The Enduring Challenges of Traffic Analysis

FloCon 2015

Wednesday, Jan 14th, 2015, 11:30-12:30

Portland Hilton

Joe St Sauver, Ph.D. (stsauver@fsi.io)

Distributed System Scientist

Farsight Security, Inc.

https://www.stsauver.com/joe/

[My Odd Slide Style]

- If you're not familiar with my slide style, let me get that "out of the way" right up front: I write **detailed slides.** This style drives some people **crazy,** so let me explain why I do it.
- I've tried the more-typical 3-4 bullets/slide with ~15 slides for an hour long talk model, but I find myself getting **sidetracked, rambling/running over,** or I end up **missing/skipping stuff.**
- I also deal with **complex issues**, and I HATE to be misquoted.
- My slide style prevents a lot of those problems, and means that **you don't need to try to take notes.**
- That said, I'm **not going to read my slides word-for-word for you.** You don't need to try to do so, either, although they are a sort of "closed captioning" if you're deaf or hard-of-hearing.
- I also write detailed slides to help people looking at them after the fact, and for indexing by web search engines.

Disclaimer

- I also wanted to share an explicit disclaimer with you. Because I wear (and have worn) a lot of different hats, there's sometimes the potential for confusion with respect to whose position I'm articulating during a given talk.
- To keep this straightforward today: **my remarks represent solely my own opinion, and do not necessarily represent the opinion of any other party,** to include FloCon and its organizers, nor any other organization or entity.
- I may mention specific protocols or products or services by way of example; **mentioning specific options is not meant to discourage consideration of other alternatives** (I just don't have time to talk about every possibility given our limited time)

I. Introduction

To Any Visitors, Welcome to Oregon!

- I hope you'll have time to explore at least a little of Oregon while you're here. You won't be able to see "all" of it even if you had a lifetime to do so – at 98,378 miles^2, Oregon's the 9th biggest state in the U.S. (and bigger than the entire U.K.'s 93,800 sq miles).

- Just my home *county*, Lane County, Oregon (4,722 miles^2), is roughly the size of the *state* of Connecticut (5,543 miles^2).

- Given how easy it is to fall in love with Portland (and Oregon's other wonderful cities), visitors sometimes overlook the Oregon outdoors. Please don't miss the Willamette Valley, the Columbia Gorge, the Cascades, the high desert, and the Oregon Coast, too. No time to tour? Check out O.P.B.'s *Oregon Field Guide* (http://watch.opb.org/program/oregon-field-guide/)

A Little About My Background

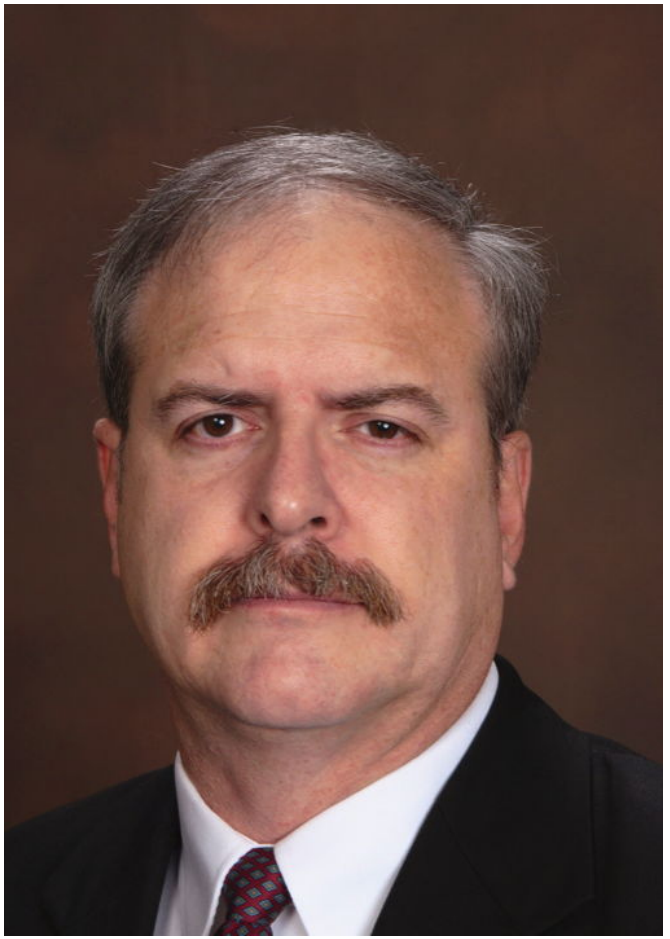

- I worked for ~28 years for the UOregon Computing Center, but in Nov. 2014, I moved to Farsight Security, Inc. (FSI).
- Prior to that time I served as Internet2's Nationwide Security Programs Manager (under contract though UO). I also ran the InCommon SSL & PKI Certificate Program, and InCommon's Multifactor Program. I've also been a member of the Research and Education Network ISAC (REN-ISAC) TAG.
- I'm one of six senior technical advisors for M3AAWG (the Messaging, Mobile and Malware Anti-Abuse Working Group), and I work with numerous other cyber security groups.

Netflow and Me

- I've had a bit of an "ongoing relationship" with Netflow:
- In June 2000, I conducted a legislatively-mandated Netflow-based audit of Oregon's public higher education network.
- Roughly five years later, in March 2005, I briefed MAAWG on "Spam Zombies and Inbound Flows to Compromised Customer Systems," describing flow-based approaches to identifying and combating bot spam (those approaches work great)
- At TNC2007 in Denmark, I delivered a CALEA talk, "Upcoming Requirements from the US Law Enforcement Community to Technically Facilitate Network Wiretaps"
- And then back in the US, I talked about the "Unidentified Half of Netflow" and how to use the Internet2 Netflow Data Archives.
- I also helped draft Internet2's IPv6 Netflow anonymization policy
- So Netflow and related topics are "old and valued friends" ☺

Picking A Topic For Today

- It's a great honor to have been asked to talk today. **Thank you!**
- **But what to talk about?** A keynote deserves a "big topic," and is not really the right place to do a deep dive into a highly technical area, especially not right before lunch, so I fought that urge. No equations and no snippets of code today, sorry.
- Having recently transitioned to Farsight Security, Inc., I would normally have talked about some of the work Farsight's doing, but my new boss, Paul Vixie, was your keynote speaker *last* year (I promise, Farsight's not trying to hijack your conference, it's just a coincidence – I'd agreed to talk here prior to moving to FSI).
- Since keynotes should ideally deliver content that's a bit provocative/challenging, I finally decided to focus on some of the work that **M3AAWG**'s been doing in the anti-pervasive monitoring area -- and some work that still need to be done.

What's M3AAWG?

- M3AAWG is the Messaging, Mobile and Malware Anti-Abuse Working Group. Its 33rd general meeting will be in San Francisco in February; 34th in Dublin in June; 35th in Atlanta in October.
- Chatham House Rules apply ("What happens at M3AAWG...")
- M3AAWG's original focus was on fighting spam, but its remit has grown over time to meet the needs of its members, most recently to include work against **pervasive network monitoring.**
- Recently published: https://www.m3aawg.org/sites/maawg/files/news/M3AAWG_TLS_Initial_Recommendations-2014-12.pdf
- List of member companies: https://www.m3aawg.org/about/roster
- Information about M3AAWG's leadership: https://www.m3aawg.org/media_center#leadership
- See also https://www.youtube.com/user/MAAWG/videos
- N.B.: I'm **NOT** speaking as a M3AAWG representative today.

TODAY'S "ASK"

- **It may sound odd, but I'm actually here to ask for your HELP**

- **Law-abiding users of the Internet need technical solutions that will let them effectively avoid pervasive metadata collection and traffic analysis, regardless of who may be targeting them. (Yes, I know this is hard)**

- At the same time, ISPs and criminal law enforcement agencies also need to be able to sustain robust traffic analytic techniques for closely tailored purposes:
  - -- these techniques are needed by providers for appropriate self-defense and for anti-abuse purposes, and
  - -- by LEOs for narrowly-targeted and court-approved lawful intercepts needed to combat online criminal activity.

Tensions Between Those Objectives

- Let me be very candid that I unquestionably get that there is definite tension between those two objectives:

  -- on the one hand, I want help making routine Internet traffic robust *against* traffic analysis for pervasive monitoring (hard, in and of itself);

  -- on the other hand, I also want the community to work on being better able to continue to perform, or even better perform, carefully targeted traffic analyses (also hard, in and of itself)

- Asking for BOTH of those things together? REALLY hard.

- This is, in many ways, directly parallel to what we see in the cryptographic world.

II. "Sneaking Up On Traffic Analysis" By Starting With The "Easy" Issue First: Countering Eavesdropping With Encryption

Deterring Eavesdropping Through Use of Encryption

- For a long time, most email traffic and most web traffic passing over the Internet was **unencrypted.** It was therefore vulnerable to **eavesdropping.**
- **Email end-to-end privacy tools** (such as S/MIME and GNU PrivacyGuard) have long been widely available, but generally have been **"too tricky"** for most "mere mortals" to routinely use.
- SSL/TLS is another cryptographic tool, but for a long time it was pretty badly technically flawed, and normally it was something that was only used to protect credit card numbers & login information and for a few other very limited use cases.
- ***Bottom line:*** **most Internet traffic content was broadly vulnerable to passive network monitoring.**
- Many of us *suspected* that monitoring of unencrypted Internet traffic was taking place, but few knew for sure until June 2013.

Users Suddenly KNEW That They Were Under Surveillance; Providers Took Steps to Harden Their Services

- Once users knew that the contents of their communications were being monitored, they wanted protection from **eavesdropping.**

- **Encryption** became an exceptionally "hot" topic and great progress was made in finding and fixing flaws, and in expanding cryptographic protections, particularly for email.

- A prime example of this can be seen in Google's "Gmail Email Transparency Report" shown on the next slide.

- **Virtually all outbound email from Gmail to top destinations worldwide is now encrypted in transit.**

https://www.google.com/transparencyreport/saferemail/

Who supports encryption in transit

Below is the percentage of email encrypted for the top domains in terms of volume of email to and from Gmail, in alphabetical order.

Select Region World

Top domains by region, inbound

Domain	%
From: amazon.{...} via amazonses.com	99.9%
From: amazonses.com	99.9%
From: ed10.net via ed10.com	0%
From: facebookmail.com via facebook.com	99.99%
From: groupon.{...}	99.99%
From: grouponmail.{...}	0%
From: linkedin.com	99.99%
From: sailthru.com	0%
From: twitter.com	99.99%
From: yahoo.{...}	99%

Top domains by region, outbound

Domain	%
To: aol.com	99.99%
To: comcast.net	100%
To: craigslist.org	99%
To: hotmail.{...}	100%
To: live.{...} via hotmail.{...}	100%
To: mail.ru	99.99%
To: msn.com via hotmail.{...}	100%
To: orange.fr	100%
To: rediffmail.com via akadns.net	100%
To: yahoo.{...} via yahoodns.net	100%

Saturday, January 3, 2015

All Those 100%'s and 99.99%'s? Those Numbers Represent A Bit of a Miracle...

- Few security technologies have *ever* successfully deployed at Internet scale.

- **PGP/GPG?** Great, but only used by a tiny subset of all users.
- **IPSec?** Never deployed (except for some *ad hoc* VPN usage)
- **DNSSEC?** Deployment of DNSSEC still trails
- **RPKI?** Another security technology that's had a slow start.

- But ***encryption of email in transit*? THAT's** an example of a security technology that **HAS** deployed at scale. We've gone from 30-40% opportunistic encryption of outbound email from Google a year ago to fully 80% in just a year. See the graph on the next slide.

% of Outbound Gmail Encrypted With STARTTLS

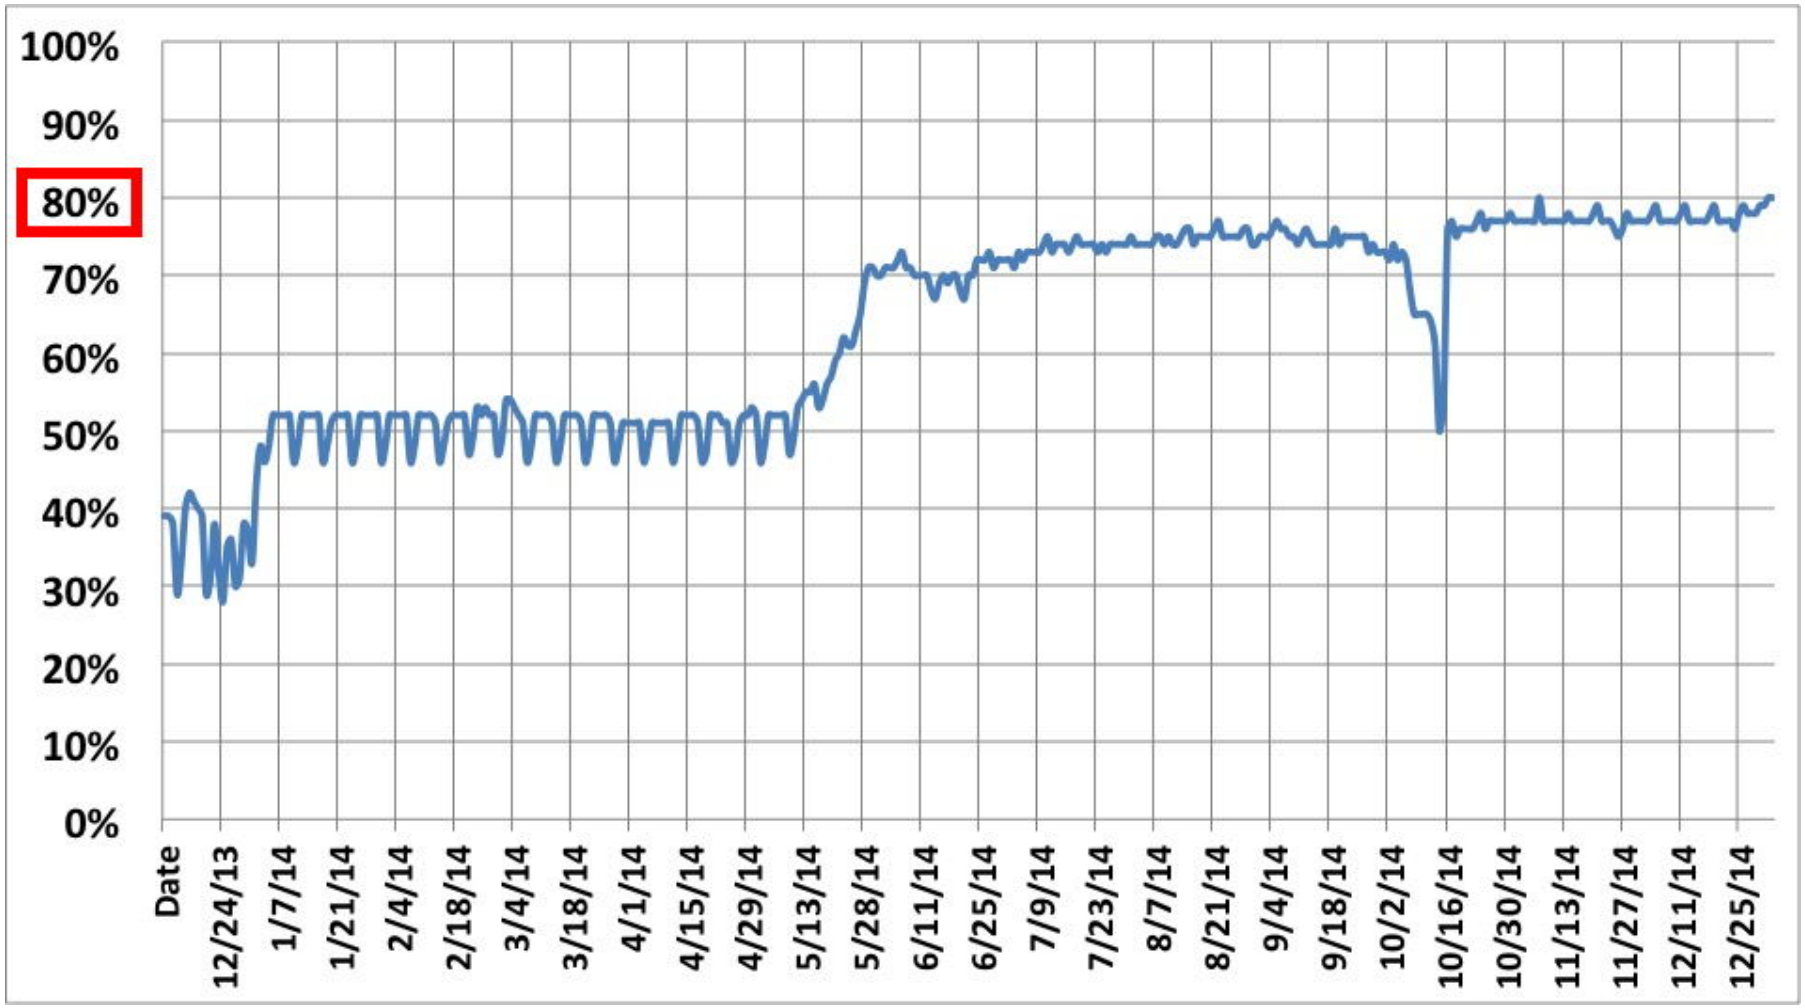

https://www.google.com/transparencyreport/saferemail/google-starttls-percentages.csv

This Does Not Mean That Gmail Is "Going Dark"

- "Going dark" is "short hand" for "law enforcement agencies will no longer be able to conduct court-ordered lawful interceptions." It is the basis for federal "push back" against encryption.

- You might think the preceding graph is an example of "going dark" what with 80% of outbound Gmail now encrypted in transit. It isn't. That 80% protection refers to email on the network *in transit*. Law enforcement is still free to obtain a court order for access to the email of a specific user on the ISP's *email servers*.

- So why bother encrypting in transit? Answer: It becomes far harder for foreign and domestic intelligence agencies, and any hacker/crackers that may be sitting on the wire, to potentially vacuum up EVERYONE's SMTP traffic indiscriminately.

Perfect Example of "Threading The Needle"?

- This is, perhaps, a perfect example of "threading the needle" or balancing apparently conflicting objectives:

  -- widespread use of encryption during transit deters pervasive surveillance efforts on the network

  -- legitimate carefully-targeted and court-authorized access has been preserved (at least as long as users don't choose to use end-to-end encryption). That is, law enforcement can still get access with appropriate paperwork for mail stored on email providers' mail servers, if they have probable cause.

- Oh, and average users don't need to become crypto experts to get reasonable protection.

More Cryptographic Work Remains To Be Done

- At least **20% of all email outbound remains unencrypted** in transit even from Gmail. We need to keep whittling away at that.
- The protection of email with STARTTLS is still imperfect. We need to keep working on helping sites move to **stronger crypto** deployments. We need to **find and patch flaws in OpenSSL, GNUTLS and other crypto implementations** (nice example? The "Frankencerts" paper from UTexas, see https://www.cs.utexas.edu/~shmat/shmat_oak14.pdf)
- We need **stronger keys** (AES 128→256, RSA 2048→4096 bit)
- We need to deploy **elliptic curve** cryptography using cryptographically **safe curves**! [see http://safecurves.cr.yp.to/]
- We need to work on deployment of **Layer 1/2/3 crypto**.
- And we need **end-to-end (not just hop-by-hop) crypto usage.**
- We should also look at **hardening ssh,** too.

Layer 1/2/3 Crypto

- Most security architectures endeavor to deploy **security in depth**: deploying overlapping layers of protection means that even if there's a flaw or compromise in one layer, redundant protection at other layers still delivers protection.
- Much of the focus to-date has been at the application layer, particularly on **SSL/TLS**.
- The time has come to remember that crypto can also be done at other layers, too. Some might assume that this means doing **IPsec in tunnel mode at layer three,** and you could certainly try that, but it would be painful if possible at all, particularly at **major-provider-and-carrier-relevant speeds (10Gbps or 100Gbps).**
- My suggestion is that providers should probably be looking at pervasively enabling encryption in **optical transport systems at layer one**, AND in Ethernet switching infrastructures at **layer two (IEEE 802.1AE, aka "MACsec" or "LinkSec")**, instead.

What About End-to-End Crypto Efforts?

- We also need to renew our effort to deploy **end-to-end crypto.**
- In the case of email, this might mean things like S/MIME and GNU PrivacyGuard, but I'd actually suggest that the community NOT focus on email for end-to-end crypto, at least not at first.
- **I think the "low hanging fruit" for end-to-end crypto is in the area of voice and instant messaging on smartphones/tablets.**
- I'm sure you've seen some of the voice and IM crypto options that have hit the market over the last few years (including solutions from **https://silentcircle.com/** and **https://whispersystems.org/**), but there are literally dozens of other options to also consider, including products from companies located outside the U.S. (such as **https://www.seecrypt.com/en/**).
- The biggest challenge we face in encouraging adoption of encrypted voice and IM is the lack of **interoperability:** most solutions are proprietary and can't talk to other vendors' products.

Hardening ssh

- A final example of an area where additional cryptographic work is required is ssh. People have spent lots of time and effort hardening SSL/TLS, but ssh has largely been overlooked.

- ssh is used at many sites for mission-critical purposes, including access to core routers and essential servers, but ssh is often not configured to be as operationally strong as it can be.

- Fortunately, people are beginning to work on improving this, too. See for example the recommendations in "Secure Secure Shell" https://stribika.github.io/2015/01/04/secure-secure-shell.html

III. Metadata

Bringing This Talk Back Around to Metadata and Traffic Analysis

- As discussed in the preceding section of this talk, **we've made *huge strides* when it comes to deploying encryption** to deter pervasive eavesdropping on traffic content at Internet scale.

- **Unfortunately, the Internet has made virtually NO PROGRESS when it comes to dealing with its metadata exposures, and when it comes to dealing with traffic analysis attacks.**

- **In fact, most people don't even understand what metadata is, or what traffic analysis is, or why they're major issues worthy of our attention.** As long as people don't understand metadata and traffic analysis, the challenges they pose will never get addressed.

Many of YOU DO Already "Get" Traffic Analysis

- I realize that I'm talking to an audience largely compromised of Netflow experts, with many of you in fact spending all day doing Netflow-based traffic analyses. That's great – this may be one of the few audiences that explicitly gets the **power of traffic analysis.** On the other hand, I also know that there are some of you who *aren't* familiar with traffic analysis:

  -- you may be just getting started with Netflow,
  -- you may use Netflow for other purposes (such as billing), or
  -- maybe you're looking at this talk after the fact and you're not really a Netflow person at all
  -- or you might know Netflow, but may not be familiar with traffic analysis of non-Netflow metadata

- For these and other reasons, let's talk a little about metadata.

Metadata: Data About Data

- Trained photographers know that most every photo taken by a smart phone (or high end digital camera) has metadata by default:
  -- *what camera/smartphone was used*? (brand, model number, etc)
  -- *when was the photo taken*? (date and time)
  -- *where was the photo taken*? (cameras include GPS receivers to help supply this information)
  -- *how was the camera configured*? (shutter speed, aperture, etc.)

  This data can be routinely helpful to photographers.

- It can *also* be key to criminal investigations. A recent example: http://www.smh.com.au/technology/technology-news/hacking-cases-body-of-evidence-20120411-1wsbh.html (a taunting photo supplied by a computer intruder contained GPS metadata, metadata which allowed identification and arrest of that intruder)

Metadata Isn't Just a "Photographic Thing"

- Metadata is something that exists for most digital objects, including:

  -- Network traffic flows
  -- Email messages
  -- Telephone calls
  -- etc.

What Metadata Is Normally Available For Network Flow-based Traffic Analyses?

- Most of you know this part...
- "Source" and "destination" IP addresses and corresponding port numbers
- Traffic start and stop times and traffic volume in octets
- Other technical traffic characteristics not related to message content (for example, packet type, TCP flags, ASNs, etc.)
- We may have all this for all traffic flows, or for some sampled subset of flows (perhaps 1 in 100 or 1 in 1000 flows, etc.)

Do Network Traffic Sources/Destinations Map Closely To Individuals? *Sometimes...*

- For example, a static IP address may be persistently used by just one person. The identity of that user can often be determined by issuing paperwork to the party responsible for that IP range, or other techniques.
- Other static IPs may be for servers shared by many users, as is the case on web servers.
- Addresses may be multiplexed across multiple users, either:
  -- shared at the same time (e.g., NAT/PAT)
  -- or serially shared (DHCP-assigned dynamic addresses).
  Those cases are harder to directly attribute -- typically only the operator of the firewall doing the NAT/PAT translation, or the operator of the relevant DHCP server, can translate IP address+time stamp+port info to a definitive customer identity.
- *Hold this thought, you'll see this information again later.*

What Metadata is Available for Email?

- Metadata for email normally includes most (but **not** all) of the information in the email message "headers"
- Stuff that IS normally considered metadata includes the contents of the "From:", "To:", "CC:", and "Date:", headers, and the multiple "Received:" headers showing the message's routing info, among other headers.
- It normally does **NOT** include the contents of the email message's **"Subject: header"** (even though that too is a header), because it contains "information concerning the substance, purport, or meaning of that communication" (see 18 U.S.C. 2510(8))
- Excellent discussion of the Subject: header in "The Content/Envelope Distinction in Internet Law," http://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2478285_code705039.pdf?abstractid=1123304&mirid=1

Telephony Metadata

- Metadata isn't limited to just Internet traffic or email. Metadata also exists in a telephony environment, too. In the telephony space, metadata is defined in the United States to be:

  "comprehensive communications routing information, including but not limited to session identifying information (e.g., originating and terminating telephone number, International Mobile Subscriber Identity (IMSI) number, International Mobile station Equipment Identity (IMEI) number, etc.), trunk identifier, telephone calling card numbers, and time and duration of call. Telephony metadata does not include the substantive content of any communication, as defined by 18 U.S.C. § 2510(8), or the name, address, or financial information of a subscriber or customer." (See for example https://www.eff.org/files/2013/11/06/mooredeclexh.pdf)

Telephony Metadata Has Been, And *Is Being*, Collected

- "Given that legislation [transitioning the Section 215 metadata collection program] has not yet been enacted, and given the importance of maintaining the capabilities of the telephony metadata program, the government has sought a 90-day reauthorization of the existing program, as modified by the changes the President directed in January. [...] the government filed an application with the FISC to reauthorize the existing program for 90 days, and that the **FISC issued an order approving the government's application. The order issued on December 4, 2014, expires on February 27, 2015."**

  http://www.dni.gov/index.php/newsroom/press-releases/198-press-releases-2014/1147-joint-statement-from-the-odni-and-the-u-s-doj-on-the-declassification-of-renewal-of-collection-under-section-501-of-the-fisa-12-14?tmpl=component&format=pdf

Why Not Collect Full Traffic Contents?

- There can be many reasons why "traffic content" isn't available to an analyst.
- Sometimes traffic may be protected with strong encryption. As a result, there may be no technical ability to access things like the body of email messages, or actual telephone conversations.
- Other times, there may be policy/administrative constraints that preclude access to traffic contents. For example, a court may not have authorized a Title III full contents lawful intercept order.
- Or perhaps storage is the binding constraint. Given fast links and a finite storage archive, there may be a trade off between:
  -- a relatively brief archive of "full content" traffic vs.
  -- a far longer window of summarized flow-level traffic.
- Sometimes "less" [detail] really does mean "more" [longer data retention window]

Domestic Bulk Metadata Collection Has Been Going On For A Long Time – And Was Revealed – Long BEFORE Dislosures Were Made In June 2013

- While everyone may assume that domestic bulk metadata first became a public issue with the disclosure in June 2013, that's a misperception. Credit for raising the metadata issue should actually go to *USA Today*. On May 10th, 2006 it published:

  "NSA has massive database of Americans' phone calls,"
  http://usatoday30.usatoday.com/news/washington/2006-05-10-nsa_x.htm

- That was **SEVEN YEARS** before the June 2013 revelations.

- That was **FIVE YEARS** after concerns about Internet communication channels increased, immediately after 9/11.

"Pen Register" & "Trap and Trace" Orders

- Traffic analytic approaches are, in fact, **strongly associated with telephony.** Law enforcement officers have used telephony "pen registers" and "trap and trace" techniques for a long time as part of their criminal investigations.
- A "pen register" records the *outgoing* calls made from a phone (in the old days of rotary pulse dialing, this meant literally tracing out the pulses made by the phone dial as it clicked around).
- "Trap and trace devices," on the other hand, focus on capturing the origin of *incoming* calls received by a phone.
- Normally BOTH pen register AND trap and trace data is collected, not just one OR the other.
- Even as investigations have moved away from telephony and toward Internet TCP/IP data, those old (now unquestionably antiquated) terms have "stuck" (even if they're a mouthful).

The Other Side of the Coin: Title III ("Full Content") Intercepts

- Pen registers and trap and trace devices do NOT provide access to message contents (normally you can't even use a pen register/trap and trace order get access to URLs, see http://www.justice.gov/usao/eousa/foia_reading_room/usam/title9/7mcrm.htm#9-7.500)
- **Title III intercepts,** on the other hand, provide the **"whole communication."** For example, in a telephony context, you'd get to hear the whole conversation. In an Internet context, you'd get the contents of an email message, not just header information.
- Because of the invasiveness of a full content intercept, requests for full content intercepts were historically given strict scrutiny, and were hard to obtain. Onesie-twosie pen registers/trap and trace orders, however, were relatively easily obtained, in part because there was little judicial appreciation for their true power.

IV. Traffic Analysis

"You can observe a lot by just watching."
Yogi Bera

Metadata Drives Traffic Analysis

- So far we've been talking about metadata. That's the "what."
- Now let's talk about the "how," aka traffic analysis.

Traffic Sources and Destinations Can Sometimes Be More Than Enough

- For example, if you observe an employee who works at a sensitive defense industrial site attempting to surreptitiously communicate with a representative of a foreign intelligence service, the exact details of what's being said are (to a first approximation) irrelevant.

- The simple fact that any such conversation is being held or attempted should be more than enough to send up a red flag (unless undertaking that communication was directed and approved at senior levels, etc.)

Ru-Roh, Scooby...

US Government Employee In Sensitive Role

Representative of A Foreign Intelligence Service

Changes In Traffic Patterns Can Also Be Key

- Assume we're monitoring traffic levels between a foreign military headquarters and its bases. Because that traffic's encrypted, we don't know what's being said, but over time, we've come to know what normal traffic looks like, e.g., we've got a "traffic baseline."
- Suddenly, out of the blue, traffic from HQ to those bases begins to run 10X or even a 100X normal levels. Something is happening. (This is an example of the "elevated level of chatter" you'll sometimes hear mentioned by the news media.)
- Alerted to this reality, monitoring authorities might decide to task other assets (such as satellite imagery or human intelligence sources) in an effort to figure out exactly what's going on. For example, is the foreign power preparing to launch an attack?
- A sudden drop in traffic can be equally concerning: is this "radio silence" prior to an attack? Was our monitoring detected and somehow circumvented?

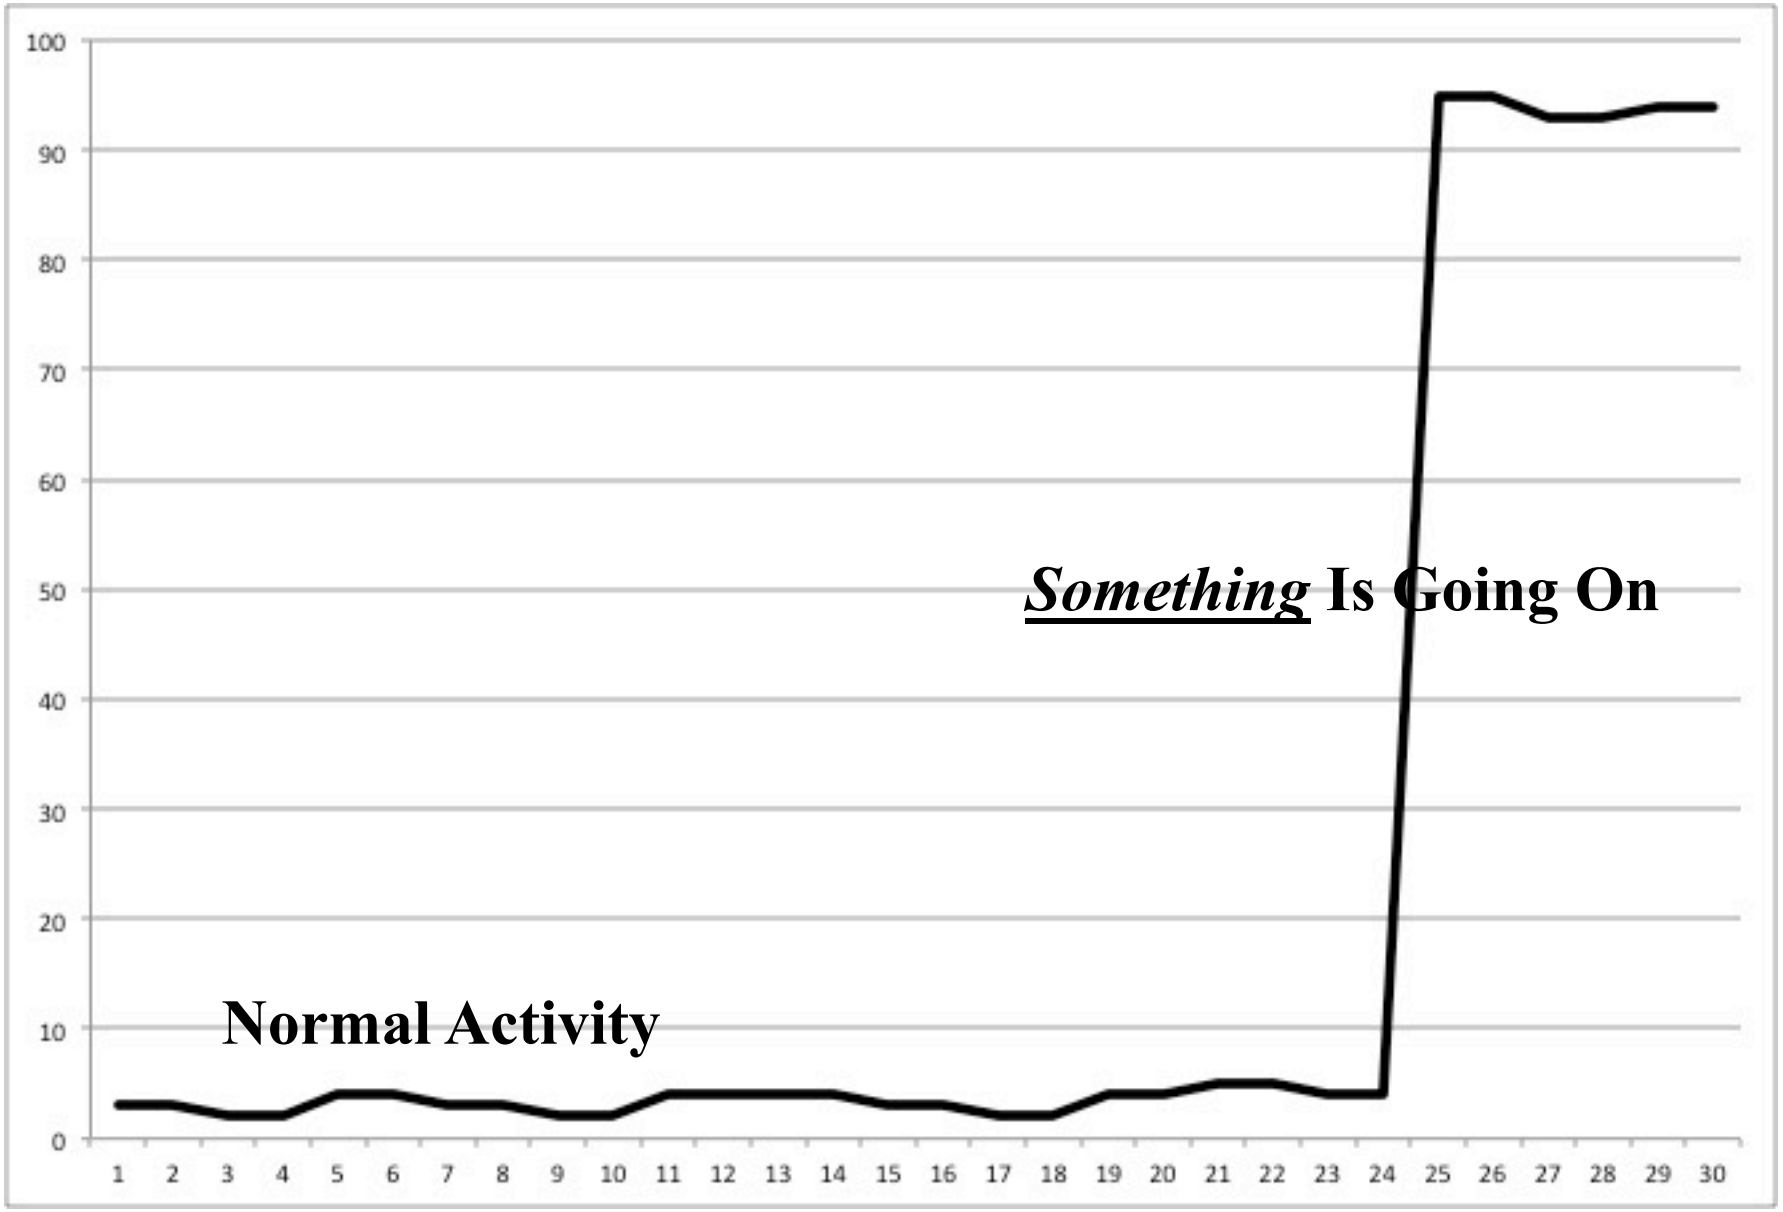
100
90
80
70
60
50
40
30
20
10
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30
Something Is Going On
Normal Activity

Traffic Sequencing Can Also Be Quite Revealing

- As another example, assume encrypted communications between parties A, B, and C are being monitored.
- Two messages (each roughly of the same size, and in close proximity time-wise) are seen. One is sent from A to B, and another sent from A to C.
- Shortly after those two messages are sent, B is observed sending a message of roughly the same size to ten additional recipients (B1-B10). C does likewise for another dozen recipients (C1-C12).
- From those observations, we might hypothesize a hierarchical communication or command-and-control structure: A directs B and C. B commands B1-B10. C commands C1-C12.
- Being able to infer these sort of relationships can crucial if you know that some of B1-B10 and C1-C12 are known drug dealers, and B and C are suspected drug distributors. Is A the "king pin?"

Communication/Control Structure Inference Example

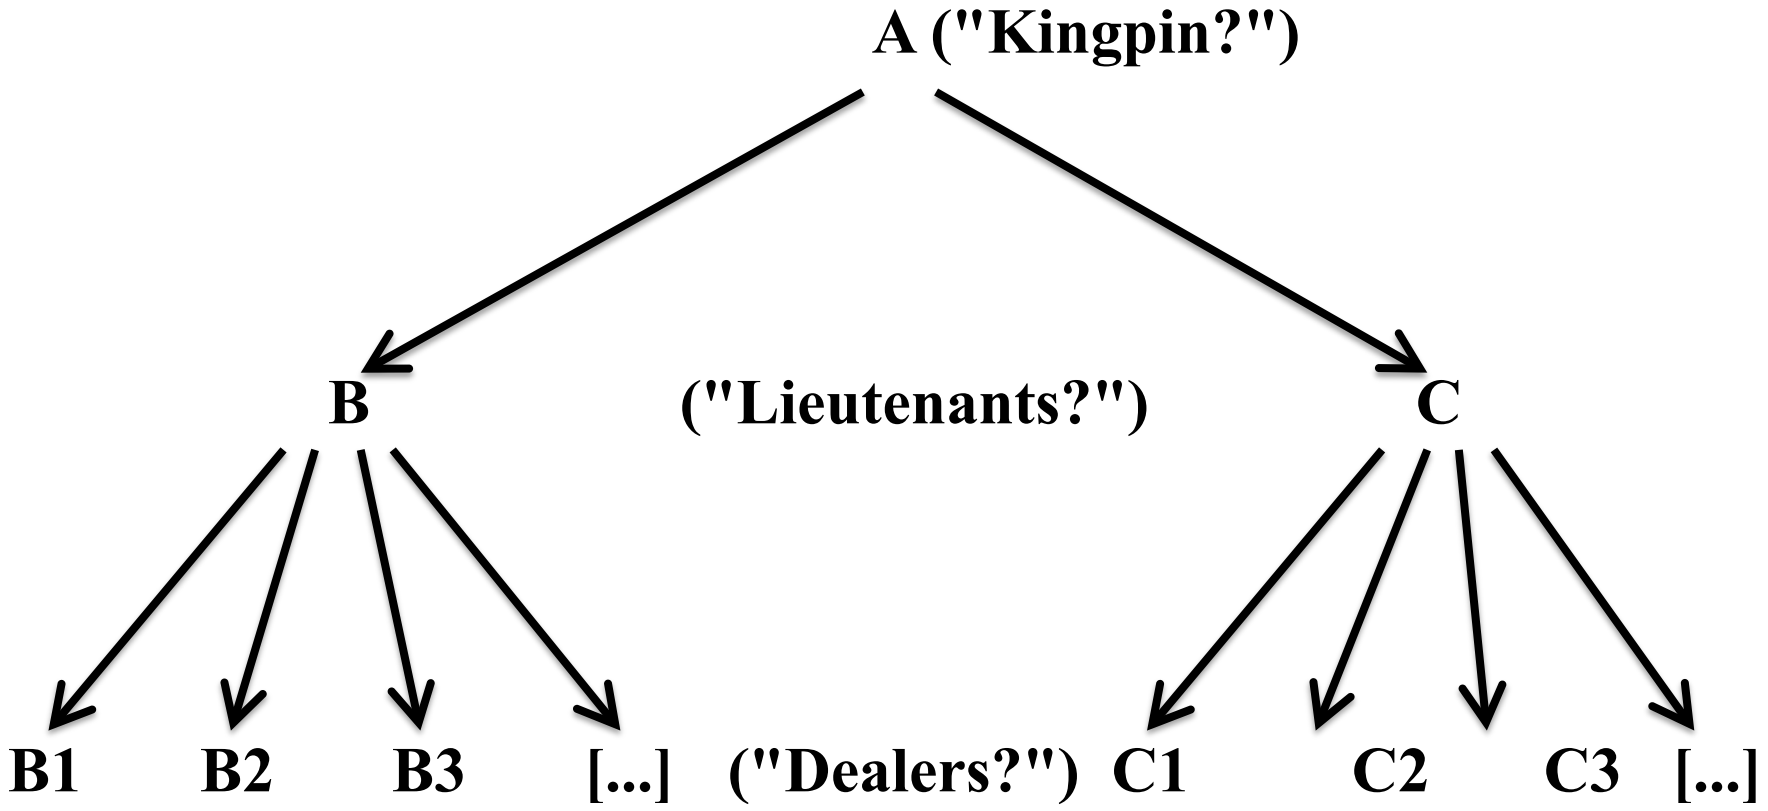

Geolocation of Specific Traffic Sources

- As a final example, assume that a kidnapper contacts the parents of a kidnapped child from his cell phone.
- The kidnapper may not know that the location of cell phones can be determined via GPS (if enabled on the cell phone), or at least by "cell phone tower triangulation" (the azimuth from two or more cell towers to the phone)
- GPS location information, or cell phone tower triangulation data, would allow law enforcement officers to see where the kidnapper is located or may be traveling, thereby perhaps also finding where a kidnapped child is being held.
- Geolocation can also be important for "E911" emergency calls (including finding a person who has called 911 after suffering a heart attack, stroke or suffered some other life-threatening emergency)

Geolocation Triangulation: Two or More Bearings: The Cell Phone's Located At the Point of Intersection

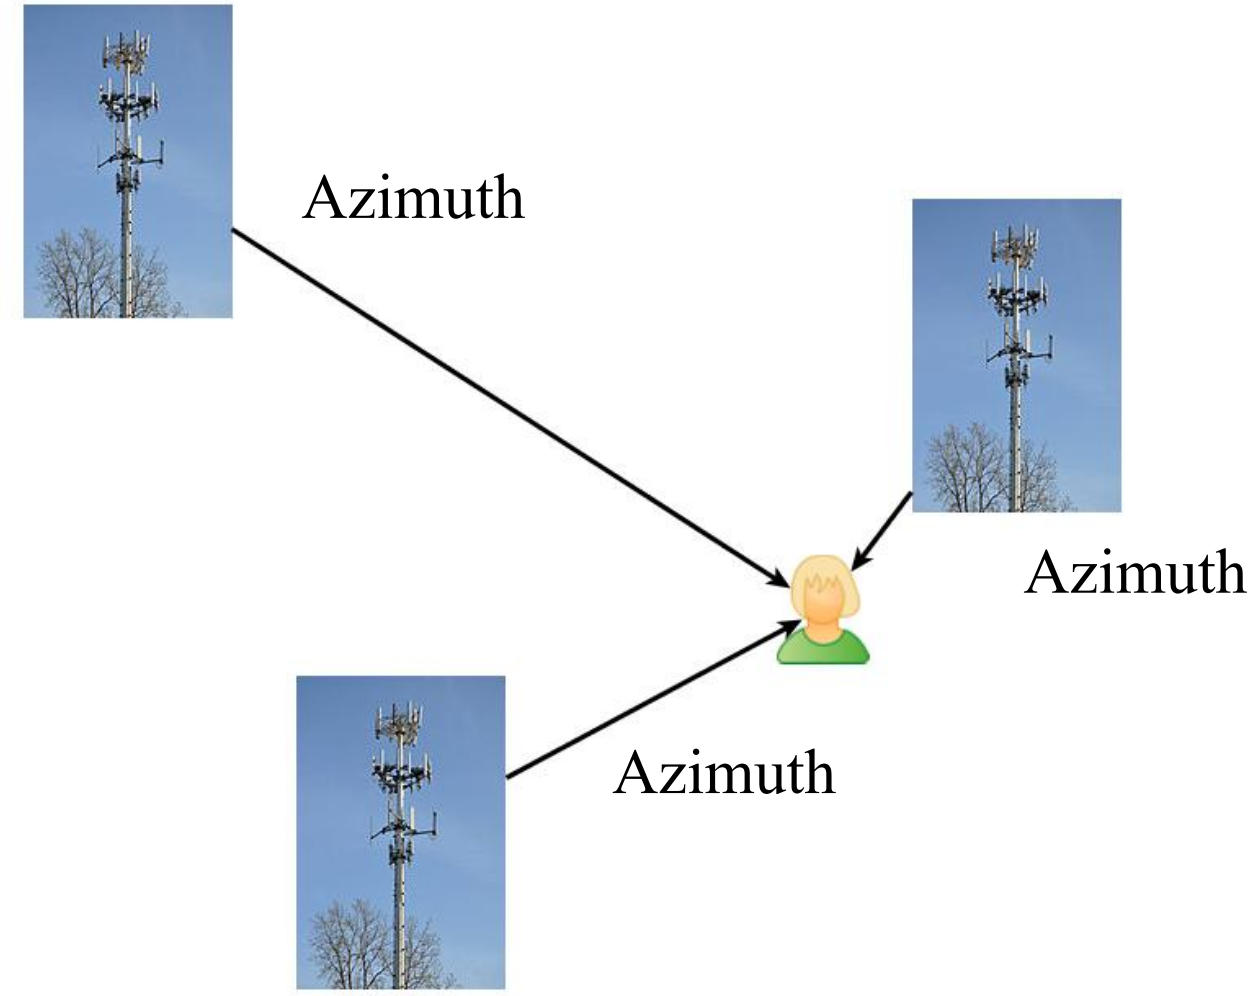

Raw cell tower image credit: Joe Ravi [CC BY-SA 3.0] via Wikimedia Commons
http://commons.wikimedia.org/wiki/File:Cell_Phone_Tower.jpg

The Preceding Gov't Uses Cases For Traffic Analysis

- The preceding examples are examples of commonly accepted governmental uses for traffic analytic approaches:

  -- Counterintelligence surveillance of a sensitive government employee (where there's never any expectation of privacy)

  -- Monitoring of a foreign military power for the purpose of detecting an attack and ensuring an appropriate national defensive response

  -- Criminal law enforcement investigation of an illegal drug ring

  -- Recovery of a kidnapping victim/arrest of a kidnapper

- These are some examples, but you could easily name others.

ISPs May ALSO Routinely Do Traffic Analysis

The following are some (but not necessarily all) examples of appropriate ISP traffic analysis:

- For **routine network operations** (as necessary to run the network, detect faults and equipment failures, plan for required expansion, negotiate peering, perform usage-based billing, etc.)

- To **protect service provider assets and services** (e.g., for intrusion detection, DDoS mitigation, fraud prevention, etc.)

- Other uses as **contractually agreed to** between provider and customer. Examples: delivery of value-added security monitoring services as an extra-cost service at the customer's request, or for research (after appropriate anonymization), etc.

Criminal Surveillance

In broad terms (and not necessarily covering every corner case)...

- Limited to specific statutorily-designated serious crimes (e.g., kidnapping, racketeering, murder-for-hire, etc.)
- Based on probable cause
- Narrowly targeted and of limited duration
- All collections are carefully minimized
- Last resort (all less-intrusive alternatives have been exhausted)
- Reviewed and approved at a senior level within the law enforcement agency requesting the lawful intercept
- Authorized and carefully supervised by an appropriate court
- Under seal only as long as necessary (NOT effectively "forever")
- Usage annually reported (see http://www.uscourts.gov/Statistics/WiretapReports/wiretap-report-2013.aspx)

Contrast With Pervasive Monitoring

- The collection is **DOMESTIC** (e.g., it includes U.S. Persons)
- The collection is **UNTARGETED** (data is collected about effectively EVERYONE in the United States), with no individualized suspicion required
- The collection is **ONGOING**
- There is no requirement for **PROBABLE CAUSE**
- No data **MINIMIZATION** takes place.
- **DATA RETENTION** is perpetual.

What Is A "US Person?"

- That definition is more inclusive than you might think.

- “United States person” means a citizen of the United States, an alien lawfully admitted for permanent residence (as defined in section 1101 (a)(20) of title 8), an unincorporated association a substantial number of members of which are citizens of the United States or aliens lawfully admitted for permanent residence, or a corporation which is incorporated in the United States, but does not include a corporation or an association which is a foreign power, as defined in subsection (a)(1), (2), or (3) of this section."

  50 U.S.C. 1801(i)

Why "U.S. Persons Matter" #1: FISA

- 50 U.S. Code § 1842 - **Pen registers and trap and trace devices for foreign intelligence and international terrorism investigations**

"Notwithstanding any other provision of law, the Attorney General or a designated attorney for the Government may make an application for an order or an extension of an order authorizing or approving the installation and use of a **pen register or trap and trace device** for any investigation to obtain foreign intelligence information **not concerning a United States person** or to protect against international terrorism or clandestine intelligence activities, provided that **such investigation of a United States person** is not conducted solely upon the basis of activities protected by the first amendment to the Constitution which is being conducted by the Federal Bureau of Investigation under such guidelines as the Attorney General approves pursuant to Executive Order No. 12333, or a successor order." [emphasis added]

Why "U.S. Persons Matter" #2: The USA PATRIOT Act "Section 215" "Business Records" Provisions

- **50 U.S. Code § 1861 - Access to certain business records for foreign intelligence and international terrorism investigations**

- (a)(1) Subject to paragraph (3), the Director of the Federal Bureau of Investigation or a designee of the Director (whose rank shall be no lower than Assistant Special Agent in Charge) may make an application for an order requiring the production of any tangible things (including books, records, papers, documents, and other items) for an investigation to obtain foreign intelligence information **not concerning a United States person** or to protect against international terrorism or clandestine intelligence activities, **provided that such investigation of a United States person is not conducted solely upon the basis of activities protected by the first amendment to the Constitution.**

V. Is Metadata Collection and Traffic Analysis Really That "Big Of A Deal?"

Real Harms From Routine Pervasive Monitoring

- The biggest problem with pervasive monitoring is the **chilling effect of that activity, once known or reasonably suspected.**
- An activist may worry that attending a peaceful political meeting or exercising her right to engage in non-violent protest will render her subject to invasive monitoring
- Citizens may become reluctant to openly support political candidates or advocate for political causes in electoral contests
- Attorneys may find it difficult or impossible to candidly advise their clients
- A cancer patient, concerned about her privacy, may be reluctant to seek relevant information about her condition online, or even to discuss her treatments options over the phone with her doctor

What Does The IETF Say? They're Blunt. They Say Pervasive Monitoring Is An "Attack"

"Pervasive Monitoring Is a Widespread Attack on Privacy

"Pervasive Monitoring (PM) is widespread (and often covert) surveillance through intrusive gathering of protocol artefacts, including application content, or protocol metadata such as headers. Active or passive wiretaps and traffic analysis, (e.g., correlation, timing or measuring packet sizes), or subverting the cryptographic keys used to secure protocols can also be used as part of pervasive monitoring. PM is distinguished by being indiscriminate and very large scale, rather than by introducing new types of technical compromise.

"The IETF community's technical assessment is that PM is an attack on the privacy of Internet users and organisations. The IETF community has expressed strong agreement that PM is an attack that needs to be mitigated where possible, via the design of protocols that make PM significantly more expensive or infeasible."

http://tools.ietf.org/html/rfc7258

"But This Is Being Done To Combat Terrorism!"

- I understand that. What's being done is being done with the best of intentions and in an effort to keep us all safe. I *totally* get that.
- I hate terrorists as much as anyone. They **must** be found and they **must** be held accountable.
- But we must remember the Constitution, and our right to privacy, notwithstanding our efforts to deal with potential terrorist threats.
- Most terrorists (except those who may have actual access to weapons of mass destruction), are limited in the damage they can directly cause. Most terrorists need to rely on a sort of "insurgent jujitsu" -- counting on the over-reaction of others -- in order to make a ***real*** impact.
- If we let terrorists goad us into ignoring the Constitution, the terrorists will *truly* have succeeded, accomplishing far more than they could ever accomplish directly.

Effectiveness of Bulk Metadata Against Terrorism?

- Any steps we take in the war against terror must also be **effective.** It's not at all clear that the benefits from the bulk domestic metadata collection program justify its existence:

> An analysis of 225 terrorism cases inside the United States since the Sept. 11, 2001, attacks has concluded that the bulk collection of phone records by the National Security Agency **"has had no discernible impact on preventing acts of terrorism."**
>
> "NSA phone record collection does little to prevent terrorist attacks, group says," *Washington Post,* Jan 12, 2014.

To see the actual report, go to
http://securitydata.newamerica.net/nsa/analysis

Alternatives To Government Bulk Collection of Domestic Phone Metadata Exist

- Phone metadata can stay with the carrier unless/until subpoenaed, and doesn't have to automatically go to the government
- Government subpoenas are required to get specific items
- Such program appears to meet law enforcement's needs while avoiding bulk domestic metadata collection by the government itself
- Data can be fresh, and responses timely
- This is a viable alternative that deserves serious consideration.

The Privacy and Civil Liberties Oversight Board

- "The PCLOB is an independent agency within the executive branch established by the Implementing Recommendations of the 9/11 Commission Act of 2007."
- "The bipartisan, five-member Board is appointed by the President and confirmed by the Senate."
- **"The PCLOB's mission is to ensure that the federal government's efforts to prevent terrorism are balanced with the need to protect privacy and civil liberties."**
- You can read the biographies of the PCLOB membership at http://www.pclob.gov/about-us/leadership.html
- The PCLOB's *Report on the Telephone Records Program Conducted under Section 215 of the USA PATRIOT Act and on the Operations of the Foreign Intelligence Surveillance Court*, from January 23, 2014 is available at http://www.pclob.gov/library/215-Report_on_the_Telephone_Records_Program.pdf

The PCLOB Said...

- **"Recommendation 1: The government should end its Section 215 bulk telephone records program. The Section 215 bulk telephone records program lacks a viable legal foundation under Section 215, implicates constitutional concerns under the First and Fourth Amendments, raises serious threats to privacy and civil liberties as a policy matter, and has shown only limited value. As a result, the Board recommends that the government end the program.**
  Without the current Section 215 program, the government would still be able to seek telephone calling records directly from communications providers through other existing legal authorities. The Board does not recommend that the government impose data retention requirements on providers in order to facilitate any system of seeking records directly from private databases. Once the Section 215 bulk collection program has ended, the government should purge the database of telephone records that have been collected and stored during the program's operation[...] subject to limits on purging data that may arise under federal law or as a result of any pending litigation." [emphasis added]

Oregon's Own Senator Wyden (Who Serves On The Senate Select Committee on Intelligence) Also Believes Bulk Metadata Collection Should End

We believe the way to restore Americans' constitutional rights and their trust in our intelligence community is to immediately end the practice of vacuuming up the phone records of huge numbers of innocent Americans every day and permit the government to obtain only the phone records of people actually connected to terrorism or other nefarious activity. We support your March 27, 2014, proposal to achieve these goals, but we also view ending bulk collection as an imperative that cannot wait. We urge you to implement your proposal with reasonable haste to protect both our national security and the personal rights and liberties of U.S. citizens.

Sincerely,

Mark Udall

Ron Wyden

Martin Heinrich

http://www.scribd.com/doc/230579555/Wyden-Udall-Heinrich-Urge-President-to-End-Bulk-Collection-While-Congress-Works-to-Pass-Real-Surveillance-Reform

Even The President Has Agreed That Domestic Bulk Data Collection by the NSA Must End

Obama to Call for End to N.S.A.'s Bulk Data Collection

By CHARLIE SAVAGE MARCH 24, 2014

Email

Save

WASHINGTON — The Obama administration is preparing to unveil a legislative proposal for a far-reaching overhaul of the National Security Agency's once-secret bulk phone records program in a way that — if approved by Congress — would end the aspect that has most alarmed privacy advocates since its existence was leaked last year, according to senior administration officials.

Under the proposal, they said, the N.S.A. would end its systematic collection of data about Americans' calling habits. The bulk records would stay in the hands of phone companies, which would not be required to retain the data for any longer than they normally would. And the N.S.A. could obtain specific records only with permission from a judge, using a new kind of court order.

The headquarters of the National Security Administration in Fort Meade, Md.
Jim Lo Scalzo/European Pressphoto Agency

http://www.nytimes.com/2014/03/25/us/obama-to-seek-nsa-curb-on-call-data.html

And Yet...

www.nytimes.com/2014/11/19/us/nsa-phone-records.html

Search

POLITICS | Bill to Restrict N.S.A. Data Collection Blocked in Vote by Senate Republicans

Bill to Restrict N.S.A. Data Collection Blocked in Vote by Senate Republicans

By CHARLIE SAVAGE and JEREMY W. PETERS NOV. 18, 2014

Email

Save

WASHINGTON — Senate Republicans on Tuesday blocked a sweeping overhaul of the once-secret National Security Agency program that collects records of Americans' phone calls in bulk.

Democrats and a handful of Republicans who supported the measure failed to secure the 60 votes they needed to take up the legislation. The vote was 58 to 42 for consideration.

Senator Patrick J. Leahy, the Vermont Democrat who drafted the bill, blamed what he said was fear-mongering by the bill's opponents for its defeat. "Fomenting fear stifles serious debate and constructive solutions," he said. "This nation deserves more than that."

Senator Mitch McConnell of Kentucky before the vote. Mr. McConnell led opposition to the bill. J. Scott Applewhite/Associated Press

Be Careful When Interpreting That Vote...

www.oregonlive.com/mapes/index.ssf/2014/05/blumenauer_bonamici_defazio_ba.html

Oregon Democratic Reps. Earl Blumenauer, Suzanne Bonamici and Peter DeFazio all co-sponsored a bill aimed at curtailing government surveillance activities.

But each of them voted against it this week after the **USA Freedom Act was revamped in last-minute negotiations** with the Obama administration. "Original bill ... was good but final bill was watered down and & filled w/ loopholes," DeFazio tweeted after the vote Thursday.

The measure still **handily passed the House on a 303-121 vote**, with most of the yes votes coming from Republicans, including Rep. Greg Walden of Oregon. Rep. Kurt Schrader, D-Ore., another co-sponsor of the original bill, also supported it on the House floor.

The measure, which now goes to the Senate, would limit the National Security Agency's bulk collection of telephone records, a practice that became controversial following last year's massive leaks from NSA contractor Edward Snowden.

Sen. Ron Wyden, D-Ore., one of the most prominent critics of government spying, **charged Friday that the bill had been so weakened** that it "fails to protect Americans from suspicionless mass surveillance."

What About Judicial Remedies?

- "[...] **I will grant Larry Klayman's and Charles Strange's requests for an injunction** and enter an order that (1) bars the Government from collecting, as part of the NSA's Bulk Telephony Metadata Program, any telephony metadata associated with their personal Verizon accounts and (2) requires the Government to destroy any such metadata in its possession that was collected through the bulk collection program. **However, in light of the significant national security interests at stake in this case and the novelty of the constitutional issues, I will stay my order pending appeal.**"

  Klayman v. Obama, Civil Action No 13-0851 (RJL), U.S. District Court for the District of Columbia, https://www.documentcloud.org/documents/901810-klaymanvobama215.html at page 67, dated Dec 16th, 2013.

That Judicial Process Continues ...

- Just a few months ago, in November 2014, the Federal appeals court heard arguments in *Klayman v. Obama*. (see for example https://www.aclu.org/blog/national-security-technology-and-liberty/yep-uncle-sam-still-wants-log-your-calls).

- A decision is now pending; a further appeal to the Supreme Court from whichever party does not prevail is likely.

- **Bottom line: no conclusive judicial resolution should be expected for many months -- if not years.**

- Other pending and past surveillance cases can be seen at: http://projects.propublica.org/graphics/surveillance-suits

The State Secrets Privilege?

- WHY am I **NOT** sanguine about this case proceeding through to normal adjudication? Answer: the **state secret privilege.** (see http://en.wikipedia.org/wiki/State_secrets_privilege)

- If it appears that the government is going to end up losing in *Klayman v. Obama,* the government may assert the state secrets privilege, and at that point the case will terminate (the judiciary routinely defers to the executive branch on national security-related questions).

What We're Left With… TECHNICAL MEANS

- Just as the community came together to tackle domestic eavesdropping with widespread deployment of encryption – a technical solution – the time has come for the community to similarly tackle traffic analysis exposures via technical means, limited though current options may be.

- This is not a step that any of us want to have to take, but continued bulk collection of domestic metadata is deeply troubling to many.

- This threat will be countered by **technical means** if no other options exist.

VI. Technical Approaches To Dealing With Traffic Analysis

What Is the "Traffic Analysis Analog" To The Use of Encryption To Defeat Eavesdropping?

- Our portmanteau of user-based anti-traffic analysis options, such as it is, is limited, currently consisting of:

  (1) non-attributable endpoints

  (2) VPNs

  (3) Tor ("onion routing")

- Yes, other options exist, but they're so obscure as to be virtually unused, or so complex as to be impractical for average users.

1) Non-Attributable Endpoints

- '[former NSA and CIA chief Michael] Hayden, who helped build the intelligence agency's response to the digital age, was pretty clear about how he viewed it, saying "**the problem I have with the Internet is that it's anonymous."** ' http://www.washingtonpost.com/blogs/the-switch/wp/2013/10/05/the-nsa-is-trying-to-crack-tor-the-state-department-is-helping-pay-for-it/ [emphasis added]
- "[...] officials surveyed by the [Office of the Inspector General] identified **pre-paid calling cards and pre-paid cell phones** as the **top two threats affecting their ability to conduct electronic surveillance.**" See "The Implementation of the Communications Assistance for Law Enforcement Act," Audit Report 06-13, March 2006 , Office of the Inspector General http://www.justice.gov/oig/reports/FBI/a0613/exec.htm [emphasis added]

Unauthenticated (Open) Network Access

- There continue to be many unauthenticated or widely available/ nearly-open wireless access points, including ones at **coffee shops** or **fast food restaurants**, free **hotel wireless** networks, etc.
- These access points may sometimes provide less-attributable access that may be valuable for those seeking to casually avoid attribution, however often these access points are subject to abuse by spammers or others who seek to inject unwanted traffic, or are heavily filtered to damp down those complaints.
- Other "open access" wireless access points may actually be outright **malicious, capturing virtually all network traffic seen.** As a random user, your expectations for security and privacy on any wireless access point you just "stumble upon" should be nil.
- Use of "inadvertently insecure" wireless access points (rather than intentionally-available access points) may also serve as the basis for claims of **computer intrusion,** a felony in some jurisdictions.

Telephonic Non-Attributable Endpoints: Prepaid Cell Phones

	Prepaid Cell Phone	Contract Cell Phone
Registered?	Not to a personal identity	Yes
Financially tied?	No (if anonymous phone cards are used to 'top up')	Yes (if personal credit card is provided)
Tied to a user's email account?	Often no	Normally yes (incl. backups to "cloud")
Address book?	Minimal or none	Often extensive
Features?	Typically few ("just a simple cheap cell phone")	Often a smart phone with camera, GPS, micro SD cards, apps
Persistently used?	No (cheap; new phones routinely purchased and not directly linked to old phone/phone number)	Yes (~two year life, with old phone chaining to new one upon replacement)
Attributability?	Minimal	Extensive

Expect America To Eventually Go The Way of Africa (and Many Other Areas) And Require Registration

Green countries are those in Africa that do NOT require SIM registration as of 2/2014:

Cape Verde
Lesotho
Mauritania
Namibia
Somalia
Swaziland

See http://firstmonday.org/ojs/index.php/fm/article/view/4351/3820

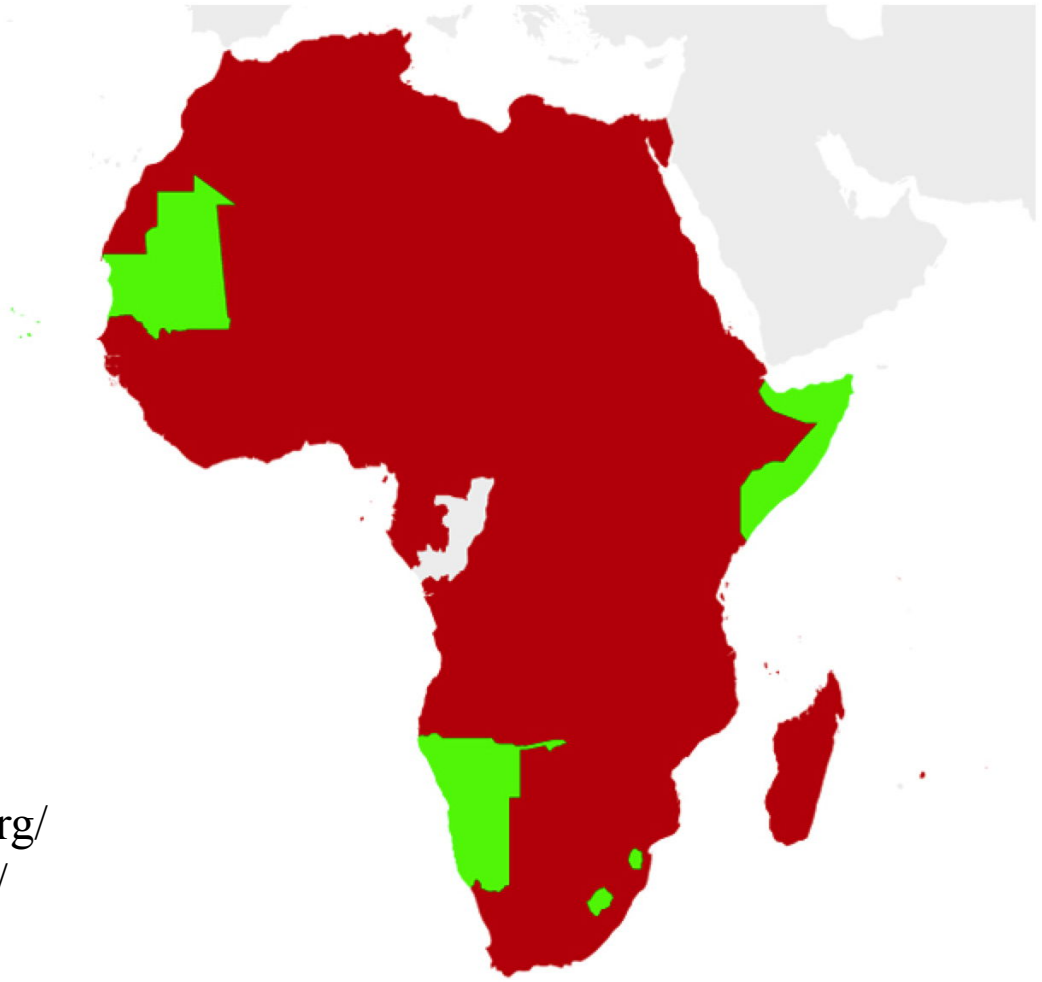

Maybe There Are Legitimate Reasons For Requiring Registration, As Many African Countries Have Done?

- Perhaps doing so would hinder fraudulent activity/misuse?
- "[...] to date there is **no evidence** that mandatory registration leads to a reduction in crime. [...**] the United Kingdom, the Czech Republic, Romania and New Zealand, have considered mandating prepaid SIM registration but concluded against it.** [...] In **Mexico, mandatory SIM registration was introduced in 2009 and repealed three years later** after a policy assessment showed that it had not helped with the prevention, investigation and/or prosecution of associated crimes"

  http://www.gsma.com/publicpolicy/wp-content/uploads/2013/11/GSMA_White-Paper_Mandatory-Registration-of-Prepaid-SIM-Users_32pgWEBv3.pdf [emphasis added]

Alternatives to Using A Cell Phone

- **Use a pay phone, instead** (but note that there are now less than half a million pay phones remaining in the U.S. according to the American Public Communications Council, see http://www.apcc.net/i4a/pages/index.cfm?pageid=40)
- **Consider using a prepaid one-way numeric pager** (these are simple one-way-only devices that receive broadcast pages, so they can't easily be tracked, although messages sent to pagers are obviously not private). N.B.: *two-way* pagers have the same issues as cellphones!
- **Go without a phone** – believe it or not, yes, you can survive without carrying a phone (but the sheer fact that you're choosing to do without may make you "stand out" as an abnormality)
- Be sure to consider the impact of cell phone usage by **family members, too,** if you choose to "go without a cell phone" yourself (their phones may become a proxy for geolocating you)

2) VPNs

- VPNs are "virtual private networks." **Inbound corporate VPNs** are routinely used to allow remote workers to securely access corporate resources while working "away from the office."
- **Outbound commercial VPN** providers also exist. These outfits, offering VPN service to any person willing to pay, are often suggested as a "solution" to overcoming traffic analysis exposure.
- If your traffic analysis threat model focuses around **local site monitoring (e.g., perhaps by your school or your employer)**, using an outbound VPN may allow you to tunnel past local traffic inspection points. **Use of a VPN for other purposes may be less effective.**
- Fundamentally, when you use a VPN, your traffic will "exit" from an alternative location, perhaps somewhere in the EU or Latin America. **You have no way to knowing if the operator of your service is trustworthy,** or routinely monitoring everything.

Virtual Private Networks (continued)

- When VPN traffic gets routed overseas, it will appear to come *from there*. **International traffic MAY be presumed to NOT be associated with a U.S. Person, and MAY therefore lose some protection from U.S. monitoring. Now add in any local host country monitoring that may be happening, too... ugh.**
- VPNs normally mix your traffic with that of other VPN users. While you may be using a VPN for laudable reasons, **other users of that same VPN service may be unsavory** (e.g., at least some of your fellow VPN users may be using a VPN in an effort to hide unlawful activities). **Your innocent traffic (and your innocent *identity*) may end up comingled and entangled with theirs.**
- Traffic from known VPN exit nodes may also be treated as untrustworthy/unwelcome by at least some mainstream sites.
- All in all, VPNs can be a bit of a "mixed bag" for the average user.

3) Tor

- If you were to ask technical people to mention one way to avoid classic traffic analysis attacks, the most common thing you'd probably hear mentioned is **Tor (The Onion Router)**.
- If you want to try Tor, it can be downloaded for free for **Windows, Mac and Linux** from https://www.torproject.org/ ; the Guardian Project has even ported it for **Android.**
- If you're a less technical person and just want to "buy hardware" in an effort to leverage Tor, see hardware offerings such as:
  -- https://pogoplug.com/safeplug , or the discussion at
  -- "Now Everyone Wants to Sell You a Magical Anonymity Router. Choose Wisely,"
  http://www.wired.com/2014/10/anonymity-routers/
- But note! You need to do **more than just install software** (or **more than just run a box) to mitigate your traffic analysis exposure.**

Tor Is Not (And Cannot Be) A "Magic Pill"

- Tor tries really hard, but if you fail to practice **strict operational hygiene**, your traffic may end up still being easily attributable (see http://www.wired.com/2014/12/fbi-metasploit-tor/)
- If a **bug arises and is exploited**, your traffic may also end up being attributable (see for example http://www.wired.com/2013/08/freedom-hosting/)
- Untrustworthy exit node operators may taint executables downloaded through their systems by **adding malware** (http://threatpost.com/researcher-finds-tor-exit-node-adding-malware-to-binaries/109008 , October 24th, 2014).
- Tor directory servers may be be targeted and attacked/seized (http://pando.com/2014/12/21/so-it-begins-operator-of-large-tor-exit-node-cluster-reports-he-has-lost-control-of-his-servers/)
- Tor's an ACTIVE focus of official attention right now, I suspect.

Tor Was/Is At Least Partially Federally Funded

- **Tor was originally a product of the Office of Naval Research and DARPA** (see http://www.onion-router.net/Sponsors.html)

- **Much of Tor's funding continues to come from the federal government, including the U.S. State Department.**
  See https://www.torproject.org/about/sponsors.html.en
  This is true, notwithstanding reported grumpiness about Tor from members of the intelligence community

- One emerging Tor alternative that's worth noting:
  https://geti2p.net/en/ ?)

So What Will ISPs Do?

- They need an **architecture** that will scale to **Internet-size audiences** and provide **reasonable protection** against traffic analysis for **average users** when they do **average stuff** with **minimal hassle for users or their providers.**

- Users can't be expected to fix the metadata/traffic analysis issue themselves. Available options are too limited, or too complex. ISPs need to protect their users from traffic analysis.

- Providers who want to protect their users need a non-disruptive solution that they can easily provision without requiring huge expense, or heroic measures.

End-User Broadband Network Providers

- End-user broadband provider networks should ensure that they're using many-to-one NAT/PAT (with many users per public IP address), DHCP with short leases, and minimal or no logging.
- In a NAT/PAT environment, users connected from behind a shared public IP. DHCP is used to dynamically assign IP addresses from a shared pool. To a first approximation, the only one who knows who's on a dynamically assigned DHCP address behind a NAT/PAT gateway is the ISP operating that network. (We'll disregard things like cookies for this initial discussion)
- If providers don't keep any DHCP or NAT/PAT logs, it will be difficult or impossible for external parties to readily map normal wide area traffic to individuals at scale.
- Unlike the European Data Retention Directive, **the US has no mandatory data retention directive** (but this is not legal advice; ISPs should check with their own legal team for legal advice).

What About Web Hosting Providers?

- Web hosting companies also have options.
- They might put as many different web site domains on a single IP address as possible, and all servers could be protected with SSL/TLS.
- Loading a large number of domains onto a single IP address may be done either on the web server itself (e.g., using regular virtual hosting) or through use of a reverse proxy front end.
- Why would loading many domains onto each IP help with pervasive monitoring? Well, recall that per DOJ policy, with only a few exceptions, web URLs are treated as "content," not "metadata, and as such require a Title III full contents intercept order, not just a pen register/trap and trace order, see http://www.justice.gov/usao/eousa/foia_reading_room/usam/title9/7mcrm.htm#9-7.500

What About ISP Anti-Abuse Efforts? And What About Lawful Intercepts by LEOs?

- **The provider** will still have the ability to identify abusers based on internal network traffic monitoring and analysis, **done from within the NAT/PAT boundary**, should they need to do so.
- **Law enforcement officers** can likewise still identify a persistently problematic user, they'd just need to serve the ISP appropriate legal paperwork and **work inside the NAT boundary.** This might not be fun or easy, nor scale to hundreds of millions of users, but it would be an option if/when its really needed.
- **IMPORTANT:** Use of NAT/PAT and DHCP without logs, and the practice of hosting many web domains on each IP could also obviously be **revisited** if/when the current bulk domestic metadata collection program gets administratively re-scoped.

Have We Threaded the Needle Again?

- So just as with deployment of encryption for email in transit, use of NAT/PAT, DHCP, and heavily shared web hosting appear to represent an example of a deployable solution to hinder bulk metadata collection and traffic analysis attacks, simultaneously ensuring:

  -- Average law-abiding users get some protection from bulk pervasive domestic metadata collection.
  -- ISPs can inexpensively protect their customers while still being able to deal with problematic abuse if it arises.
  -- LEOs can still get what they need to deal with the bad guys who truly deserve to be investigated, arrested, tried and punished.

Thanks for the Chance to Talk Today!

- Are there any questions?