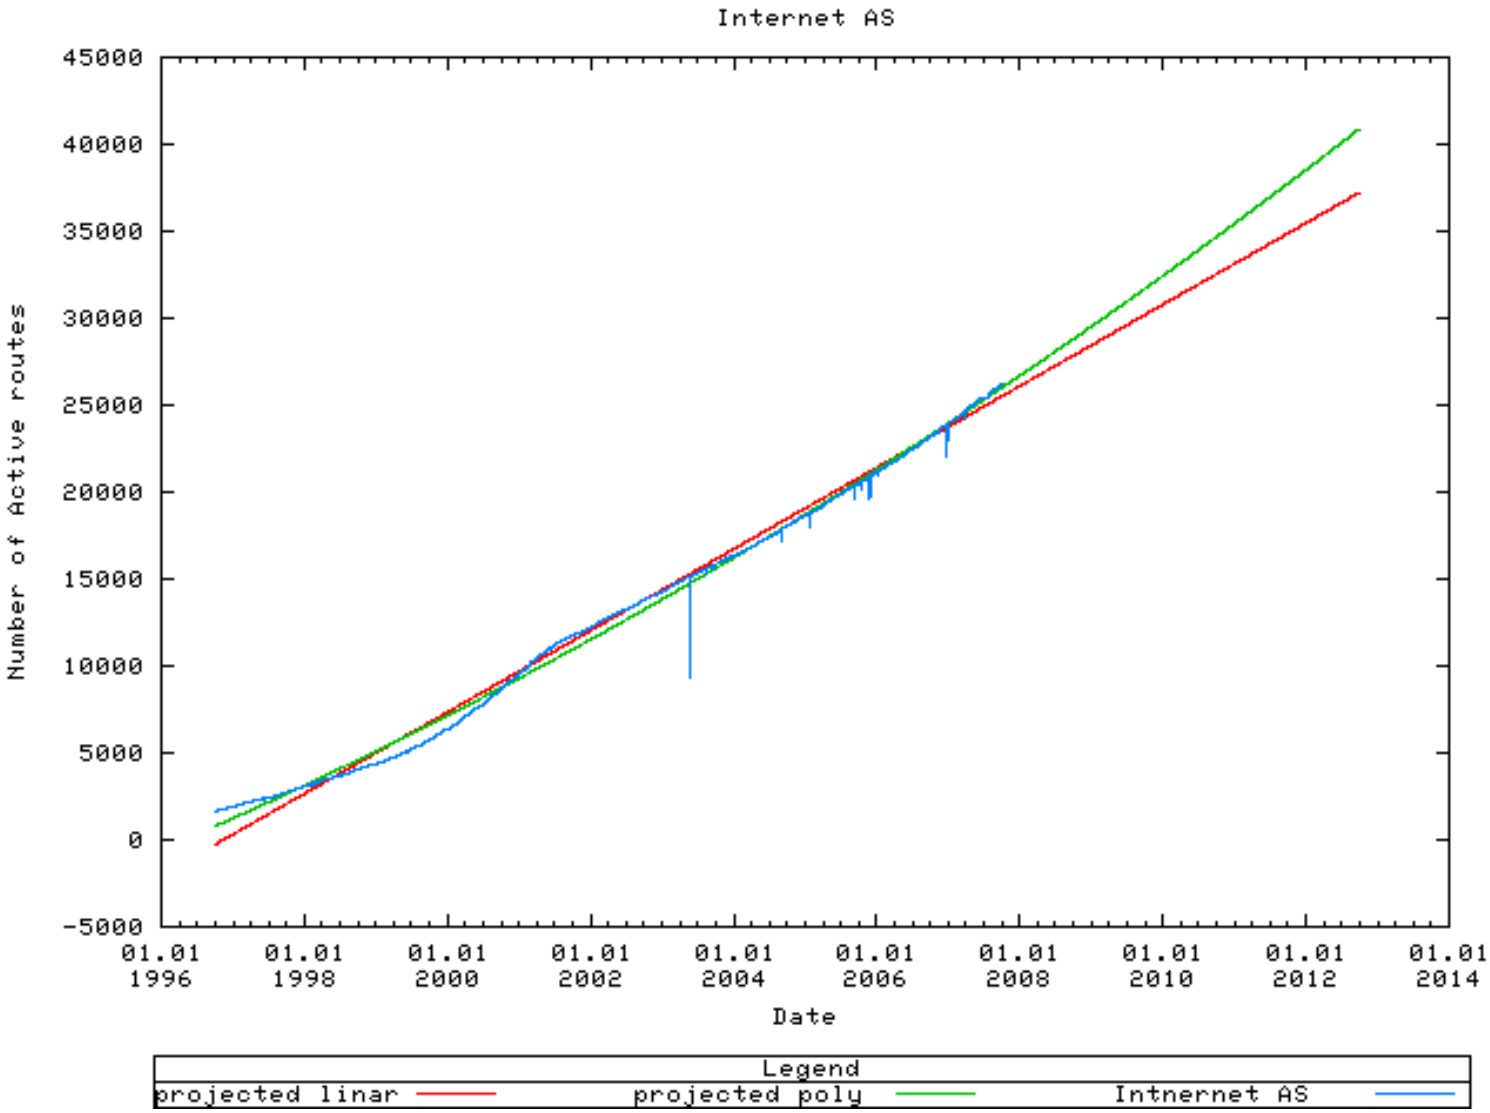

Future Projection of IPv6 Internet Growth
(IPv4 Intentional De-aggregates + Active ASes)

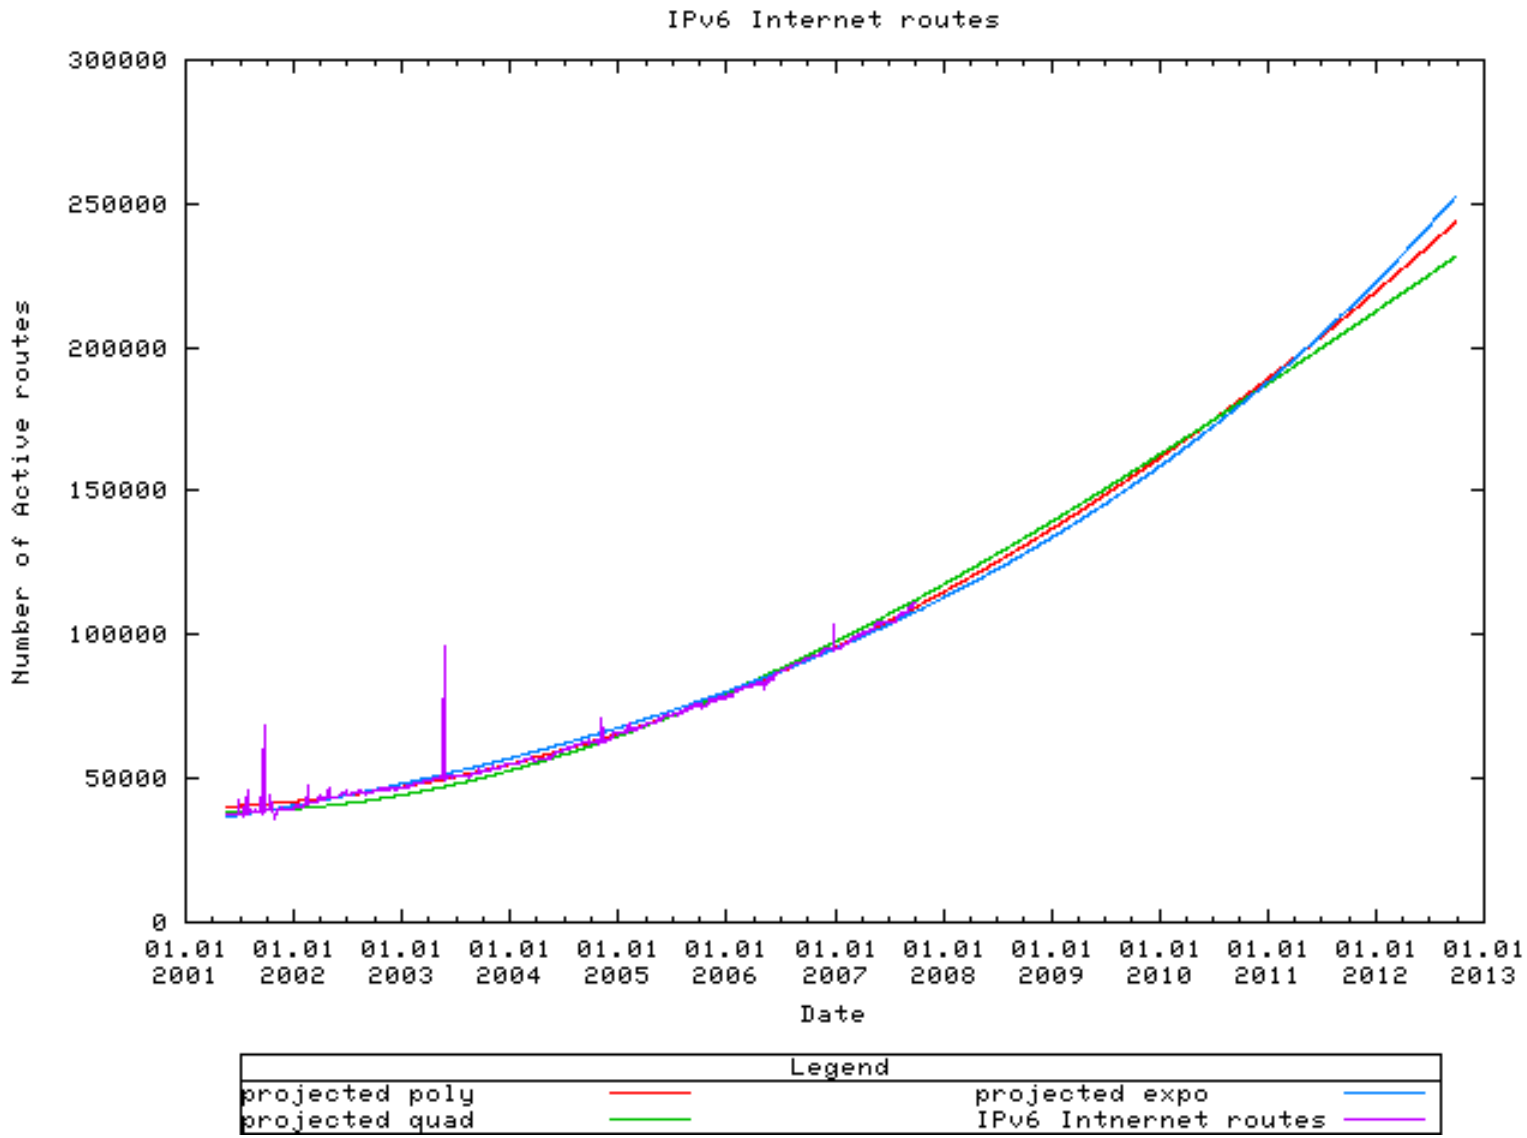

Future Projection of Combined IPv4 and IPv6 Internet Growth

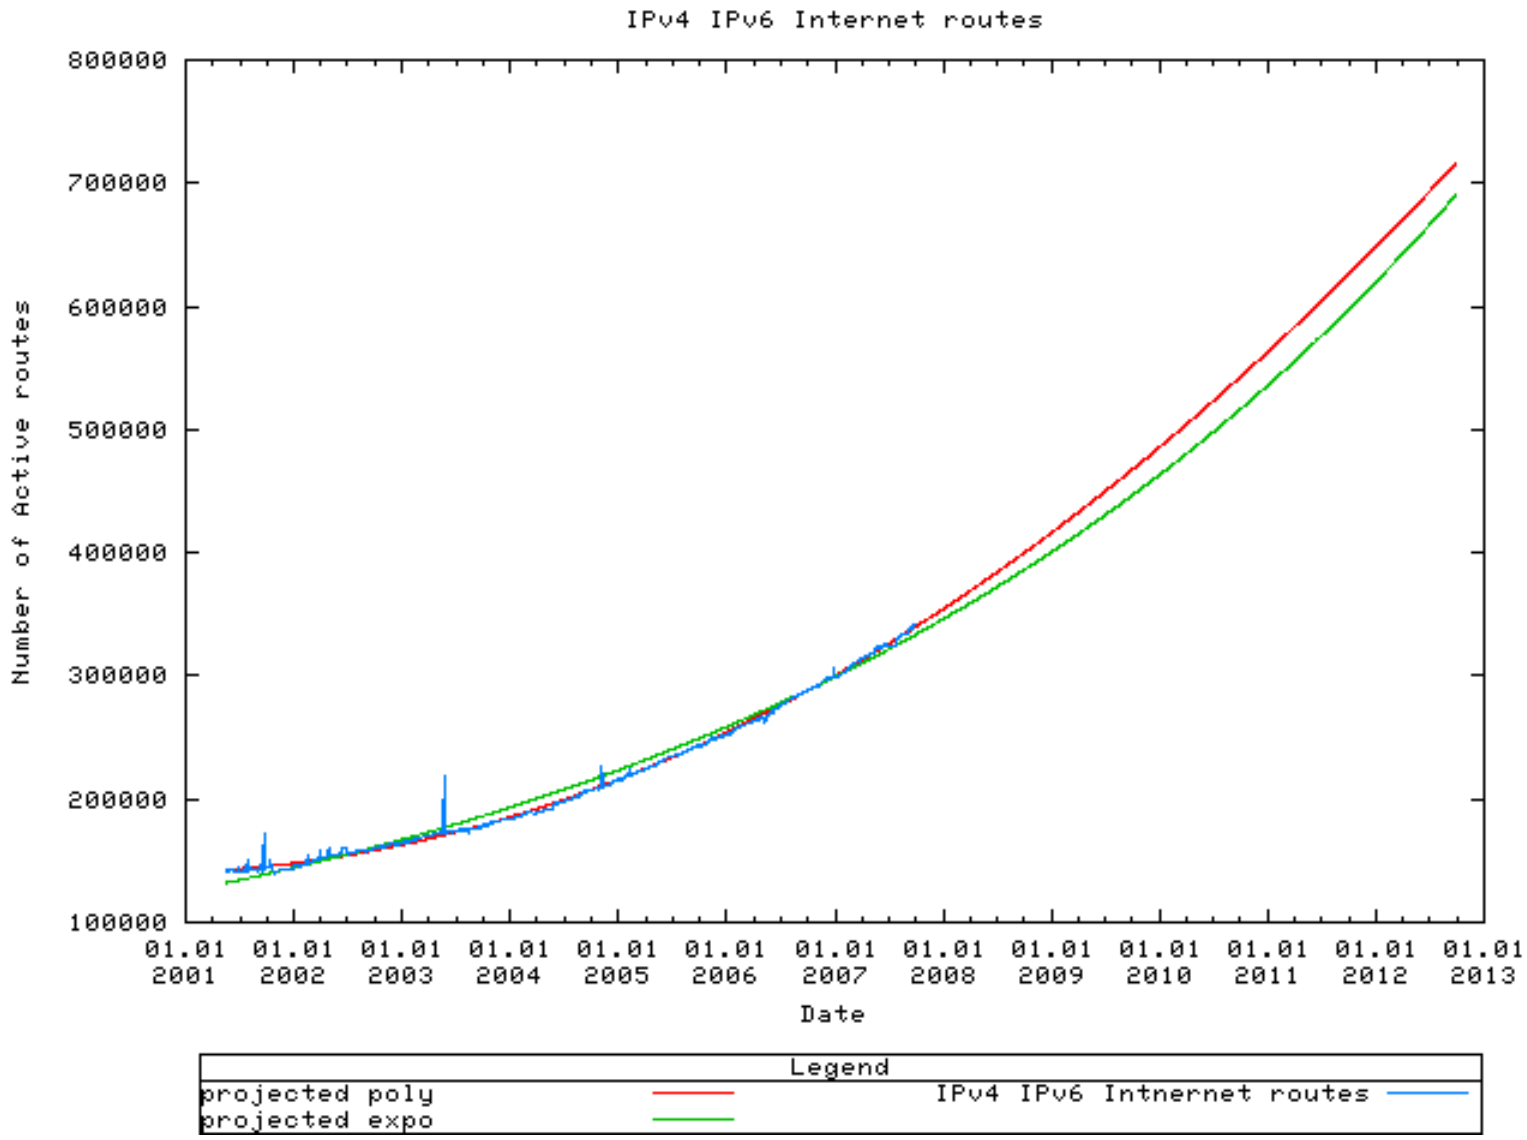

verizon

Future Projection Of Tier 1 Service Provider IPv4 and IPv6 Routing Table (03/07)

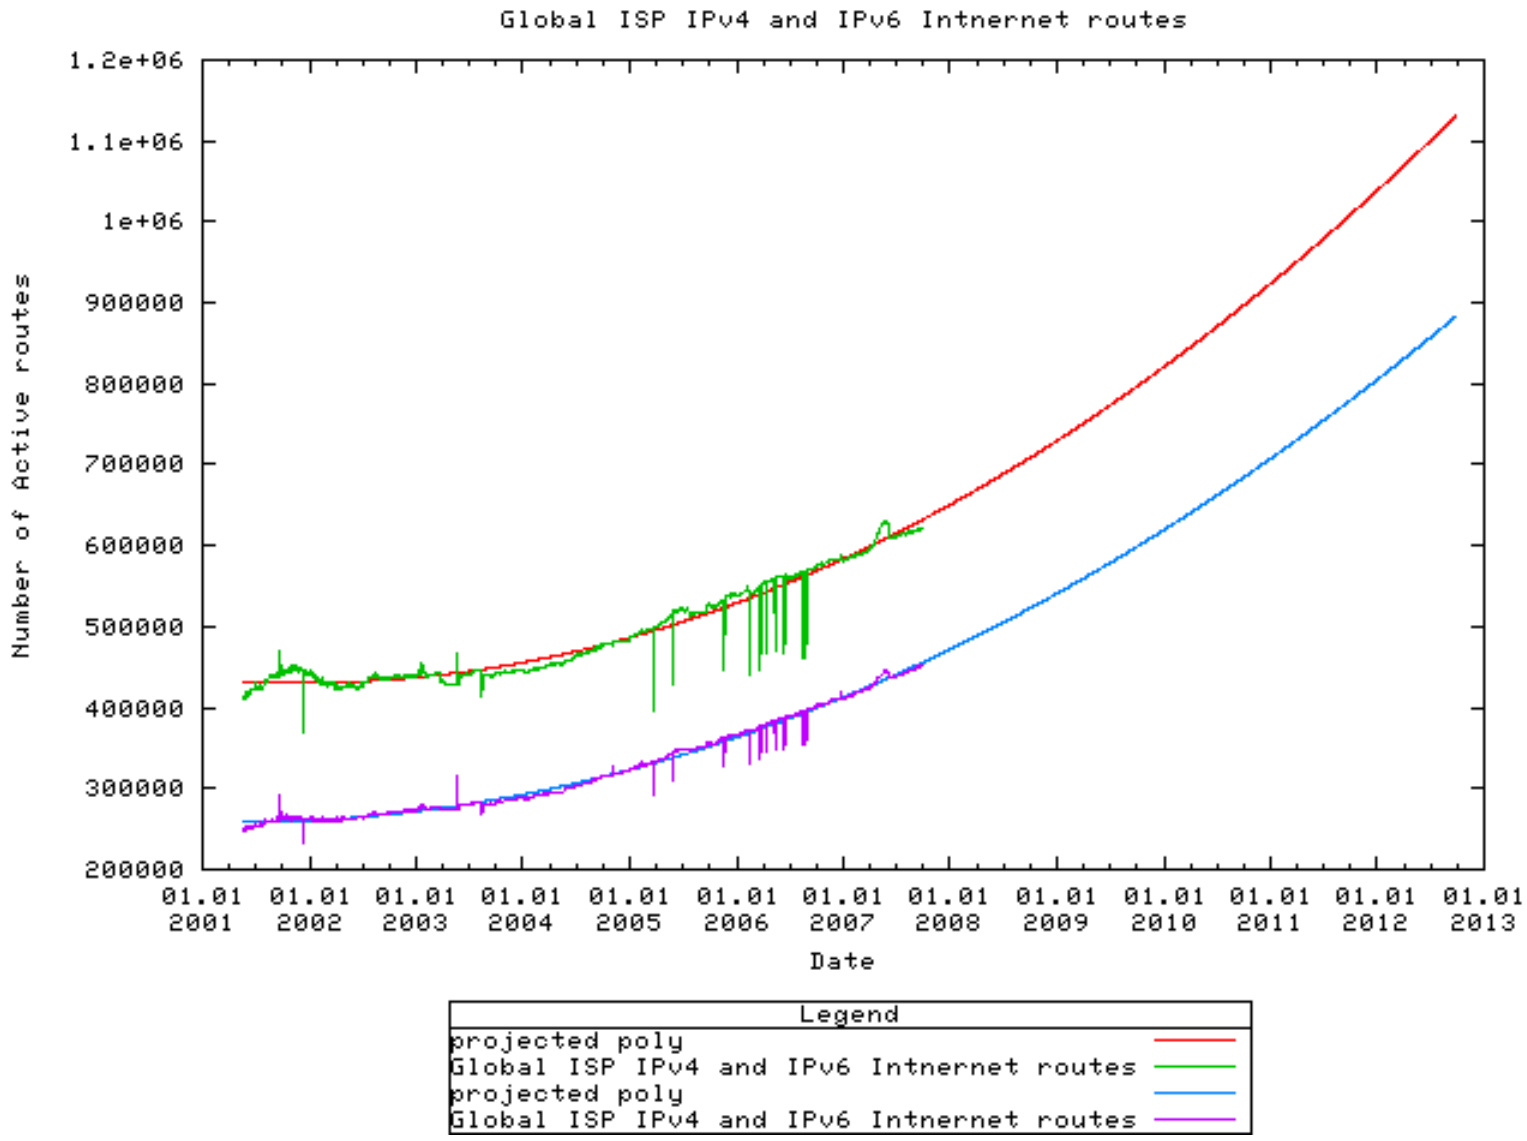

Table of big numbers (10/01/2007)

Route type	now	5 years	7 years	10 Years	14 years
IPv4 Internet routes	231,538	508,837	689,740	1,057,136	1,758,990
IPv4 CIDR Aggregates	146,939				
IPv4 intentional de-aggregates	84,599	200,462	265,335	382,324	574,014
Active Ases	26,240	37,225	41,908	48,939	58,280
Projected IPv6 Internet routes	110,839	237,651	307,243	431,263	632,294
Total IPv4/IPv6 Internet routes	342,377	746,488	996,983	1,488,399	2,391,284
Internal IPv4 (low number)	54,566	79,021	93,517	121,612	170,685
Internal IPv4 high number	137,552	199,199	235,742	306,563	430,269
Projected internal IPv6 (low)	43,653	63,217	74,814	97,289	136,548
Projected internal IPv6 (high)	110,041	159,359	188,594	245,251	344,215
Total IPv4/IPv6 routes (low)	440,596	888,726	1,165,314	1,707,300	2,698,517
Total IPv4/IPv6 routes (high)	506,984	984,868	1,279,094	1,855,261	2,906,184

Solving the Problem in Hardware

Current equipment purchases

- Assuming wide spread IPv6 adoption by 2012
- 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 2017
- Assuming equipment purchased in 2013 should last in the network for 5 years
- Vendors should be prepared to provide equipment that scales to 1.9M routes

Solving the Problem in Hardware

- Can vendors plan to be at least five years ahead of the curve for the foreseeable future or will refresh cycles compress?
- How do operator certification and deployment plans lengthen the amount of time required to be ahead of the curve?
- How do vendor time lines 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?
- 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?
- What will the impact of scaling the routes be on:
  - Cost
  - Power
  - Cooling
- What if the growth curve gets steeper?

Common Comments About the Scale Problem

- Vendors will build a box large enough if it is needed
- This problem only affects a small number of ISPs… maybe the Internet would be better without them
- There are lots of smart people working for the large ISPs, they will be affected first, and will solve the problem for everyone
- Moore's law will make this problem go away. It will be a while before we have the problem, but then the equipment will be bigger and faster

verizon

Conclusions

- Marketing literature specs currently available hardware just behind the 7 year projection
- Capacity is likely worse when considering IPv6 and firewall filters
- Vendors are optimistic that they can keep up but lack firm implementation details
- Growth is unbounded and could get worse
- Vendors unable to adapt to sudden changes in routing growth
- Need to investigate architectural solutions

verizon