Transitioning to IPv6

Feasibility, Cost, and Planning

Rocky Mountain Cisco User's Group
15 March, 2007

Jeff Doyle
Principal
Jeff Doyle and Associates, Inc.

Why the sky *was* falling, but really isn't anymore, but on the other hand might begin again Any Time Now

Rocky Mountain Cisco User's Group

15 March, 2007

Jeff Doyle
Principal
Jeff Doyle and Associates, Inc.

Why the sky *was* falling...

- IP developed 1973 - 1977
  - 2^{32} = 4.3 billion addresses
  - More addresses than anyone could possibly use!
- Who'd have guessed?

> "Over the last several years, the number of network numbers assigned and the number of network numbers configured into the Merit NSFnet routing database have roughly doubled every 12 months. This has led to estimates that, at the current allocation rate, and in the absence of corrective measures, *it will take less than 2 years to allocate all of the currently unassigned Class B network numbers*."
>
> -- RFC 1380
> November, 1992

...but isn't any more...

- Long term solution: IPng ==> IPv6
  - 2^{128} addresses
  - More addresses than anyone could possibly use!
- Short term solutions:
  - Dynamic address allocation (DHCP)
  - Network Address Translation (NAT)
  - Private IPv4 addresses (RFC 1918)
  - Classless Interdomain Routing (CIDR)

...but might begin again Any Time Now

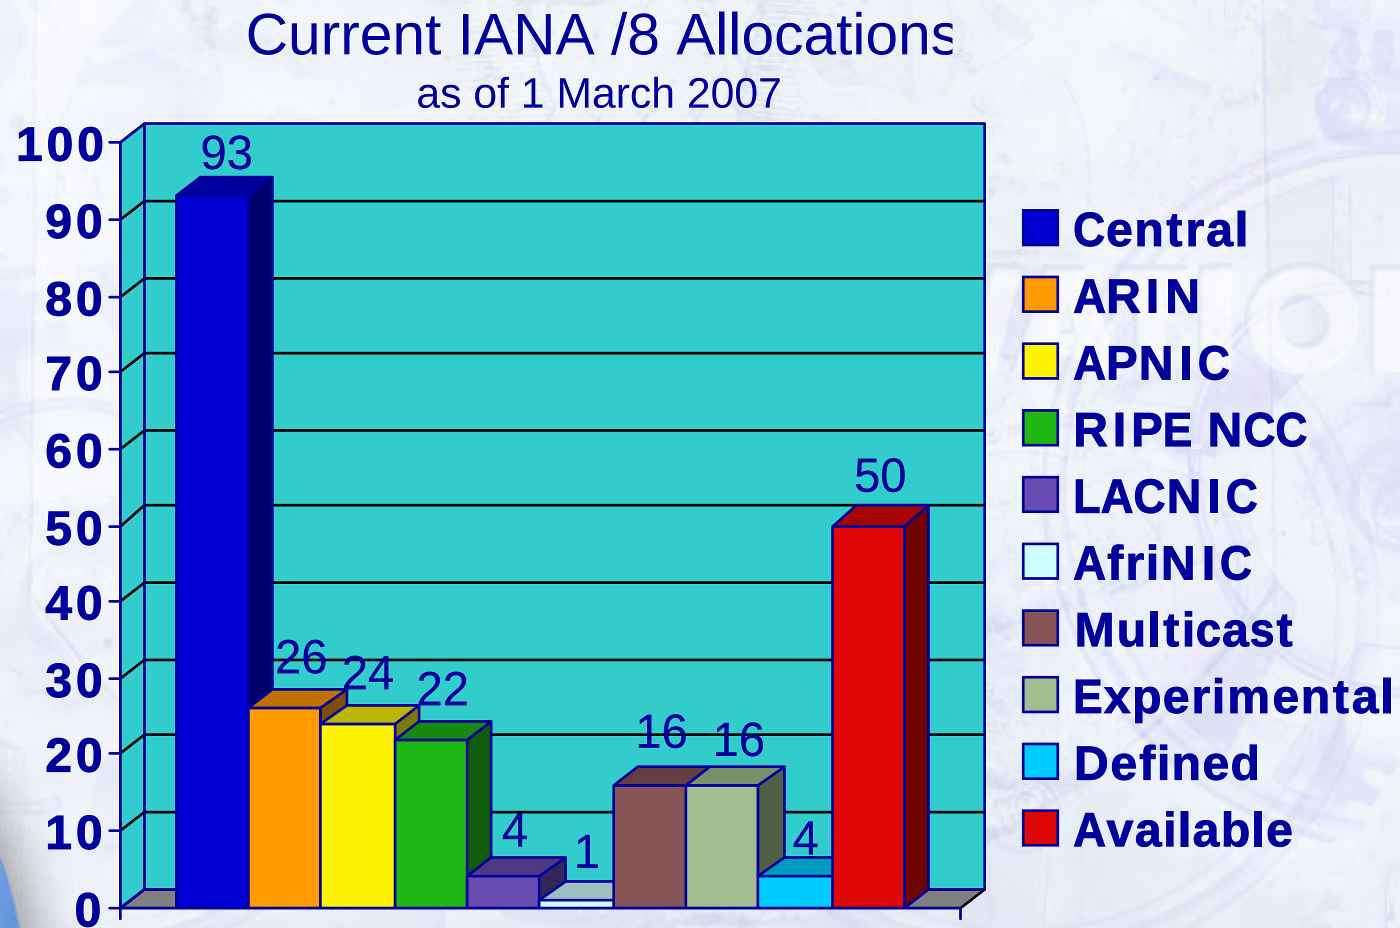

Source: www.iana.org/assignments/ipv4-address-space

How are IPv4 addresses allocated?

- Internet Assigned Numbers Authority (IANA)
  - 256 total /8 IPv4 blocks
  - /8 = 16,777,216 individual IPv4 addresses
- IANA ==> Regional Internet Registries (RIRs)
  - RIRs assume 18 month supply
  - New /8 assigned whenever:
    - Available pool falls below 50%, or
    - Available pool no longer sufficient for further 9 months
    - Allocation made to restore 18 month supply
- RIR ==> Local Internet Registries (LIRs)
  - LIRs are typically service providers
  - Might also be large company, government agency, educational institution…
- LIR ==> Customers

IANA /8 Allocations to RIRs

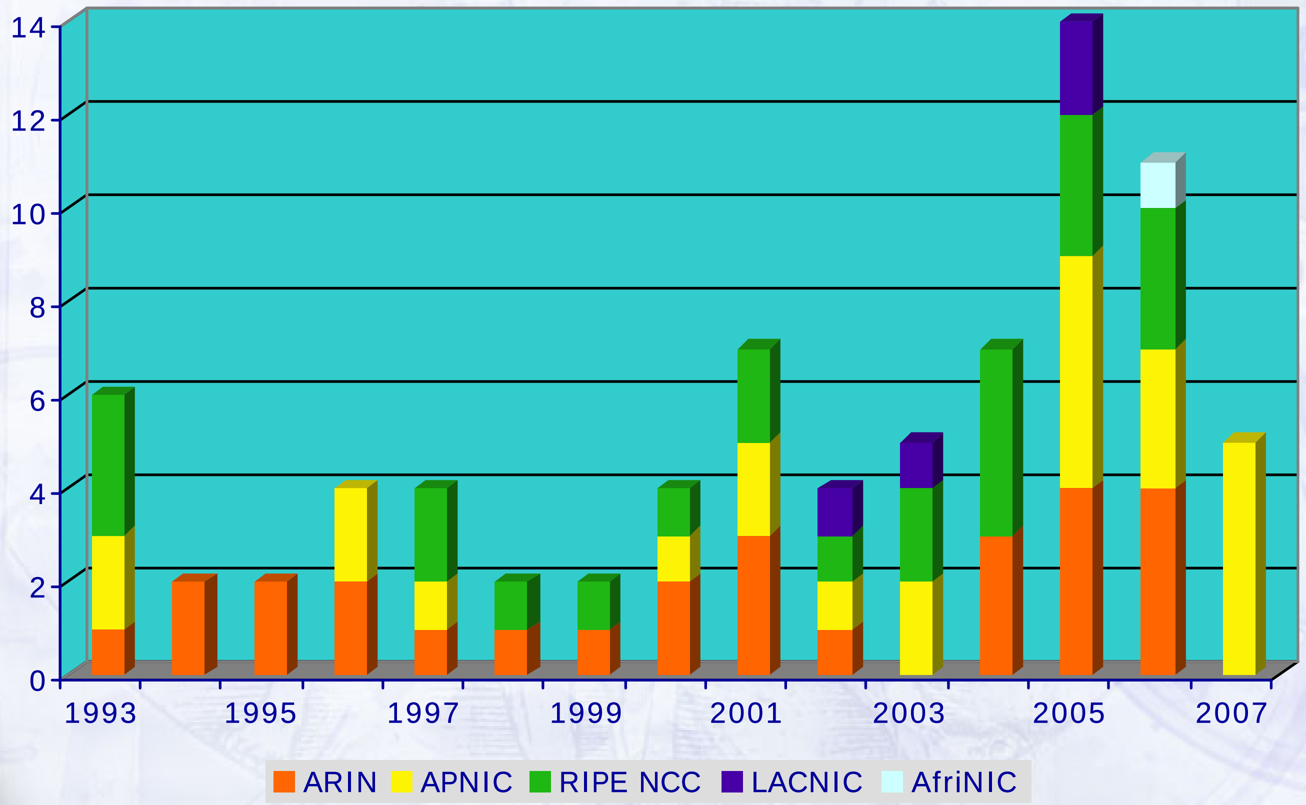

Source: www.iana.org/assignments/ipv4-address-space

What Can Be Deduced from the Allocation History?

- Prediction 1:
  - 2000 - 2006 average allocation = 7.43 /8s per year
  - Remaining /8s depleted in 6.7 years (2014)
- Prediction 2:
  - End of 2004: 80 /8s in pool ==> 31.25%
  - End of 2006: 55 /8s in pool ==> 21.48%
  - 2004 - 2006 average allocation = 4.88%
  - Remaining /8s depleted in 6.4 years (2013)
- Both predictions are clumsy
  - Assume flat rate and flat market
  - No account of growing economies

Example:
Assume 20% Internet penetration in China + India
Assume 3 person per address ratio
206 /8s required to serve this market

Geoff Huston's 2003 Study

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Source: www.potaroo.net/ispcol/2003-08/ale.html

IANA /8 Allocations to RIRs, Revisited

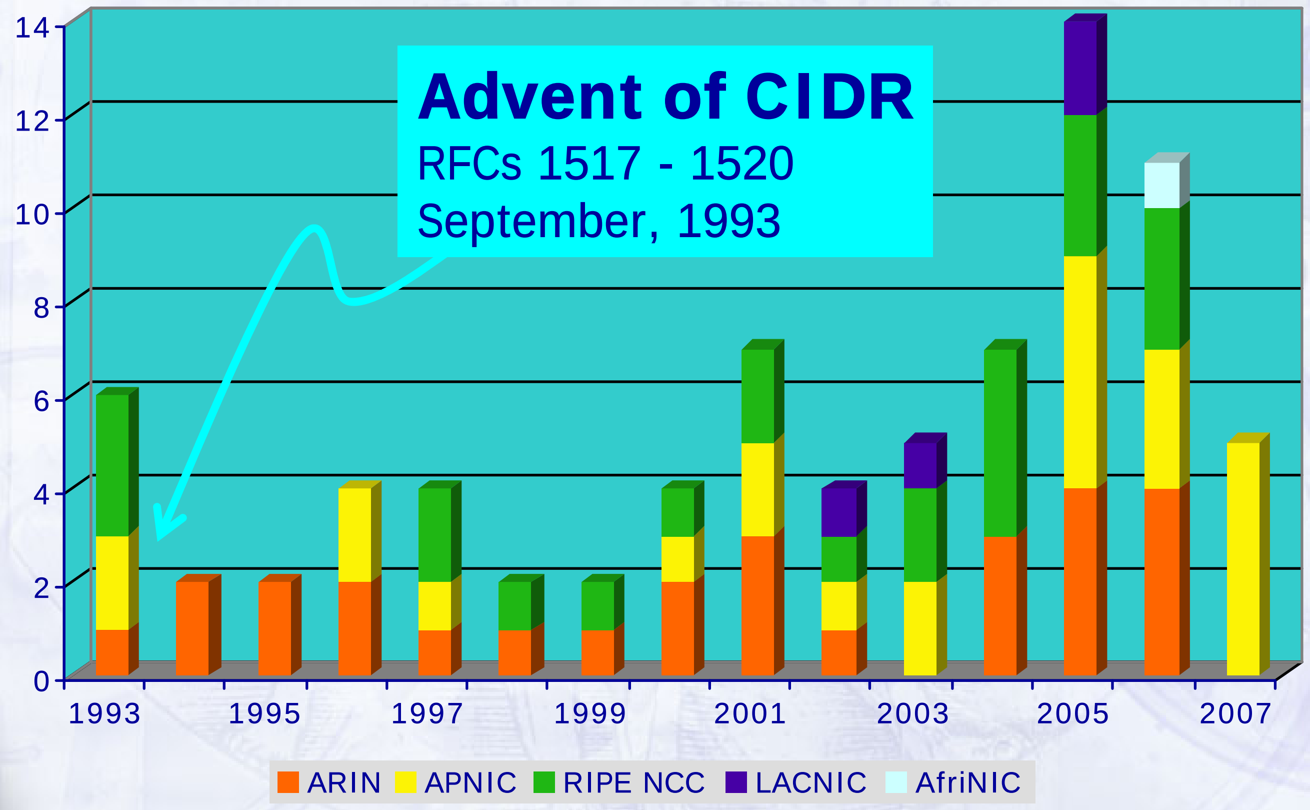

Source: www.iana.org/assignments/ipv4-address-space

Tony Hain's 2005 Study

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Source: www.cisco.com/web/about/ac123/ac147/archived_issues/ipi 8-3/ipv4.html

Geoff Huston's 2005 Study

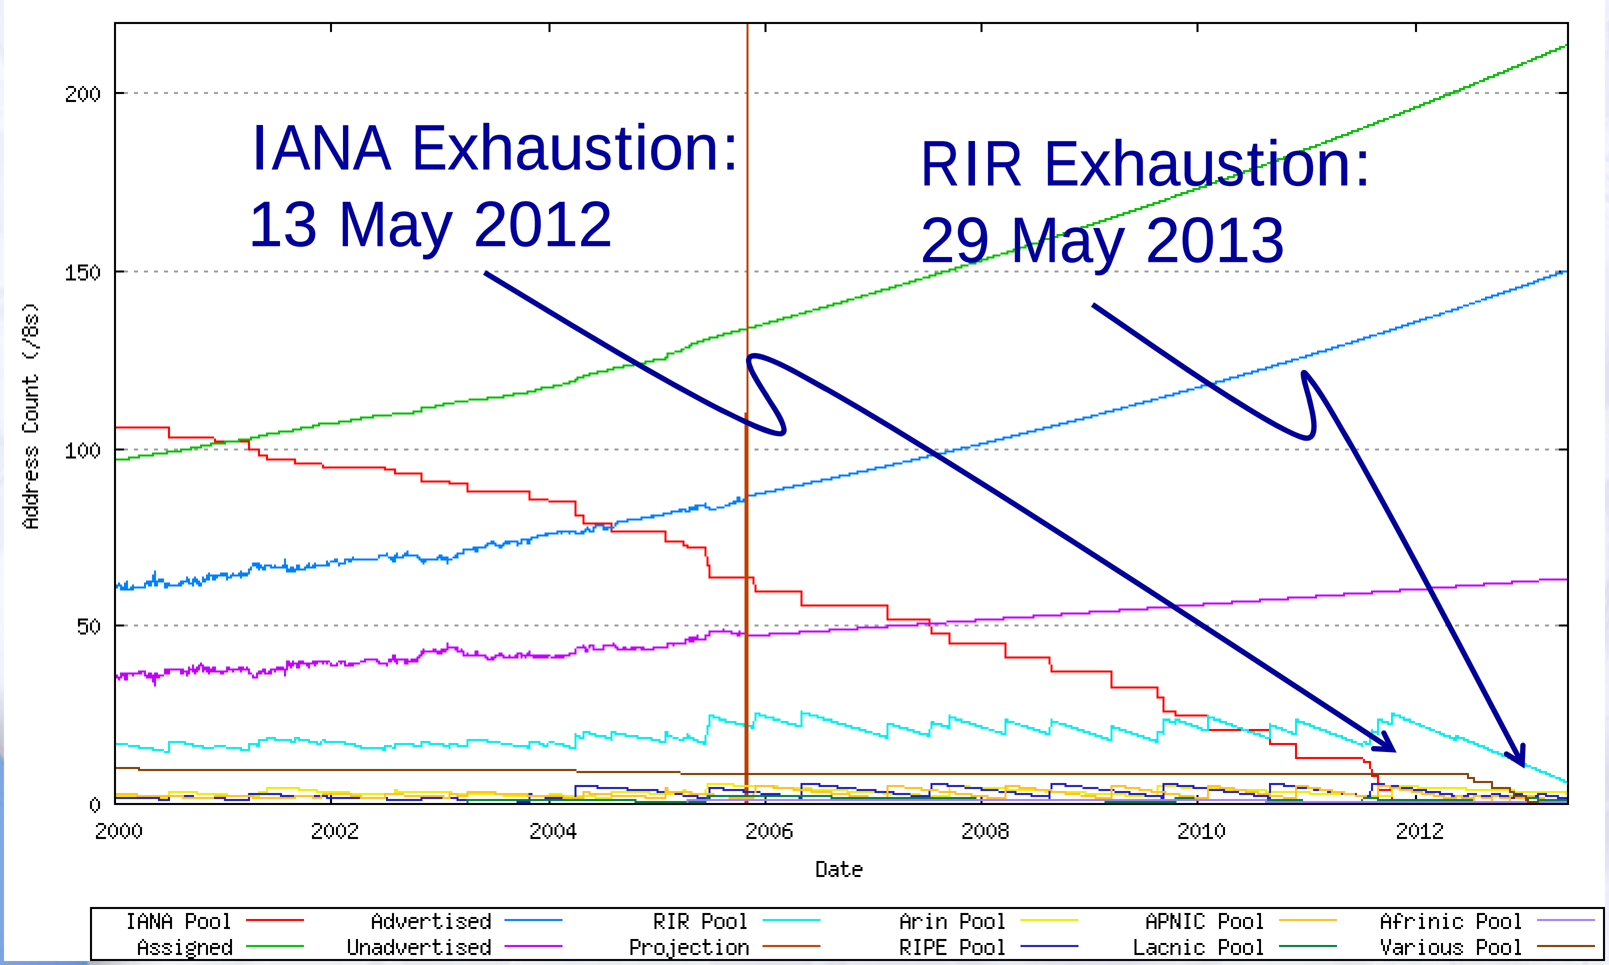

Source: www.isoc-au.org.au/ipv6summit05/abstracts.html#gh

Tony Hain's 2006 Study

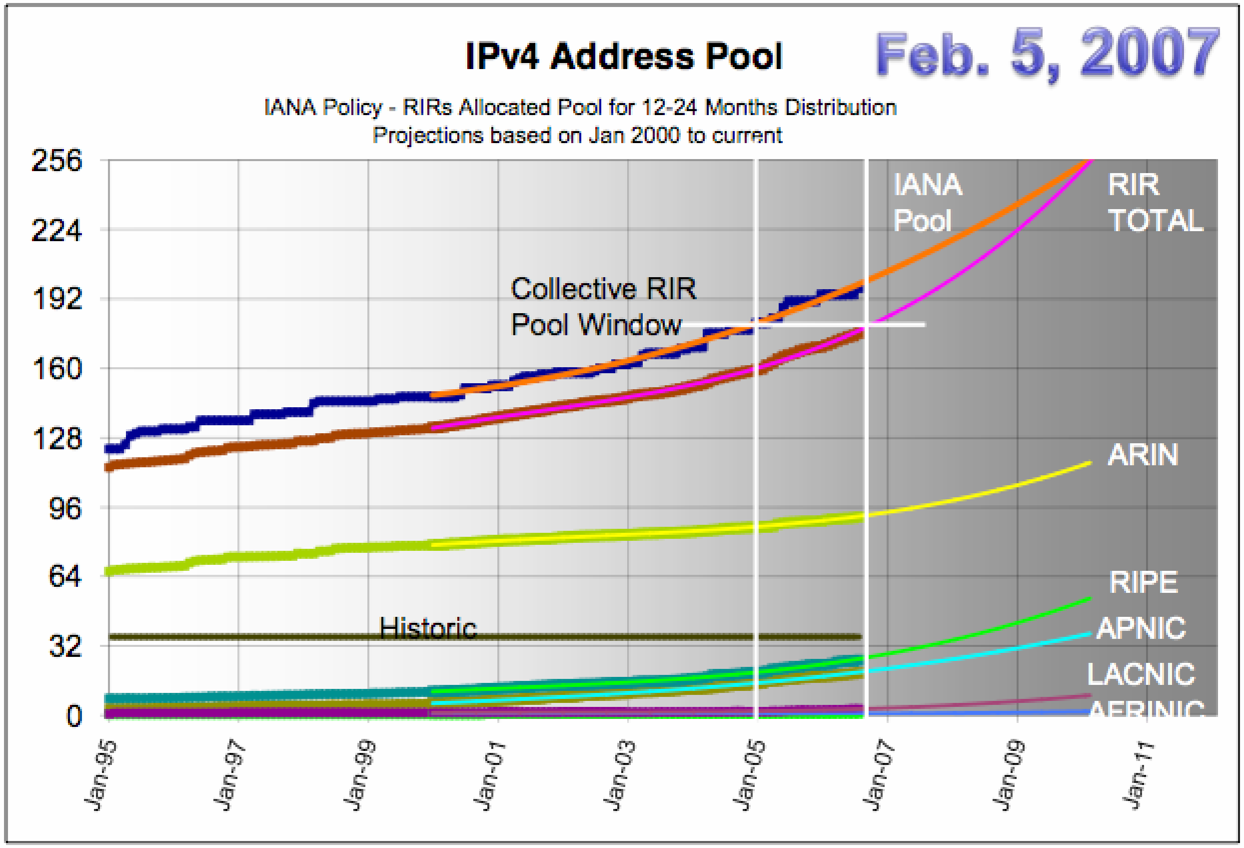

Update to: http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipj_8-3.pdf
Exhaustion of the central IANA pool - orange
Exhaustion of the collective RIR pools - magenta
Relative distribution rates between the RIRs
Time depth of collective RIR pools on pub date - white
Time depth between exhaustion events - diff between orange & magenta

Tony Hain

Source: www.tndh.net/~tony/ietf/ipv4-pool-combined-view.pdf

Possible Factors Influencing Depletion Rate

- Stricter IANA allocation rules
  - Hard to predict extent of effect
  - Will meet strong political resistance
- “Run on the bank”
- Sharing of remaining allocations among RIRs
  - Unlikely to have significant effect
- General IPv6 deployment

Proposals for Slowing Depletion Rate

- Open Experimental Class to public consumption
  - Involves significant software upgrades to existing systems
  - Only temporary relief
  - Effort better spent upgrading software for IPv6
- Expand the RFC 1918 address space
  - Larger topologies behind NATs without layering
  - Just exacerbates existing NAT limitations
  - Objections already raised
- Aggressive IPv4 reclamation project
  - Economic incentive or legal challenges
  - Either likely to be prohibitively expensive

Conclusions

- IANA IPv4 pool exhaustion circa 2010 - 2014
- RIR IPv4 pool exhaustion 12 - 24 months later
- Existing LIR and large-organization pools will last significantly beyond this
  - Huston predicts up to 5 years
- Earliest and most severe effects will be felt by:
  - Smaller companies and organizations
  - Developing nations
  - Greenfield network service companies

A Word About Cost

"Everyone knows transitioning to IPv6 is EXPENSIVE!"

- Not with a good transition plan
  - Well thought through
  - Generous timeframe
  - Insure IPv6 in natural network evolution
- IPv6 is increasingly common in IP stacks
  - Should be considered a standard part of IP
- Make IPv6 a requirement of all new hardware/software evaluations

The Feasibility Study

- Problem Statement
  - What problem am I trying to solve?
- Change Assessment
  - What needs to be changed to support IPv6?
- Standards Assessment
  - Are the standards established?
- Vendor Assessment
  - Real solutions or just slideware?
- Risk Assessment
  - I’ve got *alot* to say about this...
- Cost Rough-In
  - Rounded to the nearest million
- Value Assessment
  - Is $400,000 expensive or a good investment?
- Timeframe Rough-In
  - Time influences cost and risk!
- Conclusions

The feasibility study provides inputs to a transition plan

The Transition Plan

- High-Level Design
- Inventory
- Milestones
- Vendor Evaluation and Selection
  - Support Testing
  - Interoperability Testing
- Training
- Cost Estimates
- Methodologies
  - Don’t forget backout plans!
- Execution
- Post-Mortem (Lessons Learned)

QUESTIONS?

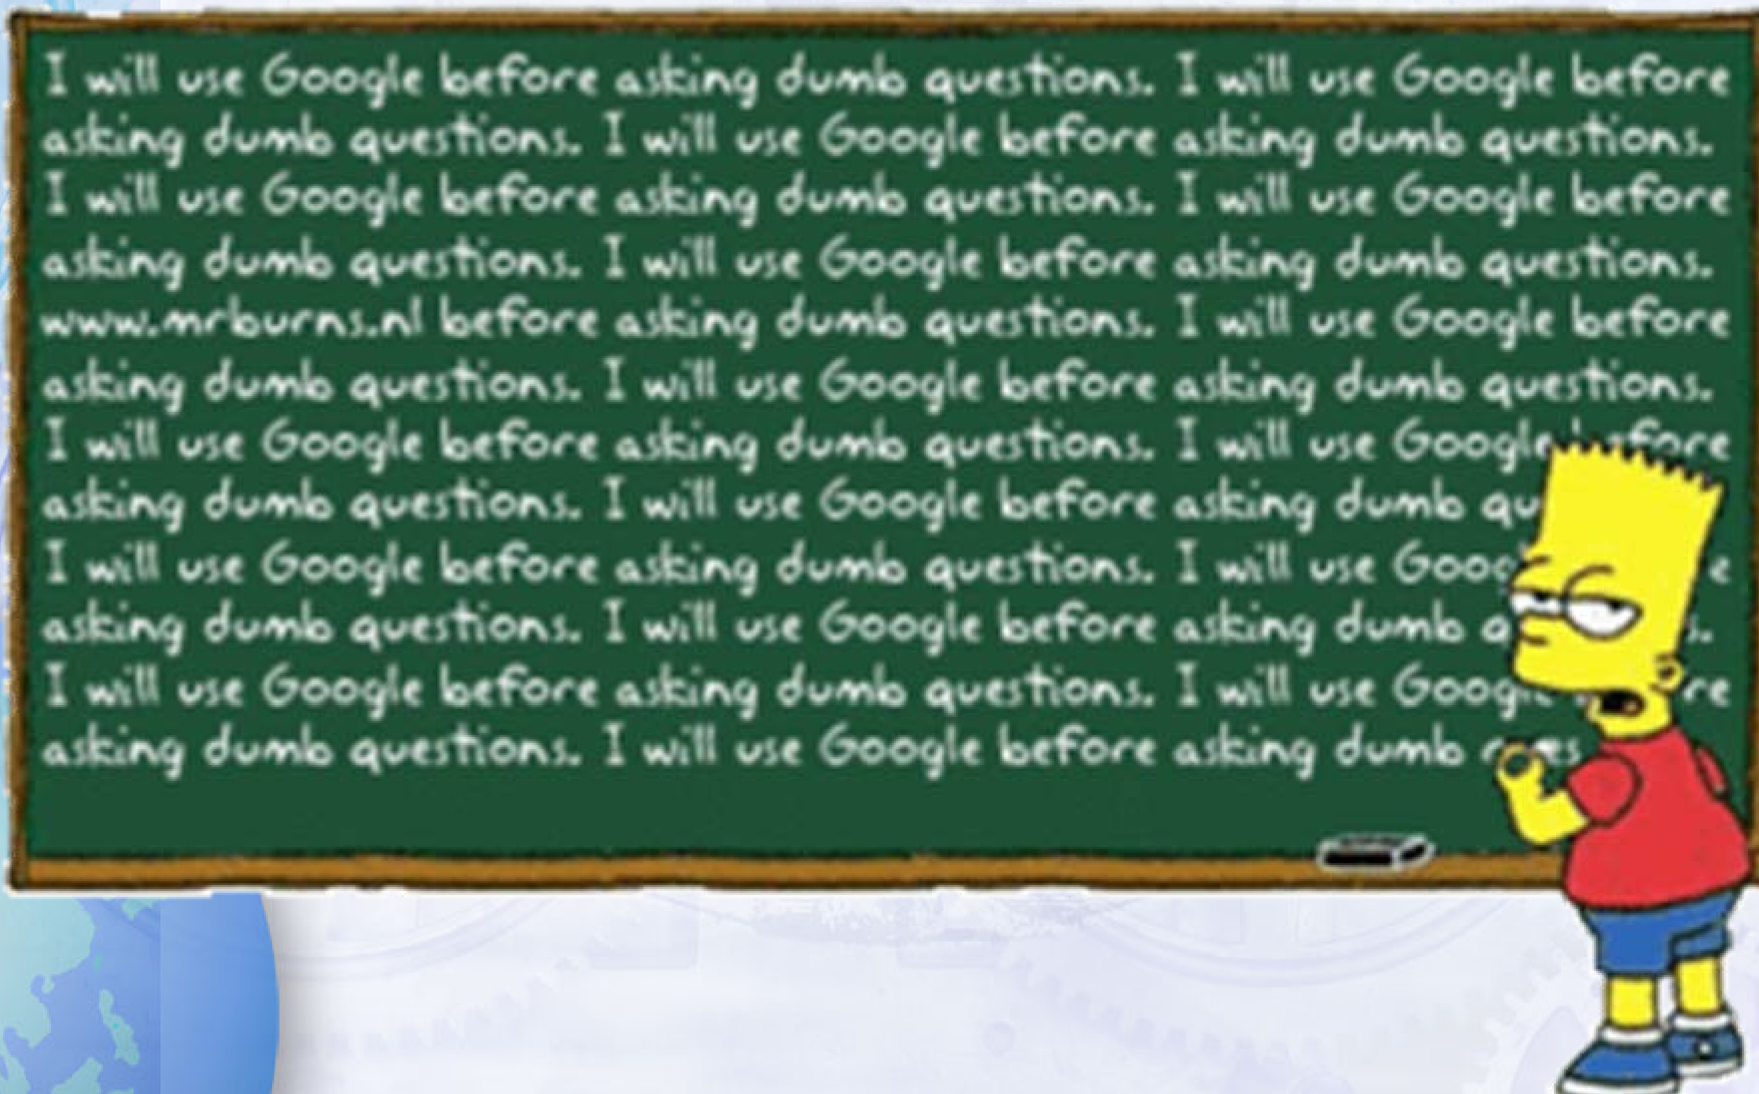
I will use Google before asking dumb questions. I will use Google before
asking dumb questions. I will use Google before asking dumb questions.
I will use Google before asking dumb questions. I will use Google before
asking dumb questions. I will use Google before asking dumb questions.
www.mrburns.nl before asking dumb questions. I will use Google before
asking dumb questions. I will use Google before asking dumb questions.
I will use Google before asking dumb questions. I will use
asking dumb questions. I will use Google before asking dumb
I will use Google before asking dumb questions. I will use
asking dumb questions. I will use Google before asking dumb
I will use Google before asking dumb questions. I will use
asking dumb questions. I will use Google before asking dumb

Thank You!

jdoyle@doyleassociates.net

(303) 428-4680

www.doyleassociates.net

Case Study: Comcast

- Uses IP to manage cable modems
- 10.0.0.0/8 space = 16.8 million addresses
- 2005: Comcast exceeded 20 million customers
  - 10.0.0.0/8 no longer sufficient
  - This is just for HSD customers
- 2007: Comcast offering triple-play bundled services
  - IP used to manage HDTV set-top boxes
  - 2 IP addresses per box
  - Average 2.5 boxes per household
  - = 100 million IP addresses across current 20 million customers
- Add VoIP into mix
  - Average 2 eMTAs per subscriber
  - = another 40 million IP addresses across current customer base
- ...and, these numbers do not account for expected growth

Projected Comcast Address Trends, 2005-2009

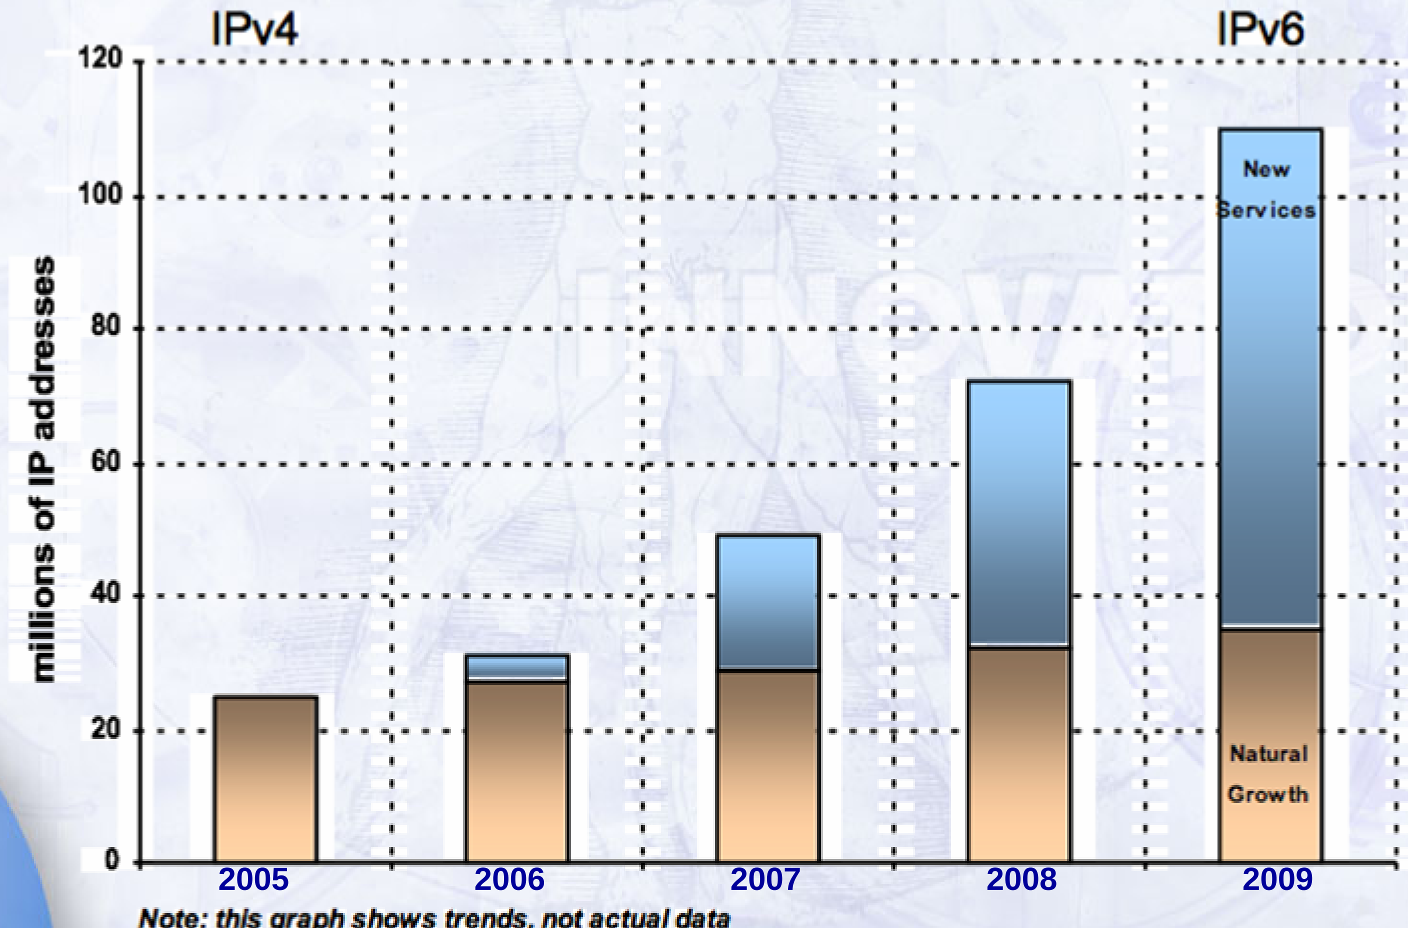

Source: Alain Durand, Comcast IPv6 Architect, NANOG 37 Presentation

Comcast Transition Plan

- Began migrating to IPv6 in 2005
- Incremental, hierarchical transition

1. Begin with backbone
   - Already completed
2. Backoffice systems
   - Not all backoffice systems need to know IPv6
3. Comcast Regional Access Networks (CRANs)
4. CMTS platforms
5. Home Systems
   - STBs
   - MTAs
   - Routers/Hosts
6. Public access

First Ping on Comcast Backbone...

```
ping ipv6 2001:558:0:f501::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:558:0:f501::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/5/14 ms

Jun  2 09:31:49.589, len: 162, hits: 1, i/p i/f: TenGigE0/7/0/0
00146a7d 29810014 6ac4dd08 86dd6000 0000006c 3a3c2001 05580000 f5010000
00000000 00022001 05580000 f5010000 00000000 00018000 50fae0da 00004480
3e53000f 062b0809 0a0b0c0d 0e0f1011 …

Jun  2 09:31:53.533, len: 162, hits: 1, i/p i/f: TenGigE0/0/0/0
00146ac4 dd080014 6a7d2981 86dd6000 0000006c 3a402001 05580000 f5010000
00000000 00012001 05580000 f5010000 00000000 00028100 4ffae0da 00004480
3e53000f 062b0809 0a0b0c0d 0e0f1011 …
```

Advantages of the Comcast Transition Plan

- Initially, IPv6 domain completely private
  - Customers do not “touch” it
  - Translation, tunneling not an issue
- Incremental, hierarchical transition
- Allows extended training and operational experience period before “going public”
- Time to research effects
- Time to pressure vendors as needed

Host Operating Systems Supporting IPv6

- Microsoft
  - Windows XP
  - Windows Server 2003
  - Windows Vista
  - Longhorn Server
- Apple OS X
- Linux
- FreeBSD
- NetBSD
- HP-UX
- Solaris
- IBM AIX
- Tru64
- OpenVMS

Routing IPv6

Versions of all IPv4 routing protocols are available for IPv6

- Static routes
- RIPng
- Cisco Systems EIGRP
- OSPFv3
- IS-IS
- BGP

Underlying features to look for:

- Neighbor authentication
- Bidirectional Forwarding Detection (BFD)
- Graceful Restart
- Routing policy capabilities
  - Source and destination address matching
  - Next-hop matching
  - Prefix filtering
- CoS capabilities
  - BA and MF classifiers
  - Unicast RPF
  - DHCP marking
  - Policing

Routers Supporting IPv6

- 6Wind
- AddPac Technologoes
- ALAXALA Technology
- AlaxalA Networks
- Alcatel
- Allied Telesis
- Allied
- Telesyn
- Alpha Networks
- Beijing Jiaxun Feuihong Networks
- Billion Electric
- BITWAY Networking Technology
- BSD (KAME)
- Cisco Systems
- Delta Networks
- Digital China Networks
- D-Link
- DrayTek
- Extreme Networks
- Fiberhome Networks
- Fujitsu
- FugtureSoft
- Harbour Networks
- Hitachi
- Huawei
- IP-infusion
- Juniper Networks
- LG-Nortel
- Linux (USAGI)
- Mercury Corporation
- Microsoft
- NEC
- Netgenetech
- NextHop
- Nokia
- Nortel
- Novell
- Omron
- Panasonic
- Ruijie Networks
- Samsung
- Shanghai Baud Data Communication
- Tsinghua Unisplendor Bitway Networking
- Xoprt
- Yamaha
- Zebra
- ZTE Corporation
- ZyXEL

IPv6 Core Support

Any IPv6 stack *must* support these functions, at a minimum:

- Link IPv6 address configuration
- Global IPv6 address configuration
- ICMPv6
- Neighbor Discovery Protocol
  - Neighbor discovery
  - Router discovery
  - Duplicate address detection
  - Neighbor unreachability detection
- Stateless address autoconfiguration
- Stateful address autoconfiguration
- Path MTU discovery