Networking in These Crazy Days: Stay Calm, Get Secure, and Get Involved

Joe St Sauver, Ph.D. (joe@uoregon.edu)

Merit Networking Summit

Ann Arbor, MI

December 12th, 2013

http://pages.uoregon.edu/joe/merit-networking/

Thanks, Technical Level and a Disclaimer

- I’d like to begin by thanking the summit program committee and Merit for the chance to talk with you today. I'm particularly honored by this invitation given that **Merit is the world's oldest regional network**, dating all the way back to 1966. (Terrific Merit history at http://www.merit.edu/about/history/index.php , BTW)
- Technical level of this talk: Given that this is a keynote for a mixed technical/managerial networking audience, **I've generally tried to hit an intermediate technical level**, but given the nature of some issues, I've also tried to provide backfill where it may be helpful. Sorry if I've provided too much or too little info at times.
- I also want to be sure to remind folks that **any/all opinions expressed this morning represent solely my own perspective,** and do NOT necessarily represent the opinion of the program committee members, Merit, Internet2, InCommon, the University of Oregon, or anyone else.

Format of This Talk

- Don't let my odd slide format shake you up. :-) I promise I'm not going to read my slides to you, and you don't need to try to read them, either. I write my slides this way:
  -- for those who may look at them later (maybe even you!)
  -- for those who aren't native English speakers
  -- to ensure accessibility for any who may be hearing impaired
  -- for ease of search engine indexing.
- I also provide detailed slides because:
  -- I tend to cover a lot of material in a relatively short period of time, and I need to be "well scripted" to stay on track
  -- I like to cite my sources/provide lots of links so you can "dig in"
  -- I don't want you to have to try to take notes today
  -- I hate to be misquoted.
- Let's handle questions, if there are any, at the end of the talk.

1. IPv4 Address Space Exhaustion

"Heavenly Father, bless us,
And keep us all alive.
There's ten of us to dinner,
And not enough for five.

"Hodge's grace, Anonymous, 1850"

"Slaughtered to extinction: Passenger pigeons in Michigan,"
http://www.detroitnews.com/article/20120318/METRO/203180301

IPv4 Address Space Exhaustion

- Large portions of the world are out (or nearly out) of available IPv4 address space. Without adequate IPv4 address space, it will be hard for the Internet to continue to grow.

- Somehow, mind bogglingly, this is NOT "front page news."

- I suppose watching global IPv4 address exhaustion it is a bit like being a lobster in a pot of water that's slowly brought up to a boil: given gradual-enough changes, you may not notice what's happening until it's too late.

- Or maybe running out of IPv4 address space really isn't all that big a deal compared to all the other stuff that's currently a problem for the world?

Global Risks 2013

- If you haven't already seen it, check out the World Economic Forum's *Global Risks 2013 Report*, www3.weforum.org/docs/WEF_GlobalRisks_Report_2013.pdf
- In particular, see Figure 2, their "Global Risks Landscape," which shows roughly fifty **major** global risks (spanning economic, environmental, geopolitical, societal and technological areas).
- The X axis is the *likelihood of occurrence* over the next decade; the Y axis showing the *impact* if the risk were to happen. Each risk also shows whether it is *increasing or decreasing* in likelihood. (It's really quite the graphic, Edward Tufte would probably like it!)
- I suppose that if/when you compare IPv4 exhaustion to "major systemic financial failures," "water supply crises," "diffusion of weapons of mass destruction," and similar major global risks, maybe IPv4 exhaustion really doesn't measure up.

But You Know, I Think The Internet Just Might Be Important...

- Notwithstanding all the other craziness that's out there, if the Internet actually is important -- and I think it is -- we'll likely want it to continue to work smoothly.
- To continue to work smoothly, **the Internet needs to be able to scale.** New users and companies need to be able to come online, and existing users and companies need to be able to grow.
- Making the Internet work smoothly also means that **Internet users need to continue to have freedom to create new applications and innovate,** not just limp along doing the same old sorts of stuff.
- Unfortunately, **it's hard to attain those goals if you're running out of IPv4 address space** and no one deploys IPv6.

The IPv4 Situation in the Americas

- At the time I prepared these slides in early December 2013, ARIN, the regional Internet registry (RIR) covering North America and the Caribbean, had just 1.57 /8s worth of IPv4 space left (a /8 netblock has 2**24=16,777,216 IPv4 addresses).

  You can see the actual "live" ARIN available space counter at https://www.arin.net/resources/request/ipv4_countdown.html

- At the current rate of consumption, ARIN (and LACNIC, the regional Internet registry for Latin America) will be down to just their last /8 in three to six months (1Q14/1H14). See Geoff Huston's run out projections as shown on the next slide. Note where the projections cross the red horizontal line.

http://www.potaroo.net/tools/ipv4/plotend.png

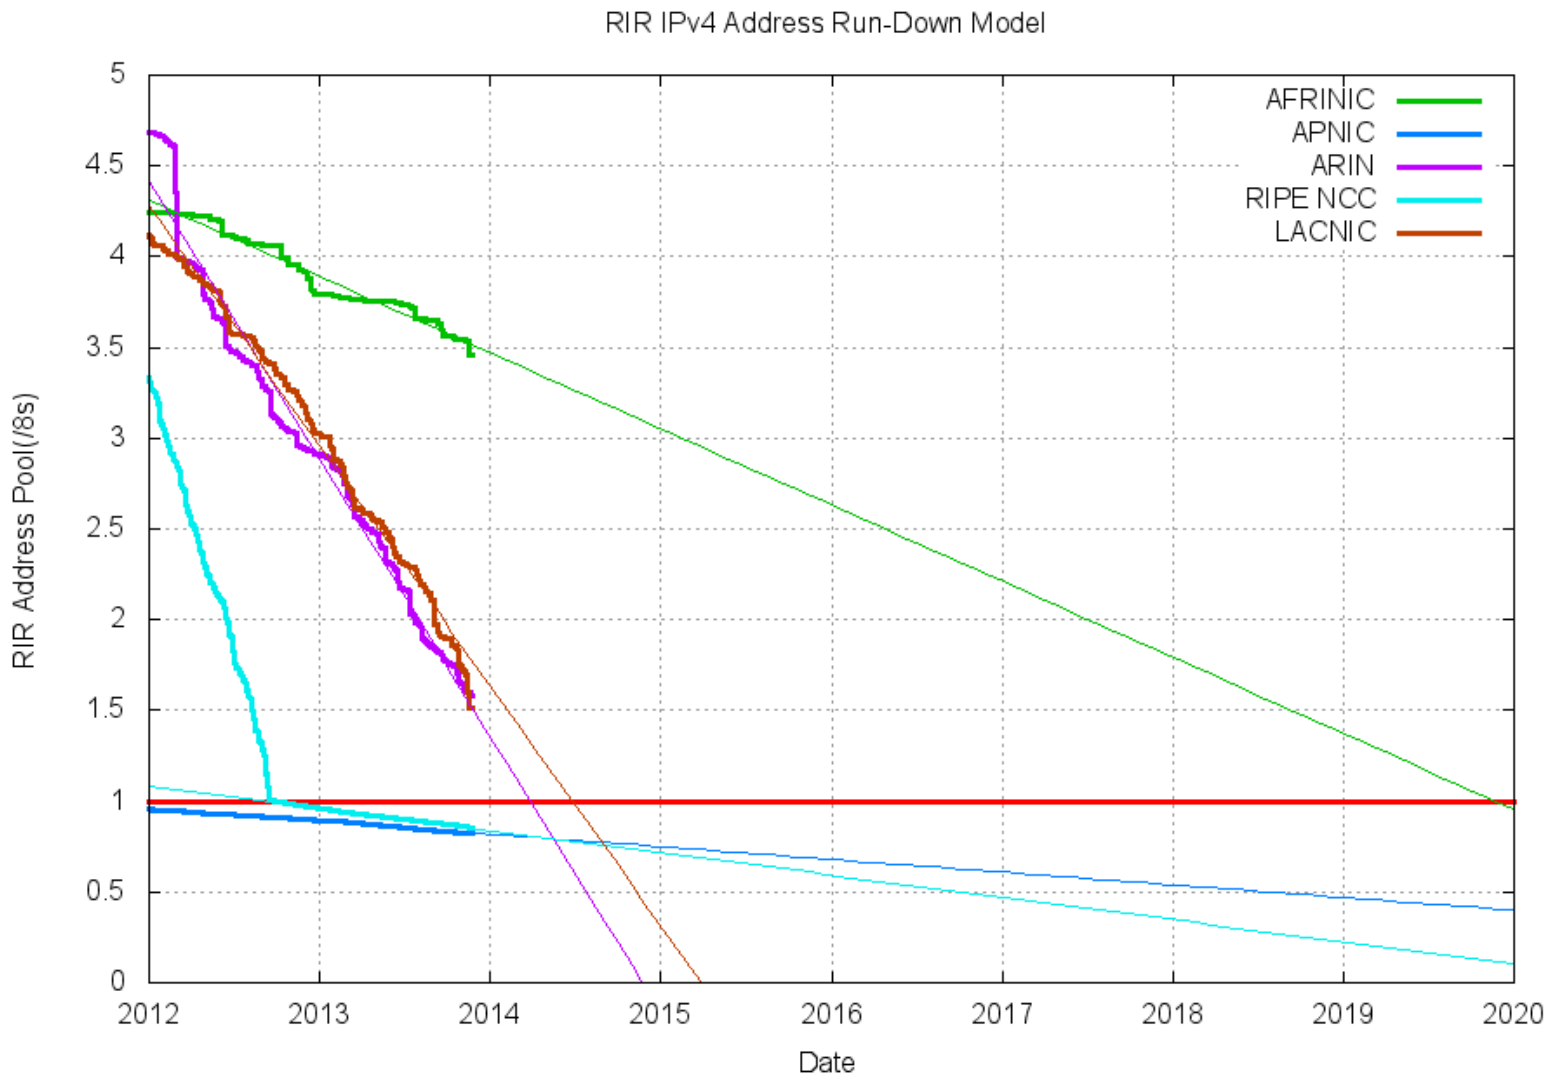

Two Implications of Impeding IPv4 Run Out

- **Implication One:** IF your site is in the Americas and you have a LEGITIMATE and JUSTIFIABLE need for more IPv4 address space, NOW is the time to ask for it. **DO NOT procrastinate.** Once ARIN and LACNIC are down to their last /8 it will be too late for you to get more IPv4 space you may legitimately need -- even if you can fully and convincingly document that request. (That said, please do NOT embarrass yourself or your organization by asking for space you don't actually need and can't compellingly justify.)

- **Implication Two:** At the same time you review your IPv4 address requirements, it is critical that ***everyone*** get IPv6 deployed.

- By the way, bad as things are in the Americas, the IPv4 situation is WORSE in Europe and Asia.

The IPv4 Situation in Europe and Asia

- Europe and Asia are each *already* down to their last /8, and RIPE and APNIC have begun rationing address space from those final blocks.

- To understand what "rationing" means, assume you're a new site in Europe or Asia. You want IPv4 address space to start your business or connect your campus. When you approach RIPE or APNIC, as of now they're only be able to give you **a single /22 (1,024 IPv4 addresses)** -- no matter how great your need and no matter how good your justification for more address space might be.

- That's too small a netblock to route globally, and just enough IPv4 address space for you to sort of "limp along" with Large Scale NAT while you (and the rest of the world) hopefully get IPv6 deployed.

Large Scale NAT (aka Carrier Grade NAT)

- As most tech folks know, when NAT (technically, "PAT") gets used, a public IP (or a small pool of public IPs) gets shared across multiple devices behind the NAT box. Devices behind the NAT box typically use private RFC1918 (e.g., 10.0.0.0/8) addresses locally. You probably use NAT on your home wireless "router" today.
- LSN is similar to NAT as used on your home wireless router, but LSN may use RFC6598 address space (100.64.0.0/10), instead.
- LSN has many drawbacks when used as an ISP-scale technology, most notably: (a) not working for public web servers (and other Internet facing server infrastructure), (b) causing a loss of end-to-end transparency (see RFC2775 and RFC4924), and (c) potentially limiting the throughput that can be achieved.
- If you're thinking about LSN, I also urge you to read RFC7021, "Assessing the Impact of Carrier-Grade NAT on Network Applications" (a new RFC from just September 2013)

Large Scale NAT and Security

- From a **security** point of view, the two biggest drawbacks to large scale NAT are:

  (a) Misbehavior by ONE customer sharing a public IP address will negatively affect the IP reputation of ALL the users on that IP

  (b) Use of LSN complicates abuse reporting: mapping an abuse report to a specific customer will be impossible unless reports include **source port** data as well as the usual IP-plus-time-stamp-with-time-zone-info. For what it may be worth (and as your security team probably already knows all-too-well), abuse reports SELDOM IF EVER include source port information.

- You may also want to see http://www.maawg.org/sites/maawg/files/news/M3AAWG_Carrier_Grade_NAT_BP.pdf

Universities, LSN and Legacy IPv4 Space

- Most universities don't tend to think much about large scale NAT.

- Why? Well, higher education's user base is far smaller than that of major ISPs (just a total of 21.8 million students nationally) and most universities have "lots" of IPv4 address space, having received legacy address space way back in the '90s.

- When most colleges and universities got address space in the 1990s (during the pre-CIDR era) your choices were basically:
  -- a "class A" address block (a /8 with 16,777,216 addresses), or
  -- one or more class B address blocks (/16s, with 65,536 addresses),
  -- one or more class C address blocks (/24s, with 256 addresses).

- **How much space do you currently have? How much is used?**

Getting Involved

- There are many opportunities to get involved with IPv4 run out issues, ranging from local to regional to national to global.

- A local opportunity: **talk about IPv4 run out when you go home.** Is your campus community broadly informed? Has your school reviewed your own IPv4 usage and address space requirements?

- Nationally, consider getting involved with the ARIN Policy Development Process (PDP). One way to do that is by signing up for the ARIN **Public Policy Mailing List (PPML)**, see https://www.arin.net/participate/mailing_lists/

- Globally, the public can also participate in ICANN; for ICANN-related opportunities, see http://www.icann.org/en/about/participate

2. IPv6 (Non) Deployment

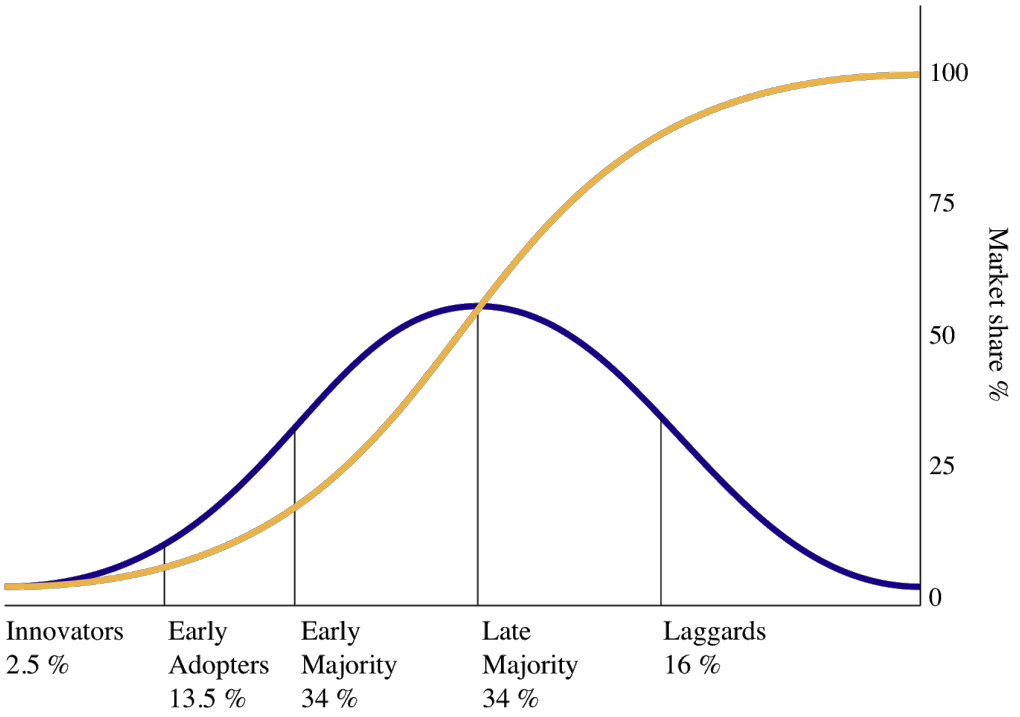

http://en.wikipedia.org/wiki/File:Diffusion_of_ideas.svg

See also "Diffusion and Adoption of IPv6 in the United States,"
http://www.cs.indiana.edu/pub/techreports/TR661.pdf

IPv6

- While we're talking about address space, let's take a minute or two to talk about IPv6 (non) deployment.

- Even though the Internet's just about out of IPv4 addresses, every one's very busy putting out a million other fires.

- Thus many sites haven't done much to get ready to actually use IPv6.

- For example, consider the basic act of simply **acquiring IPv6 address space** so you can begin to deploy IPv6. Unlike the shortage of IPv4 address space we just talked about, pretty much any organization can get abundant IPv6 address space. Only a relatively small number of sites have done so, so far...

Some Merit-Related Sites With IPv6 Address Space

AS237:

- Merit itself.... 2001:48A8::/32
- Regional Ed. Media Center 4, Muskegon: 2604:380::/32
- Michigan State: 2605:dd00::/32
- Wayne State: 2606:9700::/32
- Andrews University: 2620:0:2be0::/48
- Davenport University: 2620:a8::/48

[See http://bgp.he.net/AS237#_prefixes6]

Other Autonomous Systems:

- Mott Community College (AS1638): 2620:114:5000::/40
- U Michigan (AS36375): 2607:f018::/32

Non-Uptake Within Merit Isn't Merit's "Fault"

- Merit has been very active for many years when it comes to encouraging/supporting adoption of IPv6, e.g., see for example:

  -- "IPv6: Time To Get Started" by Andy Rosenzweig
  http://www.merit.edu/connections/Jan2012/ipv6.php
  -- "IPv6 Workshop," Aug 2-3, 2011
  http://www.merit.edu/events/archive/specialevents/ipv62011/
  -- "IPv6 Workshop," Nov 11-13, 2009
  http://www.merit.edu/events/archive/specialevents/ipv62009/
  -- "IPv6 Workshop," Apr 17-18, 2007 (at Merit Networks)
  http://ipv6.internet2.edu/merit/
  -- etc.

- **Have y'all heard what Merit's been trying to tell you?**

So Why *Aren't* Sites Deploying IPv6?

- If you talk to sites that haven't deployed IPv6 address space, you'll hear many reasons why not. Some may legitimately point out: "We're busy! We've got plenty of IPv4 space, and there's hardly any IPv6 traffic. So why bother getting IPv6 address space, and enabling IPv6 connectivity, and working to make workstations and servers dual stack, if IPv6 won't get used?"

- Of course, **if *everyone* remains reluctant to deploy IPv6, one wouldn't *expect* there to be much IPv6 traffic, right? And if there *isn't* much traffic, then there's *not* much point to deploying IPv6...** This is a classic circular dependency: chicken, egg; egg, chicken.

- We need your help to break that cycle, deploying IPv6 even if you don't "need it" and even if demand for IPv6 is still just ramping up.

IPv6 Traffic On I2 Is Already Non-Negligible

- On Internet2, the ratio seems to run roughly 10:1, IPv4:IPv6:

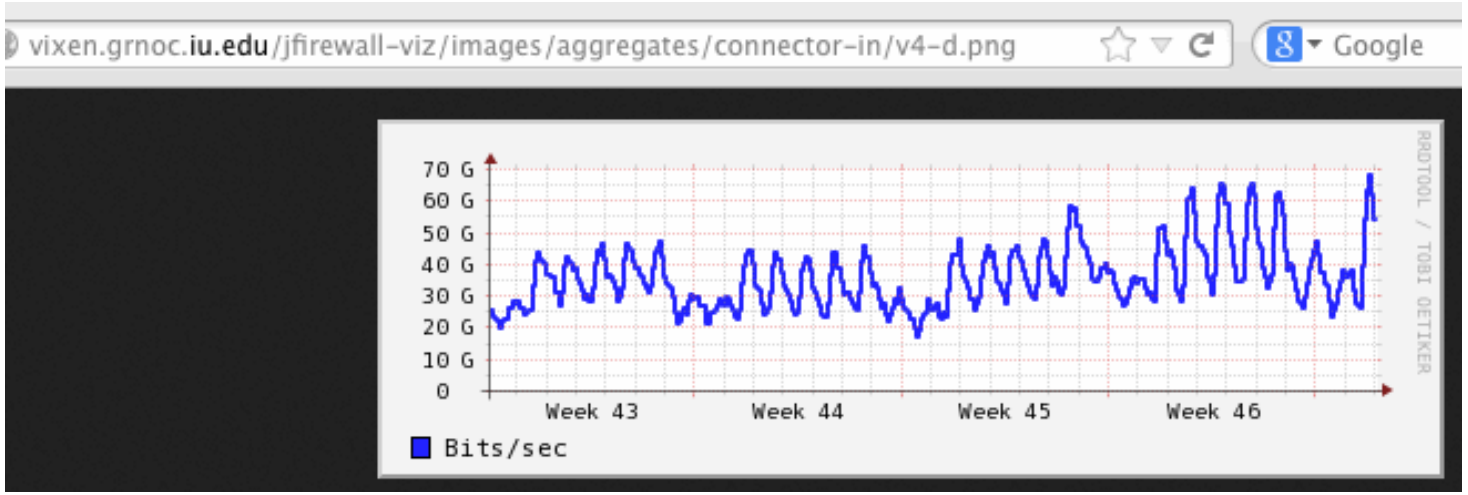

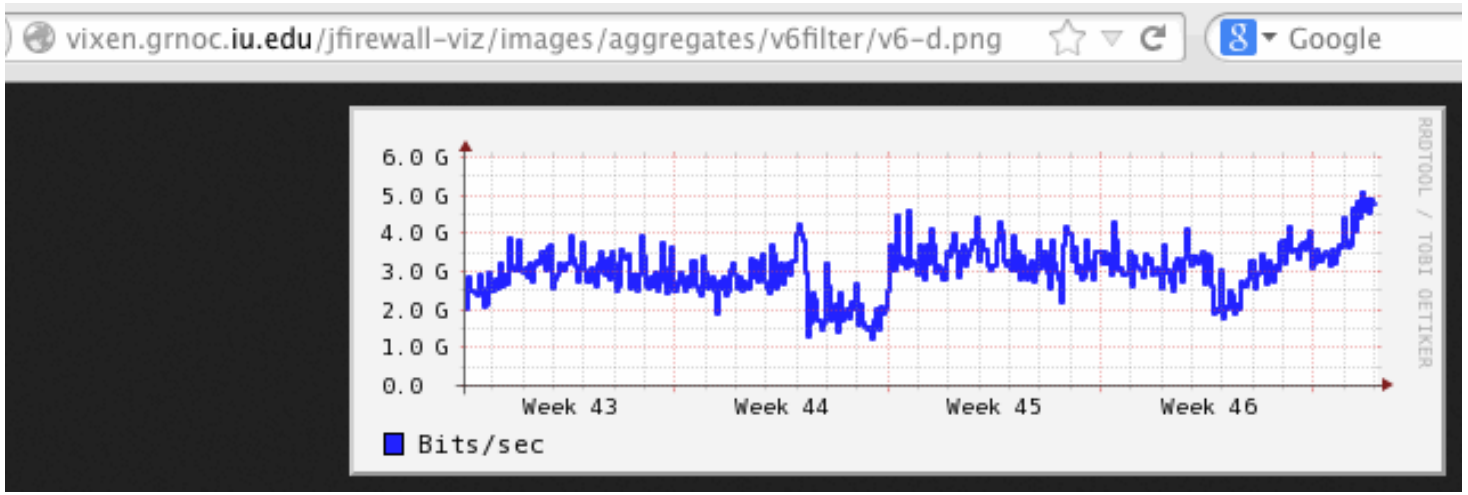

Google's IPv6 Traffic Is Ramping Up, Too...

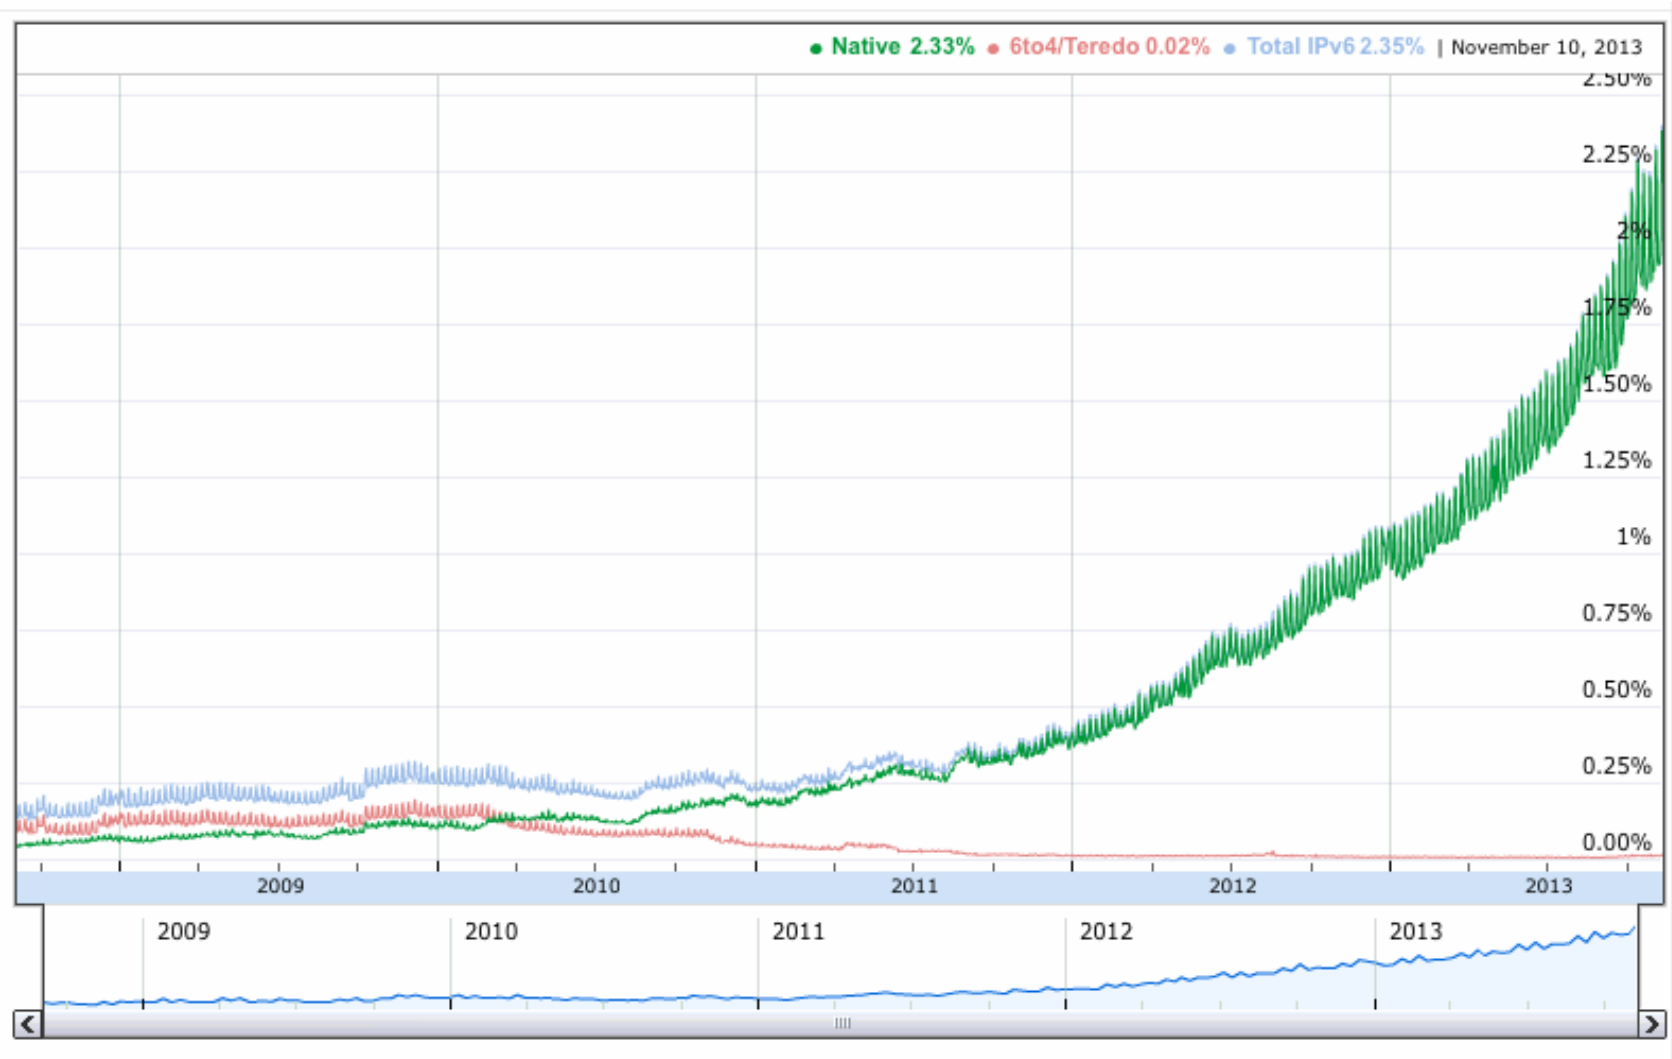

Source: http://www.google.com/ipv6/statistics.html

Does That Google Graph Show Failure or Success?

- I could see how some might look at that preceding graph and be *discouraged* ("Gee, Google IPv6 traffic is only at ~2% currently")
- Personally, I think we should all be *ecstatic* (IPv6 traffic levels appear to be more-or-less doubling annually, year over year):
  - 2012, IPv6 traffic was roughly ½ of 1%
  - 2013, IPv6 traffic was roughly 1%
  - 2014, IPv6 traffic will likely be well over 2%
  - 2015 → 4%?
  - 2016 → 8%?
  - 2017 → 16%?
  - 2018 → 32%?
  - 2019 → 64%?
- That's probably a conservative (low) estimate for the actual rate of growth (but the curve could also plateau). We don't yet know.

Technically Enabling IPv6 Success

- In order for traffic to flow over native IPv6, the entire path needs to be IPv6 enabled "end-to-end." This means that:
  -- **You need IPv6 address space**
  -- **Networks need to route IPv6**, both locally and over the wide area (both Merit and Internet2 already route IPv6 traffic today)
  -- **Network middleboxes** (firewalls, load balancers, etc.) **need to stay out of the way of IPv6 traffic** (and most now do so)
  -- **DNS servers** (authoritative name servers and recursive resolvers) **need to at least support quad A resource records**, and should ideally also have dual stack transport -- do yours?
  -- **Application servers need to be configured to use IPv6**
  -- **Laptops and other end user devices need IPv6 connectivity**
  -- **Applications need to be IPv6 aware**, and prefer IPv6 addresses over IPv4 if the target server is dual stack
- What do we see if we look at the IPv6 status of some servers?

Servers and IPv6: www.mrp.net/ipv6_survey/

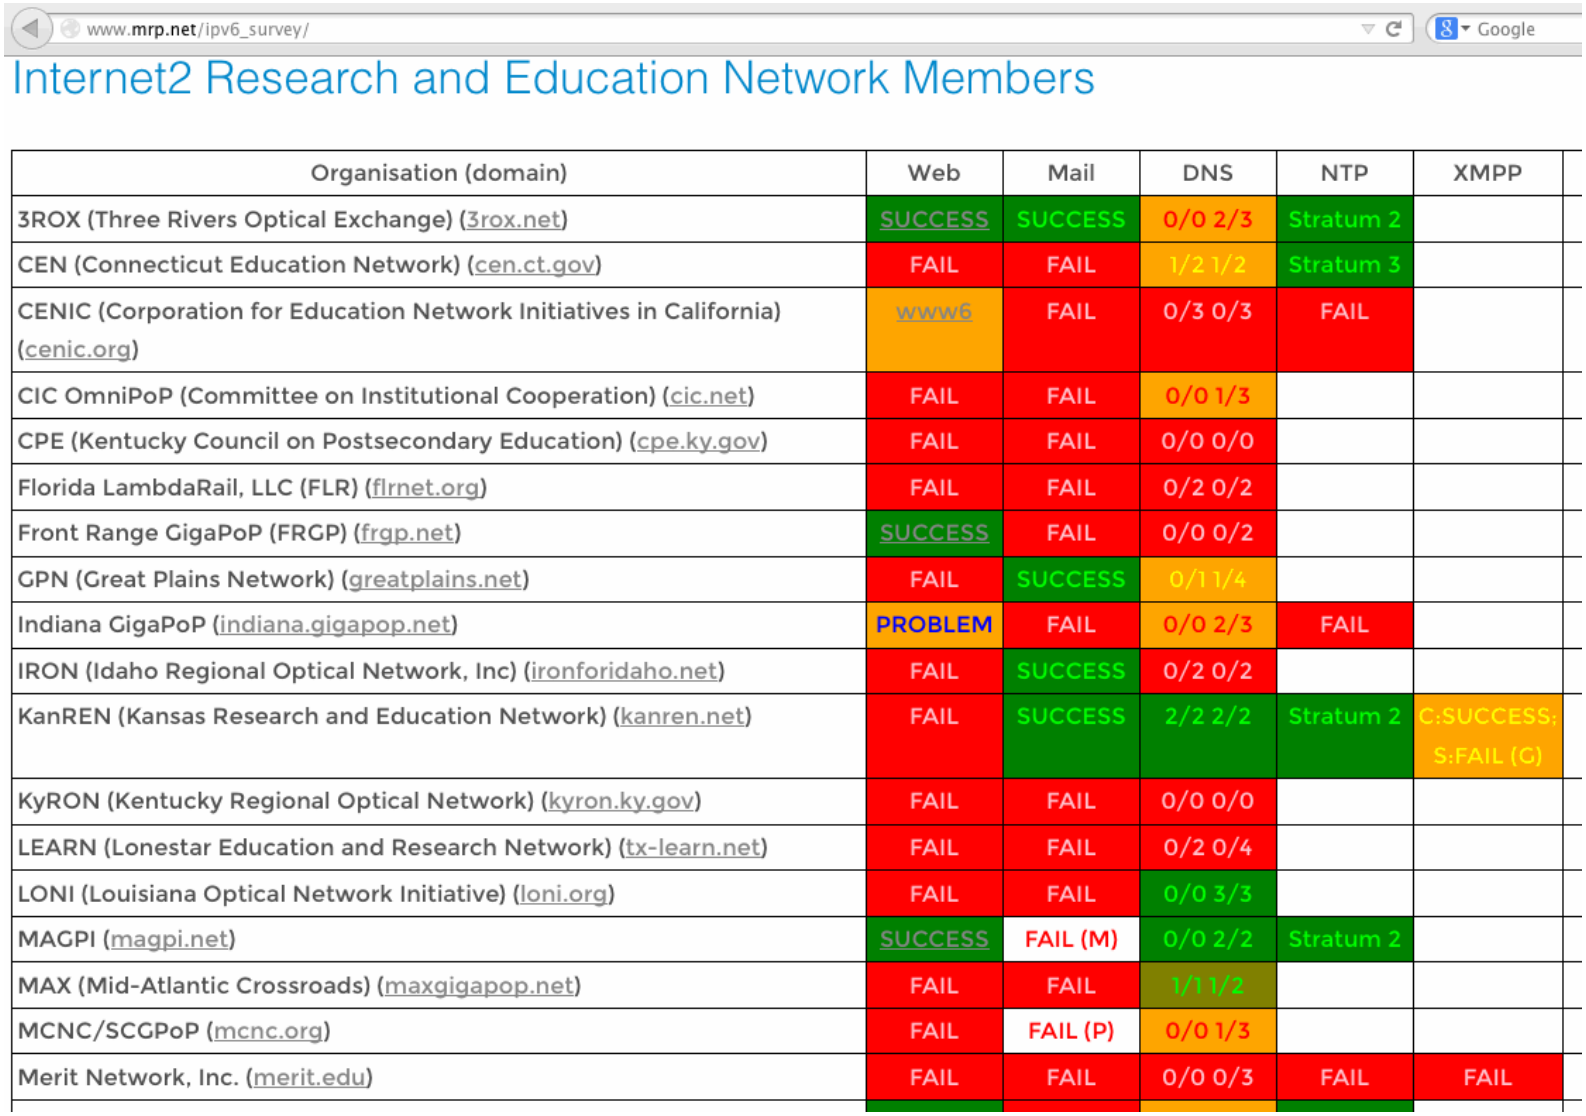

www.mrp.net/ipv6_survey/

Google

Internet2 Research and Education Network Members

Organisation (domain)	Web	Mail	DNS	NTP	XMPP
3ROX (Three Rivers Optical Exchange) (3rox.net)	SUCCESS	SUCCESS	0/0 2/3	Stratum 2	
CEN (Connecticut Education Network) (cen.ct.gov)	FAIL	FAIL	1/2 1/2	Stratum 3	
CENIC (Corporation for Education Network Initiatives in California) (cenic.org)	www6	FAIL	0/3 0/3	FAIL	
CIC OmniPoP (Committee on Institutional Cooperation) (cic.net)	FAIL	FAIL	0/0 1/3		
CPE (Kentucky Council on Postsecondary Education) (cpe.ky.gov)	FAIL	FAIL	0/0 0/0		
Florida LambdaRail, LLC (FLR) (flrnet.org)	FAIL	FAIL	0/2 0/2		
Front Range GigaPoP (FRGP) (frgp.net)	SUCCESS	FAIL	0/0 0/2		
GPN (Great Plains Network) (greatplains.net)	FAIL	SUCCESS	0/1 1/4		
Indiana GigaPoP (indiana.gigapop.net)	PROBLEM	FAIL	0/0 2/3	FAIL	
IRON (Idaho Regional Optical Network, Inc) (ironforidaho.net)	FAIL	SUCCESS	0/2 0/2		
KanREN (Kansas Research and Education Network) (kanren.net)	FAIL	SUCCESS	2/2 2/2	Stratum 2	C:SUCCESS; S:FAIL (G)
KyRON (Kentucky Regional Optical Network) (kyron.ky.gov)	FAIL	FAIL	0/0 0/0		
LEARN (Lonestar Education and Research Network) (tx-learn.net)	FAIL	FAIL	0/2 0/4		
LONI (Louisiana Optical Network Initiative) (loni.org)	FAIL	FAIL	0/0 3/3		
MAGPI (magpi.net)	SUCCESS	FAIL (M)	0/0 2/2	Stratum 2	
MAX (Mid-Atlantic Crossroads) (maxgigapop.net)	FAIL	FAIL	1/1 1/2		
MCNC/SCGPoP (mcnc.org)	FAIL	FAIL (P)	0/0 1/3		
Merit Network, Inc. (merit.edu)	FAIL	FAIL	0/0 0/3	FAIL	FAIL

What About Laptops? Is Yours Ready For IPv6?

- Check by visiting the web site http://test-ipv6.com/ with Chrome

- Note that if you visit with Firefox, even if you're dual stacked (e.g., have both IPv4 and IPv6 connectivity), **Firefox will normally prefer IPv4 by default.** Therefore, if you're trying to check your IPv6 status, be sure to check with Chrome, instead.

- Bottom line, if you're NOT able to routinely use IPv6, you really want to look into why, and make it a priority to get that fixed.

- **The general rate of IPv6 non-deployment in higher education -- at least given the simultaneous rate of global IPv4 exhaustion – is another example of something that I find totally *crazy*.**

The Sort of IPv6 Report You Want To See

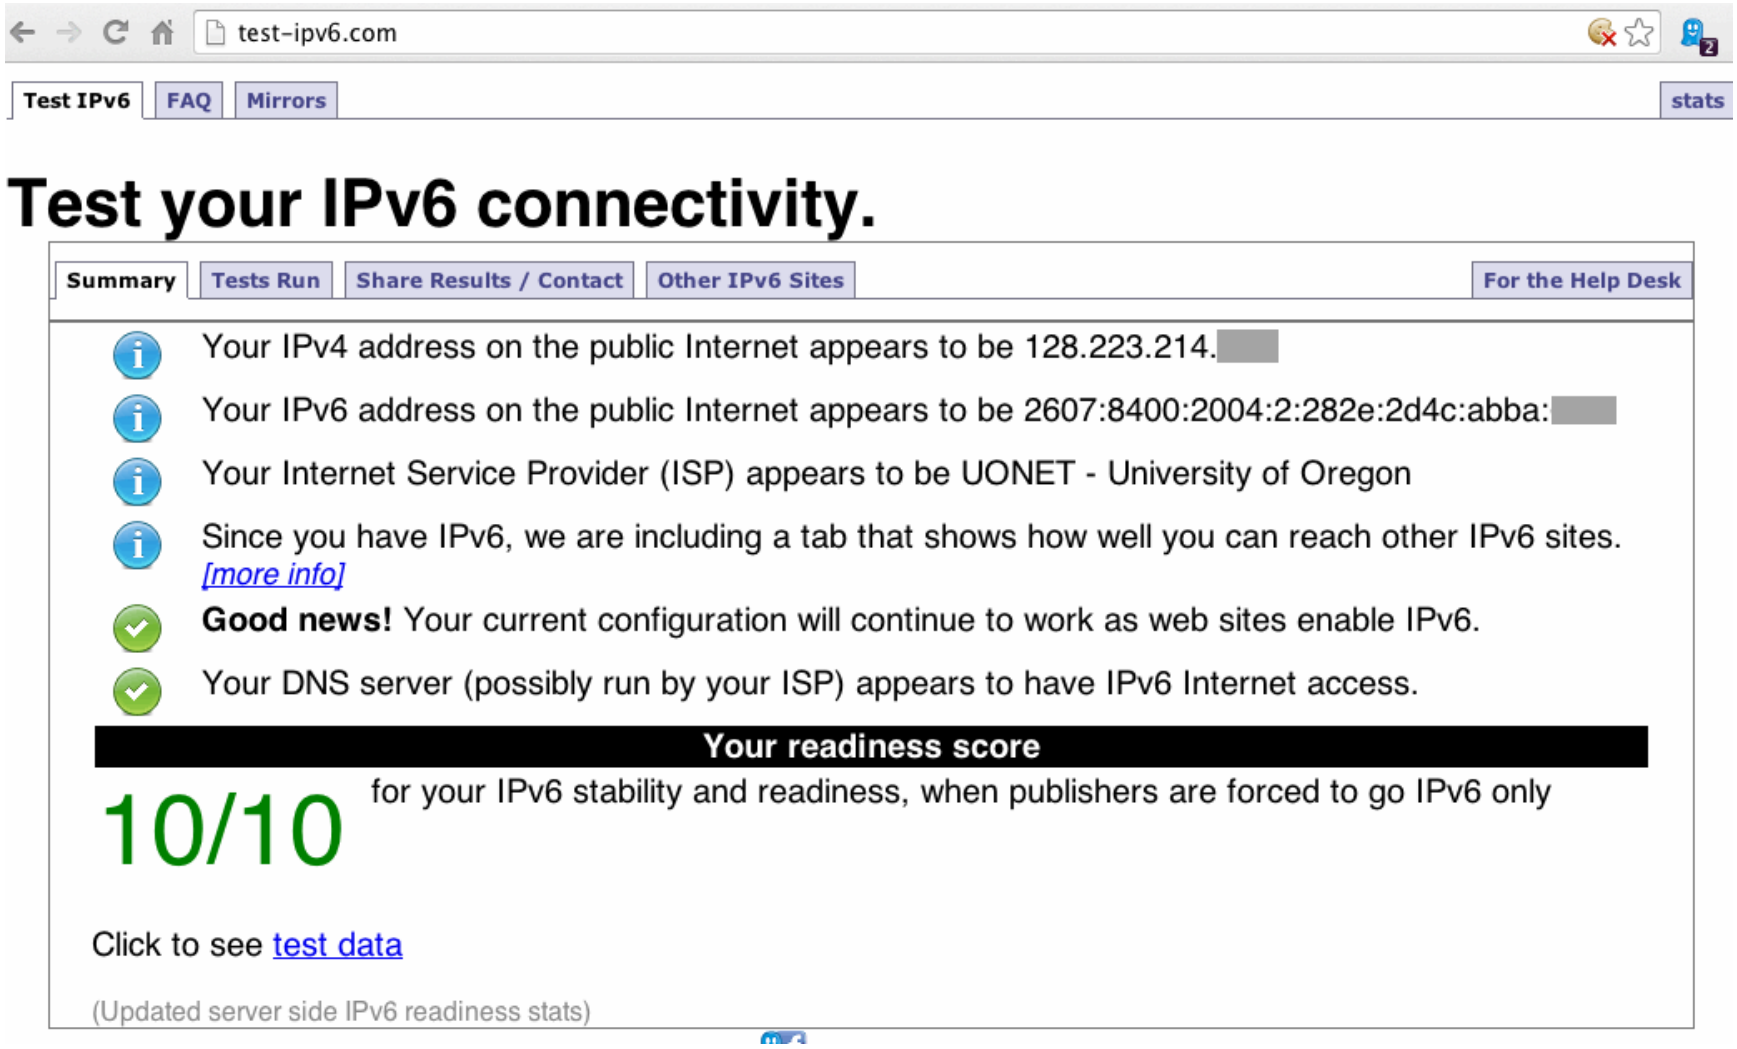

An IPv6 Aside: The Merit IPv6 Darknet Project

- If you look at AS237 IPv6 routes, you'll see that Merit announces huge "covering" IPv6 /12 routes from each of the five regional Internet registries (see http://software.merit.edu/darknetv6/ for the letters of authorization from each of the RIRs okaying this).
- The net effect of those broad IPv6 covering routes is that any IPv6 traffic destined for unassigned IPv6 space (e.g., random scans) will end up being collected (by default) by the Merit IPv6 darknet.
- The good news is that (at least based on preliminary results now reported by Merit), there isn't much "background radiation" in IPv6 space right now (~1Mbps), and what is there is typically the result of misconfiguration rather than malicious scanning (see www.merit.edu/research/pdf/2013/ipv6_darknet_paper_r6098.pdf)
- So, if you've been holding off on IPv6 because you've been worried that it's some sort of online cesspool of cyber badness, relax.

Speaking of IPv6 and Security...

- In other talks I've discussed the relationship between *IPv6 and security* ("IPv6 and the Security of Your Network and Systems," see http://pages.uoregon.edu/joe/i2mm-spring2009/). Short form: IPv6 is neither a security "magic bullet," nor is it an impossible and intractable security morass. IPv6 is something that you'll eventually need to do one way or the other, so you might as well get some experience now while it's still relatively new -- thereby "future proofing" your network (and you!) against obsolescence.

- One exception: I'd currently proceed with restraint when it comes to deploying SMTP over IPv6. Management of unwanted wide area mail traffic from unauthenticated IPv6 sources is still technically tricky. See "MAAWG IPv6 Training for Senders and Others," http://pages.uoregon.edu/joe/maawg-senders-ipv6-training/maawg-senders-ipv6-training.pdf

Getting Involved/Encouraging Campus Use of IPv6

- Are you up to speed on how IPv6 works? If not, get trained!

- When it comes to *getting involved*, **every campus is going to need "IPv6 champions,"** including folks in the university's top executive leadership, grassroots users, and everyone in between. When you go back to your home institutions, what will you tell your colleagues about the need for IPv6 deployment, eh?

- **Many of the key people needed to make IPv6 a reality at Merit schools are sitting in this room today.** Maybe Merit could consider standing up an "IPv6 Deployment Leadership Council," if such a group doesn't already exist?

- Don't forget to talk to your home cable/DSL ISP providers, too!

3. Domain Names

Sticks and stones will break my bones,
but names will never hurt me.

19th century English nursery rhyme

Domain Names

- We've been talking about IP addresses (e.g., numbers) but domain names also play a critical role on the Internet.
- When it comes to domains, higher ed's focus has always been on dot edu, because that's where our primary "home" has always been (surprisingly, we have relatively little direct control over how that gTLD gets administered, and yes, I say that with full awareness of http://net.educause.edu/edudomain/policy_public_forum.asp)
- The rest of the Internet, however, has long lived with generic top level domains (gTLDs) such as .com, .net, .org, .biz, .info, etc., plus country code top level domains (ccTLDs) such as .br, .ca, .cn, .de, .es, .fr, .jp, .mx, .ru, .uk, etc.
- There are hundreds of TLDs already in use, but we have some sense from root node name server data about which domains actually currently "matter", modulo DNS caching effects.

stats.l.root-servers.org/cgi-bin/dsc-grapher.pl?window=86400&plot=qtype_vs_all_tld&server=L-root

Google

Servers/Nodes

L-root
› Africa
› AsiaPacific
› Europe
› LatinAmericaCaribbean
› NorthAmerica
L-root-Africa
L-root-AsiaPacific
L-root-Europe
L-root-LatinAmericaCaribbean
L-root-NorthAmerica
as112
b.ip6-servers.arpa
b.in-addr-servers.arpa
ns.icann.org

Plots

By Node
Qtypes
Rcodes
Classification
Client Geography
TLDs
› valid
› invalid
› numeric
2nd Level Domains

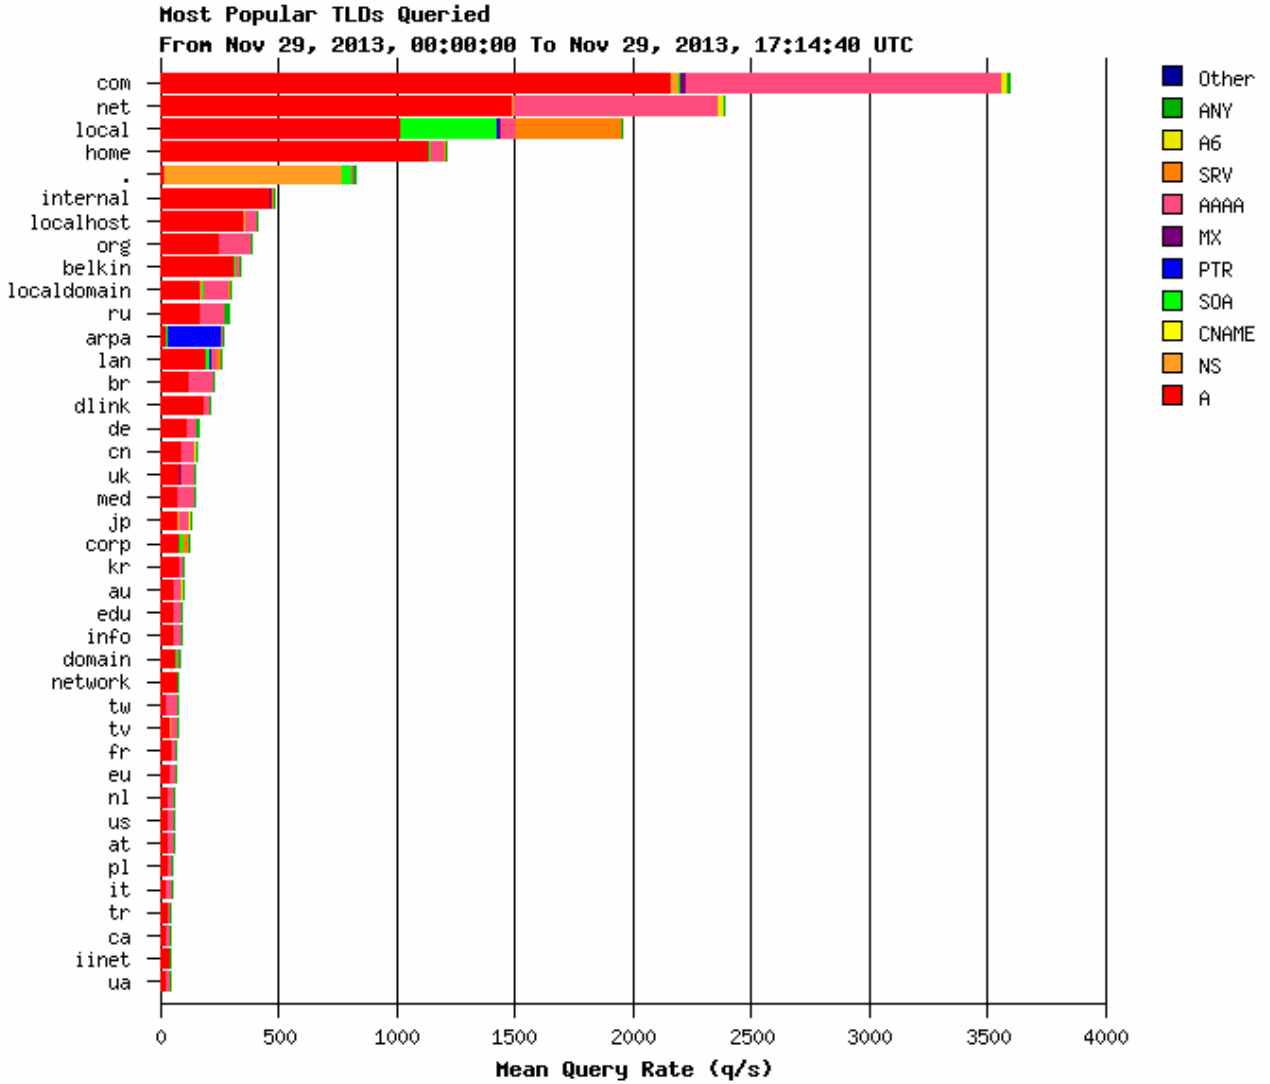

ICANN's New Top Level Generic Domain Names

- Given that it has largely been a "dot com/dot net" world to-date, in June 2011, ICANN authorized the launch of ICANN's new gTLD program, meant to "enhance competition" and increase "consumer choice," while also enabling innovation, particularly in the area of Internationalized Domain Names (IDNs). Many questions about the new gTLD program are answered in the (long) ICANN new gTLD guidebook, see http://newgtlds.icann.org/en/applicants/agb
- 1,970 applications for new gTLD were received by ICANN. After resolution of duplicate requests, etc., that's now down to "just" 1,400 or so new gTLDs.
- Some of those gTLDs have now begun to be delegated. To see a list of what's already been delegated so far, check out http://newgtlds.icann.org/en/program-status/delegated-strings or see the next couple of slides...

DATE	STRING	REGISTRY
23 November 2013	みんな (xn--q9jyb4c) -- Japanese for "everyone"	Charleston Road Registry, Inc.
19 November 2013	.DIAMONDS	John Edge, LLC
19 November 2013	.TIPS	Corn Willow, LLC.
19 November 2013	.PHOTOGRAPHY	Sugar Glen, LLC
19 November 2013	.DIRECTORY	Extra Madison, LLC
19 November 2013	.ENTERPRISES	Snow Oaks, LLC
19 November 2013	.KITCHEN	Just Goodbye, LLC
19 November 2013	.TODAY	Pearl Woods, LLC
14 November 2013	.PLUMBING	Spring Tigers, LLC
14 November 2013	.GRAPHICS	Over Madison, LLC
14 November 2013	.CONTRACTORS	Magic Woods, LLC
14 November 2013	.GALLERY	Sugar House, LLC
14 November 2013	.SEXY	Uniregistry, Corp.
14 November 2013	.CONSTRUCTION	Fox Dynamite, LLC
14 November 2013	.TATTOO	Uniregistry, Corp.
14 November 2013	.TECHNOLOGY	Auburn Falls
14 November 2013	.ESTATE	Trixy Park, LLC

DATE	STRING	REGISTRY
14 November 2013	.LAND	Pine Moon, LLC
14 November 2013	.BIKE	Grand Hollow, LLC
06 November 2013	.VENTURES	Binky Lake, LLC
06 November 2013	.CAMERA	Atomic Maple, LLC.
06 November 2013	.CLOTHING	Steel Lake, LLC
06 November 2013	.LIGHTING	John McCook, LLC
06 November 2013	.SINGLES	Fern Madison, LLC
06 November 2013	.VOYAGE	Ruby House, LLC
06 November 2013	.GURU	Pioneer Cypress, LLC
06 November 2013	.HOLDINGS	John Madison, LLC
06 November 2013	.EQUIPMENT	Corn Station, LLC
23 October 2013	شبكة (xn--ngbc5azd) -- Arabic for "web/network"	International Domain Registry Pty. Ltd.
23 October 2013	онлайн (xn--80asehdb) -- Cyrillic for "online"	CORE Association
23 October 2013	сайт (xn--80aswg) -- Cyrillic for "site"	CORE Association
23 October 2013	游戏(xn--unup4y) -- Chinese for "game(s)"	Spring Fields, LLC

Will Those New gTLDs *Really* Matter?

- Maybe. Each applicant that just spent $185,000 per domain (plus substantial additional implementation costs!) sure thinks so...
- Let me share one example of why I think the new gTLDs matter: **over a hundred of the new gTLDs will be Internationalized Domain Names,** using 12 different scripts (Arabic, Cyrillic, etc.)
- While the ICANN Board approved non-Latin character sets for ccTLDs in October 2009, four years later **only 2% of the world's domains are non-Latin.** (see "World Report on IDN Deployment, 2013," http://www.eurid.eu/files/publ/insights_2013_idnreport.pdf)
- IDN uptake will hopefull be better in some of the new gTLDs... remember, **nearly 2.8 billion people use major non-Latin scripts** (see http://en.wikipedia.org/wiki/List_of_writing_systems). For comparison, there are fewer than 1/3 billion US residents, and just over half a billion people in all of the European Union.

"So What Should *We* Do Re the New gTLDs?"

- In the case of the new gTLDs, there's not really much that you actually can do -- at this point, the new gTLDs are being rolled out, whether you or I like them/need them/want them or we hate them.

- So what if a user from your campus comes up to you and says:

  "We've been hearing about all these new ICANN domains! What *should* we do? Should we be thinking about doing defensive registrations in all of them to protect all of our school's brands?"

- To answer that, begin by reviewing what you currently do when it comes to your brands and traditionally available top level domains. For example, what about domain names for the label "merit"?

gTLD Registrations for The String "merit"

- merit.edu: Merit Network Inc.
- merit.net: Merit Network Inc.
- merit.org: Merit Network Inc.
- merit.asia: Marie-Louise van Dijk, Hoofdorp NL
- merit.biz: Evone Farha, Los Angeles CA
- merit.com: Merit Medical Systems, Inc.
- merit.info: Gabriele Hoefer, Aurolzmuenster AT
- merit.mobi: Philip Morris Brands S.A.R.L., Neuchatel CH
- merit.name: blocked by defensive registration
- merit.pro: Andrei V Titushkin, Moscow RU
- merit.tel: Philip Morris Brands S.A.R.L., Neuchatel CH
- merit.travel: Merit Holdings Inc., Toronto CA
- merit.xxx: blocked by defensive registration

More "Merit" Domains (these in selected ccTLDs)

- merit.ca: Merit Travel Group Inc.. Toronto CA
- merit.cn: 北京国网信息有限责任公司
- merit.cz: Merit Group, a.s., Olomouc CZ
- merit.de: Michael Ruppert, Rossbach/Wied DE
- merit.dk: A/S Mirit-Glas, Herlev DK
- merit.ee: AS Merit Tarkvara
- merit.fi: Haastattelukeskus Merit Oy, Helsinki FI
- merit.fr: Merit Li-Lin Europe SA, Eragny sur Oise FR
- merit.hk: Merit Company Limited, HK
- merit.in: Philip Morris Brands S.A.R.L., Neuchatel CH
- merit.co.uk: Magazine Subscription Ltd, Darlington UK

(BTW, that's just a sampling of the ccTLD "merit" domains...)

The New gTLDs Trademark Clearinghouse

- If you're not **already** trying to exhaustively register your marks in all currently available gTLDs and ccTLDs, should you start doing so now for all the new gTLDs? Is it really worth it? Would trying to do so even be *feasible* given the number of new gTLDs that have been created?

- Fortunately, the new gTLD process includes a new Trademark Clearinghouse, which potentially simplifies the process of at least protecting your trademarked brands, if you're worried. See http://www.trademark-clearinghouse.com/ for more information.

- **The <u>real issue</u> probably isn't all the new domains and cybersquatting risks, it's how (or IF!) you will maintain control of the domains you already have and rely on.**

Secure Critical Domain Registrations

- Some may not appreciate how easily some domain name can be hijacked, even including those for major online properties.
- Recent hijacks have involved the domains of the NY Times, Twitter, and the Huffington Post (http://www.theregister.co.uk/2013/08/27/twitter_ny_times_in_domain_hijack/), among others.
- If a hacker/cracker can gain access to a domain owner's domain administration panel at the domain owner's registrar, the hacker/cracker can totally control the domain, including doing things like changing the domain name servers to point to their hosts.
- If the domain name servers can be changed to name servers of a hacker's choice, the hacker can then hijack/eavesdrop on email for the domain, or they can direct visitors to look-alike sites that will drop malware on the visitor, steal their PII, etc.
- This is obviously bad.

Which Registrar Are You Using?

- You can review a relatively long list of registrars at http://www.icann.org/registrar-reports/accredited-list.html
- Some registrars specialize in offering cheap and easy domain name registrations for vanity domains/small business owners. Other registrars may specialize in bulk domain registrations for speculators who routinely register 1000's of domains per day. Only a few registrars specialize in securely managing high value corporate domains so that they don't get hijacked.
- If you check whois to see what registrars get used by the top domains on the Internet, you may be surprised to see how many of those domains all use one of a small number of registrars.
- When it comes to your dot edu domain, you have no choice (Educause is your only option), but when it comes to your other domains, are you using the "right" registrar? Or did you just you one that happened to be cheap and well known? Check your whois!

What Registrars Do Top U.S. Alexa Domains Use?

1.	google.com	Mark Monitor	18.	huffingtonpost.com	Melbourne IT
2.	facebook.com	Mark Monitor	19.	cnn.com	Corp. Domains
3.	youtube.com	Mark Monitor	20.	paypal.com	Mark Monitor
4.	yahoo.com	Mark Monitor	21.	wordpress.com	Mark Monitor
5.	amazon.com	Mark Monitor	22.	instagram.com	Mark Monitor
6.	wikipedia.org	Mark Monitor	23.	msn.com	Mark Monitor
7.	linkedin.com	Mark Monitor	24.	apple.com	Corp. Domains
8.	ebay.com	Mark Monitor	25.	netflix.com	Mark Monitor
9.	twitter.com	Corp. Domains	26.	imgur.com	ENOM
10.	bing.com	Mark Monitor	27.	aol.com	Melbourne IT
11.	craigslist.com	Network Sol'n	28.	walmart.com	Corp. Domains
12.	pinterest.com	Ascio Tech	29.	reddit.com	Gandi
13.	blogspot.com	Mark Monitor	30.	imdb.com	Mark Monitor
14.	go.com	CSC Global	31.	weather.com	Network Sol'n
15.	live.com	Corp. Domains	32.	yelp.com	Mark Monitor
16.	espn.go.com	CSC Global	33.	microsoft.com	Mark Monitor
17.	tumblr.com	Mark Monitor	34.	bankofamerica.com	Mark Monitor

Facilitating The Tracking of Registrar Reputation

- A quick sidebar: a perennial issue is abusive domain names with concealed/anonymized contact information. While ICANN has conducted multiple whois studies (some excellent, others with almost comically bad analyses), cyber criminals will likely be able to continue to hide behind private/proxy domain registrations, and investigators will suffer from unduly restrictive whois rate limits.
- Way back in 2008, I demonstrated that a small number of registrars were disproportionately associated with abusive domain names, see http://pages.uoregon.edu/joe/maawg12/domains-talk.pdf
- I'm therefore very happy to see some people beginning to talk about making it easier to systematically obtain domain registrar data at scale... There certainly shouldn't be any privacy concerns when it comes to that information! See part 2A of Paul Vixie's note: http://mm.icann.org/pipermail/itipanel/2013-November/000017.html

Multifactor Auth For Your Domain Admin Panel

- If your institution has critically important domain names (and is there any modern business that doesn't?), at a minimum, wouldn't it be nice to use a registrar that doesn't relies on just plain passwords?

- There ARE now registrars that do offer **multifactor authentication** to protect your domain registration admin panel access, including:

  -- http://www.101domain.com/domain_name_security.htm
  -- http://blog.dnsimple.com/protect-your-dnsimple-account-with-two-factor-authentication-from-authy/
  -- http://www.dynadot.com/domain/security.html
  -- http://wiki.gandi.net/en/contacts/login/2-factor-activation
  -- https://www.name.com/services/namesafe
  -- http://community.namecheap.com/blog/2013/10/08/two-factor-authentication/

Getting Involved: Keep Track of Your Domains!

- If your institution has a substantial portfolio of domain names, perhaps created by a handful of different administrators and departments, is there anyone keeping an eye on all of them? Do you even know all the domain names your school uses?

- Are all those domains registered to your university? Or are some registered to third parties, such as contractors or individual employees? (If they're hidden behind private/proxy registrations, are you SURE you know who's registrant of record for them?)

- Are all the details (such as email contacts) up-to-date? Or do some refer to former employees and now-retired email systems?

- And when will each of your domains expire? Any expiring soon?

4. DNS

"Running a nameserver is not a trivial task.
There are many things that can go wrong [...]"

David Barr, Penn State University, February 1996,
RFC1912, "Common DNS Operational and Configuration Errors"

DNS: A Crucial (If Often-Neglected) Service

- We've talked about IPv4 and IPv6 addresses and domain names. Now let's now talk a little about the glue that ties them together, the domain name system (DNS).
- DNS transparently and efficiently maps names, such as www.merit.edu, to IP addresses such as 207.75.117.26
- DNS can also (ideally) do the reverse, mapping an IP address (such as 207.75.117.26) to a domain name, although *sometimes* that isn't properly and symmetrically configured (even though it *should* be).
- Without DNS, the Internet would be a real pain to use, wouldn't scale very well, and wouldn't be very flexible. DNS is important. Given its importance, it's surprising how often it's neglected.
- I've previously talked about this, see "Securing DNS: Doing DNS As If DNS Actually Mattered," http://pages.uoregon.edu/joe/secprof10-dns/secprof10-dns.pdf

Starting With The Basics

- Given how important and useful DNS can be, it would be great if all DNS servers were correctly configured and operating well.

- **You should periodically check your domain's DNS using a free DNS checker such as the one that's at http://dnscheck.iis.se/**

  Do all your domains get a "clean bill of health?" Checking at least some Merit-related university domains, I'm seeing some domains that look great, but others where DNS errors exist.

- **For a service this mission critical, that's crazy.**

- Here's the sort of thing you SHOULD be seeing...

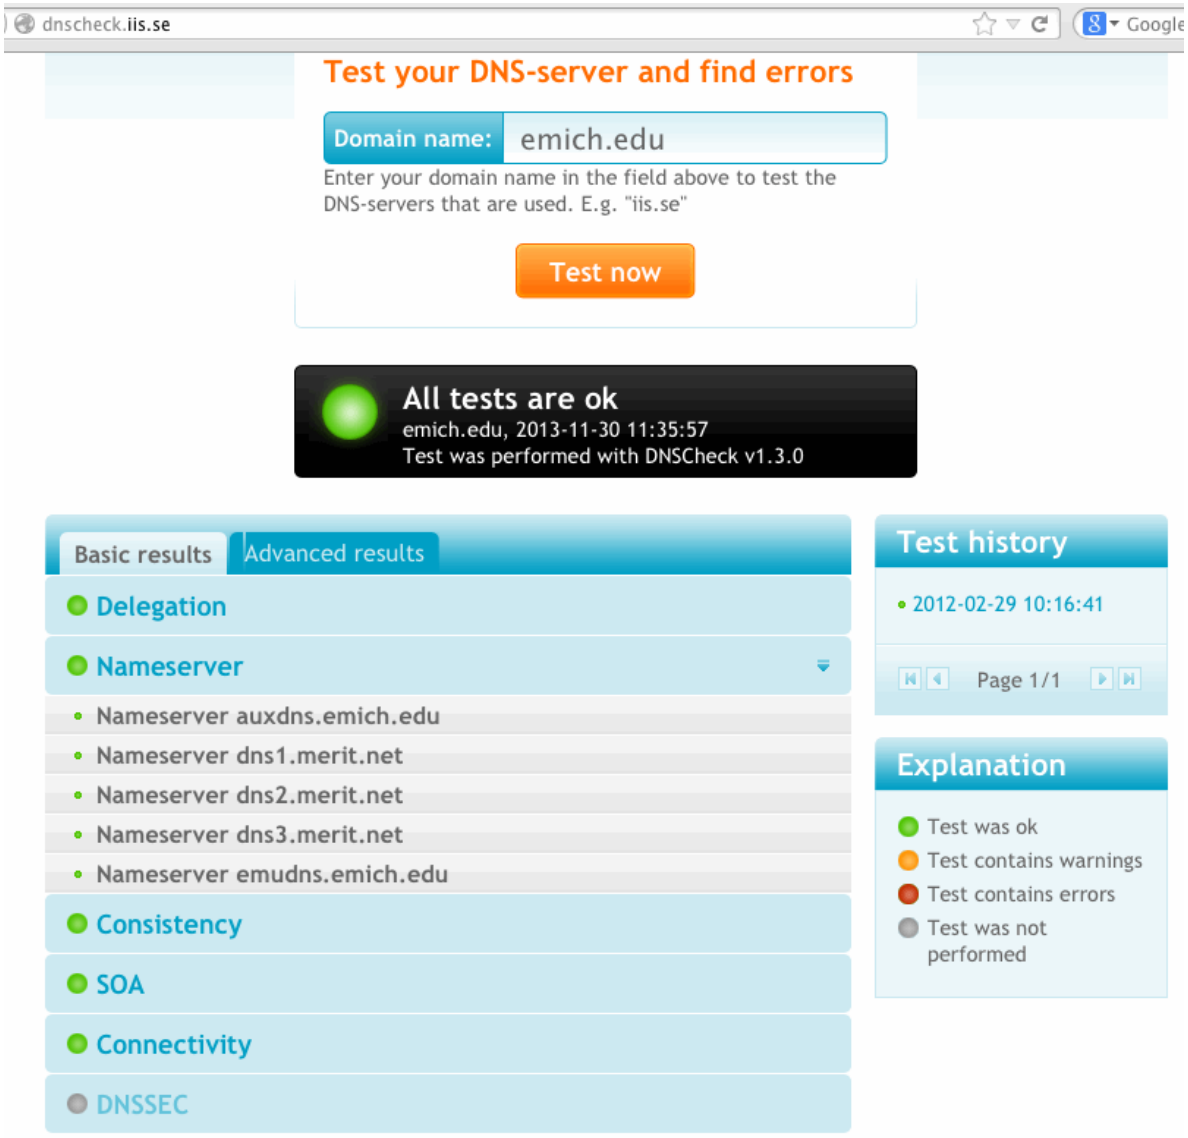

Unfortunately, not all "members of the class" get an equally clean report!

Just Having Correctly Configured DNS Isn't Enough: You Need DNSSEC, Too

- You can have perfectly configured and fully functional DNS servers, and yet still have important DNS-related work left to do.
- For example, what about DNSSEC? If you're not currently using DNSSEC to cryptographically protect your DNS, your infrastructure is vulnerable to cache poisoning attacks.
- Because DNS serves the critical function of mapping names to IP addresses, you need DNS to be a "trustworthy guide" and not an unfaithful servant that may pretend to take you where you want to go only to actually drop you off in a dangerous neighborhood.
- DNSSEC helps cryptographically protect that mapping process -- IF sites bother using it. Most sites don't. If sites don't use DNSSEC, an attacker can replace real IP addresses with alternative addresses of their choosing. That can be disastrous.

Signing Your Own DNS Zones and/or Validating Domains That Others Have Signed

- When it comes to DNSSEC, there are **two things** you can choose to do: you can sign your own authoritative zones (as Internet2 and InCommon currently do), or your recursive resolvers can validate the domains that others have signed (as the University of Oregon currently does). Ideally, **you should do both.** If you want to start slowly, unless your domains are potential high value targets for hijacking, I'd suggest beginning by validating the DNSSEC signatures of other sites' DNSSEC signed domains.
- Fortunately that's easily accomplished if you're using BIND (technically, it's literally a matter of adding a couple of lines into BIND's config file and then restarting BIND):

```
dnssec-enable yes;
dnssec-validation yes;
```

A Caveat: DNSSEC Can Increase DNS Fragility

- If you do decide to validate DNSSEC-signed domains, one thing to be aware of: if a site has signed its domains but lets its keys expire or otherwise "screws things up," DNSSEC **will** "perform as designed," and those DNSSEC-secured domains **won't resolve** for you even though those domains may still resolve just fine for all the other sites that **aren't** validating DNSSEC signatures. Some refer to this as "increased DNS fragility" because even innocent/accidental DNSSEC crypto errors can end up knocking entire domains offline.

- If you do decide to try enabling DNSSEC validation, you need to know about DNSviz, a DNSSEC zone checking tool written by Casey Deccio of Sandia National Laboratory. DNSviz is the "go-to" tool for debugging DNSSEC-signed zones if/when something goes wrong, and you need to understand what. See http://dnsviz.net/

Open Recursive Resolvers And DDoS Attacks

- While DNS is an essential service, it has a number of properties that make it particularly prone to being exploited for distributed denial of service attacks IF not correctly configured.
- **It is critical that only authorized local users should be able to use your recursive resolvers to resolve arbitrary names.**
- **Unfortunately ~28 million sites run with their recursive resolvers open to any user.** That list includes ***some Merit schools.*** When recursive resolvers are open to anyone, it is common for them to be used as part of a DNS amplification attack, DDoS'ing innocent victims. For example, Spamhaus was hit with a **300Gbps DDoS** via open recursive resolvers. **THAT'S REALLY CRAZY.**
- The Internet really needs you to make sure your recursive resolvers have been appropriately locked down. See http://www.team-cymru.org/Services/Resolvers/ for details.

Rate Limiting Authoritative Name Servers

- At the same time you fix any open recursive resolvers at your site, be sure to also **check your authoritative name servers. They should be rate limited** so that they can't be exploited as yet another DNS-based DDoS attack tool.

- For more information about rate limiting authoritative name servers to prevent abuse, see:

  http://www.redbarn.org/dns/ratelimits

A Third Critical Bit: BCP38/BCP84

- In addition to securing your open recursive resolvers and rate limiting your authoritative name servers, the other thing the Internet urgently needs you to do is to filter traffic with spoofed source IPs.
- The principle behind BCP38/BCP84 is really pretty simple: your network shouldn't be emitting traffic with source addresses pretending to be from someone else's address range. For example, UO's network address block is 128.223.0.0/16. Give that, there's no reason why devices in that range should be emitting traffic that appears to be from someone else's IP addresses.
- Many (but not all) networks currently do BCP38/BCP84 filtering. See summary stats at: http://spoofer.cmand.org/summary.php
- Does your network do BCP38/BCP84 filtering? It should!
- With the permission of your local network admin, check and see if your network does using http://spoofer.cmand.org/software.php

DNS As A Policy Enforcement Mechanism

- Because most users access sites either by clicking on a link that contains a domain name or by typing in a domain name manually, domain names also potentially represent a way to "enforce policy."
- For example, Response Policy Zones (RPZ) can be beneficially used to block access to the C&C domains used by at least some bots, worms, and other malware. That process works because people consensually *want* to be protected from that sort of badness. (For a nice introduction to RPZ, see https://dnsrpz.info/)
- On the other hand, some governments, including our own Congress, have considered trying to use the domain name system to block access to some content: remember SOPA and PIPA, back in December 2011? If you'd like to know why SOPA technically wouldn't have worked if Congress had tried to "impose it from above," see http://pages.uoregon.edu/joe/managers-amendment/sopa-amended-version.pdf (pwd: "final")

Getting Involved

- When it comes to DNS-related work, there are also many opportunities to participate.

- If you're profoundly interested in DNS, perhaps as a researcher or implementer, you probably already know about DNS-OARC, but if not, see https://www.dns-oarc.net/

- If you're a DNS operator, or a member of the cyber security/anti-abuse community, a nice opportunity to get involved with DNS-related work is through participating in the Farsight Security SIE (Security Information Exchange), see https://archive.farsightsecurity.com/ for more.

5. BGP Security

"Revealed: The Internet's Biggest Security Hole"
http://www.wired.com/threatlevel/2008/08/revealed-the-in/

"The New Threat: Targeted Internet Traffic Misdirection,"
http://www.renesys.com/2013/11/mitm-internet-hijacking/

BGP

- Most end users have no idea how Internet traffic gets routed from their ISP to its destination. They simply have no idea what's "happening under the hood." "It's magic." "It just happens."

- Network engineers, however, can tell you that BGP (the "Border Gateway Protocol") is the key underlying magic (technically, the "exterior routing protocol") that helps packets get where they need to go. Relevant BGP RFCs can be seen at http://www.bgp4.as/rfc

- Given the size of the Internet and the basic simplicity of BGP, the fact that BGP works and scales as well as it has is really quite impressive.

"Oh Noes, Mr. Bill!"

- Unfortunately, BGP is potentially subject to a variety of intentional (and accidental) attacks.
- One such attack is known as "route injection" or "BGP hijacking."
- In a route injection attack, a site "injects" or "advertises" unauthorized routes via BGP for part or all of someone else's address space.
- When that happens, particularly if the injected route is "more specific" than the normally-advertised-route, network traffic that should be going via the authorized route to its real destination instead gets misrouted ("hijacked") via the evil competing route.
- This has been well described by *Wired* as "the Internet's biggest security hole," and has recently been seen getting exploited (see the Renesys write up mentioned on the title slide for this section)

Yes, This Vulnerability *Is* Getting Exploited

- Route injection attacks have been observed many other times on the Internet. One of the most famous route injection incidents occurred in 2008, when Pakistan, in an attempt to domestically limit access to the video sharing site YouTube, accidentally leaked routes for YouTube's address space worldwide. The route monitoring company Renesys has a nice summary of this incident, see "Pakistan Hijacks YouTube," http://www.renesys.com/2008/02/pakistan-hijacks-youtube-1/
- Besides those sort of accidental incidents, as the Internet comes increasingly close to exhausting its supply of IPv4 address space, we'll see more and more address space hijacking attacks by spammers and other miscreants.
- BGP route injection can also be exploited by intelligence services. They can temporarily reroute selected traffic, eavesdrop upon it, and then silently reintroduce it for "normal"-appearing delivery.

Deterring BGP Hijacking

- Multiple approaches have been tried over the years to prevent these sort of vulnerabilities, most of them only partially successfully.

- A minimum standard of care entails providers checking whois to verify assignment of any provider-independent address block that an ISP gets asked to route for a customer, and requiring customers to provide a letter of authorization if the provenance of a particular netblock is at all clouded.

- See, e.g., http://business.comcast.com/enterprise/services/internet/ethernet-dedicated-internet/edi_tech_specs at section 5.9.

Routing Registries

- Another approach is the use of "routing registries" (see http://www.irr.net/) with "Routing Policy Specification Language" (RPSL). A nice intro to "Using RPSL in Practice" can be seen in RFC2650, by Dave Meyer et. al. In a nutshell, routing registries allow ISPs to describe the routes they originate and the ASs ("Autonomous Systems") that should be announcing them.
- If everyone was conscientious about documenting their routes in routing registries, and all network service providers built their operational routing filters directly from routing registry data, it would be difficult for a third party to accidentally hijack another site's address space.
- An example of an ISP that requires customers to use a routing registry can be seen at http://www.us.ntt.net/support/policy/routing.cfm#RR
- Merit, of course, also runs the Merit RADB, see http://www.ra.net/

RPKI

- There has also been growing community interest recently around RPKI. RPKI uses cryptographically-verifiable certificates, known as "Route Origin Authorizations," or "ROAs", to specify what ASNs are authorized to originate a particular prefix. ROAs are normally issued by the RIRs (ARIN, RIPE, APNIC, etc.), see for example https://www.arin.net/resources/rpki/

- Unfortunately, as noted in http://tools.ietf.org/id/draft-ietf-sidr-origin-ops-22.txt at section 7, RPKI is **not** intended to deal with malicious/intentional route injection, but only inadvertent incidents (such as the Pakistan YouTube incident).

- Uptake of RPKI to-date has been somewhat limited to-date, too: http://rpki.surfnet.nl/validitytables.html reports just over 20,000 valid routes at the time this talk was prepared.

BGPSEC

- Yet another stream of work-in-progress involves BGPSEC, see http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-overview-03 and http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-07 (draft expired August 25th, 2013).
- BGPSEC builds on RPKI, but endeavors to secure the chain (or ASPath) of autonomous systems that should be originating each authorized prefix.
- A discussion of the threat model underlying BGPSEC can be seen in http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-threats-07
- Unfortunately, there appears to be little current momentum around BGPSEC, see for example the discussion in Section 6 of http://www.cs.stonybrook.edu/~phillipa/papers/CCR2014.pdf (preprint of a paper to be published in the January 2014 issue of the ACM Computer Communications Review).

Detecting BGP Hijacking

- While this all gets sorted, one alternative focuses on at least **detecting** route hijacking if/when it does happen.

- Detecting route hijacking typically depends on the availability of routing data from ISPs all around the world, since "every routing table is different," and some hijacked routes may not propagate globally. (One such repository of routing data is the Oregon Routeviews Project, http://www.routeviews.org/ ; another resource is http://www.ripe.net/data-tools/stats/ris/ris-raw-data).

- Some companies also offer productized route monitoring, see for example http://www.bgpmon.net/services/route-monitoring/ (free for up to five prefixes)

- **Are you monitoring YOUR routes? You should be!**

Route Deaggregation

- If the risk of route injection isn't intelligently managed via one of the preceding options, some providers may attempt to minimize their risk of experiencing route injection via deaggregation, or the announcement of multiple more specific netblocks (rather than using maximally-aggregated routes).

- When this happens, **every** border router on the Internet ends up getting penalized due to having to carry all those additional routes.

- See for example http://www.cidr-report.org/as6447/#Aggs

- Surprisingly, some universities show up on the report I checked on December 2nd, 2013...

Top ASNs Advertising More Specific Routes

More Specifics	Total Prefixes	ASnum	AS Description
5381	5494	AS4538	ERX-CERNET-BKB CN Education & Research Network
4395	4535	AS7029	WINDSTREAM - Windstream Communications Inc
3665	6082	AS3	MIT-GATEWAYS - Massachusetts Institute of Technology
3601	4552	AS4	ISI-AS - University of Southern California
3421	3428	AS28573	NET Serviços de Comunicação S.A.
3192	3272	AS577	BACOM - Bell Canada
3004	3041	AS6389	BELLSOUTH-NET-BLK - BellSouth.net Inc.
2848	2930	AS4766	KIXS-AS-KR Korea Telecom
2747	2962	AS4323	TWTC - tw telecom holdings, inc.
2696	2709	AS17974	TELKOMNET-AS2-AP PT Telekomunikasi Indonesia
2672	2673	AS10620	Telmex Colombia S.A.
2201	2748	AS2	UDEL-DCN - University of Delaware
2144	2212	AS22773	ASN-CXA-ALL-CCI-22773-RDC - Cox Communications
2040	2113	AS7545	TPG-INTERNET-AP TPG Telecom Limited
2033	2052	AS18566	MEGAPATH5-US - MegaPath Corporation
1992	2268	AS174	COGENT Cogent/PSI
[etc]			

[Hmm. AS2, AS3, AS4. Hmm. Do I detect a pattern?]

A Few BGP-Related Involvement Opportunities

- If you're an engineer, are your routes in a routing registry -- and up-to-date? Are you also be monitoring your university's prefixes? Have you gotten RPKI ROAs for all your netblocks?
- If you're into cyber security research and want to push yourself, this would be a wonderful area to choose. Of ALL the areas we've talked about, the security of BGP is in by far the worst shape and most in need of your contribution.
- A couple of excellent FCC CSRIC resources for more information:
  -- "BGP Security Best Practices,"
  http://www.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG4_Report_March_%202013.pdf
  -- "Secure BGP Deployment,"
  http://www.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG6_Report_March_%202013.pdf

6. Malware

"This quarter McAfee Labs cataloged 18.5 million new malware samples, bringing the total McAfee “zoo” to more than 147 million unique pieces of malware."

http://www.mcafee.com/us/resources/reports/rp-quarterly-threat-q2-2013-summary.pdf

"Sophos' findings is that approximately 1 in 36 Mac systems (2.7 percent) were found to have true OS X-based malware on their systems. Of these systems, the majority were infected with the recent Flashback malware [...].

http://reviews.cnet.com/8301-13727_7-57420025-263/one-in-five-macs-infected-with-malware-is-inaccurate/

The Malware Problem in General

- There's a lot of malware that has been created to-date: note McAfee's assertion on the preceding slide that their malware collection had **147 million variants** as of the middle of this year.

- Malware installation has also become such a sophisticated science that a PC user doesn't need to do anything except visit a perfectly normal/routine site to become infected.

- Another key fact: miscreants can create new malware (or new variations on old malware) faster than antivirus venders can produce new A/V signatures to identify, block and remove that malware. Signature-based antivirus is increasingly less than satisfactory.

Most Malware Targets MS Windows and Android

- Virtually all malware targets MS Windows on the laptop/desktop, or Android in the mobile space. If you're not using those operating systems, your risk of getting infected with drops dramatically.
- And yet, what do most people use? Well, given that systems running MS Windows or Android have the software they want, and may cost less than half what some alternatives cost, people often discount the malware issue entirely.
- But if you worry about malware -- and **you should** -- **does it make sense to use the operating systems that most of the bad guys specialize in attacking? I don't think so.**
- **BTW, what's your plan to get Windows XP off your campus networks by April 8th, 2014, when it goes EOL? If any operating system isn't supported, you can't keep it patched and safe! (See: http://www.microsoft.com/en-us/windows/enterprise/endofsupport.aspx)**

The Helper Application Problem

- Another fact: a lot of malware leverages vulnerabilities in Java, Adobe Flash or Adobe Reader.

- Those helper apps are very popular, and hard to live without, but some estimates are that those three products may collectively account for 2/3rds of all exploited vulnerabilities. (See for example http://www.av-test.org/en/news/news-single-view/artikel/adobe-java-make-windows-insecure/)

- Basic step: encourage your users to run **PluginCheck** to get helper apps up-to-date: http://www.mozilla.org/en-US/plugincheck/

- Another terrific tool for consumers is **Secunia PSI**, see: http://secunia.com/vulnerability_scanning/personal/

Online Ads, Potential Malvertising and Trackers

- Online advertising supports many popular online sites. You may even find that some advertising tells you about intriguing offers.

- In general, however, because every ad represents a potential source of malware, you should block all advertising in your browser by default. One good tool for this is **Adblock Plus** (see https://adblockplus.org/)

- Since you're not viewing advertisements, there's also no real point to letting marketers track your online activities, right? One of the best tools for blocking this sort of thing is **Ghostery**, see http://www.ghostery.com/ (note that you may be initially surprised to see how heavily instrumented some popular sites can be!)

One Specific Recent Malware: CryptoLocker

- "CryptoLocker is Trojan horse malware which surfaced in late 2013. A form of ransomware targeting computers running Windows, a CryptoLocker attack may come from various sources; one such is disguised as a legitimate email attachment. When activated, the malware encrypts certain types of files stored on local and mounted network drives using RSA public-key cryptography, with the private key stored only on the malware's control servers. The malware then displays a message which offers to decrypt the data if a payment (through either Bitcoin or a pre-paid voucher) is made by a stated deadline, and says that the private key will be deleted and unavailable for recovery if the deadline passes. If the deadline is not met, the malware offers to decrypt data via an online service provided by the malware's operators, for a significantly higher price in Bitcoin.

  "Although CryptoLocker itself is readily removed, files remain encrypted in a way which researchers have considered infeasible to break. Many say that the ransom should not be paid, but do not offer any way to recover files; others say that paying the ransom is the only way to recover files that had not been backed up. [...]

  "Symantec estimated that 3% of users infected by CryptoLocker chose to pay. Some infected users claimed that they paid the attackers but their files were not decrypted."

  http://en.wikipedia.org/wiki/CryptoLocker

Some Implications of CryptoLocker

- **If you're not routinely backing up your system, you're nuts**
- If you do have backups, you can remove CryptoLocker and then restore your files from backups -- but many people don't have clean recent backups.
- Without backups, you need to think about whether or not you'd be willing to "pay the CryptoLocker ransom" to get your files back if you got infected/encrypted. Arguably, the reason cybercriminals have deployed this malware is that it works as a way to earn money. If people refused to pay, miscreants might stop using this strategy. But are you willing to "throw away" irreplaceable files?
- Some backup strategies (such as mirroring your files) will protect you against hard drive failures, but mirroring your files will **not** provide adequate protection against CryptoLocker. You should be saving write protected copies offline to ensure that your backups don't end up encrypted by CryptoLocker, too.

7. "Snowdonia"

"Nobody does the right thing."

Marie Kreutz, *The Bourne Identity,* 2002

Obumbrata et velata, michi quoque niteris
[In the dark and secretly, you work against me]

"O Fortuna," *Carmina Burana,* 13th century poem

The Online World: Upside Down As Of 6/2013

- In June 2013, Glenn Greenwald published an article in *The Guardian* revealing that the NSA had been vacuuming up phone records for millions of American customers who use Verizon. The Internet suddenly tilted. See http://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order
- The next day, the online world turned completely upside down when the *Washington Post* subsequently reported on PRISM, see "U.S., British intelligence mining data from nine U.S. Internet companies in broad secret program," www.washingtonpost.com/investigations/us-intelligence-mining-data-from-nine-us-internet-companies-in-broad-secret-program/2013/06/06/3a0c0da8-cebf-11e2-8845-d970ccb04497_story.html
- A seemingly endless stream of additional Snowden revelations continue to be released through various newspapers worldwide.

The Seemingly Endless Stream of Revelations

- Those front page revelations are perfect examples of just how crazy things have become today – it really wasn't routine to see papers publishing stories about pervasive domestic Internet monitoring.
- Crazy as things already have become, **things will probably only get crazier in the days ahead.** Why? Well, only **1%** of all the documents Snowden gave to journalists have been published so far. (http://world.time.com/2013/12/03/guardian-editor-says-paper-only-published-1-of-snowden-nsa-leaks/)
- Only Snowden, his immediate colleagues, and the intelligence community itself know precisely what may be coming in all those other documents, but just based on what has *already* been publicly released, 2013 clearly marks the dawn of a new era for the Internet and its security and privacy.
- The best we can do right now is probably to try understand why folks are doing what they're doing...

Why Did Snowden Become a Whistleblower?

- Based on what has been disclosed, **Edward Snowden seems to have thought that given what he knew, he had no moral option but to become a "whistleblower," informing the public about what was being done in their name, and what was being done to them, notwithstanding his secrecy oath.**

- In doing so, he's already paid a tremendous personal price: he's now living in self-exile in Russia. He may yet be returned to the U.S., tried, and put in prison for life, or even put to death.

- Doing so would be a major mistake: if executed, Snowden would become a martyr, and and his death might trigger the wholesale release of all the documents he reportedly took and then cached in "button down mode," as a sort of insurance.

Why Did The NSA Do All These Spy Programs?

- The National Security Agency's actions, even including some fairly astonishingly programs, were undertaken with the best of intentions: **the NSA genuinely wanted to keep Americans safe.**

- As an American, I think that's a good objective to pursue.
  I want to be safe. I'm just not willing to accept potentially unlimited encroachments on my privacy to obtain that goal.
  I want the government to operate strictly within the Constitution, and in a proportionate way.

- Whether the NSA went too far or used unacceptable means in pursuing the various programs that Snowden has revealed is a matter that will ultimately be decided by the court system and Congress, not by you or me.

<u>Why</u> Did The Media Publish the Documents?

- Some officials have been appalled at what **the media** has chosen to publish.

- To understand the media's actions, you need to remember that the media views itself as a pseudo-"4th branch of government," responsible for **educating the public and providing a final check and balance against the otherwise unlimited power** of the Executive Branch, the Congress, and the Judiciary. In the reporting I've seen, the media has exercised considerable self-restraint, including self-redacting content before releasing copies of documents they'd been given.

- Any attempt at prior restraint would be wholly inconsistent with 1st Amendment protections here in the United States.

Why Were Civil Liberty Groups Irate?

- **Numerous civil liberty & privacy groups** are also understandably incensed by what has been disclosed about the actions of the NSA -- merely knowing that what you say might be getting monitored can have a significant chilling effect on political or religious speech, or private attorney/client discussions, etc.

- A representative list of organizations that want the NSA to stop snooping on users of the Internet includes the ACLU, the American Library Association, the Association of Research Libraries, the Center for Democracy and Technology, the Electronic Frontier Foundation, the Free Software Foundation, the Internet Archive, and many others -- see the list at https://optin.stopwatching.us/

- We should also note that the **U.S. State Department** has also gone on record opposing pervasive Internet surveillance and censorship...

Hillary's "Remarks on Internet Freedom"

- "Some countries have erected electronic barriers that prevent their people from accessing portions of the world's networks. They've expunged words, names, and phrases from search engine results. **They have violated the privacy of citizens who engage in non-violent political speech.** These actions contravene the Universal Declaration on Human Rights [...] And we must also grapple with the issue of anonymous speech. Those who use the internet to recruit terrorists or distribute stolen intellectual property cannot divorce their online actions from their real world identities. **But these challenges must not become an excuse for governments to systematically violate the rights and privacy of those who use the internet for peaceful political purposes."**

  Secretary of State Hillary Clinton, January 21, 2010, http://www.state.gov/secretary/rm/2010/01/135519.htm [emphasis added]

The Perspective of American Businesses

- Another party impacted by the NSA monitoring consists of **American businesses**, as a result of foreign customers changing networking plans due to a loss of trust in American providers:

-- "**Cisco has seen a huge drop-off in demand for its hardware in emerging markets, which the company blames on fears about the NSA using American hardware to spy on the rest of the world.** Cisco chief executive John Chambers said on the company's earnings call that he believes other American technology companies will be similarly affected. Cisco saw orders in Brazil drop 25% and Russia drop 30%."
[http://qz.com/147313/ciscos-disastrous-quarter-shows-how-nsa-spying-could-freeze-us-companies-out-of-a-trillion-dollar-opportunity/ (emphasis added)]

-- "Earlier this month The Information Technology & Innovation Foundation (ITIF) published a prediction that the U.S. cloud computing industry stands to lose up to $35 billion by 2016 thanks to the National Security Agency (NSA) PRISM project, leaked to the media in June. **We think this estimate is too low and could be as high as $180 billion or a 25% hit to overall IT service provider revenues in that same timeframe."**
[http://blogs.forrester.com/james_staten/13-08-14-the_cost_of_prism_will_be_larger_than_itif_projects (emphasis added)]

Eight Leading Internet Companies Urge Reform

- On December 9th, 2013, AOL, Apple, Facebook, Google, LinkedIn, Microsoft, Twitter, and Yahoo issued an open letter to President Obama and to Congress urging Washington to reform government surveillance (see http://reformgovernmentsurveillance.com/).

- While acknowledging that governments need to take action to protect their citizens safety and security, those leading Internet companies called for Washington to adopt five key common-sense principles to shape and govern surveillance efforts moving forward:

  -- Limit Governments' Ability to Collect Users' Information
  -- Oversight and Accountability
  -- Transparency About Government Demands
  -- Respecting the Free Flow of Information
  -- Avoiding Conflicts Among Governments

Snowden's Impact on *Other* Members of the IC?

- The NSA is just one of 16 publicly acknowledged American intelligence community (IC) agencies. See http://www.dni.gov/index.php/intelligence-community/members-of-the-ic
- Because of the NSA's actions and Snowden's disclosures about them, all parts of the IC are likely to run into more (and stronger) encryption, potentially interfering with their ability to do critical work. In other cases, some companies collaborating with the IC may reconsider continuing to do so.
- Non-NSA members of the IC may also be becoming more potentially *over-collection-prone* than they otherwise might have. See the logo for a recent classified NRO satellite launch...

What Do Average People Think?

- And then there's the perspective of the **average person.** Poll data tells us that Americans are nearly evenly divided. According to the Huffington Post (citing an Angus Reid Global online poll):

  -- 51% of Americans said the NSA leaker was “something of a hero who should be commended for letting the public know that our governments are running electronic surveillance programs that threaten people’s privacy,"
  -- 49 % labeled him “more of a traitor who should be condemned for publicizing security activities and threatening western intelligence operations."

  [http://www.huffingtonpost.com/2013/10/30/edward-snowden-poll_n_4175089.html]

- Nice example of how well settled all our perspectives can be, eh?

Legalities

- While there seems little question that Snowden violated his secrecy agreement, his disclosures *have* raised legitimate questions about the legality of some of the NSA programs he exposed, e.g.:

> **Attorney General Eric Holder** claimed Tuesday that members of the Obama administration had concerns about the extent of the National Security Agency's surveillance operations before former NSA contractor Edward Snowden leaked details to the press. In an interview on CNN, Holder singled out the NSA's controversial program to collect records on all U.S. phone calls. [...] **Sen. Patrick Leahy** (D-Vt.), the chairman of the Senate Judiciary Committee, and **Rep. James Sensenbrenner Jr.** (R-Wis.), the original author of the Patriot Act, have introduced legislation to end the NSA's bulk collection of phone records.
>
> See "Holder Questions NSA Phone Data Collection," Nov 5, 2013, http://thehill.com/blogs/hillicon-valley/technology/189343-holder-questions-vast-nsa-phone-data-collection [emphasis added]

The Problem With Bulk Collection

- I share the concerns of Attorney General Holder, Senator Leahy, and Congressman Sensenbrenner.

- Specifically I worry about the constitutionality of bulk collection of call records for all domestic phone calls. A traditional criminal search warrant, if issued with similarly sweeping scope, would fail appellate review for a **"lack of particularity,"** much like a blanket search warrant authorizing law enforcement to search "any or all people, premises or documents located in the state of Michigan" for unspecified illegal activity.

- BTW, do you remember learning about colonial **"writs of assistance"** in your high school history classes? If those lessons have faded, see http://en.wikipedia.org/wiki/Writ_of_assistance

The NSA Might Have Assumed There Actually *IS* Precedent For Bulk Collection

- If you spend time looking for a reason why the NSA might have assumed it was okay to "bulk collect" "metadata" (e.g., source and destination addresses plus call timing details, etc.), you don't need to look very hard to find examples of other surprisingly sweeping surveillance programs.
- For example, on August 2nd, the NY Times reported on the U.S. Postal Service's "Mail Isolation and Tracking Service:"
  "Postal Service Confirms Photographing All U.S. Mail,"
  http://www.nytimes.com/2013/08/03/us/postal-service-confirms-photographing-all-us-mail.html?_r=0
- I could see an NSA person arguing that if it's okay to collect the delivery & return address from all postal mail, why shouldn't it also be permissible to collect source & destination addresses from all email messages, or calling & called numbers from phone calls?

Drilling Holes In the Bottom of Our Own Boats

- Another disturbing revelation was that the NSA may have **intentionally weakened** or c**ompromised the strength and technical integrity** of some cryptographic protocols as part of the standards development process. We all need to be able to rely on these protocols to secure confidential information online, but we can't if they've been intentionally weakened or compromised.
- A specific example of this: **RSA** has now publicly told its customers to stop using the NSA-influenced Dual_EC_DRBG random number generator that had been used for key parts of some RSA products (see http://arstechnica.com/security/2013/09/stop-using-nsa-influence-code-in-our-product-rsa-tells-customers/)
- **NIST** has also noted the loss of trust that the NSA protocol manipulation efforts have caused. In an effort to regain some of that lost trust, NIST has begun a review of its cryptographic standards development process: http://csrc.nist.gov/groups/ST/crypto-review/

"But What About The Fight Against Terrorism?"

- The rationale for many intelligence community activities is that they "protect Americans from terrorism and other security threats."
- A closer review of the terrorism problem makes it clear that except for certain (thankfully rare) exceptions (such as use of nuclear, chemical, biological or radiological weapons), most terrorist attacks -- while unquestionably awful and completely reprehensible -- don't *directly* have much of a national impact.
- To have a real national impact, terrorists need to count on either (a) **media amplification** or (b) **official over-reaction**. If terrorists can get extensive media coverage or government officials to over-react, then they can force us to become scared and to take self-defeating actions as a result. For example, the NY Times estimates that by 2011 we had spent **$7 million dollars for every one dollar** that Al Qaeda spent planning and executing the terrible 9/11 attacks. (http://www.nytimes.com/interactive/2011/09/08/us/sept-11-reckoning/cost-graphic.html)

Risk Management; Proportionate Responses

- Our goal should be the prudent management of terrorism-related risks, and, if terrorist incidents do occur despite those reasonable efforts, measured and proportionate response against those perpetrators and those supporting them. Put another way, "How much is enough?" Just because "we can" doesn't mean "we should." We cannot have "perfect security" in anything resembling a free society. We cannot eliminate all risks to our security...
- Trying to attain that standard of perfection would come at too great a cost to the country, requiring us to forgo many of the Constitutional freedoms that make this country what it is. If we try to do so, we will have won a Pyrrhic victory ("a victory with such a devastating cost, it is tantamount to defeat").
- Don't let the terrorists win by pushing us into a totalitarian system of pervasive monitoring where privacy for law-abiding citizens becomes a "luxury" of some dimly-remembered past.

Side Effect: Damage To Relations With Key Allies?

- If we do go overboard in our use Internet monitoring as a tool against terrorism, we risk damaging crucial alliances.
- You may have heard about the "Five Eyes" intelligence sharing alliance of English-speaking countries (the US, Canada, Great Britain, Australia and New Zealand), but what about all our *other* allies? In most cases, apparently every other country is "fair game for surveillance," see http://www.theguardian.com/world/2013/dec/02/nsa-files-spying-allies-enemies-five-eyes-g8
- A specific example: the United States allegedly tapped German Chancellor Angela Merkel's phone; see http://www.theregister.co.uk/2013/11/26/merkel_phone_tapped_by_5_countries/
- If we treat key allies (like Germany) this way, how long can we count on them remaining our allies? America needs all the friends abroad that we can get these days...

Good for the Goose, Equally Good for the Gander?

- Speaking of other countries, we also need to assume (as a purely pragmatic matter) that if America runs surveillance programs targeting them, those other countries will reciprocate with similar programs that *target us*. [See, e.g., "Uproar over French plan to extend online spying," http://www.thelocal.fr/20131126/france-surveillance-privacy-internet-online-snowden-nsa]
- Thus, even if you're completely comfortable trusting the U.S. government to protect our constitutional right to privacy, would you be *equally* comfortable trusting every *other* country to do the same?
- As Nicholas Weaver said in *Wired*, "We now live in a world where, if we are lucky, our attackers may be every country our traffic passes through except our own." www.wired.com/opinion/2013/11/this-is-how-the-internet-backbone-has-been-turned-into-a-weapon/
- As we do unto others, so will they do unto us.

PRISM and Higher Education Email

- PRISM is the "number one source of raw intelligence used for NSA analytic reports," accounting for 91% of the NSA's Internet traffic acquired under FISA section 702. It is reportedly based on messages collected directly from the servers of **Microsoft, Yahoo, Google, Facebook, PalTalk, AOL, Skype, YouTube and Apple.** [See www.washingtonpost.com/wp-srv/special/politics/prism-collection-documents/ and https://www.eff.org/sites/default/files/filenode/fisc_opinion_-_unconstitutional_surveillance_0.pdf]

- **Many** schools currently outsource their student email, if not *all* their email, to providers chosen from that set of providers.

- FWIW, http://www.merit.edu/services/meritmail/ has **NOT** been mentioned as an email services involved with PRISM collection.

No ISP Is Going To "Take a Bullet" for You

- As you think about your email and its privacy, let's be pragmatic: it isn't realistic to expect *any* provider to refuse to obey a lawful court order. No ISP will defy a judge and go to jail (or be forced to pay huge contempt of court fines) just to protect your privacy.
- Thus, if you, as an end user, really want to protect your email privacy, you'll need to use an **end-to-end email encryption** technology such as **PGP, GNU Privacy Guard**, or **S/MIME**, where only YOU control the private key needed to decrypt those messages. (For a tutorial on S/MIME, see http://pages.uoregon.edu/joe/secprof2012/sec-prof-2012-client-certs.pdf).
- Let me also note that using PGP, GPG or S/MIME won't be painless: you will sacrifice some convenience for your privacy.
- Even if the contents of your email are fully protected by strong encryption, you should also know that your email messages will still be potentially vulnerable to traffic analysis.

"Huh? Traffic Analysis?"

- An intelligence analyst doesn't always need to be able to see the plain text contents of a communication in order to be able to infer important information about a conversation.
- Sometimes just knowing that any conversation is happening between two particular parties is enough to convey significant information. A classic example: a government official in a sensitive role begins exchanging non-official messages with a known agent of a foreign power. If that was noted, a big red flag might appropriately go up for the counterintelligence types.
- Other times, the frequency/volume of communications may signal that something's afoot. For instance, normal traffic volume might be sporadic and brief, but during an incident, message count and message volume might ramp up dramatically (this is the classic "increased chatter" phenomena you nay hear reported in the news).
- Social graphs can also be used to ID the relationships of parties.

Countering Email Traffic Analysis

- **The normal approach to defeating email-based traffic analysis is through use of an anonymous remailer**, see a description of them at http://en.wikipedia.org/wiki/Anonymous_remailer

- Some may even combine use of PGP/GPG encrypted messages with an anonymous remailer, gatewaying the output to the Usenet newsgroup alt.anonymous.messages

- Again, this needs to be done carefully to preserve the user's privacy.

  There's an excellent report you may want to read at: http://ritter.vg/blog-deanonymizing_amm.html

What About The Privacy of Web Traffic?

- We know from CAIDA statistical data (see http://www.caida.org/data/passive/trace_stats/chicago-A/2013/) that roughly 2/3rds of all Internet traffic (measured by bits/sec) consists of HTTP.
- Some Internet web traffic, such as login credentials and credit card numbers, routinely gets encrypted to protect that sensitive data from potential eavesdropping.
- The rest of the Internet's web traffic is unencrypted plain text, and can potentially be eavesdropped upon, whether that's by a foreign or domestic intelligence service, hacker/crackers, or someone else.
- The "obvious" response to this vulnerability is to encrypt all web traffic, much as (encrypted) ssh has virtually completely replaced (unencrypted) telnet for terminal sessions. People are already taking steps to encrypt more web traffic.
- Before we try to run, however, let's at least walk.

A Pretty Good Qualys SSL/TLS Server Report

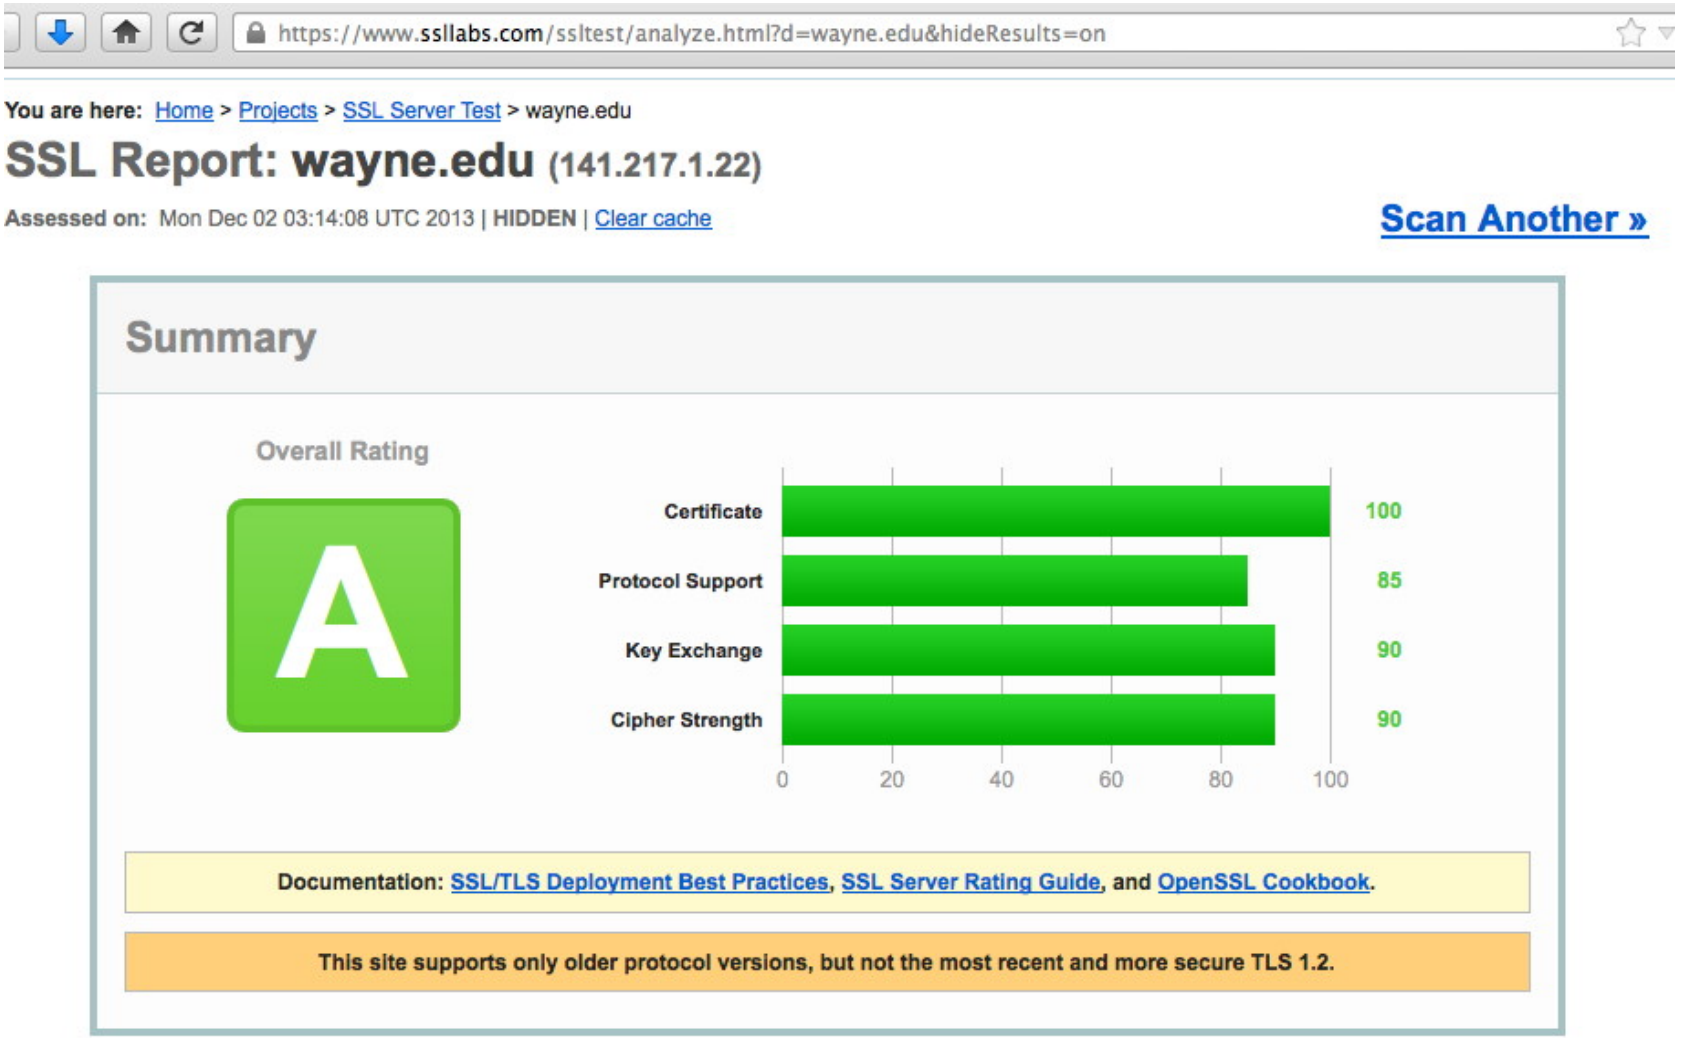

Sadly, not all Merit-related domains I checked did this well. Please check (and if necessary, fix) *your* site(s).

Some Industry Wide Summary Results

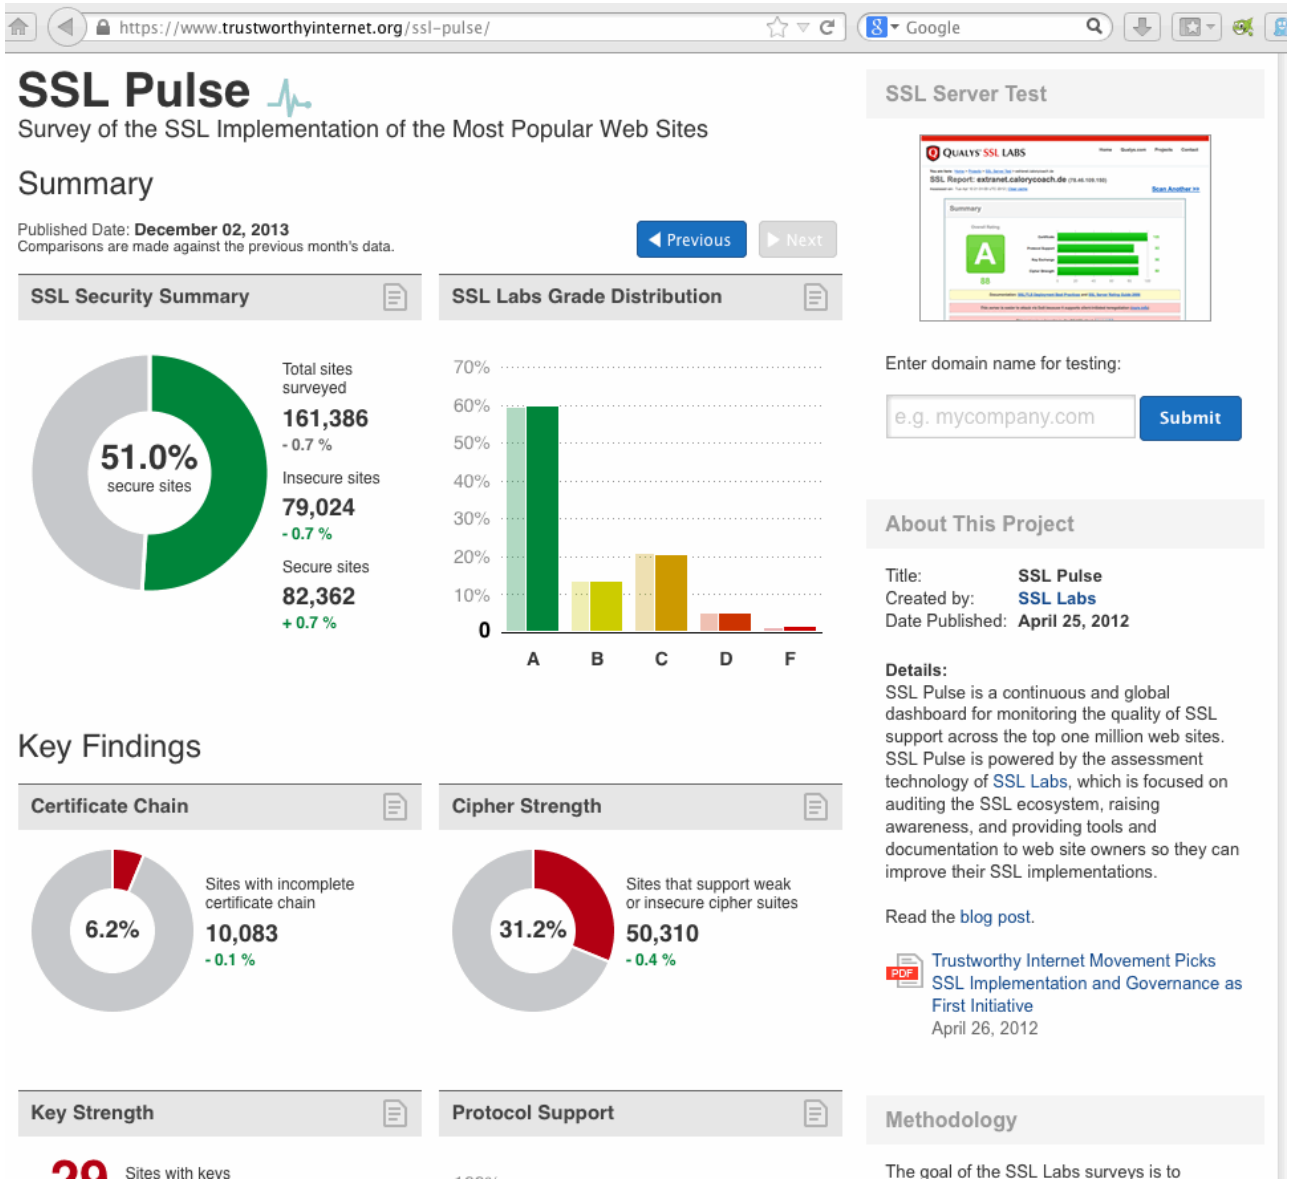

Add'l Specific Crypto Configuration Guidance

- Besides the excellent guidance available as documentation links from the Qualys tester output, you may also want to see:

- ENISA's "Recommended Cryptographic Measures," http://www.enisa.europa.eu/activities/identity-and-trust/library/deliverables/recommended-cryptographic-measures-securing-personal-data/at_download/fullReport (34 page PDF document)

- Another work in progress: "Applied Crypto Hardening is a project to define a reasonably safe copy & paste-able set of recommendations for sys admins on which crypto settings they should use on their systems. [...]"

  See http://www.bettercrypto.org/

What's The EFF Eyeballing, Web Crypto-Wise?

https://www.eff.org/deeplinks/2013/11/encrypt-web-report-whos-doing-what

	Encrypts data center links	Supports HTTPS	HTTPS Strict (HSTS)	Forward Secrecy	STARTTLS
amazon	undetermined	limited	✗	undetermined	✗
Apple	undetermined	✓ (iCloud)	✗	undetermined	✗ (me.com, mac.com)
at&t	undetermined	undetermined	✗	undetermined	✗ (att.net)
comcast	undetermined	undetermined	✗	undetermined	✗ (comcast.net)
Dropbox	✓	✓	✓	✓	✓
facebook	✓ in progress	✓	✓ planned	✓	✓ (in progress, facebook.com)
foursquare	undetermined	✓	✓	undetermined	✗
Google	✓	✓	in progress for select domains, see notes	✓	✓
LinkedIn	✗ contemplating	✓ planned 2014	✓ planned 2014	✓ planned 2014	✗ contemplating
Microsoft	✗	✓	✗	undetermined	✗ (outlook.com)
myspace	undetermined	✓	✗	undetermined	✗

HSTS

- HTTP Strict Transport Security is defined in RFC6797 (November 2012), and provides a way for domains to specify that **ALL web connections to a given domain should ONLY happen via encrypted (https) connections.**

- HSTS also eliminates (as an additional protocol feature) any possibility of users mistakenly "clicking through" any SSL/TLS errors associated with a domain's cert. Once HSTS is live, either the cert's right, or you aren't given the option to inappropriately trust it when you shouldn't!

- **Should your campus be considering adoption of an HSTS policy for your campus domains?**

HTTP 2.0 and Ubiquitous SSL/TLS

- Mark Nottingham, chair of the IETF httpbis working group (charged with developing the long awaited 2nd major version of the HTTP protocol), stated on Nov 13, 2013 that:

I believe the best way that we can meet the goal of increasing use of TLS on the Web is to encourage its use by only using HTTP/2.0 with https:// URIs.

This can be effected without any changes to our current document; browser vendors are not required to implement HTTP/2.0 for http:// URIs today. However, we will discuss formalising this with suitable requirements to encourage interoperability; suggestions for text are welcome.

To be clear - we will still define how to use HTTP/2.0 with http:// URIs, because in some use cases, an implementer may make an informed choice to use the protocol without encryption. However, for the common case -- browsing the open Web -- you'll need to use https:// URIs and if you want to use the newest version of HTTP.

"Moving forward on improving HTTP's security"
http://lists.w3.org/Archives/Public/ietf-http-wg/2013OctDec/0625.html

TLS, Certificates and MITM Risks

- While having all HTTP traffic encrypted by default would greatly reduce the risk of traffic being routinely monitored, this requires us to be able to trust sites secured with certificates, specifically, SSL/TLS certificates issued by trusted certificate authorities (CA).
- Unfortunately, as currently implemented, ANY trusted CA has the technical ability to issue a trusted cert for ANY site.
- Sites don't currently have any generally deployed way to say, "Hey, the certificate you SHOULD be seeing for our site is certificate number *foo* from certificate provider *bar*. If you see a certificate from anyone else, it's fake, so don't trust it!"
- If the IETF's DANE work (https://ietf.org/wg/dane/charter/) gets widely adopted, that will help tremendously, but that's still "work in progress," and work that will require sites to do DNSSEC (as urged, earlier in this talk).
- In the mean time, maybe try using CertPatrol in your browser?

Compelled Disclosure of Private Keys

- Another risk to relying on SSL/TLS is the judicially-compelled disclosure of private keys. **We now know that this has actually happened, at least to Ladar Levison, owner of the ISP "Lavabit."** There was an excellent Q&A session on this at NANOG, Oct 9th: http://www.youtube.com/watch?v=uo9-0So2A_g
- Under threat of contempt, Ladar was compelled to provide a copy of his SSL/TLS private keys, keys that protected the data of over 400,000 customers. With a copy of those private keys, the security of *all* Lavabit users could have been completely undercut. After being forced to surrender his keys under seal, Ladar took the only action he felt was morally left to him: he shuttered his business.
- When Godaddy learned that their customer's private keys may have been compromised, they revoked his certificate: www.forbes.com/sites/kashmirhill/2013/10/09/godaddy-pulls-lavabits-security-creds-because-the-government-got-ahold-of-its-encryption-keys/

Precluding Compelled Private Key Disclosure

- Given the importance of SSL/TLS private keys, **those who rely on public key cryptography should consider protecting those keys through use of a hardware security module (HSM).**
- In protecting your private keys secured with an HSM, you're wrestling with two mutually conflicting requirements:
  -- On the one hand, private keys are operationally critical. If you lose them, they can't be recreated. Thus, you'd really like to be able to back them up (thereby protecting them from loss)
  -- On the other hand, if it *is* possible for you to backup (and restore) private keys from your HSM, there's also the possibility that you could be compelled to restore those keys onto a 2nd HSM that could then be used by others for surreptitious monitoring.
- As long as your private keys are only used for transport security (and thus could easily be replaced if your HSM fails), creating the key pair on the HSM, non-exportably and w/o backup, may be best.

An Alternative Approach: Forward Secrecy

- We now also know that under some circumstances, even traditionally-encrypted network traffic may end up getting vacuumed up and retained (disk is cheap). If the collecting entity does eventually obtain the appropriate private key, they can then retrospectively decrypt everything they'd previously harvested.
- This retrospective decryption can be prevented by using "ephemeral" cipher suites that offer forward secrecy, such as ephemeral Diffie-Helman (DHE) or ephemeral Elliptical Curve Diffie-Helman (ECDHE).
- While there's a lot of interest in forward secrecy, if done imperfectly, **it can result in weak encryption being used.** See https://community.qualys.com/blogs/securitylabs/2013/06/25/ssl-labs-deploying-forward-secrecy and "How to Botch TLS Forward Secrecy," https://www.imperialviolet.org/2013/06/27/botchingpfs.html

The Web and Traffic Analysis

- Just as email can be subject to traffic analysis even if individual messages are "perfectly encrypted," the same is true of web traffic.
- To try to avoid web traffic analysis, **consider using Tor** (see https://www.torproject.org/), but note that it, too, has been targeted (albeit with limited success) by the NSA. See for example "NSA and GCHQ target Tor network that protects anonymity of web users," http://www.theguardian.com/world/2013/oct/04/nsa-gchq-attack-tor-network-encryption
- New hardware Tor implementation: https://pogoplug.com/safeplug
- BTW, in pushing users toward anonymity networks such as Tor, authorities may have helped incent the creation of an entire new world of content available only from Tor, including the infamous "Silk Road" anonymous marketplace (http://en.wikipedia.org/wiki/Silk_Road_%28marketplace%29) operated within the .onion pseudo-domain (see http://en.wikipedia.org/wiki/.onion).

Getting Involved

- Stay informed. I know that there's almost overwhelming amounts of press coverage relating to Snowden and the NSA programs, but that's because a lot is going on. **Do your best to try to keep up.**
- Second, recognize that this issue isn't just a technical one. **Regardless of how you feel, let your legislators know your opinion -- that's how a participatory democracy should work!**
- If you're a technical person and concerned about privacy and confidentiality, **begin using end-to-end encryption (such as PGP/GPG or S/MIME), anonymous remailers, and Tor**, and help teach others to do so, too.
- Work on improving campus web privacy and security, too.
- There are also lists discussing many of the issues from this section, including the IETF Pervasive Passive Monitoring ("Perpass") list, the IETF Privacy list, and Stanford's "liberationtech" mailing list, just to mention a few of many.

8. Internet Governance

"So Who's In Charge of the Internet, Anyhow?"

- It is probably fitting to conclude these remarks with a few comments on Internet governance. The preceding parts of this talk certainly highlight some of the challenges we currently face.
- When facing those challenges, it may be tempting to want some authority figure to be "in charge" of making it all work. But who?
  -- Our Internet Service Providers? (Even the biggest of those are "just businesses," after all...)
  -- The U.S. Government? (If so, what about Europe, Asia, South America, Africa, and Australasia? Should one country, even a country as cool as ours, shoulder all that responsibility?)
  -- The United Nation's International Telecommunications Union? Sadly, that doesn't appear to be working so well anymore. Consider the schism between the West and the rest of the world reflected in this voting map for a recent ITU treaty... That's not the answer, either.

Who Signed The ITU WCIT Treaty, Dec 2012?

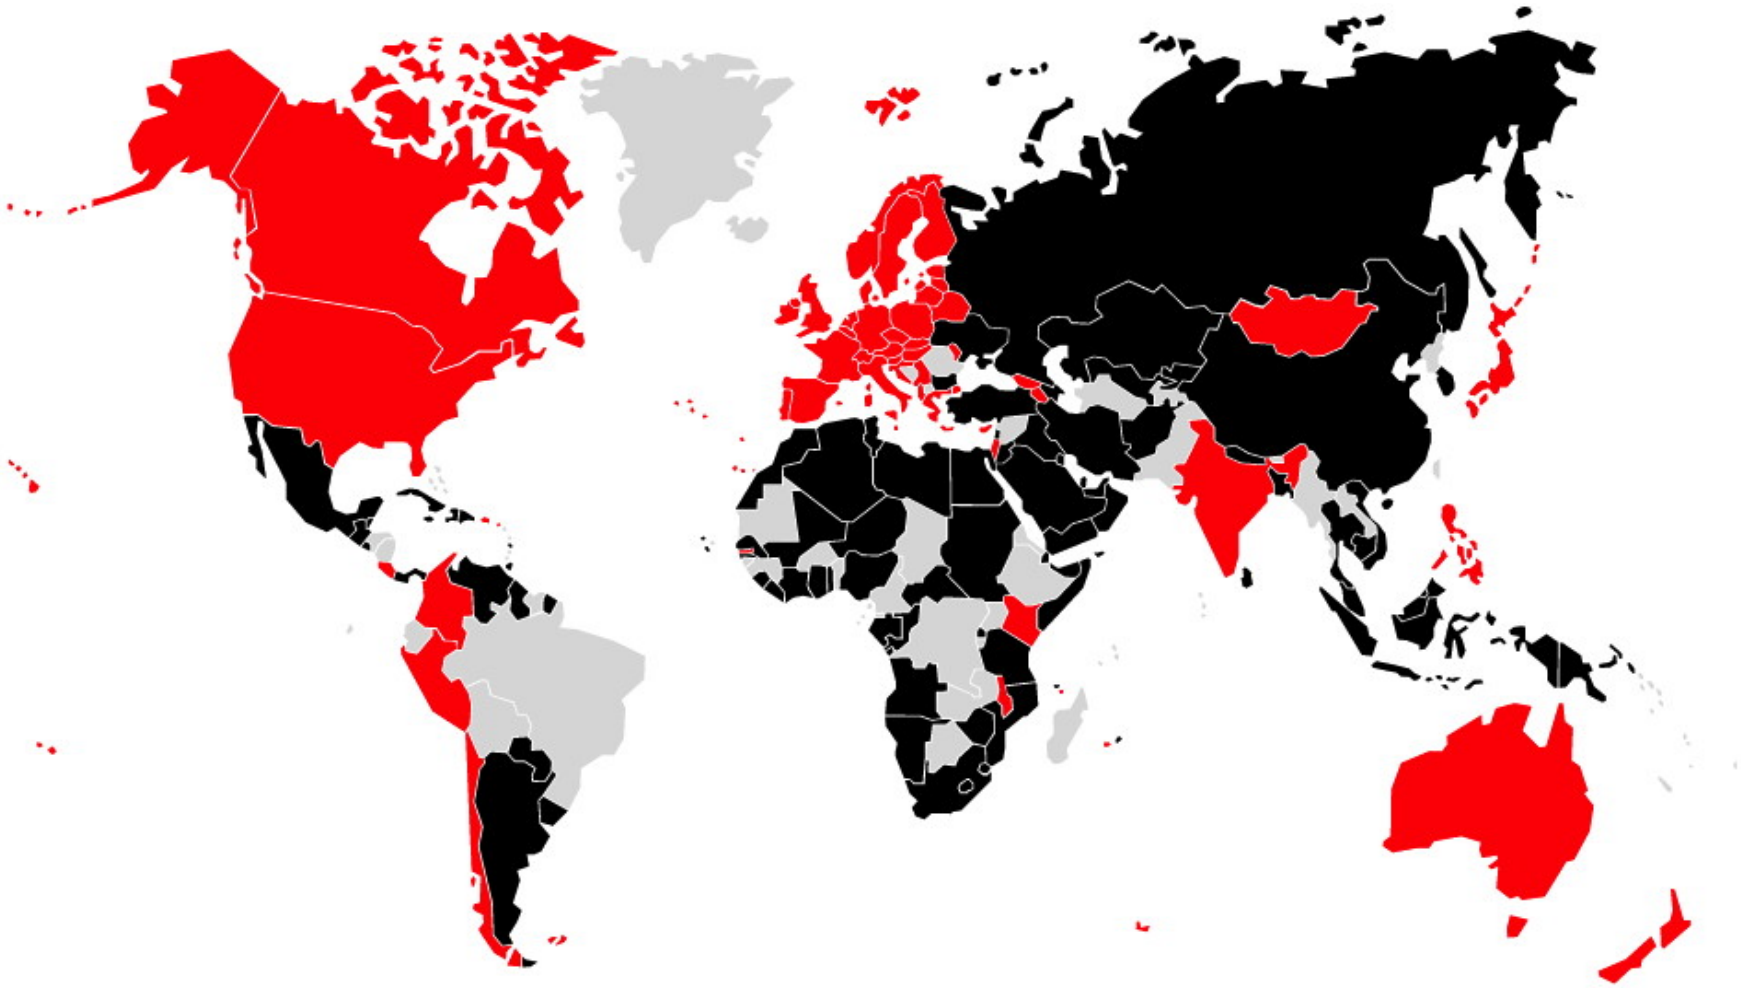

http://www.ipv.sx/wcit/

Ultimately, YOU Are In Charge of the Internet

- You, and all the other users of the Internet, are the ones who run it -- at least if you choose to get involved.
- There are so many opportunities for you to do so, whether you're an engineer or an operator or a coder or a policy person or someone focused on a particular country or region.
- **Please participate! Join working groups. Comment on draft documents. Contribute code, or help test code that others have written. Attend meetings, in person if possible, or online.**
- There are a lot of crazy things happening right now, but you can make a difference IF you get involved, and IF you participate. I hope you all will! Ultimately the Internet is an experiment, and one that's still a work in progress.
- The choices you and I make will help to determine if the Internet continues to be "crazy cool," or ends up being just plain nuts.

Thanks for the Chance to Talk Today

- You can download a copy of these slides in Powerpoint or PDF format from http://pages.uoregon.edu/joe/merit-networking/

- Are there any questions (if we have time)?