Messaging Anti-Abuse Working Group

MAAWG IPv6 Training for Senders and Others

Segment 2 of 3

Understanding the Drivers & Myths Behind IPv6

October 4th-6th, 2010

MAAWG 20th General Meeting, Washington DC

Messaging Anti-Abuse Working Group

Joe St Sauver, Ph.D.,

MAAWG Sr. Technical Advisor

joe@oregon.uoregon.edu

www.uoregon.edu/~joe/maawg-senders-ipv6-training/

Disclaimer: All opinions expressed in this talk are strictly my own, and do not necessarily represent the opinions of any other entity. This talk is provided in a detailed written form to insure accessibility and for ease of web indexing.

MAAWG IPv6 Training – Video Segments

Segment 1 - 21 minutes Basic IPv6 For Senders	Segment 2 – 25 minutes Understanding IPv6 Drivers & Myths	Segment 3 – 21 minutes. IPv6 Technology Primer
1-Do Mail Senders *Really* Need IPv6 Today?	Drivers: 6a-IPv4 Address Exhaustion	8a-IPv6 Addresses
2-IPv6 and SMTP	6b-Regaining Internet Transparency	8b-IPv6 Prefixes
3-Obtaining Native IPv6 Connectivity and Address Space	6c-Controlling Route-Table Bloat	8c-Types of IPv6 Addresses
4-Enabling IPv6 on Your Servers	6d-Regulatory Compliance	8d-Addresses and Systems
5-Enabling IPv6 in Your MTA	Myths: 7a-Improve "Network Security"	8e-IPv6 DNS
	7b-*Everyone's* Running Out of IPv4 Address Space	8f-Enabling IPv6 in Desktop Operating Systems
	7c-IPv6 Will Simplify Renumbering . . .	8g-IPv6 Web Browsers
	7d-Access to "Cool New Stuff"	8h-IPv6 Email Clients
	7e-Techies Will Stretch Out What IPv4 Space	8i-ssh for IPv68. (j): Web Servers
	7f-Customers Just Aren't Asking For IPv6	

PART II.
Understanding the Drivers and Myths Behind IPv6

6 (a). Understanding the Drivers for IPv6: IPv4 Address Exhaustion

IPv4 Addresses: A Scarce Resource

- There is a finite pool of available IPv4 addresses, and we're getting really, really close to running out.
- Based on the best available forecasts, see http://www.potaroo.net/tools/ipv4/index.html , the last IPv4 blocks will be allocated by IANA on 27-May-2011
- **The regional internet registries (RIRs), such as ARIN, RIPE, APNIC, LACNIC and AFRINIC will exhaust the address space they've received from IANA less than a year later, around 21-Jan-2012**
- These best estimates are based on current trends, and actual exhaustion might accelerate (or might slow) depending on what the community does (but probably not by much). **From now till 21-Jan-2012 is roughly 1 year 3 and a half months away.** That's <u>not</u> much time.

inetcore.com/project/ipv4ec/index_en.html

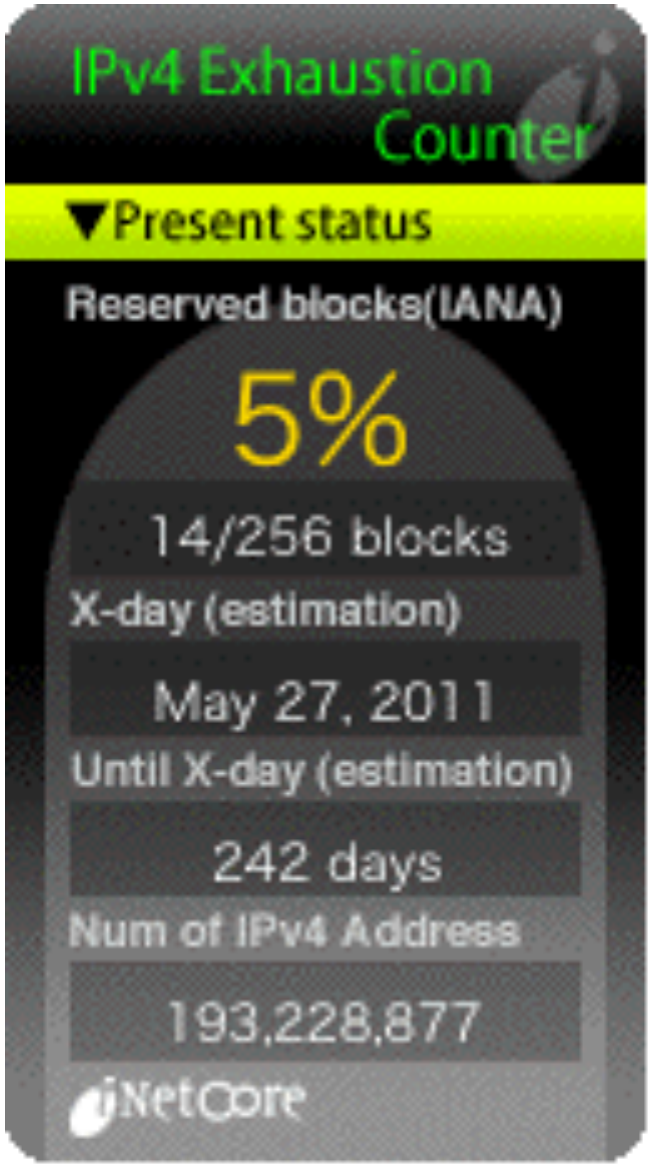

The Internet, Post-IPv4 Run Out

- Running out of IPv4 addresses isn't like running out of water in the middle of the desert, or air while SCUBA diving -- if you already have IPv4 address space, the space you already have will continue to work just fine.
- People who WILL run into problems, however, include:
  -- growing ISPs who need more IPv4 IP addresses
  -- new ISPs who need IPv4 addrs just to get started
  -- customers of existing IPv4-based ISPs who may need to access network resources available ONLY via IPv6
  -- customers behind weird/broken stopgap kludges
- Eventually, we risk the bifurcation of the Internet: part of the Internet may cling to IPv4 addressing, while the rest may end up having no choice but to use IPv6 addressing. Eventually, this will be a serious issue.

6 (b). Understanding the Drivers for IPv6: Regaining Internet Transparency

"But What About NAT?"

- While some sites (including uoregon.edu) assign each system on campus a globally routable IP address. Other sites (including many home users and many corporate sites) routinely employ network address translation (or "NAT"). NAT (actually PAT) makes it possible for multiple workstations to all use a single shared globally routable IPv4 address. If all you do is browse the web or use a web email service such as Hotmail, or Yahoo! Mail, or Gmail, NAT may superficially work fine for your needs.
- On the other hand, if you want to do Internet video conferencing, or use peer-to-peer applications, or you're trying to track down and fix malware-infested hosts connecting from behind a NAT, you may find that NAT will make your life significantly more difficult.

NAT: A (Semi) Protocol-Aware Protocol

- Some network protocols (such as H.323) embed IP addresses in the traffic generated by those protocols.
- Because NAT rewrites network addresses, it needs to know HOW each protocol embeds IP addresses in network traffic streams. That is, NAT boxes need to be "protocol aware," and thus networks using NAT are NOT "end-to-end transparent." (Packets get rewritten during transport while passing through a NAT)
- If a NAT box faces traffic of a type that it doesn't know how to handle, such as some new protocol, it can't rewrite that traffic, and as a result that application will fail when run behind a NAT. This is very commonly the case for H.323 video conferencing, for example.
- Because of this, NAT'd networks can stifle application layer network innovation, or at least make it far harder!

The "Two Port Internet"

- Because of the problems that application developers face getting past NAT boxes (and restrictive firewalls!) it is common for developers to implement new protocols over http instead of developing new native protocols. Some of my colleagues refer to this as the "two port Internet" -- in this model, virtually all user traffic is either http (port 80) or https (port 443).
- Obviously this is something of an exageration (they forgot about DNS for example :-)), but it isn't entirely an argument w/o merit. All you need to do is look at network traffic and try to identify what applications make up most of the traffic to see the problem -- you can't do it any more just based on ports.
- C.F.: "A Look at the Unidentified Half of Netflow," pages.uoregon.edu/joe/missing-half/missing-half.pdf

End-To-End Transparency

- If you'd like to read about the importance of end-to-end transparency, some excellent starting points are:

  -- RFC2775, "Internet Transparency," B. Carpenter, February 2000, http://tools.ietf.org/rfc/rfc2775.txt

  -- RFC4924, "Reflections on Internet Transparency," B. Aboba and E. Davies, July 2007, http://tools.ietf.org/rfc/rfc4924.txt

- While Internet transparence is less often mentioned than imminent IPv4 address exhaustion as a reason why we need to deploy IPv6, transparency is nonetheless a very important underlying motivation for IPv6.

6 (c). Understanding the Drivers for IPv6: Controlling Route-Table Bloat

Controlling Route Table Bloat

- Another important (if little recognized) reason for promoting use of IPv6 has been the need to control the growth in the size of the global routing table. In fact, RFC4984 (http://www.ietf.org/rfc/rfc4984.txt) states,

  "[...] routing scalability is the most important problem facing the Internet today and must be solved [...]"

The IPv4 Route Table Continues to Grow…

350,000

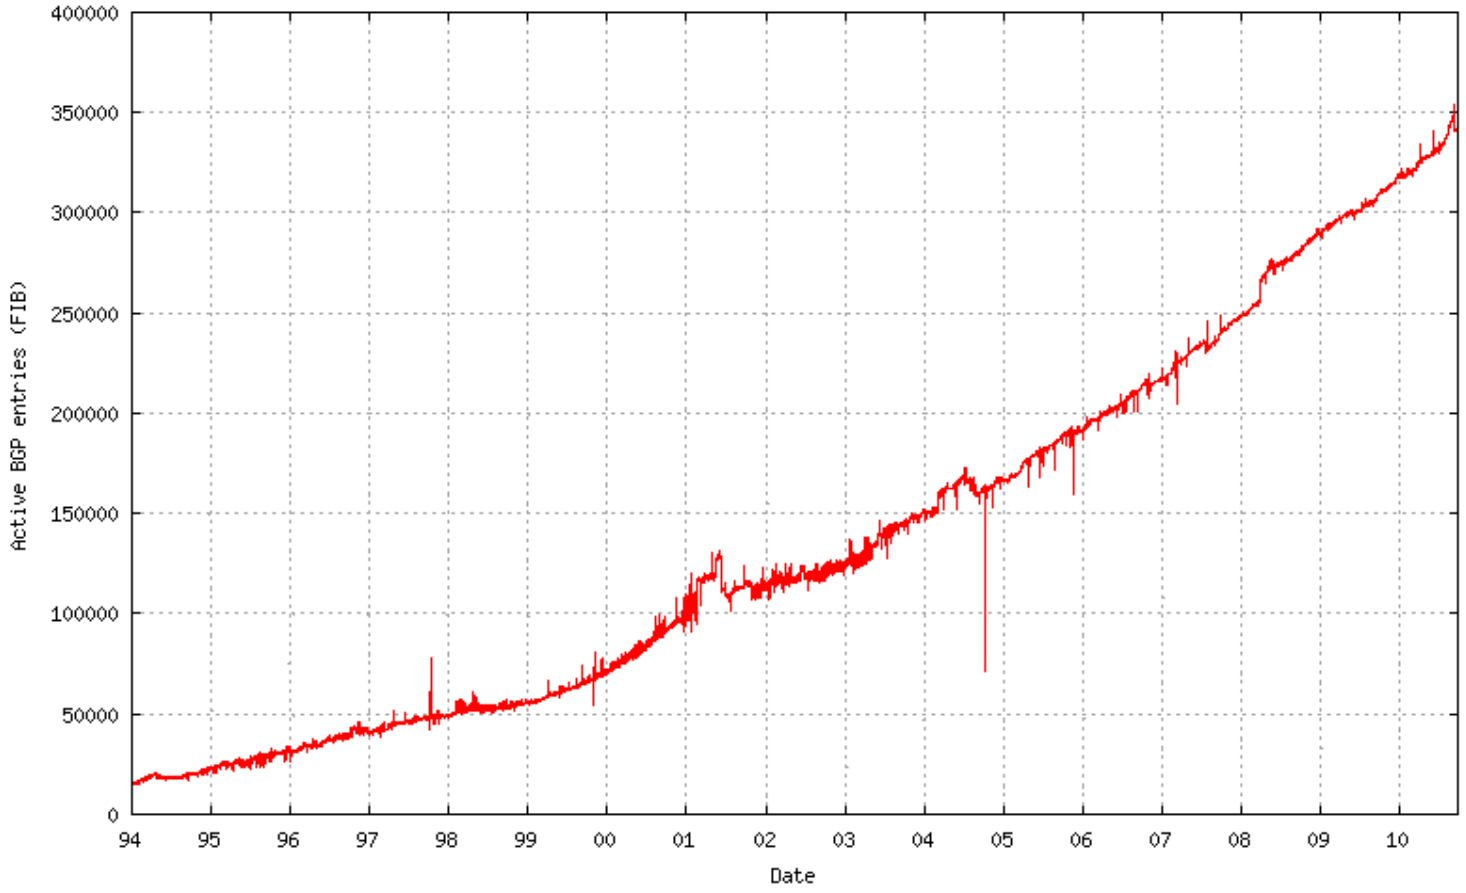

Source: http://bgp.potaroo.net/as6447/

IPv6 Was Supposed to Help Fix That

- When IPv6 was designed, address assignment was supposed to be hierarchical. That is, ISPs would be given large blocks of IPv6 address space, and they'd then use chunks of that space for each downstream customer, and only a single entry in the IPv6 routing table would be needed to cover ALL the space used by any given ISP and ALL their downstream customers (see RFC1887, "An Architecture for IPv6 Unicast Address Allocation")
- But now, let's pretend that my Internet connectivity is important to me, so I don't want to rely on just a single ISP -- I want to connect via multiple ISPs so that if one provider has problems, the other ISPs can still carry traffic for my site. This connection to multiple sites is known as "multihoming."

If I'm Multihomed, Whose Address Space Do I Use?

- When I get connectivity from sites A, B and C, whose address space would I announce? Address space from A? Address space from B? Address space from C? No...
  -- A doesn't want me to announce part of its address space via B and C
  -- B doesn't want me to announce part of its address space via A and C
  -- C doesn't want me to announce part of its address space via A and B.
- I need to either assign each host multiple addresses (e.g., one address from A, one from B, and one from C), or I need to get my own independent address space which I can use for all three ISPs, but which will then take up a slot in the global routing table.

The Original Multiple IP Approach in IPv6

- The multiple IP approach was the original "answer" to this question in the IPv6 world.
- But if I assign multiple IPs to each host, one for each upstream ISP I connect to, how do I know which of those IP addresses I should use for outbound traffic generated by each host? Do I arbitrarily assign the address from A to some traffic? The address from B to other traffic? What about the address from C?
- Which of those addresses do I map to my web site or other servers via DNS? Do I use just A's address? Just B's? Just C's? All three of those addresses? What if one of my providers goes down? Will traffic failover to just the other two providers quickly enough?

The Multihoming Reality Today

- IPv6 multihoming without use of provider independent address space is one of the unsolved/open issues in the IPv6 world today. Operationally, in the real world, ISP customers who need to multihome request their own provider independent IPv6 address space (cue Sonny and Cher: "The beat goes on, and the beat goes on...")
- Route table growth may be a critical issue facing the Internet in the long term, but for now, the community has "dropped back into punt formation," and we're doing what needs to be done (at least for now) to get IPv6 deployed in a robust way (e.g., with multihoming). The good news is that the IPv6 table is still small (so we still have time to solve the IPv6 routing table growth issue); the bad news is that the IPv6 table is still small (which means many people still haven't deployed IPv6!)

IPv6 Route Table Growth

3500

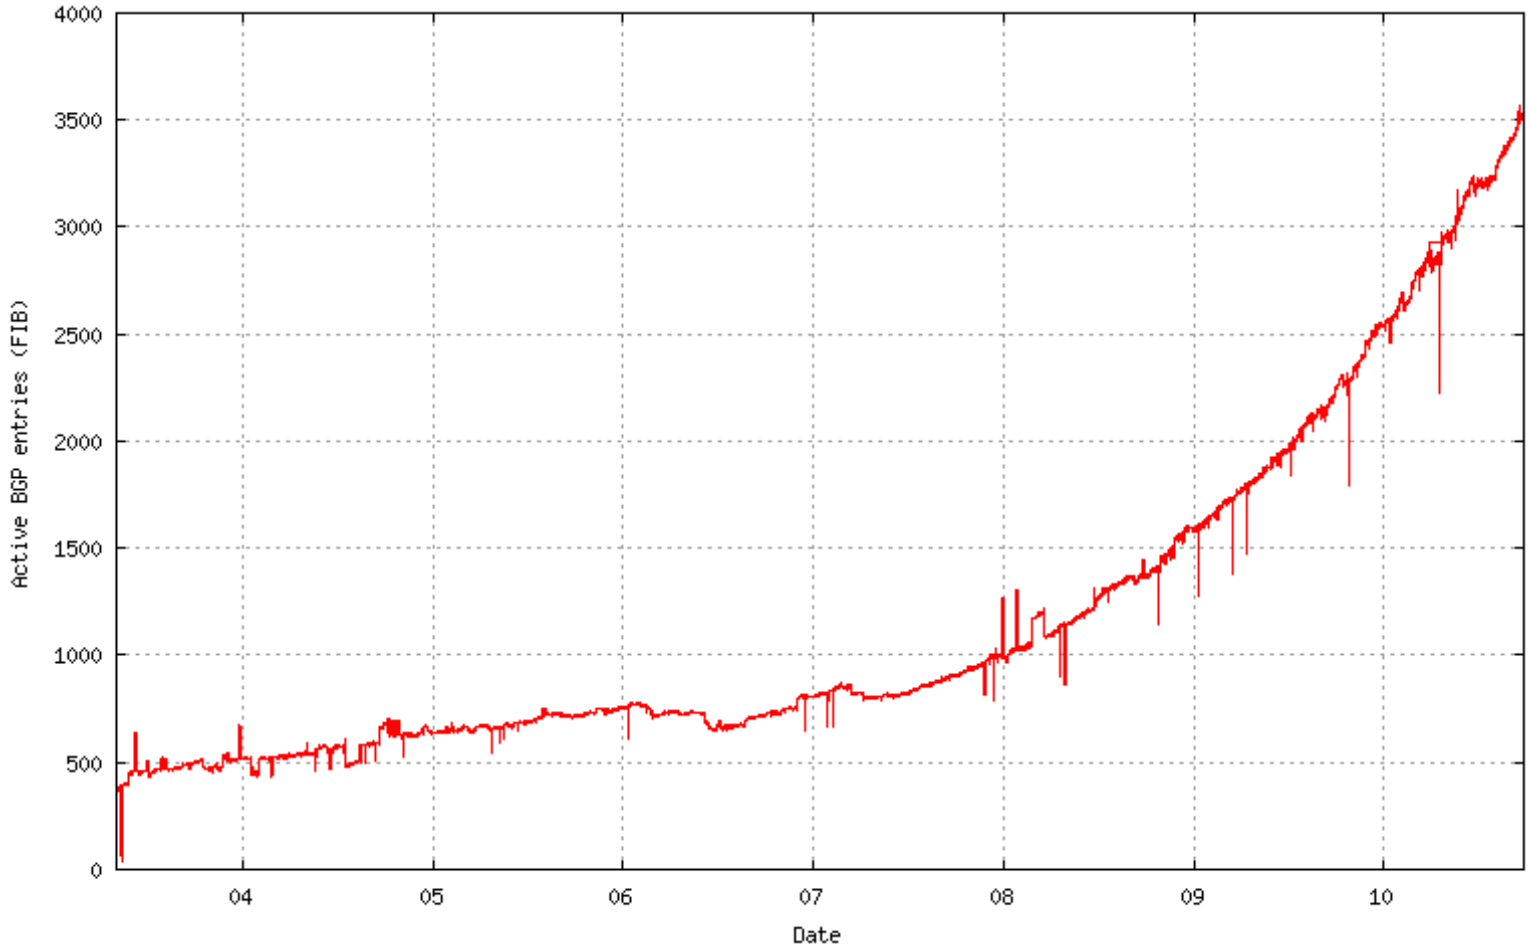

Source: http://bgp.potaroo.net/v6/as6447/

6 (d). Understanding the Drivers for IPv6: Regulatory Compliance

Federal Networks, For Example, Are Supposed to Be IPv6 Ready

M-05-22

August 2, 2005

MEMORANDUM FOR THE CHIEF INFORMATION OFFICERS

FROM: Karen S. Evans
Administrator
Office of E-Government and Information Technology

SUBJECT: Transition Planning for Internet Protocol Version 6 (IPv6)

As I stated in my testimony of June 29, 2005, before the House Committee on Government Reform, we have set June 2008 as the date by which all agencies' infrastructure (network backbones) must be using IPv6 and agency networks must interface with this infrastructure. This memorandum and its attachments provide guidance to the agencies to ensure an orderly and secure transition from Internet Protocol Version 4 (IPv4) to Version 6 (IPv6). Since the Internet Protocol is core to an agency's IT infrastructure, beginning in February 2006 OMB will use the Enterprise Architecture

Source: www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf

Theoretically All Federal Networks (At Least Temporarily) Met That Mandate, But…

- Reportedly many federal networks, having passed one IPv6 packet (and thus, however briefly, demonstrated that their backbones were IPv6 capable), promptly "re-disabled" IPv6 (ugh!)
- Check your favorite fed sites -- are they v6 accessible? See: http://www.mrp.net/cgi-bin/ipv6-status.cgi

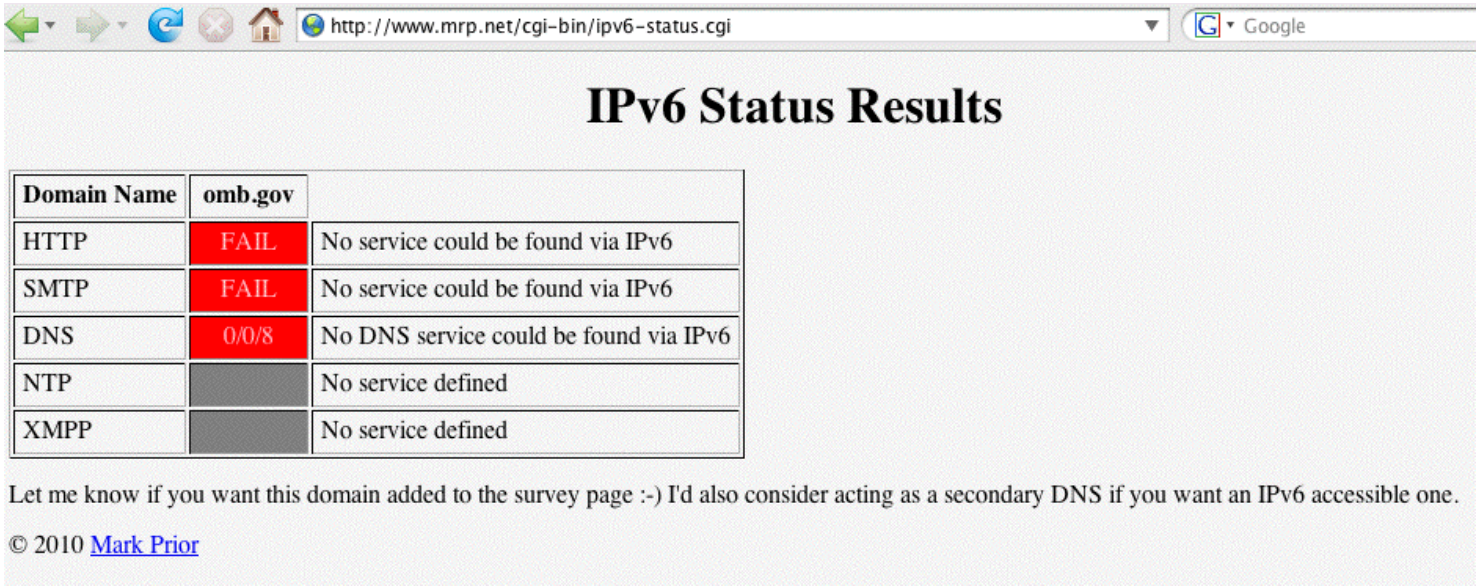

IPv6 and Fed Scorecard Network Homepages?

www.dhs.gov --> no
www.doc.gov --> no
www.dod.gov --> no
www.doe.gov --> no
www.dot.gov --> no
www.ed.gov --> no
www.epa.gov --> no
www.hhs.gov --> no
www.hud.gov --> no
www.doi.gov --> no
www.doj.gov --> no
www.dol.gov --> no
www.nasa.gov --> no
www.nsf.gov --> no
www.nrc.gov --> no
www.opm.gov --> no
www.sba.gov --> no
www.ssa.gov --> no
www.state.gov --> no
www.usaid.gov --> no
www.usda.gov --> no
www.ustreas.gov --> no
www.va.gov --> no

Or pick another federal agency of your choice: the pattern is pretty consistent I'm afraid…

"Planning Guide/Roadmap Toward IPv6 Adoption Within the US Government"

- This is a new document (ca. May 2009) from the Federal CIO Council Architecture and Infrastructure Committee, Technology Infrastructure Subcommittee, Federal IPv6 Working Group, see http://tinyurl.com/fed-cios-ipv6
- I quote: "The purpose of this document is to provide U.S. government agency leaders with practical and actionable guidance on how to successfully integrate Internet Protocol version 6 (IPv6) throughout their enterprise. [...] without a concentrated effort by Federal agencies to effectively and efficiently deploy secure IPv6 network services, the Government's technical advancement and ability to meet its mission needs will be critically impacted during the next 2 to 3 years."

And Just This Last Week…

- On Sept. 28th, 2010, the NTIA held a workshop at which Federal CIO Vivek Kundra announced a directive "requiring all U.S. government agencies to upgrade their public-facing Web sites and services by Sept. 30, 2012, to support IPv6…" and that access must be via native IPv6 rather than an IPv6 transition mechanism.
- A second deadline, Sept. 30th, 2014, applies for federal agencies to upgrade internal client applications that communicate with public servers to use IPv6.
- For more, see
  "White House Issues IPv6 Directive,"
  http://www.networkworld.com/news/2010/092810-white-house-ipv6-directive.html?page=1

A Major Potential Stumbling Block: Non-IPv6 Content Delivery Networks (CDNs)

- Many federal web sites (and key commercial web sites) use Akamai (or another CDN) in order to handle huge online audiences, deliver good performance worldwide, and to resist DDoS attacks.
- For example, www.irs.gov is actually just a cname for www.edgeredirector.irs.akadns.net; whois confirms that akadns.net actually belongs to Akamai.
- If Akamai doesn't do IPv6, will current major Akamai customers (such as Apple, Cisco, Microsoft, RedHat, the Whitehouse, etc.) be willing to deploy IPv6 for critical sites without them?
- BTW, at least one vendor, Limelight, DOES offer an IPv4 **and IPv6** CDN service…

But Speaking of Akamai, Akamai *Is* Reportedly Working On IPv6…

- I'm happy to report that Akamai is now reportedly working on IPv6-ifying its CDN infrastructure. See, for example, the coverage in:

  "Akamai: Why Our IPv6 Upgrade Is Harder Than Google's,"
  http://www.networkworld.com/news/2010/091610-akamai-ipv6.html
  September 16th, 2010

The Issue Isn't Just Web CDNs…

- A growing number of sites **outsource their email operations.** Unfortunately some email-as-a-service (and some cloud-based spam filtering services) don't support IPv6, thereby limiting the ability of their customers to integrate IPv6 into their existing IPv4-based services.
- CDNs and outsourced email and spam filtering services aren't the only reason why IPv6 adoption has been slow at some major Internet sites, but they can certainly be stumbling blocks that will need to get resolved.
- Some of these outsourced email/spam filtering services and CDNs are being actively monitored on MRP's survey page, see www.mrp.net/IPv6_Survey.html , including Akamai, Google, McAfee, MessageLabs, MacQuarie, Microsoft, Postini, ProofPoint, RedCondor and Websense.

7 (a). Understanding IPv6 Myths: 'IPv6 Will Improve "Network Security" Due to IPv6 Having "Mandatory" IPSec'

(Sorry, No)

IPv6 and IPsec

- IPsec is not new with IPv6; in fact, IPsec dates to the early 1990's. What's different when it comes to IPv6 is that support for IPsec was made "mandatory" for IPv6 (see for example "Security Architecture for IP," RFC4301, December 2005 at section 10, and "IPv6 Node Requirements," RFC4294, April 2006 at section 8.)
- **If actually used**, IPsec has the potential to provide:
  -- authentication
  -- confidentiality
  -- integrity, and
  -- replay protection
- All great and wonderful security objectives -- **IF** IPsec gets used. **Unfortunately, as we'll show you, what was supposed to be the cornerstone of the Internet's security architecture has proven in fact to be widely non-used.**

IPv6 Does NOT Require IPSec to Run...

- IPv6 can be brought up without IPSec getting enabled, and in fact this is routinely the case -- see an example on the next slide.
- More broadly, if people are doing cryptographically secured protocols of *any* sort, they inevitably run into problems -- crypto stuff just tends to be inherently complex and hard to learn to use. For example, how many of you routinely use PGP or GPG to cryptographically sign or encrypt your email, eh? How many of you are doing DNSSEC to cryptographically protect the integrity of your DNS traffic?
- Now think about how often you see people moaning about problems they're having getting IPSec to work with IPv6 -- do you EVER see that sort of thing on the mailing lists or discussion groups you're on? No, right? That's because hardly anyone is doing IPSec with IPv6.

Some IPv6 Traffic Statistics From A Mac OS X Host (Note That It's NOT Doing IPSec)

```
# netstat -s -finet6   <-- let's look just at IPv6 traffic
[snip]

ip6:
        124188 total packets received
        [snip]
        84577 packets sent from this host
        [snip]

    ipsec6:              <-- how much of it was Ipsec protected?
        0 inbound packets processed successfully
        0 inbound packets violated process security policy
        [snip]
        0 outbound packets processed successfully
        0 outbound packets violated process security policy
        [snip]
```

IPsec (Even on IPv4!) Isn't Getting Much Use

- Raw IPsec traffic (AH+ESP, protocols 50 & 51) isn't seen much on the commercial IPv4 Internet.

- For example, about a year or so ago, Jose Nazario of Arbor Networks estimated IPsec traffic at 0.9% of octets (statistic courtesy the ATLAS project).

- CAIDA (thanks kc!) also has passive network monitoring data available; see

  http://www.caida.org/data/passive/monitors/equinix-chicago.xml

  You can see the protocol distribution from a couple of CAIDA's monitors for one recent day on the next couple of slides. IPsec traffic is basically too small to even be seen for the most part.

Protocol Distribution From One of CAIDA's Passive Monitors

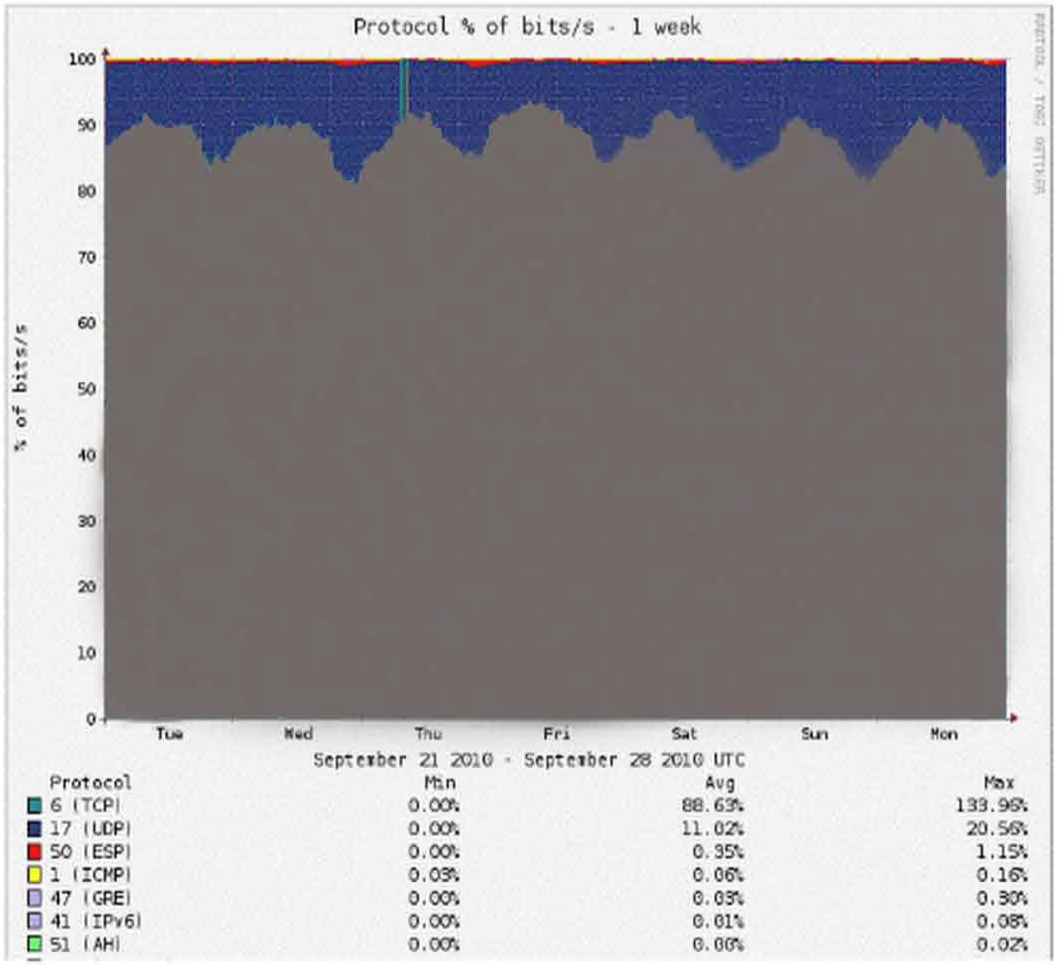

Bad News/Good News

- **It would be foolhardy to expect IPsec to provide any material improvement to your site's security since the vast majority of your aggregate traffic (including virtually <u>all</u> your IPv4 traffic) will NOT be IPsec secured.**

- On the other hand, the "good news" is that a lack of IPsec usage in the IPv6 world is substantively no worse than a lack of IPsec usage in the IPv4 world.

7. (b) Understanding IPv6 Myths: "IPv6 Will Simplify Renumbering, Improve Routing Performance By Simplifying Packet Formats, Improve Support For QoS, Facilitate Mobility, etc."

(maybe, but frankly, no one cares)

I Really Don't Mean to Sound Harsh, But…

- IPv6 may very well bring many cool new features to networking, but quite frankly, all of these incidental new features really don't matter -- they're not "make or break" drivers when it comes to adoption of IPv6.

- I'd love to hear hard evidence to the contrary, but truly, I've seen no indication that any of these other factors carry much weight in helping to shape the IPv6 go/no-go decision.

- On the other hand, there are some genuine reasons why people adamantly are NOT interested in doing IPv6.

7. (c) Understanding IPv6 Myths: *Everyone's* Running Out of IPv4 Address Space and Will Need IPv6

(No, Actually Some Sites Still Appear To Have Quite A Bit of IPv4 Space Left)

If You Have Abundant IPv4 Address Space, It May Be Hard to Get Excited About IPv6

- The definition of "abundant" will vary from site to site, but, for example, many universities have legacy /16's, and 2^16 (or roughly 65,500 addresses) can seem like "a lot" of addresses (even though they can go awfully fast when you have a campus of 20,000-25,000 people, most of whom have multiple networked devices, plus lots of printers and other networked infrastructure).
- Others, however, have /8's (2^24, or roughly 16,777,200 addresses per /8), and that may be large enough to eliminate all worries at those sites about address scarcity.
- If you haven't recently looked at the list of who has /8's, you can check http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space

Considering Just Large Federal Netblocks…

- The US government (typically as the US Department of Defense) controls a relatively large fraction of the entire Intenet IPv4 space, including a dozen slash 8's: 6/8, 11/8, 21/8, 22/8, 26/8, 28/8, 29/8, 30/8, 33/8, 55/8, 214/8, and 215/8 (to say nothing of additional /8 blocks controlled by defense industrial contractors with close/ extensive military contracts, plus miscellaneous smaller netblocks scattered hither and yon). Given that level of public IPv4 address space availability, one can understand that some federal agencies have not felt a particularly pressing need to move to IPv6 real ricky-ticky soon now.
- An aside about those large federal netblocks: even though addresses from some of those block may not show up in public routing tables, they <u>are</u> still being used, you're just not seeing them in some <u>public</u> routing tables.

Another Manifestation of the "Oh, No Worries, I've Got Plenty of Address Space" Phenomena

- While there are obviously only a comparative handful of sites which have their own public /8, many sites use RFC1918 "private" non-publicly routable address space, such as addresses from 10/8.
- If NAT currently meets your needs, as long as you have 10.0.0.0-10.255.255.255 (and 172.16.0.0-172.31.255.255, and 192.168.0.0-192.16.255.255) available, again, you may feel like your private IPv4 addressing needs are generally being well met (provided, of course, that you can also get the comparative handful of public IPv4 addresses you may also need).
- But, as we'll see later, use of NAT forecloses some of the easiest IPv6 transition mechanisms, such as use of 6to4.

7 (d). Understanding IPv6 Myths: 'If You Enable IPv6, You'll Suddenly Have Access to Lots of "Cool New Stuff" That's Not Available via IPv4'

(Nope. Anything That's Available Via IPv6, Will Also Be Equally Available via IPv4)

An Important/Subtle Point to Understand

- Sometimes folks ask, "So if I begin to do IPv6, ***what new stuff can I get at that I can't get at already?***"
- Can you imagine Google, or Amazon, or Microsoft, or CNN, or <fill in the name of an important Internet site here> making their web site or other online resources ONLY available via IPv6? No, probably not.
- Any/all important Internet resources will ALWAYS be available via IPv4, even if those resources are ALSO available via IPv6.
- Thus moving to IPv6 does NOT magically give you access to new stuff that you couldn't get to via IPv4 (well, technically there are a few IPv6-only things mentioned at http://www.sixxs.net/misc/coolstuff/ , but nothing in and of itself that's enough to justify deploying IPv6)

At Some Point In The Future, Though, The Default *Will* Change

- At some point in the future, people will eventually ask, "So if I still bother doing IPv4, ***what sort of 'old stuff' can I get at that I can't get at already via IPv6?***"

- We're still a LONG way off from that point, but it WILL eventually happen.

- Remember when people used to carry AppleTalk or IPX or DECNet on their local area networks? They sure don't anymore (or at least no one I know still does!)

7 (e). Understanding IPv6 Myths: "All The Techies Will Figure Out Some New Way To Stretch Out What IPv4 Space We've Got Left -- I'm Not Going to Worry"

(I Admire Your Confidence,
But Personally, I Don't Think
I'd Be So Sanguine)

We're Not The (Only) Ones Driving The Address Consumption Bus!

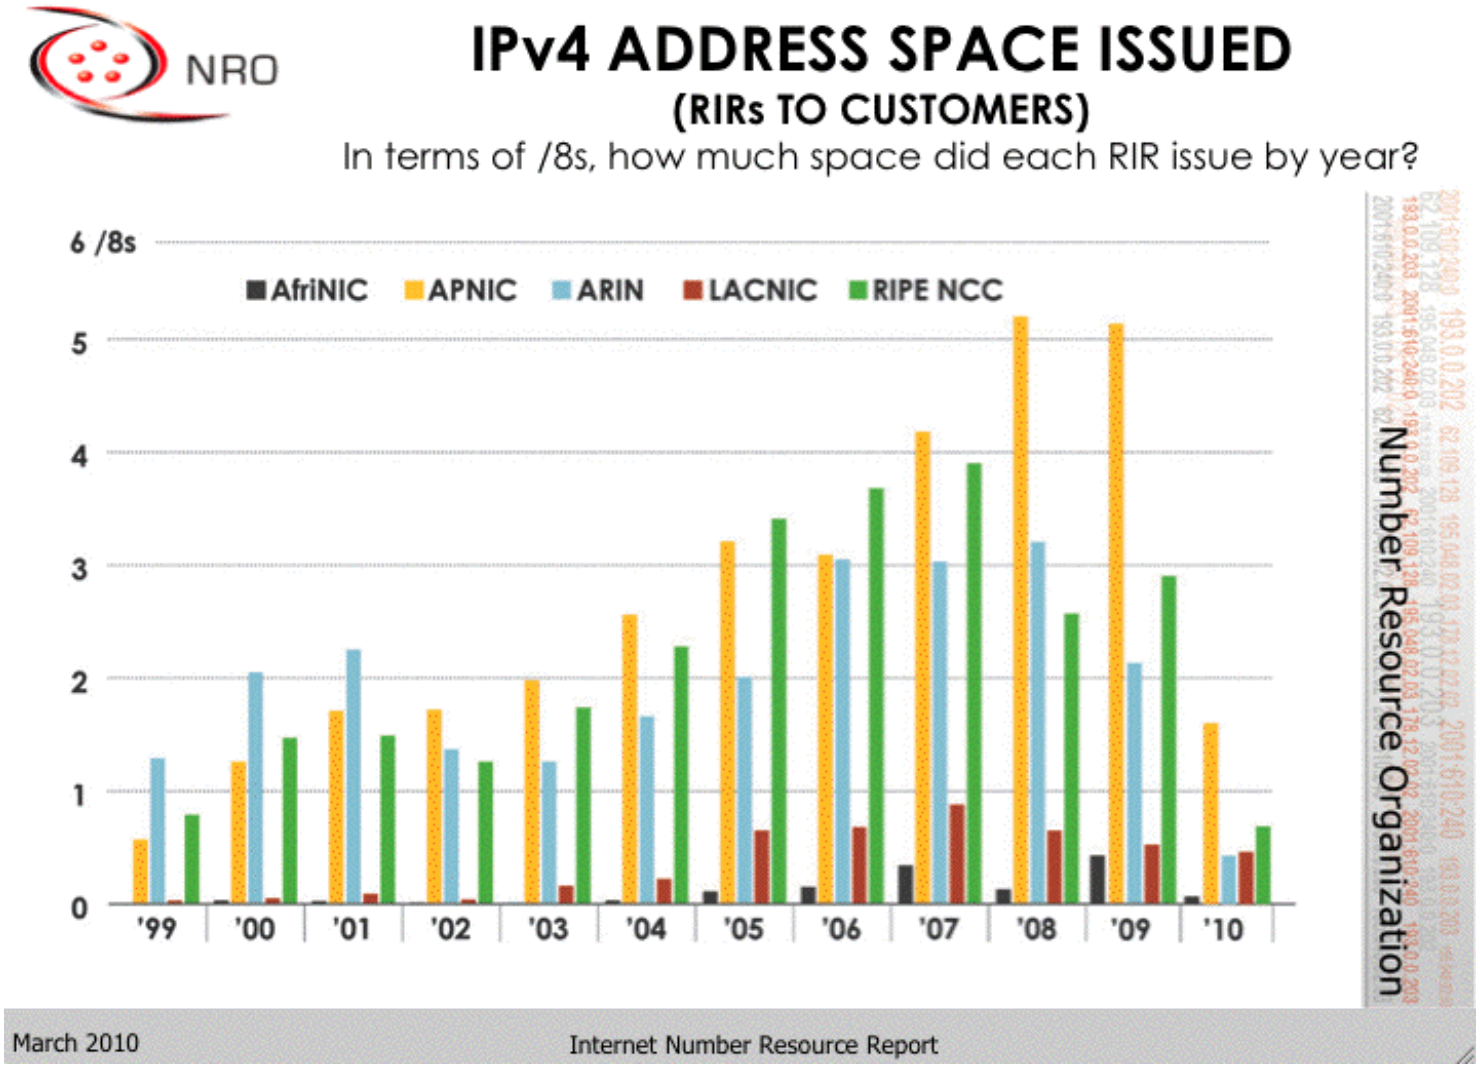

http://www.arin.net/participate/meetings/reports/ARIN_XXV/PDF/Monday/Nobile_NRO_joint_stats.pdf

A Cumulative View

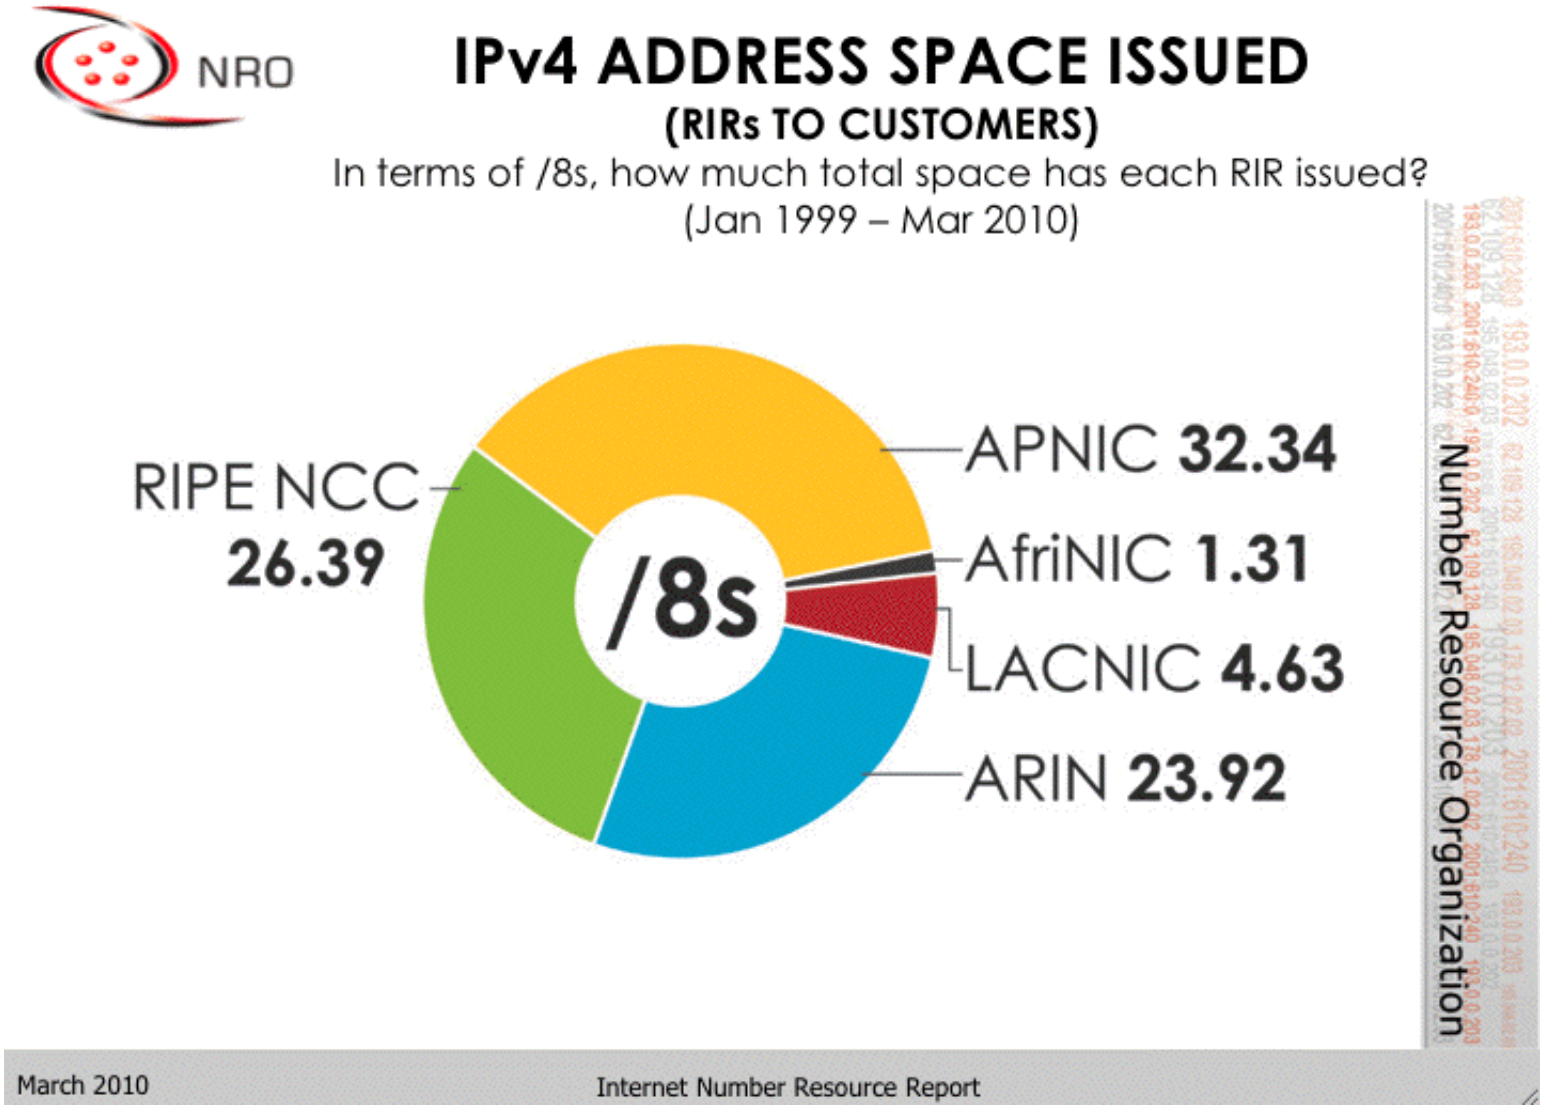

http://www.arin.net/participate/meetings/reports/ARIN_XXV/PDF/Monday/Nobile_NRO_joint_stats.pdf

What If IPv4 Address Usage Was Proportionate to Regional Population?

	Population	%	/8's	%	Ratio
• Asia:	4,121,097	60.3%	32.34	36.5%	0.605
• Africa:	1,009,893	14.7%	1.31	1.4%	0.095
• Europe:	732,206	10.7%	26.39	29.7%	**2.775**
• L. Amer.:	582,418	8.5%	4.63	5.2%	0.611
• N. Amer.:	348,360	5.1%	23.92	27%	**5.29**
• Oceania:	35,387	0.5%			
• Total:	6,829,360		88.56		

Population in thousands, mid year 2009 estimates

Note: Oceania's addresses are handled by APNIC (e.g., Asia)

http://esa.un.org/unpd/wpp2008/jpg/WPP2008_Wall-Chart_Page_1.jpg

The Future: Notice Relatively "Tiny" Africa

- Historically, Africa's IP address usage to date has been minimal, less than one and a half /8s.
- This was likely due to a variety of factors, but at least one important factor was the high cost of connectivity (thousands of dollars per Mbps per month vs. just dollars per Mbps per month in the US (for large customers)).
- Another driver was widespread use of satellite Internet connectivity, with high latency, NAT'd connections and IPs provided by the satellite operator.
- Improved fiber connectivity is changing all that. Some of the world's largest and most densely populated regions in both Africa and Asia are now coming online, and I believe the improved connectivity to those areas will result in a substantial demand for new IPv4 addresses.

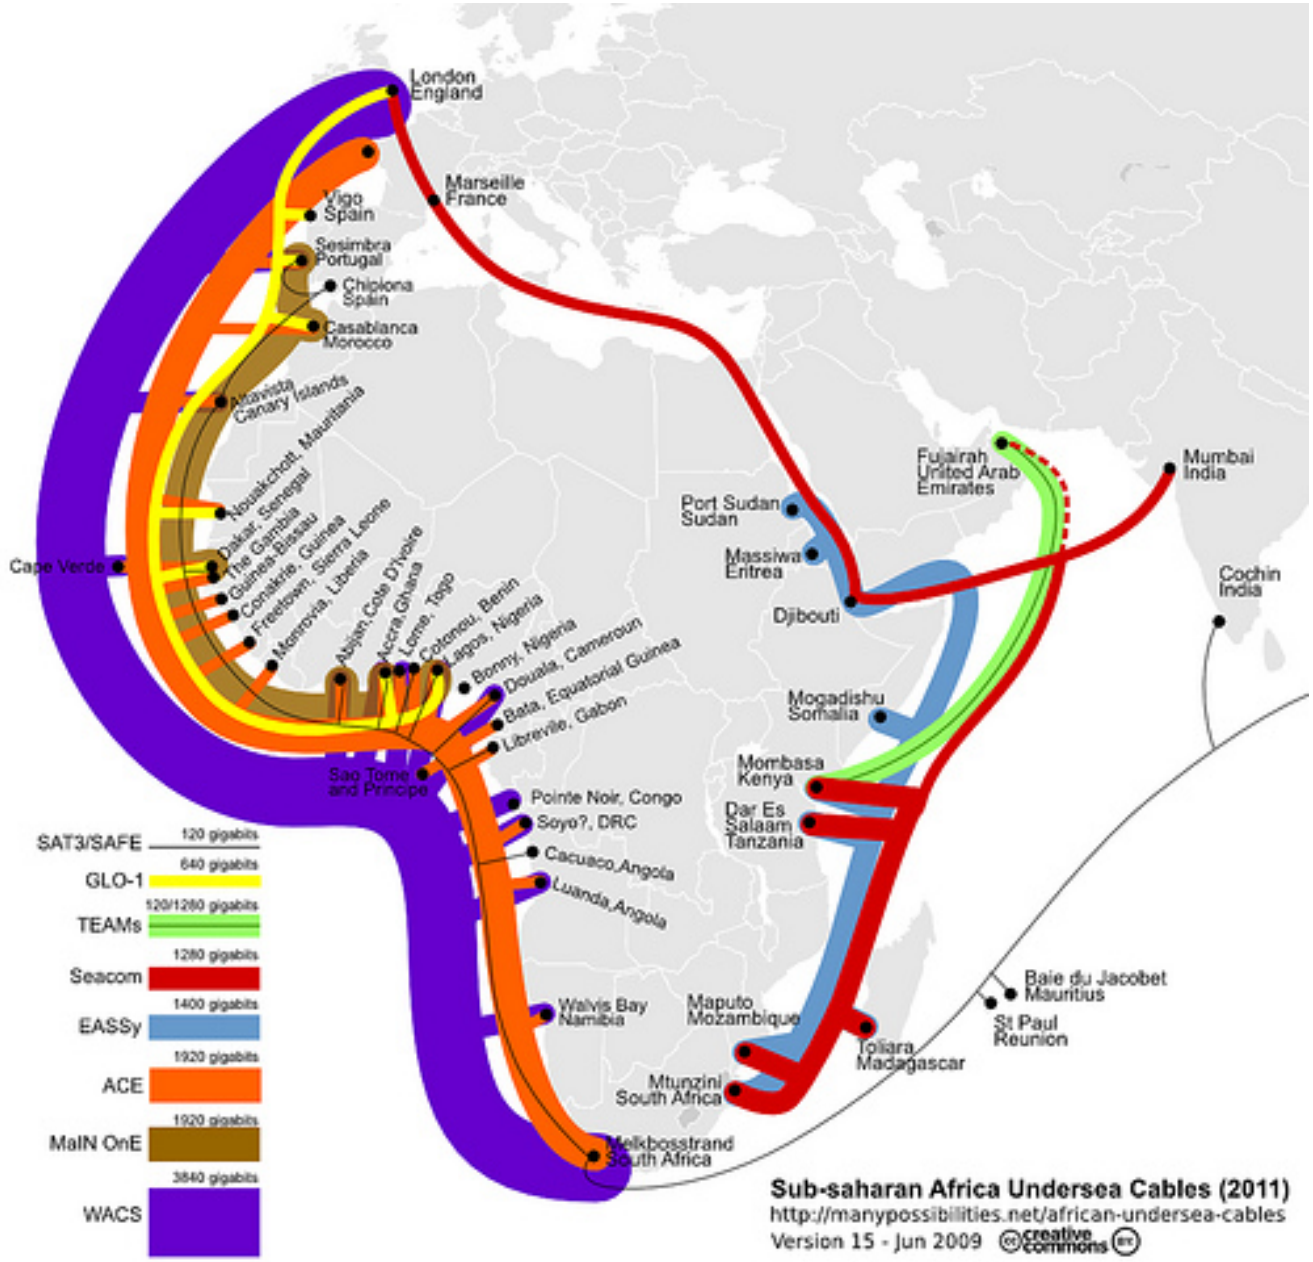

http://blog.foreignpolicy.com/files/images/090618_africa_underseas_cables.jpg

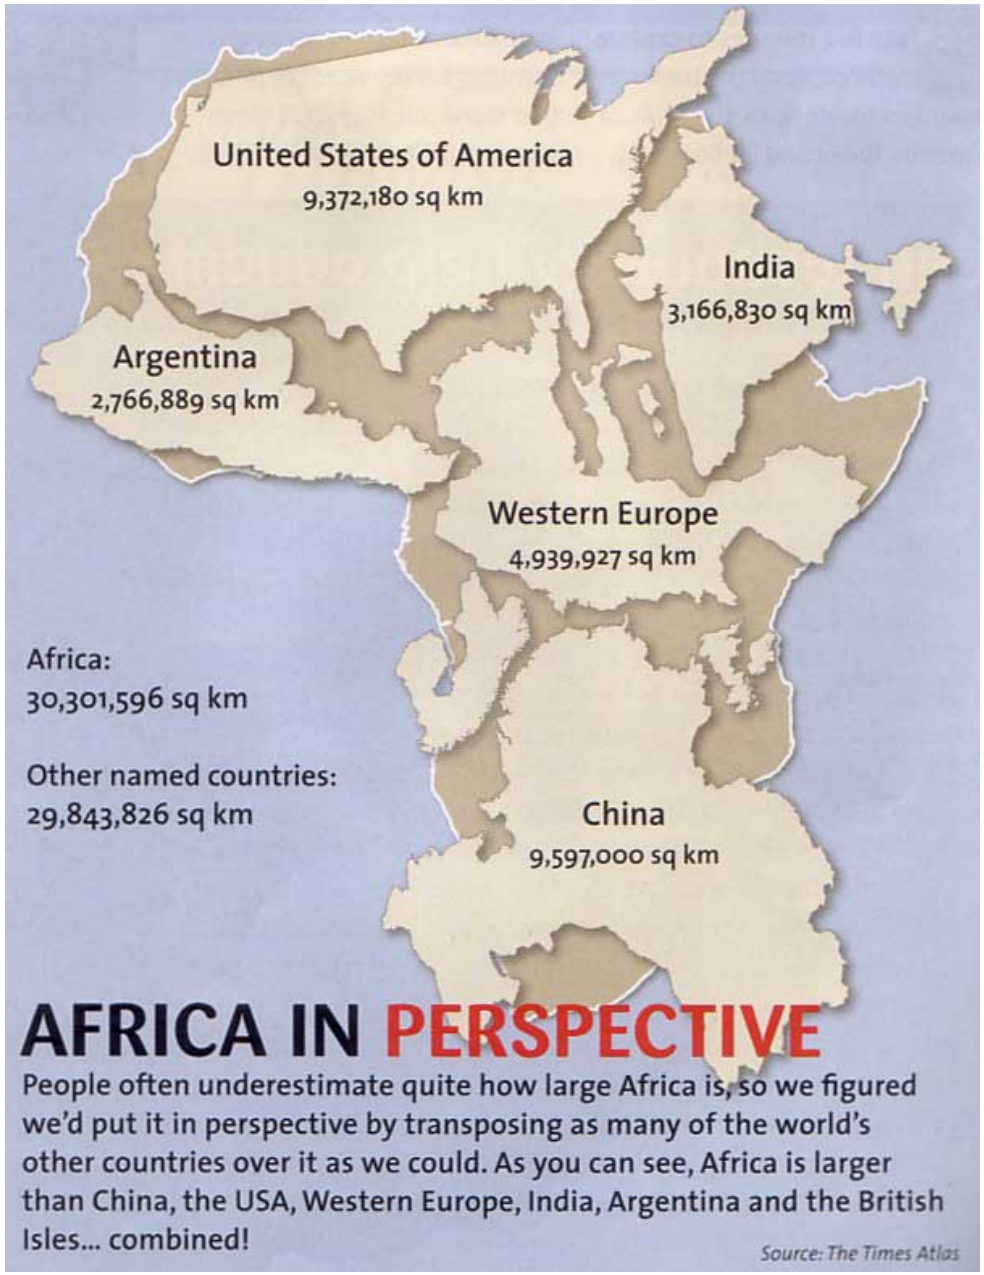

http://strangemaps.files.wordpress.com/2006/11/africa_in_perspective_map.jpg

If You Believe We Have Enough IPv4 Addresses

- Given the preceding slides, you must also believe in miracles! :-)
- In my case, I'd rather be prepared to work with IPv4 AND IPv6 :-)

7 (f). Understanding IPv6 Myths: "Customers Just Aren't Asking For IPv6 (Except You!)"

That's Not The Impression *I'm* Getting…

IPv6 Allocations Issued (/32s)

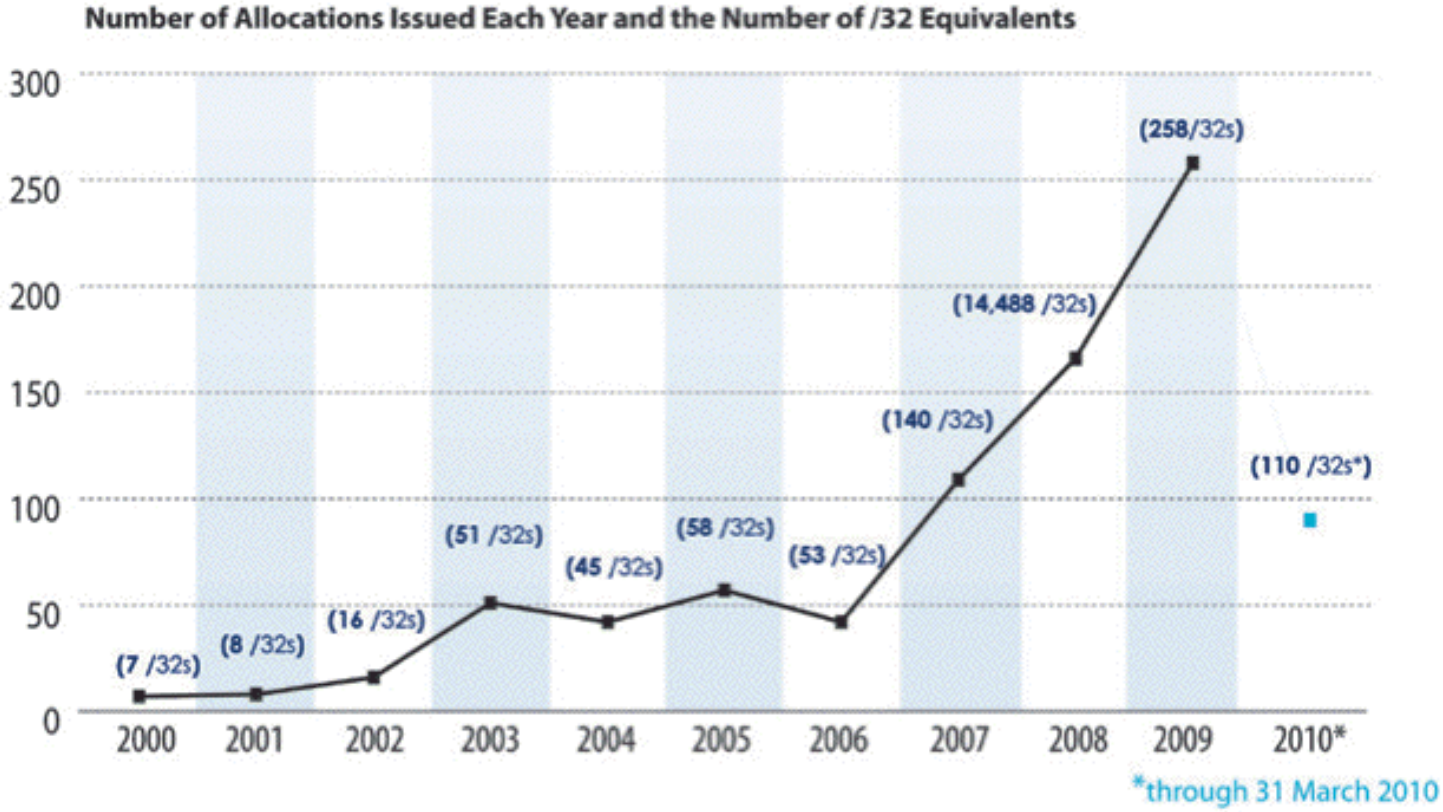

Source: https://www.arin.net/participate/meetings/reports/ARIN_XXV/PDF/Wednesday/Nobile_RSD.pdf

From the NANOG Mailing List, April 2010

- > What I heard at a recent (within the past six months)
  > conference was that "there is no customer demand for v6" so it
  > isn't on the immediate needs list. He said they had a lot of
  > inquiries about v6, but to date not having native v6 wasn't a
  > deal breaker with anyone

  I watched a vendor at one conference tell 20 people in a row that each one of them was the only one asking for IPv6. I mentioned to him that he should have his short-term memory loss checked out by a physician. At first he was confused. When I pointed out what I had just seen him do, he went from confused to embarrassed and admitted that it was the party line from his marketing department and they knew IPv6 was important, but, didn't have a story to tell yet, so, they were trying to spin for delay.

  http://seclists.org/nanog/2010/Apr/71

But The Perception Persists…

- Nobody actually cares*
  - Very little market demand
  - A "checkbox" feature
  - Few real endpoints

* Except Asia, US Government, Internet2

www.metasploit.com/data/confs/sector2008/exploiting_ipv6.pdf

The Question That's Really Being Asked Is…

- "What's the business case?"
- Vendors who ask that question obviously fail to understand the technology adoption lifecycle:

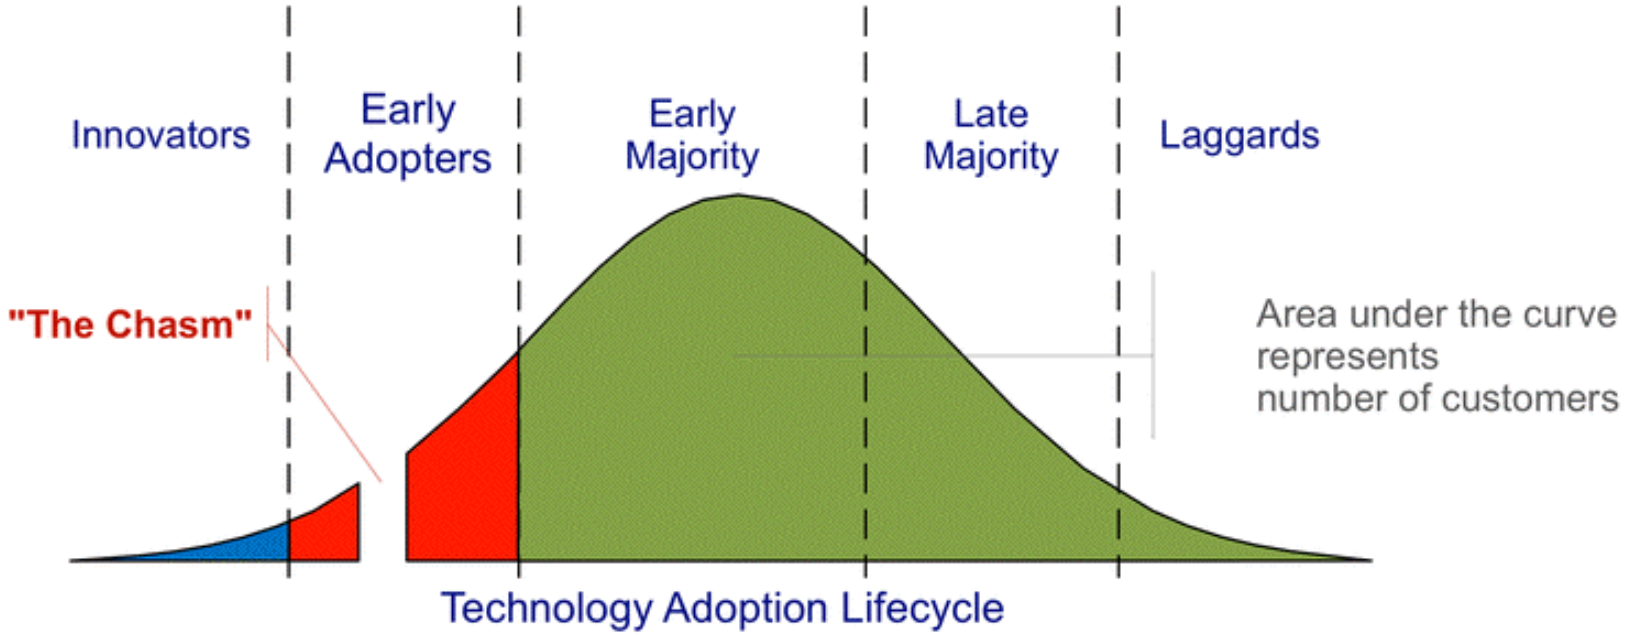

http://commons.wikimedia.org/wiki/File:Technology-Adoption-Lifecycle.png

We're Crossing "The Chasm" Right Now

- The companies that **are** ready to support IPv6 will be well positioned to support emerging customer needs and to thrive in an IPv6 environment

- The companies that **are not** ready will find their customers doing what they need to do (nothing personal, it's just business, and since "no one but me" cares about IPv6, I'm sure you won't mind that I'm taking my business elsewhere, see ya…)

A 2nd Opportunity For You To Bail Out

- So that's the end of Part II of today's IPv6 discussion; the final chunk of our training today will begin to expose you to some technical material relating to IPv6.

- If you're not a "techie" and your head already hurts and you're not interested in diving in deeper, feel free to bail out now. You will not hurt my feelings if you do so. :-)

This video is presented by the
Messaging Anti-Abuse Working Group

MAAWG IPv6 Training for Senders and Others

can be viewed in three segments from the training pages at www.MAAWG.org.
This has been part 2 of 3.

Our thanks to MAAWG Senior Advisor Joe St Sauver, Ph.D.,
for developing the materials in this training session
and allowing MAAWG to videotape it
for the benefit of professionals worldwide.