

Update on the IAB Routing and Addressing Workshop

David Meyer/Chris Morrow
IAB Plenary
IETF67
San Diego, CA

Agenda

- Why Hold This Workshop?
- Logistics
- Workshop Objectives
- Participant Perspective: Chris Morrow
- Key Findings
- Workshop Recommendations

Why Hold This Workshop?

- The Internet's routing system is facing a set of serious scaling problems, and...
- *We are the IAB, after all, and...*
 - *“A is for Architecture” -- Leslie Daigle*
- And importantly...
 - There is a shared opinion among *many backbone operators* that none of the existing IETF efforts provides a complete set of solutions

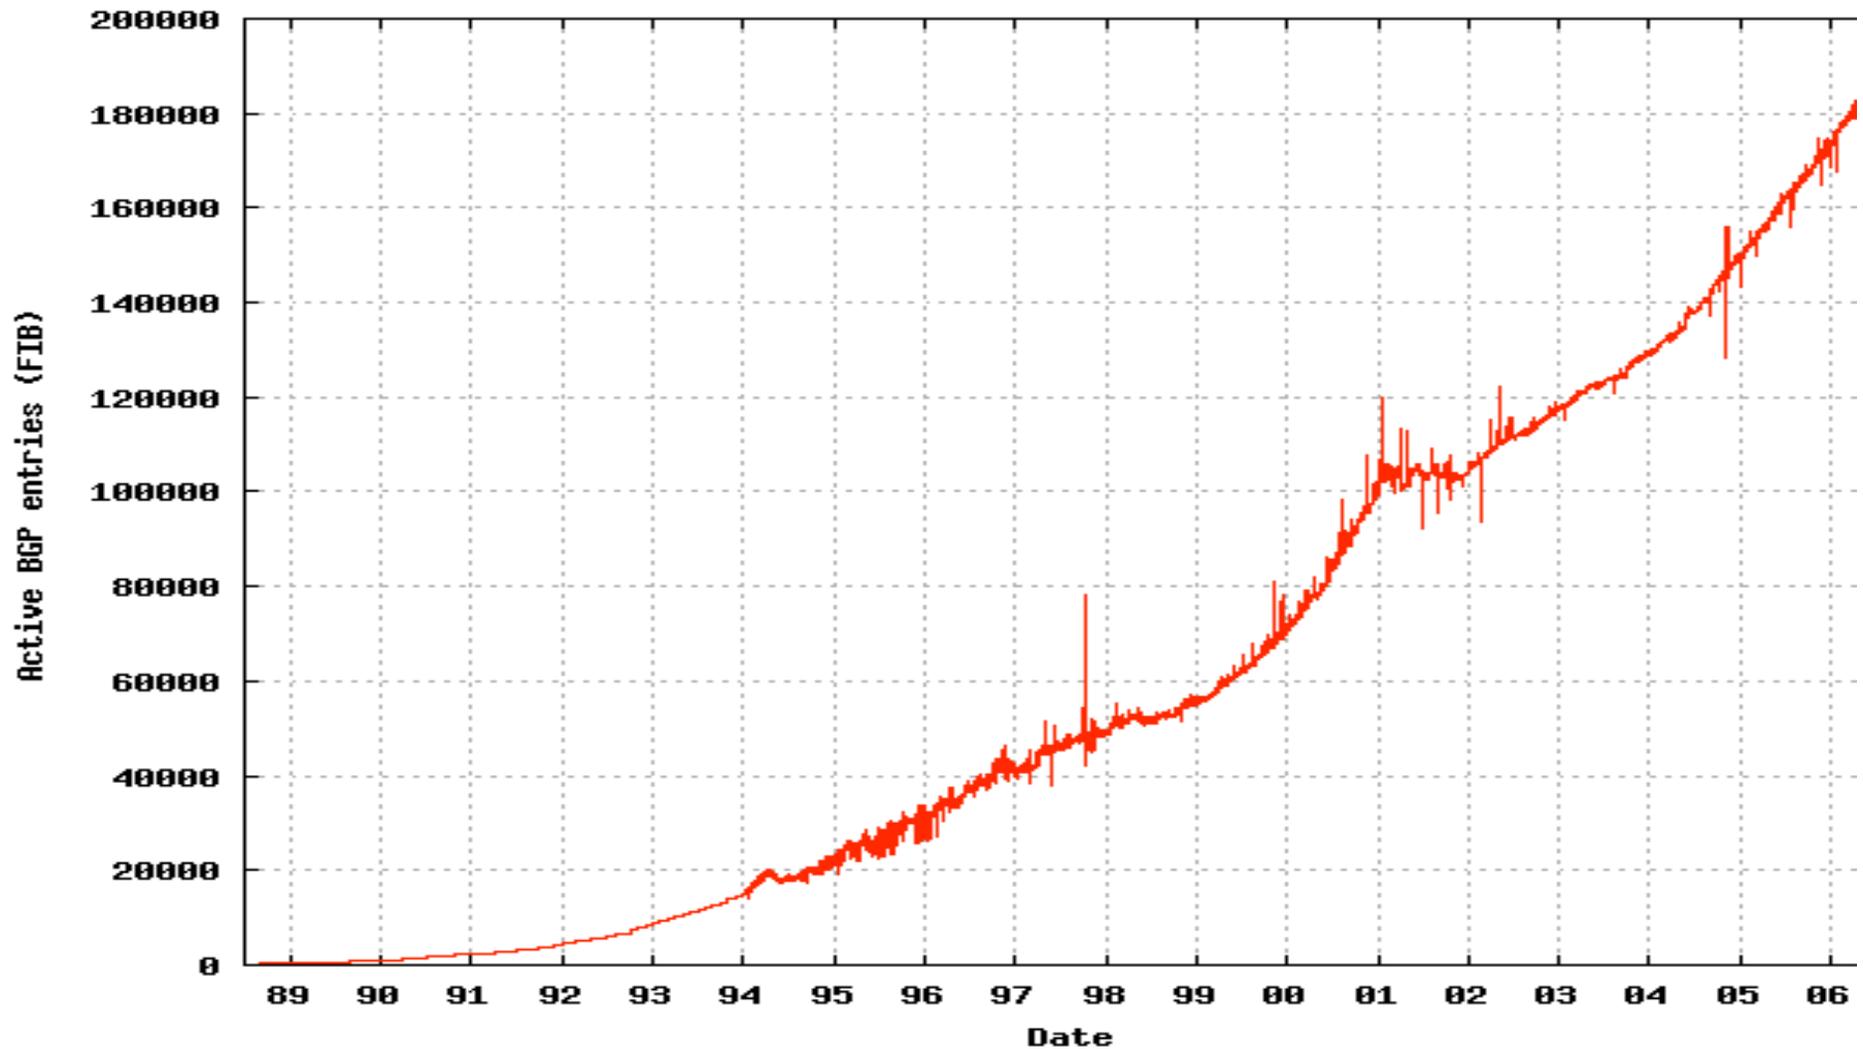
Logistics

- The workshop was held in Amsterdam, Netherlands on October 18-19, 2006
- 38 attendees
 - Focused on backbone operators
 - Also a few h/w designers, enterprise types
 - 18 (of the 38) were IESG, IAB, or IRTF
 - One scribe
- Many thanks to ISOC/RIPE NCC/NLnet Labs/Cisco
- And everyone who made the trip to help us think about these issues

Workshop Objectives

- To develop a shared understanding of the problems that operators are facing with today's routing and addressing system, and
- To use that information to inform the IETF process

Participant Perspectives: Chris Morrow



Problem Statement

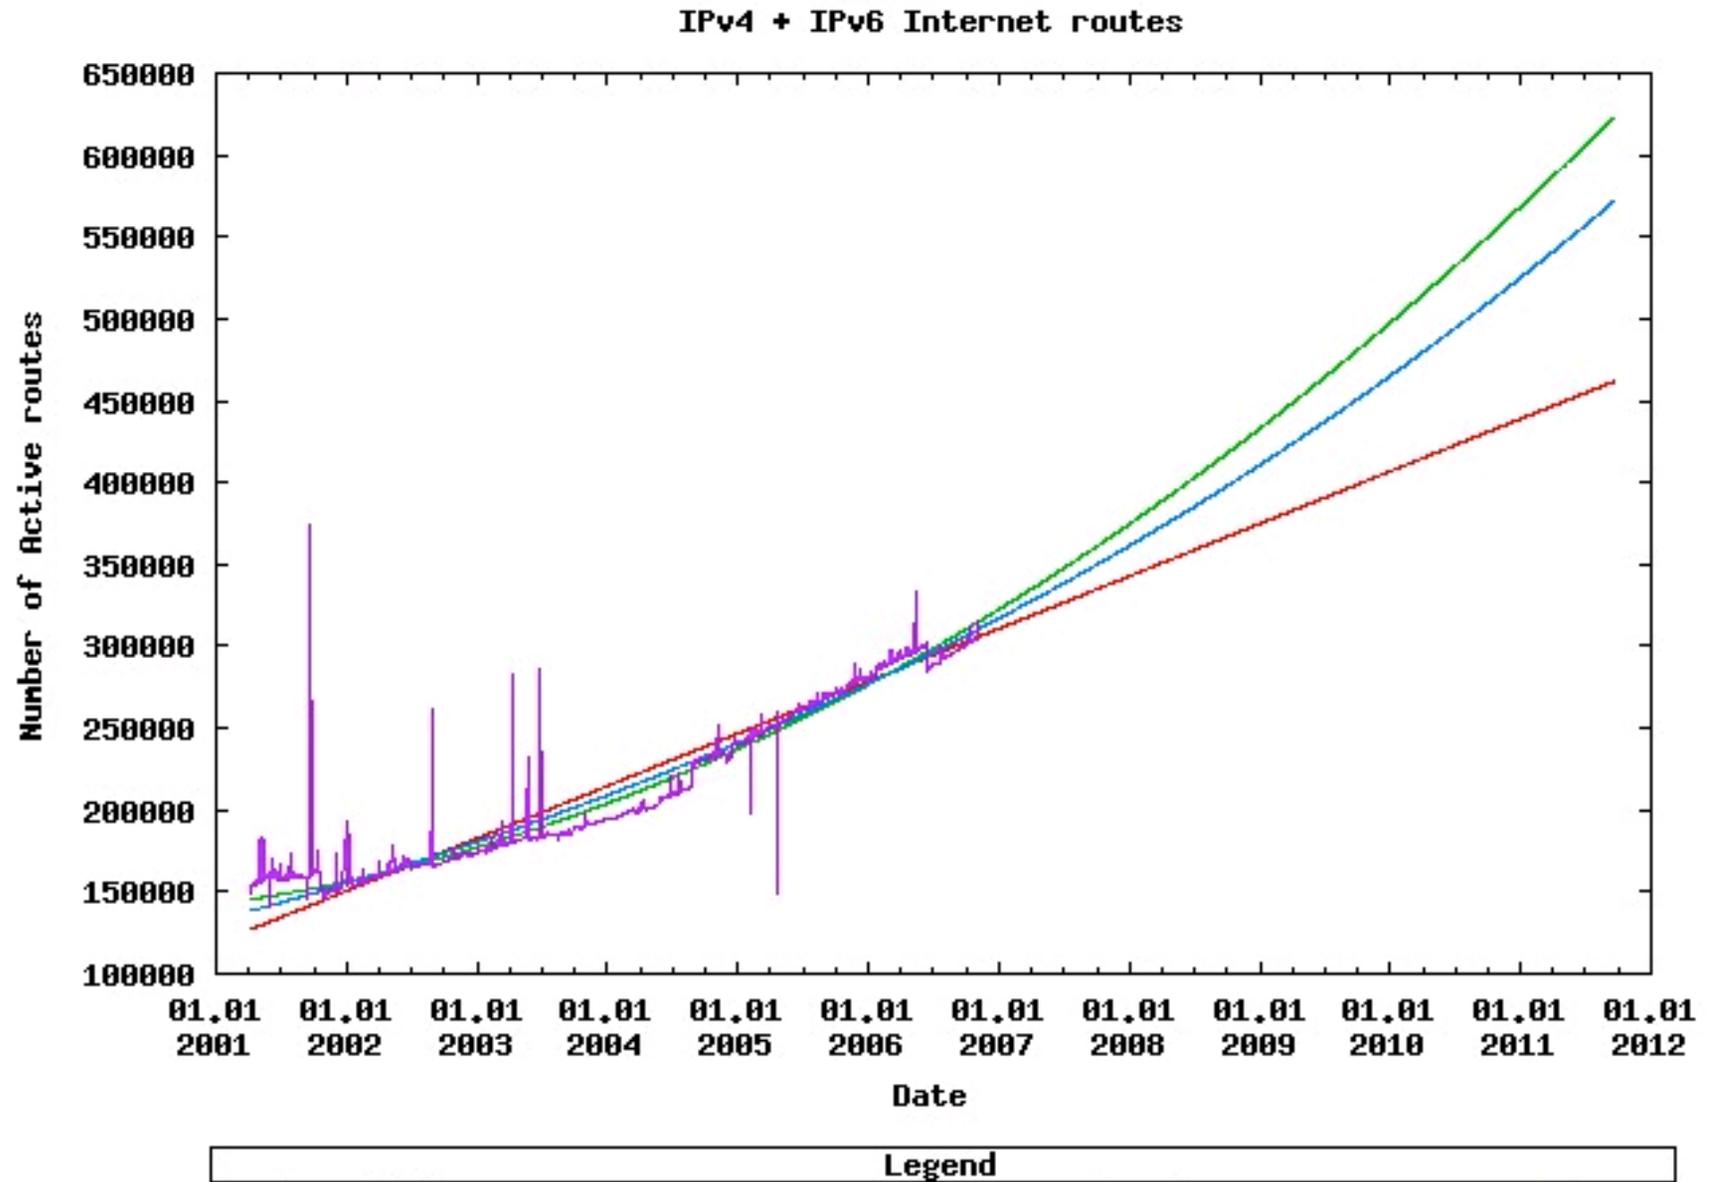
- Current trends in the growth of routing and addressing state on the global Internet are not scalable (in the long term)
- Among the major causes of this growth are multihoming and traffic engineering, which themselves are growing
- IPv6 is not significantly different than IPv4 – it shares many of the same properties and scaling characteristics

Estimated IPv4+IPv6 Routing Table (Jason Schiller, 11/06)

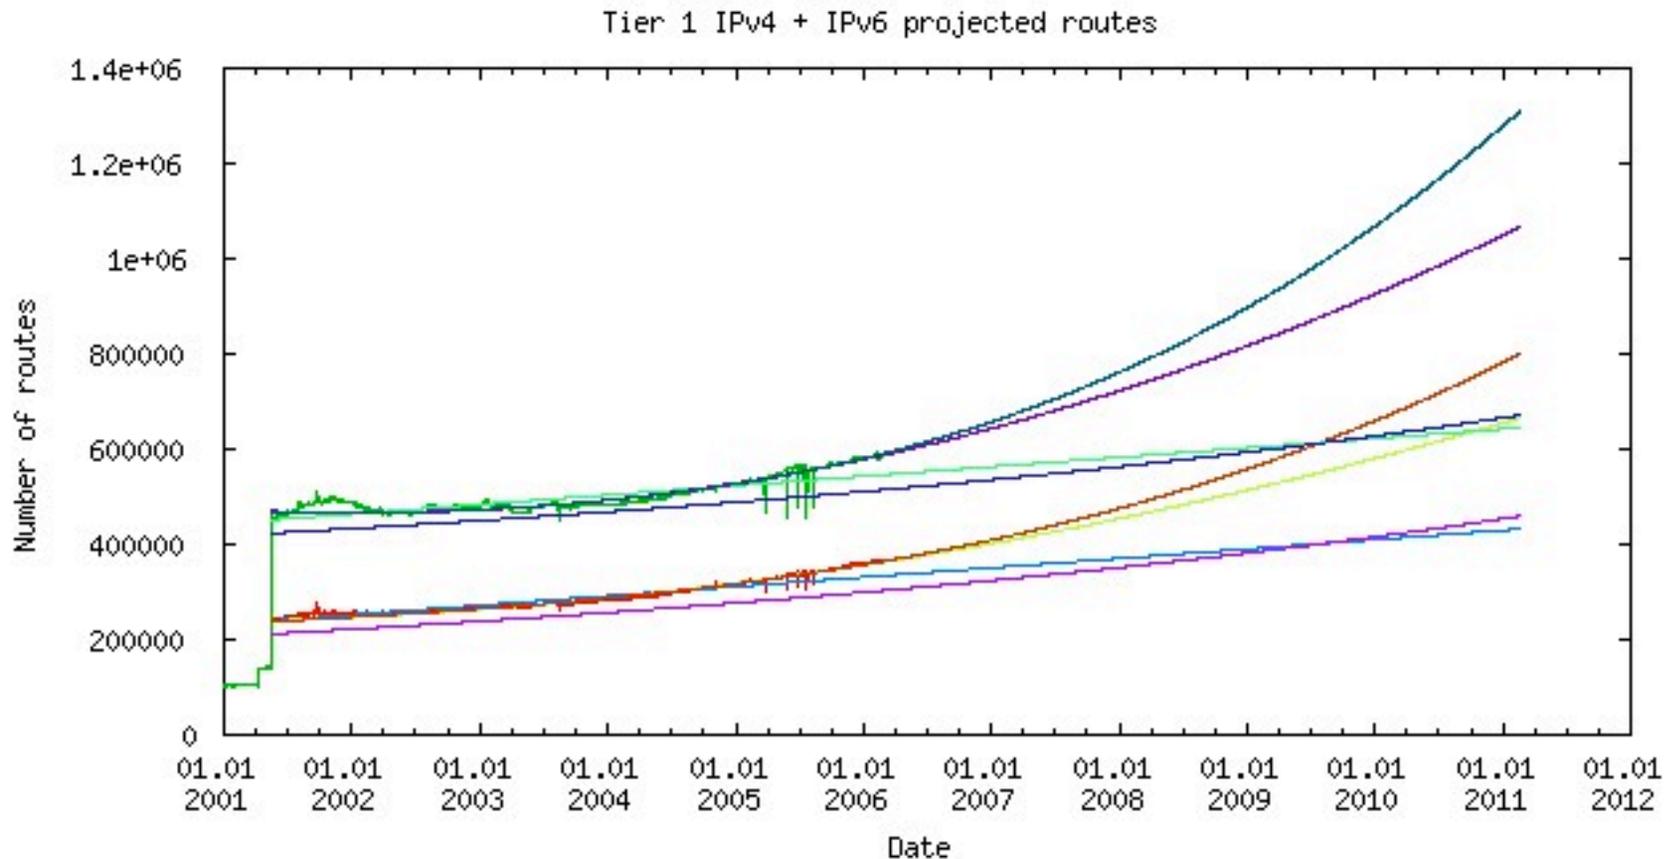
Assume that tomorrow everyone does dual stack...

Current IPv4 Internet routing table:	199K routes
New IPv6 routes (based on 1 prefix per AS):	+ 23K routes
Intentional de-aggregates for IPv4-style TE:	+ 69K 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	381K to 561K routes

Future Projection of Combined IPv4 and IPv6 Internet Growth



Inside a “tier-1” is even more “interesting”...



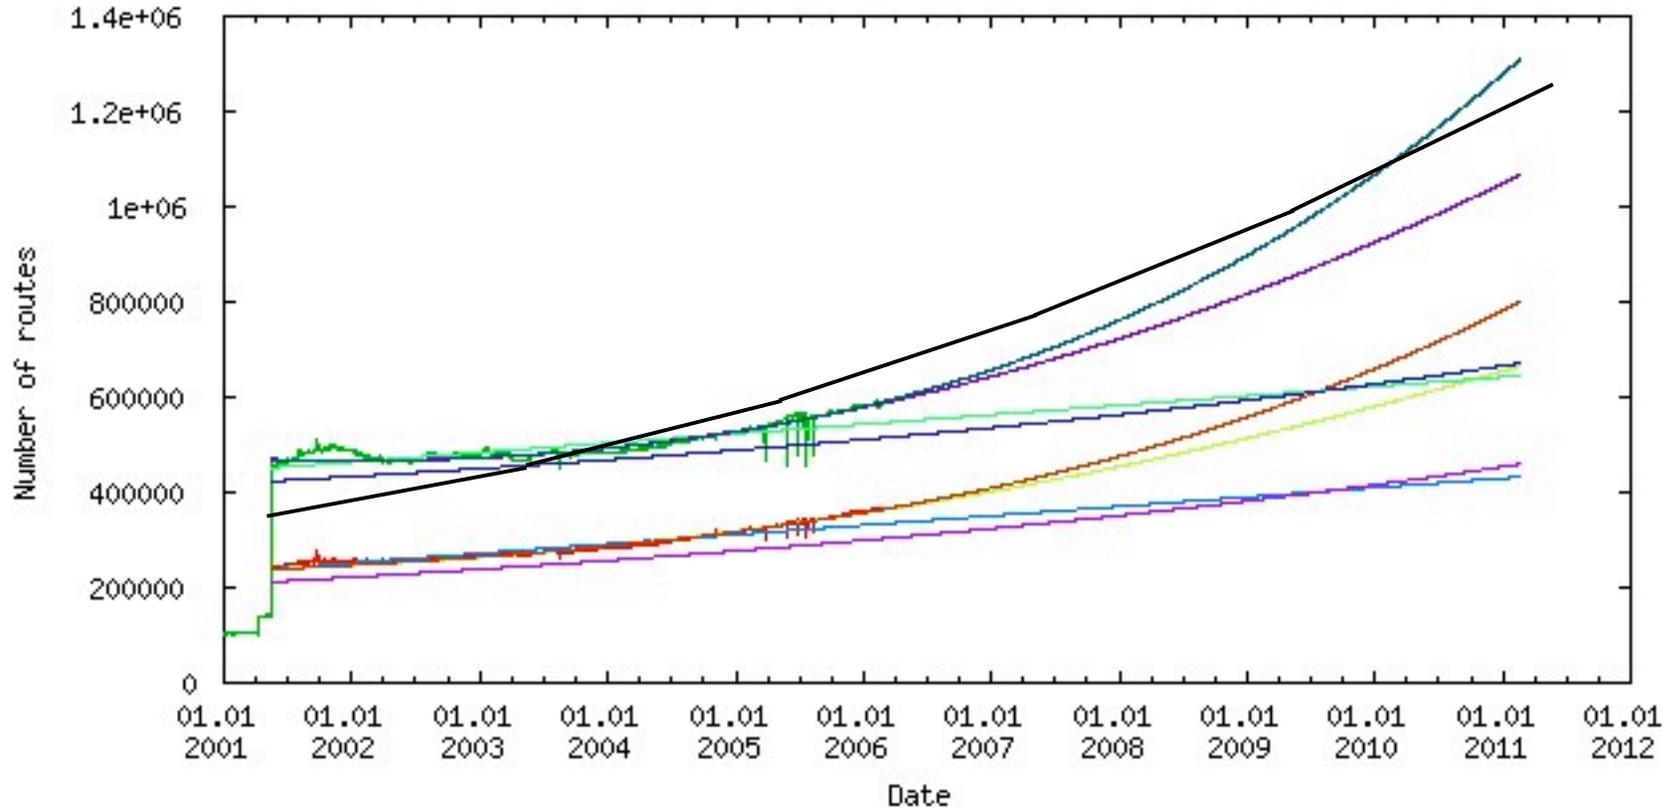
Legend	
Internal IPv4 + IPv6 routes	—
Internal IPv4 + IPv6 routes	—
projected IPv4 + IPv6 linear regression	—
projected IPv4 + IPv6 Power Regression	—
projected IPv4 + IPv6 quadratic regression	—
projected IPv4 + IPv6 cubic regression	—
projected IPv4 + IPv6 linear regression	—
projected IPv4 + IPv6 Power Regression	—

What About Moore's Law?

- Applicable to high volume components - think PC's: CPUs, main (DRAM) memories, and disk drives
- Critical router components (TCAM, SRAM) are low-volume and have much lower growth rates
- Forwarding ASICs already push limits of technology
- Memory speeds improve at about 10% per year
- Bottom line: state growth in excess of 1.3x every 2 years is problematic (translation: expensive)

Hardware growth vs. routing state growth

Tier 1 IPv4 + IPv6 projected routes



Legend	
Internal IPv4 + IPv6 routes	—
Internal IPv4 + IPv6 routes	—
projected IPv4 + IPv6 linear regression	—
projected IPv4 + IPv6 Power Regression	—
projected IPv4 + IPv6 quadratic regression	—
projected IPv4 + IPv6 cubic regression	—
projected IPv4 + IPv6 linear regression	—
projected IPv4 + IPv6 Power Regression	—

Some Interesting Numbers/Projections

Route type	11/01/06	5 years	7 years	10 Years	14 years
IPv4 Internet routes	199,107	285,064	338,567	427,300	492,400
IPv4 CIDR Aggregates	129,664				
IPv4 intentional de-aggregates	69,443	144,253	195,176	288,554	362,400
Active Ases	23,439	31,752	36,161	42,766	47,400
Projected ipv6 Internet routes	92,882	179,481	237,195	341,852	423,800
Total IPv4/ipv6 Internet routes	291,989	464,545	575,762	769,152	916,200
Internal IPv4 (low est)	48,845	101,390	131,532	190,245	238,400
Internal IPv4 (high est)	150,109	311,588	404,221	584,655	732,400
Projected internal ipv6 (low est)	39,076	88,853	117,296	173,422	219,800
Projected internal ipv6 (high est)	120,087	273,061	360,471	532,955	675,800
Total IPv4/ipv6 routes (low est)	381,989	654,788	824,590	1,132,819	1,374,400
Total IPv4/ipv6 routes (high est)	561,989	1,049,194	1,340,453	1,886,762	2,324,400

Thanks!

(darrel/dave/jason/ted/vince/vijay)

Key Workshop Findings

- **The scalability of the routing system is an urgent problem**
- Super-linear RIB growth is a great concern
 - Increased BGP convergence times and associated costs
 - RIB (UPDATE) dynamics also an issue (cf deaggregation)
- Questions about the applicability of Moore's Law to high-end routers (in particular, FIB memories)
- And of course, along with all of the various constraints e.g., no provider lock (PA/CIDR), TE, multihoming, ...
- Shared problem between IPv{4,6}
 - Larger IPv6 address space exacerbates these problems

Key Workshop Findings

- **The use of IP addresses for both ID and Locator is a problem**
 - Workshop participants felt that a solution to this overloading may solve the mobility and multihoming problems
 - Examined the tradeoffs inherent in SHIM6 and GSE
 - Long term solutions need to consider the anticipated “orders of magnitude” growth in new mobile end devices

Key Workshop Findings

- **Costs and Benefits in current practices are not aligned**
 - Canonical example: multihoming
- **Cost/Benefit curves vary by stakeholder**
 - An enterprise may have a very different view of the cost/benefit tradeoffs of a given solution set than say, content provider might

Personal Observations

- The workshop generated a nice enthusiasm, and ...
 - Everyone seemed to leave with a new energy around the problem
 - Some folks who hadn't engaged with IETF leadership recently (or ever before) re-engaged
- A lot of positive socialization occurred
- So the time for decisive action is **now**

Workshop Recommendations

- **These problems are urgent**
 - Need to start working on solutions now
- **Need to reach out to all stakeholders**
 - In addition to backbone providers, we need to reach out to the content providers, enterprises, applications folks, vendors,...

Workshop

Recommendations

- **We must develop solutions in an open & transparent way, engaging the broad community**
 - Engaging research community as well
- **Look into whether interim solutions are necessary to buy us a little time**

Workshop Recommendations

- **Need to develop a clear and coordinated approach to solutions development**
- Roadmap
- Near, intermediate, and long term steps from current state to solutions

Questions/Comments?

Thanks!