

Upcoming Requirements from the US Law Enforcement Community to Technically Facilitate Network Wiretaps

TERENA, Lyngby Denmark

Tuesday, May 22nd, 2007 4-5:30PM

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

Internet2 and the University of Oregon

A condensed version of this talk is available online at <http://www.uoregon.edu/~joe/calea-requirements/>

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

Introduction

Today's Talk

- I'm neither a lawyer nor a law enforcement person, so this talk is not meant to be legal advice, nor does it in any way express any sort of "official" opinion about CALEA.
- What I've done is to:
 - look at what law enforcement (LE) appears to want/need,
 - look at the sort of networks and systems architectures that higher education currently has planned or deployed, and
 - review some public documents relating to lawful intercept.
- Considering those requirements, facilities and documents, I've then endeavored to discuss and explain the issues which I believe may ultimately frustrate law enforcement's goals and objectives, frustrations which may (and probably should) end up driving requests by them for clarifying amendments to CALEA, and specific technical assistance.
- But first let's take a quick look at 4 important considerations.

"What You Talk About Today Will Help Law Enforcement To 'Spy' On Us!"

- Some may be concerned that today's discussion may in some way encourage law enforcement to become illegally invasive of our privacy. I don't believe this talk will do that.
- The factors which will deter any potential abuse of communication interception capabilities are our legal system, the personal ethics of the people involved, agency internal monitoring, independent oversight, structural checks & balances, and public attention (including media coverage).
- Things which this talk might help prevent include:
 - large institutional expenditures spent enabling unusable or ineffective lawful interception capabilities, and
 - the introduction of architectural weaknesses or the incidental creation of vulnerable or exploitable systems.
- **All of us want to our right to privacy to be protected.**

"What You Talk About Today Will Help The Bad Guys Avoid Surveillance!"

- Just as there are those who worry today's talk may help law enforcement do something illegal, there are also some who may worry that today's talk may somehow help drug dealers, child pornographers, terrorists or other really bad people "slip through LE's net." I wouldn't want that to be the case, and I sincerely don't think that what we talk about today will cause that to become true – in fact I've been very careful to insure that that doesn't become the case.
- I'd like to also emphasize that higher education has always been very cooperative when it comes to lawfully assisting authorities in dealing with criminal activities and helping to protect and defend our country.
- **We all want our loved ones and country to be safe.**

"CALEA's A US Topic – Why Talk About A US Topic At An EU Event?"

- When I was asked by the program committee to present on CALEA, I wondered if CALEA would indeed be a relevant topic for a largely EU audience, but then I recognized that:
 - the Internet is transnational, and many European systems or users routinely connect to US systems or users
 - with today's mobile society, many Europeans may physically travel to, and spend time in, the United States, and thus potentially become locally subject to CALEA
 - while some technologies may end up deployed first in the United States or first in the EU, similar requirements and a global marketplace mean that eventually most capabilities (including lawful intercept solutions) diffuse worldwide.
- Thus, while CALEA is nominally a US topic, **lawful intercept is actually a topic of worldwide importance and impact**⁶

"CALEA's a 'Legal Thing;' This Is a Technical Networking Conference"

- CALEA is indeed a law, but it is a law that deals at fundamental levels with very complicated technical networking issues. In fact, I would go so far as to say that **many of the current problems with CALEA come from a lack of appropriate technical input from the networking and system administration community. This talk is an attempt to begin correcting that shortcoming.**
- Much of this talk is oriented at network architects and engineers, system administrators and security professionals, although higher education leadership, government officials, policy analysis, and others may also find it accessible and perhaps of interest.
- Whew... with that out of the way, **let's begin by trying to understand what CALEA is trying to do.**

Perspectives and Objectives

What Is CALEA?

- CALEA is the United State's "Communication Assistance for Law Enforcement Act of 1994," see 47 SC 1001-1021.
- Quoting <http://www.askcalea.net/>, CALEA “defines the existing statutory obligation of telecommunications carriers to **assist law enforcement in executing electronic surveillance pursuant to court order or other lawful authorization**. The objective of CALEA implementation is to preserve law enforcement's ability to conduct lawfully-authorized electronic surveillance while preserving public safety, the public's right to privacy, and the telecommunications industry's competitiveness.” Recent FCC administrative actions (and court decisions targeting those actions), have clarified that this 1994 law includes “facilities based broadband providers,” and under some circumstances, **some higher education networks**. The statute says...

47 USC 1002

(a) Capability requirements

Except as provided in subsections (b), (c), and (d) of this section and sections 1007(a) and 1008(b) and (d) of this title, a telecommunications carrier shall ensure that its equipment, facilities, or services that provide a customer or subscriber with the ability to originate, terminate, or direct communications are capable of--

(1) expeditiously isolating and enabling the government, pursuant to a court order or other lawful authorization, to intercept, to the exclusion of any other communications, all wire and electronic communications carried by the carrier within a service area to or from equipment, facilities, or services of a subscriber of such carrier concurrently with their transmission to or from the subscriber's equipment, facility, or service, or at such later time as may be acceptable to the government;

(2) expeditiously isolating and enabling the government, pursuant to a court order or other lawful authorization, to access call-identifying information that is reasonably available to the carrier—

- (A) before, during, or immediately after the transmission of a wire or electronic communication (or at such later time as may be acceptable to the government); and
- (B) in a manner that allows it to be associated with the communication to which it pertains,

except that, with regard to information acquired solely pursuant to the authority for pen 10

47 USC 1002 (cont. #1)

registers and trap and trace devices (as defined in section 3127 of title 18), such call-identifying information shall not include any information that may disclose the physical location of the subscriber (except to the extent that the location may be determined from the telephone number);

(3) delivering intercepted communications and call-identifying information to the government, pursuant to a court order or other lawful authorization, in a format such that they may be transmitted by means of equipment, facilities, or services procured by the government to a location other than the premises of the carrier; and

(4) facilitating authorized communications interceptions and access to call-identifying information unobtrusively and with a minimum of interference with any subscriber's telecommunications service and in a manner that protects—

- (A) the privacy and security of communications and call-identifying information not authorized to be intercepted; and
- (B) information regarding the government's interception of communications and access to call-identifying information.

[see also the legislative history for this act, House Report No. 103-827, Part I, available online at <http://www.askcalea.net/docs/hr103827.pdf>]

CALEA Was Created With Limitations

(b) Limitations

(1) Design of features and systems configurations

This subchapter does not authorize any law enforcement agency or officer--

(A) to require any specific design of equipment, facilities, services, features, or system configurations to be adopted by any provider of a wire or electronic communication service, any manufacturer of telecommunications equipment, or any provider of telecommunications support services; or

(B) to prohibit the adoption of any equipment, facility, service, or feature by any provider of a wire or electronic communication service, any manufacturer of telecommunications equipment, or any provider of telecommunications support services.

(2) Information services; private networks and interconnection services and facilities

The requirements of subsection (a) of this section do not apply to--

(A) information services; or

(B) equipment, facilities, or services that support the transport or switching of communications for private networks or for the sole purpose of interconnecting telecommunications carriers.

(3) Encryption

A telecommunications carrier shall not be responsible for decrypting, or ensuring the government's ability to decrypt, any communication encrypted by a subscriber or customer, unless the encryption was provided by the carrier and the carrier possesses the information necessary to decrypt the communication. [continues]

CALEA and Law Enforcement

- Law enforcement (LE) wants to catch terrorists and criminals, and thereby help keep us safe
- The terrorists and criminals whom LE struggle with are using telecommunications to plan and conduct illegal operations
- Telecommunication technologies in use include wireline and wireless telephones, fax machines and digital pagers, **but also** Internet-based communication technologies such as voice over IP (VoIP), instant messaging (IM), video teleconferencing (H.323) and email
- Traditional wiretaps can (theoretically) be used to intercept VoIP, IM, H.323 traffic and email in transit, but...
 - traditional wiretaps may be impracticably **slow**, and
 - traditional wiretaps are labor intensive and **can't scale** to the volume of intercepts which may be potentially required

CALEA and Higher Ed

- Like everyone else, those of us in higher education want to be safe from terroristic violence and criminal schemes, but...
- We also want to insure that constitutional safeguards and critical procedural limitations are fully respected.
- Higher education administrators also worry about the budgetary impact of unfunded federal mandates imposed upon already-strained institutional budgets.
- And some may want assurance that if we're asked to invest money in a new capability:
 - you really need what you're asking for, and
 - you're going to use what we build, buy or deploy for you.
- Oh yes, if we have to do a technical project, we also want to do it in a technically sound way.
- None of this is inconsistent with what LE wants.

The Internet: We Don't Want It To Be a Sanctuary for the Bad Guys!

- When passed in 1994, CALEA was intended to modernize & preserve LE's ability to conduct lawful intercepts, particularly given advanced calling features and growing cell phone use.
- If CALEA had remained focused on telephones, we wouldn't be talking today, but technologies evolved, and the result was that CALEA needed to evolve with those technologies to encompass communications taking place over the Internet
- To understand why, note that if LE could only conduct lawful interceptions in the traditional wireline or wireless phone domain, the bad guys & gals would quickly learn that, and at least the smart crooks would go to where they'd be immune from surveillance, e.g., they'd move to the Internet (and as we all know, the Internet is a great communication tool).

CALEA Wasn't (and Isn't) Perfect

- **CALEA involves many federal agencies:** the lead agencies are the FCC and the FBI, but the DEA and other agencies may also be providing input and direction (and ironically the consensus result may be fully satisfactory to none of them).
- CALEA's evolution and extension to the Internet occurred by the FCC's **creative interpretation of an existing statute**, rather than clear and unambiguous legislative action *de novo*.
- Not surprisingly, CALEA has been the **subject of litigation**, including litigation which yielded a complex and tortured judicial decision which, when read, does little to shore up CALEA's legitimacy as applied to Internet technologies.
- CALEA has had a very **slow roll out**, in part because CALEA involves complex technical matters and required industry help in developing **appropriate technical standards**
- CALEA is also potentially very **expensive**.

CALEA and the DOJ Inspector General

- "According to the Federal, state, and local law enforcement officials we interviewed and surveyed, their agencies do not request intercepts requiring CALEA features for several reasons (i.e., the **high cost charged by carriers**, [...], or the investigation **only required a traditional wiretap**)."
- 'Law enforcement's biggest complaint regarding CALEA is the relatively high fees charged by carriers to conduct electronic surveillance. **A traditional wiretap costs law enforcement approximately \$250. However, a wiretap with CALEA features costs law enforcement approximately \$2,200** according to law enforcement officials and carrier representatives we interviewed. A law enforcement official noted that, "[w]ith CALEA, the carriers do less work but it costs approximately 10 times as much to do a CALEA-compliant tap" [emphasis added]

Additional DOJ Inspector General CALEA Report Comments...

- 'According to the FBI, Internet “hotspots” such as **cyber cafés that provide anonymity with multiple access points**, third-party calls using calling cards, and toll free numbers are a “technologically unsolvable problem.” These services can only be addressed through investigative techniques, rather than through the application of CALEA. In addition, FBI officials said that commercially available electronic **encryption** will also hinder law enforcement’s ability to collect information from electronic intercepts.'
- Lots more interesting data is in "The Implementation of the Communications Assistance for Law Enforcement Act," U.S. **Department of Justice Audit Report 06-13**, March 2006, Office of the Inspector General Audit Division; see <http://www.usdoj.gov/oig/reports/FBI/a0613/final.pdf> [emphasis added]

CALEA Also Ended Up Getting Conflated With A Controversial War on Terror

- CALEA became law in **1994**, but the CALEA of today has been strongly impacted by government reaction to the **tragic events of 9/11/2001** (e.g., counter terrorism and counter intelligence intercepts now dominate criminal intercepts)
- Because of the scale (and the scope of funding involved!), the war on **terror has ended up becoming a very partisan issue**, and thus anything **related** to the war on terror or homeland security becomes inherently contentious & political
- CALEA also lives in a sort of **damned-if-you-do, damned-if-you-don't** continuum. If CALEA were to become routinely used, LE agencies might be accused of becoming "**wiretap happy**," yet if CALEA were to be reserved for only the most dire of circumstances, critics might claim that CALEA is just an obscure, unimportant, & basically "**unneeded**" **capability**.

Lawful Intercepts by the Numbers

Nationwide, in 2006 (the most recent reporting year available):

- Intercepts authorized by federal & state courts in '06: 1,839 (461 by federal judges and 1,378 by state judges)
- State courts with the most approved intercepts: CA (430), NY (377), NJ (189), FL (98) – those four states accounted for 79% of all state intercept orders; 27 state courts reported no intercepts whatsoever.
- Average days installed wiretaps were in operation: 40 days
- Average number of people whose communications were intercepted per wiretap order: 122
- **80% of all wiretaps involved drug offenses**; racketeering and homicide/assault were the other two top offenses cited.
- Average cost of a federal intercept: \$67,044. Average cost of a state intercept: \$46,687.

Lawful Intercepts by the Numbers (cont.)

- **Wiretap requests which were for telephones: 96%**
- **Wiretap requests involving mobile devices, such as cell phones: 92%**
- **Number of federal or state intercepts encountering encryption: 0.**
- **Wiretap requests which were for "digital pagers, fax, or computers:" roughly 0.7% (13 requests out of 1,839)**
- **Source: U.S. Courts' 2005 Wiretap Report,**
<http://www.uscourts.gov/wiretap06/contents.html>
- **Note:** these numbers do **not** include FISA intercepts and some selected other categories of interceptions. Data on national security-related intercepts is included in the expanded version of this talk available via the web. (FISA and other national security intercepts dwarf criminal intercepts by law enforcement)

An Aside: So What About FISA?

- ***FISA (and other foreign intelligence-related programs) dwarf domestic criminal use of electronic intercepts...***
- 9,254 reportable National Security Letters (NSLs) involving 3,501 different US persons were issued during 2005;
47,221 total NSLs involving 9,475 US and 8,536 non-US person
- 2,074 applications for electronic surveillance to the Foreign Intelligence Surveillance Court in 2005 (2,072 approved)
- 155 applications for access to certain business records (including the production of tangible things) were made
- **Sources:**
<http://www.fas.org/irp/agency/doj/fisa/2005rept.pdf>
<http://www.usdoj.gov/oig/special/s0703b/final.pdf> (199 pages)
<http://www.usdoj.gov/oig/special/s0703a/final.pdf> (112 pages)
<http://www.eff.org/legal/cases/att/>

CALEA and The Cybercrime Wave

- CALEA has also implicitly gotten caught up in the Internet "cybercrime wave." When you look at things like:
 - nine out of every ten mail messages are now spam
 - thousands of new viruses are created per month
 - there are millions of broadband-connected botnet hosts,
 - leading businesses and critical infrastructure services are constantly being scanned, probed and compromised (or labor under DDoS attacks)so you might wonder, **"Does LE have the online investigative tools, methods and authority they need to cope with these online cyber threats?"**
- The answer **may** be "no." For example, while CALEA might help tackle some of the cyber crimes mentioned above, internal LE procedures may be so careful and methodical as to render many legitimate uses of LI impracticable.

Wiretaps Without CALEA

- Traditionally, wiretapping has been a manual process:
 - A court of competent jurisdiction would issue suitable an order authorizing a wiretap to occur,
 - The local provider (whether that's a telephone company, Internet service provider, or other entity) would be contacted by LE and asked to provide assistance
 - Legal review of LE's request for assistance would occur
 - Assuming that local legal review is positive, technical steps would be taken to facilitate the requested intercept, such as mirroring a switch port or installing an optical splitter
 - Traffic from that intercept would be minimized to insure that only traffic covered by the paperwork would be extracted
 - The minimized traffic would be delivered to LE
 - LE analysis of the intercepted traffic would then occur
- **This is not a painless, rapid, or inexpensive process.**

Traditional Wiretaps Can't Be Provisioned At "Internet Speed"

- Traditional wiretaps aren't very agile – they don't (and can't!) be provisioned at "Internet speed."
- By this I mean that in many cases a network connection may be used for just a very brief period of time, but traditional wiretaps might take days (or weeks!) to request, approve and arrange, and by that time, the subject of the interception order might be long gone, having moved on through a series of one or more other connections in the intervening time.
- In some cases it may be possible to obtain a court order authorizing a so-called "roving" or "multipoint" wiretap (but those have traditionally been uncommon), and even then, the physical mechanics of effecting the interception can be thwart the intent of the order.
- CALEA may be designed to partially begin fixing this.

Universities and Compliance

Key CALEA Resource For Higher Ed

- Educause has an excellent CALEA resource page for higher education users at

<http://www.educause.edu/calea>

and there is also a CALEA-HE mailing list for higher ed users which you can join via

<http://listserv.educause.edu/cgi-bin/wa.exe?A0=CALEA-HE>

- Many of the resources referred to in the following sections can be readily found on the Educause CALEA web site.

Deliverables and Dates

- **If (and only if)** you're an American campus network or other entity subject to CALEA, you have a number of new substantive and procedural responsibilities.
- Relevant dates for CALEA "deliverables" have included:
 - By February 12th, 2007, you should have filed FCC Form 445, "CALEA Monitoring Report for Broadband and VoIP Services" (see <http://www.fcc.gov/Forms/Form445/445.pdf>)
 - By March 12th, 2007, you should have filed the required "System Security and Integrity" ("SSI") Plan (examples available on the Educause site)
 - Finally, May 14th, 2007, was the deadline for full CALEA compliance. Full compliance required meeting the requirements of the appropriate industry technical standard(s) (see <http://www.askcalea.net/standards.html>)

“How Do We Technically Become CALEA Compliant?”

- If you need to instrument your network for CALEA, some options include:
 - You can “roll your own” CALEA solution ala Merit (see <http://www.opencalea.org/>)
 - You can purchase a commercial vendor solution (some options are listed at <http://www.educause.edu/ir/library/pdf/EPO0708.pdf>)
 - You can employ a “trusted third party” to effect CALEA compliance for you (see for example the list at <http://www.educause.edu/ir/library/pdf/EPO0707.pdf>)
- Which of those solutions makes sense for a given site may be a technical, financial or political question.

The Question of the Month: “Does My Campus Need to Be CALEA Compliant?”

- Because everyone’s circumstances will differ, and because this is a very complex issue, this is a question that your administration will ultimately need to decide after consultation with your legal staff. Subtle differences in circumstances, or in the analysis of those circumstances, may lead seemingly identical entities to radically different conclusions.
- A relatively large number of potential exemptions have been identified.
- Some of the exemptions your legal counsel may be considering include...

The Private Network Exemption

- 47 U.S.C. 1002 (b)(2)(B) exempts
"equipment, facilities, or services that support the transport or switching of communications for private networks."
- Unfortunately, "private network" is not a term explicitly defined in the Act, and because the Internet is a series of interconnected hierarchical private networks, it can sometimes be difficult to ascertain exactly where a "private network" ends and "the public Internet" begins.

The Private Network Exemption (cont.)

- Clearly, a network which exists solely within a single building or facility and which does not interconnect with any networks owned or operated by other entities would be a "private network" for the purposes of CALEA.
- That sort of physically isolated private network is rare, however, and restricting it to just that one extreme type of "private network" would be unduly and unnecessarily limiting since the **FCC has made it clear that the private network exemption potentially encompasses far more.**

See **footnote 100** on PDF page 19 of the FCC's "First Report and Order and Further Notice of Proposed Rulemaking" as adopted August 5th, 2005, FCC 05-153.

I quote...

Footnote 100

"Relatedly, some commenters describe their provision of broadband Internet access to specific members or constituents of their respective organizations to provide access to private education, library and research networks, such as **Internet2's Abilene Network, NyserNet, and the Pacific Northwest gigaPoP**. See, e.g., EDUCAUSE Comments at 22-25. To the extent that EDUCAUSE members (or similar organizations) are engaged in the provision of facilities-based private broadband networks or intranets that enable members to communicate with one another and/or retrieve information from shared data libraries not available to the general public, these networks appear to be private networks for purposes of CALEA.

"Indeed, **DOJ states that the three networks specifically discussed by EDUCAUSE qualify as private networks under CALEA's section 103(b)(2)(B)**. DOJ Reply at 19. We therefore make clear that **providers of these networks are not included as "telecommunications carriers" under the SRP with respect to these networks**. To the extent, however, that these private networks are interconnected with a public network, either the PSTN or the Internet, providers of the facilities that support the connection of the private network to a public network are subject to CALEA under the SRP."

Public Access and Interconnectivity

- Institutions interested in relying on this exemption thus need to pay attention to the extent to which their private networks end up being **publicly accessible**, and to any **interconnections** between their private network and either the **public switched telephone network or the Internet**. It is particularly worthy of note that at least in some cases a private institutional network may interconnect with a private regional network or private national network, and only with private regional or private national networks, and thus the institution may not be subject to CALEA compliance obligations. [Please see the American Council on Education (ACE)'s document “The Application of CALEA to Higher Education,” and ACE vs. FCC, U.S. Court of Appeals for the District of Columbia Circuit, No. 05-1404, June 9, 2006 particularly at PDF page 19.]

Public Access May Not Be Knowable

- Given the potential importance of that "not publicly accessible" stipulation, it may be worth noting that **most networks cannot definitively determine whether the general public has "access" to their networks or not.**
- For example, the public may end up having access to a nominally private network via a rogue insecure wireless access point installed by a student or faculty member, or there may be insecure, unadvertised, but used open network jack in a publicly accessible space, such as a classroom.
- Most would describe this as "de minimis" and inadvertent public access, distinguishing it from the intentional delivery of public Internet access (whether that's free unauthenticated Internet access provided to anyone using the campus library, or for-fee resale of ISP services to community entity).
- Upstream networks may not be able to independently identify **any of those types** of downstream public access.³⁵

Internet Gateway Compliance (Only)

- At one point there was also concern that universities would need to replace virtually all their network equipment in order to make it possible to do lawful CALEA interceptions within private networks themselves.
- That is, if you wanted to be able to lawfully intercept traffic going from one local user to another local user, with both users connecting via the private network, it would not be sufficient to just be able to intercept traffic at the Internet gateway -- traffic exchanged between two local users would remain entirely within the local private network, and since it would never touch the Internet gateway, it would not be able to be lawfully intercepted.
- In its second report and order, however, the FCC clarified that in fact private networks did in fact **only** need to be CALEA compliant at their **Internet gateway**.

Internet Gateway Compliance (2)

- See, for example, the FCC's Second Report and Order and Memorandum Opinion and Order, Adopted May 3, 2006, FCC 06-56 at page 82, which states,

"Petitioners' professed fear that a private network would become subject to CALEA "throughout [the] entire private network" if the establishment creating the network provided its own connection between that network and the Internet is unfounded. The [First Report and Order] states that only the connection point between the private and public networks is subject to CALEA. This is true whether that connection point is provided by a commercial Internet access provider or by the private network operator itself."

Internet Gateway Compliance (3)

- Thus, it is possible to envision a scenario whereby:
 - an institution's private network in turn connects to
 - a private regional network.
- Given the gateway compliance rule, CALEA compliance would only be required at the point where the private regional network interconnects with the public Internet or the PSTN, but that requirement also needs to be viewed in light of the Interconnecting Telecommunications Carriers Exemption.

Interconnecting Telecommunications Carriers Exemption

- 47 U.S.C. 1002 (b)(2)(B) also exempts "equipment, facilities, or services that support the transport or switching of communications [...] for the sole purpose of interconnecting telecommunications carriers." Thus, "equipment, facilities, or services that support the transport or switching of communications [...] for the sole purpose of interconnecting telecommunication carriers" would not be subject to CALEA.
- But what is a "telecommunication carrier?" The FCC clarified this for CALEA purposes in rules it issued, see FCC 06-56 at page 45, section 1.20002 (e)...

Interconnecting Telecommunications Carriers Exemption (2)

- ***Telecommunications carrier. The term telecommunications carrier includes:***

(1) A person or entity engaged in the transmission or switching of wire or electronic communications as a common carrier for hire;

(2) A person or entity engaged in providing commercial mobile service (as defined in section 332(d) of the Communications Act of 1934 (47 U.S.C. 332(d))); or

(3) A person or entity that the Commission has found is engaged in providing wire or electronic communication switching or transmission service such that the service is a replacement for a substantial portion of the local telephone exchange service and that it is in the public interest to deem such a person or entity to be a telecommunications carrier for purposes of CALEA.

Interconnecting Telecommunications Carriers Exemption (3)

- In considering those definitions, note that only one of two alternatives may logically be true: either an entity is a telecommunication carrier, or it isn't.
- If the entity IS NOT a telecommunication carrier, it is not subject to CALEA (e.g., see for example, Section 103(a) "Except as provided in subsections (b), (c), and (d) of this section and sections 108(a) and 109(b) and (d), a **telecommunications carrier** shall..." (emphasis added) and see also ACE vs. FCC, U.S. Court of Appeals for the District of Columbia Circuit, No. 05-1404, June 9 '06, at p.4
- Thus a private regional network which is also not a telecommunications carrier would not be subject to CALEA compliance obligations (its upstream, if a public Internet provider or PSTN provider, would be).

Interconnecting Telecommunications Carriers Exemption (4)

- If the entity **IS** a telecommunication carrier, when focusing on the Interconnecting Telecommunications Carriers Exemption, one should then ask, "Does the telecommunication carrier have equipment, facilities, or services that support the transport or switching of communications [...] for the **sole purpose** of interconnecting telecommunication carriers?" If so, then that equipment, and those facilities and services may **ALSO** not be subject to CALEA obligations.
- So what, then, of a carrier-to-carrier equipment, facilities or services which also happen to be an "Internet gateway" for downstream private networks? (I am not a lawyer, and I don't have a good answer for this particular question)

Last Mile Focus

- This issue of network hierarchy and gateway compliance is also relevant in so far as CALEA's emphasis is on so-called "last mile" connectivity, not backbone interconnections between carriers.
- Why is law enforcement **not** particularly interested in connections between backbone carriers for CALEA compliance purposes?
- Backbone carriers may lack the knowledge needed to identify network traffic that may be associated with a named lawful intercept subject of interest ("All network traffic originated by or destined for Susan Marie Anderson of 345 Elm Street, Wagonwheel, Oregon.") – the backbone carrier would simply have no idea what traffic is associated with that person of interest. E.G., only the last mile provider might know what IP address or MAC address she's using.

A Strange Potential Situation

- With that for background, now consider a scenario where:
 - the institutional private network is exempt,
 - the regional private network is exempt, and since
 - compliance need only occur at the gateway from the private network to the public Internet (or PSTN), but the "Internet gateway" might effectively end up "pushed up" to an interconnecting telecommunications carriers link, but that link may also have been exempted by CALEA (and if not, the carrier may simply not have access to the data they'd need to comply). Wild, eh?
- One more potential exemption to mention...

Retail Establishment Exemption

- A final potentially relevant exemption can be found in the so-called "**coffee shop**" exemption or "**retail establishment exemption**" described at paragraph 36 and footnote 99 on PDF page 19 of 59 of the First Report and Order, FCC 05-153 which states,
*"Finally, in finding CALEA's SRP to cover facilities-based providers of broadband Internet access service, we conclude that **establishments that acquire broadband Internet access service from a facilities based provider to enable their patrons or customers to access the Internet from their respective establishments are not considered facilities-based broadband Internet access service providers subject to CALEA under the SRP.** [footnote 99] We note, however, that the provider of underlying facilities to such an establishment would be subject to CALEA, as discussed above."* [emphasis added]

Footnote 99

- Footnote 99 reads:
'Examples of these types of establishments may include some hotels, coffee shops, schools, libraries, or book stores. DOJ has stated that it has "no desire to require such retail establishments to implement CALEA solutions," DOJ Comments at 36, and we conclude that the public interest at this time does not weigh in favor of subjecting such establishments to CALEA.' [emphasis added]
- This exemption might provide additional grounds for some schools to assert that they are exempt from CALEA compliance obligations. Note, too, that it effectively deprecates the possibility of a hierarchy of exempt private networks, since the "provider of underlying facilities to such an establishment would be subject to CALEA" apparently as an absolute matter by this finding.

One or More of Those Exemptions May Apply to Many HE Institutions

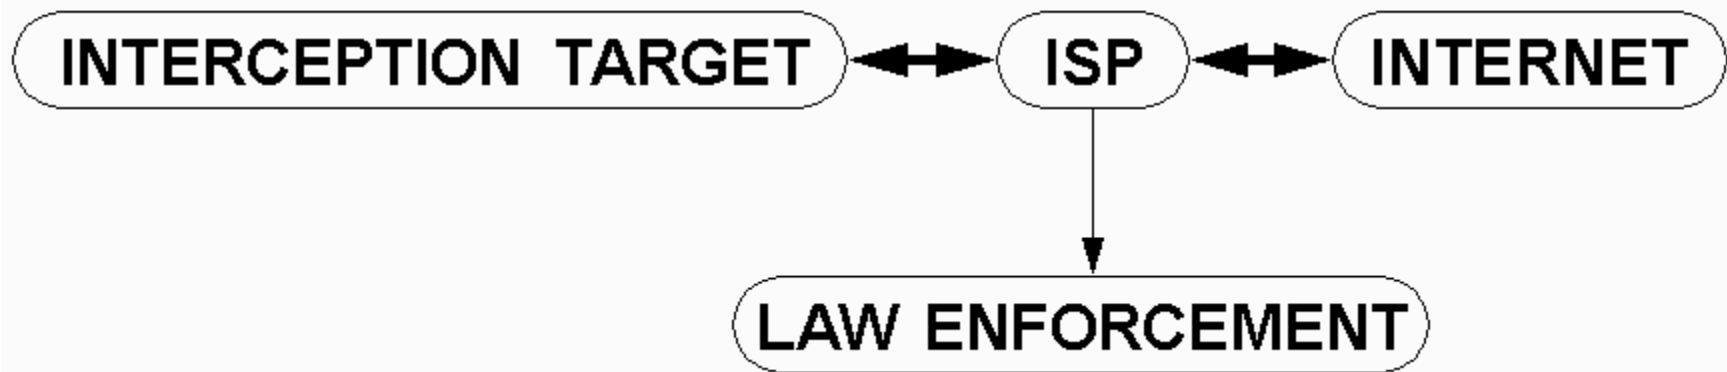
- Because one or more of those exemptions may apply to many higher education institutions, **many colleges and universities (and statewide or region-wide higher education networks) have not filed either the “CALEA Monitoring Report for Broadband and VoIP Services” reports, nor a “System Security and Integrity” (“SSI”) Plan, nor have they instrumented their network to be able to deliver CALEA-related data to law enforcement...**
- Even if all college and university campuses, and all statewide or region-wide higher education networks were specifically required to become CALEA compliant, there'd still be a lot of things that would make compliance tricky.
- To understand those things, let's look at the conceptual connection model (which often bears no relation to reality)?

The Connection Model

Customer, ISP, Internet and LE

- In what we might think of as the "conceptual connection model," the model which I think Congress had in mind when they were drafting CALEA and which the FCC had in mind when they were applying it to facilities based broadband providers, there are basically just four relevant parties:
 - the customer of the ISP/LE's interception target,
 - the ISP,
 - the Internet, and
 - law enforcement.
- This is obviously a very simple model (see the diagram on the next page).

The Simple (Wrong) Conceptual Model



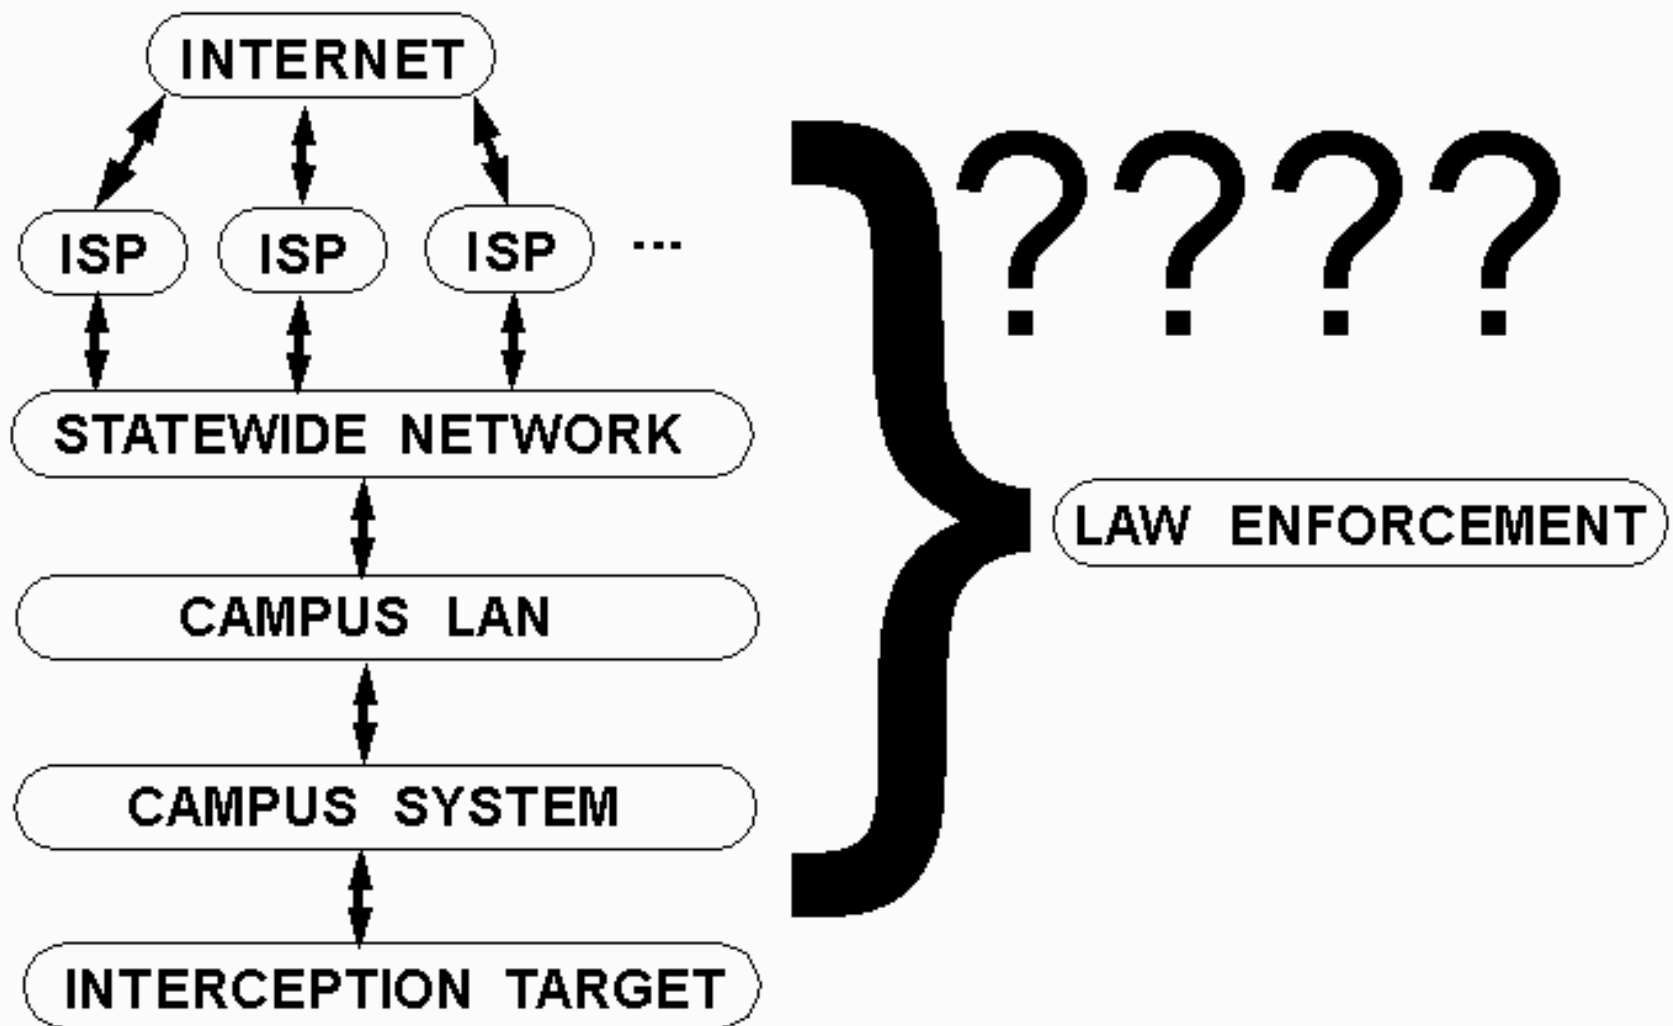
Because There Are Only Four Participants In the Standard Model...

- ... everything is pretty simple.
- If you're writing lawful intercept laws, and you've got the standard model in mind, the problem's pretty straightforward:
 - you know who's targeted for monitoring
 - you know who needs to be able to make the intercept work on a technical layer, and
 - you know who needs to receive the intercepted traffic.
- The "standard model" may even be an approximately accurate representation of how an individual user or small college connects directly to the Internet via a commercial Internet Service Provider.
- But often that simple model breaks down...

A More Realistic Model

- A more realistic model, at least for large universities, often involves additional layers (see the next slide). Rather than having a model with only four participants (intercept target, ISP, Internet and law enforcement), many more "layers" and many more participants may be involved.
- For example, some traffic may flow through a large shared university system, where it may emerge commingled with traffic from hundreds or thousands of other users.
- In other cases, network traffic may cross a private campus network or a private state wide network, perhaps being rewritten as it crosses a firewall/network address translation (NAT) device.
- When traffic finally leaves the statewide network, it may exit on any one of a number of ISPs, rather than just one.
- So who's responsible for getting LI traffic to LE?

More Common (Conceptual)
Higher Ed Network Environment



Everyone; No One; No One Knows

- Because of the complexities we've been discussing, there's a distinct possibility that **"everyone; no one; no one knows"** may be the phrase that best describes who's responsible for being able to lawfully intercept traffic in a complex university network environment.
- The campus, the statewide network, AND the upstream ISPs may ALL instrument their networks so as to be able to support CALEA, a permissible situation, but potentially expensive "overkill."
- On the other hand, a campus might expect the statewide network to take on that obligation; the statewide network might expect the campuses or the statewide network's upstream ISPs to handle it; the upstream ISPs might expect the campuses or the downstream statewide network to handle it. Ultimately no one may be ready/able to respond.⁴


The Costs and Benefits of Doing Lawful Intercept (LI) "Upstream"

- Let's assume for the sake of argument that the connections between the statewide network and a national service provider end up being determined to be the "Internet Gateway," and thus must be compliant.
- Technically, it can be tricky to do LI on high bandwidth pipes.
- Remember, too, that the state network may have no idea who's traffic they're delivering by the time it gets to the gateway between the state network and the ISP.
- Finally, at least in some cases, there may be literally millions of users "downstream" from that state network's Internet gateway, and **any** traffic within that network – as long as it **stays** within that "private" network – would not need to be able to be monitored under CALEA's Internet "gateway compliance only" provision.

Even Being Gateway Compliant at the Campus Level May Mean Missed Flows

- Even if we tighten up, and move the compliance point from the statewide network/ISP demarc to the statewide network/campus demarc, you'll still end up missing internal flows...
- See the red arc on the diagram on the following slide... only traffic that goes through the "Internet Gateway" (in this case shown as the campus border router) would be potentially able to be intercepted.
- The probability that traffic will end up passing through the Internet Gateway is partially a function of how "far upstream" the "Internet Gateway" may be – the more users who are downstream, in the "private network" region below the Internet Gateway, the greater the chance that their traffic will remain local (going to another local user of the private network), and hence be effectively unmonitorable.

Per Port vs. Gateway Compliance: Local Flows Between A&B Don't Touch The Border Router And Hence Aren't Subject to Lawful Interception



The Closer You Go To The Edge...

- The closer you go to the edge of the network, the more potential intercept points you need to manage, and the more "network aware" your intercept solution needs to become.
- For example, *which* edge interception device would you need to monitor to get traffic for a particular MAC address or IP address?
- What if a user is working via multiple subnets, perhaps on a wired subnet and on a wireless subnet, or some of her traffic is going to the Internet via a proxy server?
- Doing a lawful intercept for even a single user in an environment where instrumentation has been pushed to the edge of the network may require collecting and correlating traffic from multiple interception points, which can quickly become tedious.

The Risk of Architecting for CALEA

- There's a temptation to architect one's network for easy CALEA compliance – there are clearly steps that one could take which would make it easier for a site to become CALEA compliant, or to avoid compliance obligations entirely
- For example, look at the ACE discussion at <http://www.acenet.edu/AM/Template.cfm?Section=HENA&Template=/CM/ContentDisplay.cfm&ContentID=17276>
Given their position, one might be a tempted to stop running the campus border router, to stop working to obtain and light fiber, and to simply rely on a commercial third party to provide Internet connectivity for campus, a step which would likely exempt the campus from CALEA compliance (leaving satisfaction of those obligations to the commercial provider).
- One could take that step, but I think that taking that sort of step would be a huge step **backwards**.

Schools Should Continue to Architect, Build and Run Resilient IP Networks

- Network connectivity is a fundamental academic resource, and one which is too important to outsource in an effort to sidestep compliance obligations.
- Academia has unique requirements for networks, and needs operational experience with networks to keep its network research-related work real and relevant.
- There's also the matter of cost: academia can run its own networks at far lower cost than what can be obtained from a commercial provider. (Or, put another way, given fixed budgets, if you end up outsourcing your network you'll end up getting less capacity for the same expenditure)
- When you run your network yourself, you can also do it right. For example, you can make sure you are fully multihomed.

Multihoming

- While in a simple network model a college or university might purchase connectivity from just one upstream network provider, many higher education institutions know the importance of purchasing network connectivity from multiple upstream providers, a process known as "multihoming."
- For example, a higher education institution might purchase commodity Internet connectivity from both Sprint and Level3, thereby insuring that even if one provider were to experience an outage, the other redundant provider could still provide a path to the Internet.
- Similarly, higher education institutions might purchase commodity Internet transit service from one network service provider, and high performance network connectivity from Internet2, National Lambda Rail, or an aggregator thereof such as a gigapop or regional network.

Why Does Multihoming Matter for CALEA?

- Multihoming can matter for CALEA because network traffic can come in & go out via different network pipes. Because of that, if law enforcement wanted intercept traffic via the upstream ISPs, they'd need to work with (and correlate data from) all of the ISPs which the campus is using.
- What if LE were to instrument just one or two pipes (out of perhaps three or four)?
 - LE might not see any relevant traffic (it might all happen to go via one of the uninstrumented pipes)
 - LE might see only some relevant traffic, or
 - LE might see everything (but not know that that's the case)
 - or they might see only inbound or only outbound traffic (the traffic flows can be asymmetric, coming into the network on one pipe and leaving the network on another)₆₂

Peering at Internet Exchange Points

- Another potential traffic exit is via university peering connections at Internet exchange points (see the list of exchange points at <http://www.ep.net/ep-main.html>).
- Unlike commodity Internet transit connections, or high performance research and education network connections, peering takes place when two networks agree that it is mutually beneficial to exchange customer traffic, and only customer traffic, directly.
- Thus, if two sites exchange traffic via a peering point, traffic between those sites would be:
 - exchanged directly, and so would not touch "the public Internet" via a commodity Internet transit provider, BUT
 - I suspect that LE would still expect that traffic to be able to be lawfully intercepted... but wouldn't this would be an "interconnecting carrier"-ish CALEA-exempt situation? ⁶³

So Where Does "The Internet" Begin?

- So where does the Internet begin? Because the Internet is "just" an interconnected "network of networks," maybe:
 - the point where I begin to **pay** someone else to carry my traffic (but note that peering points would fail that test!)
 - the point where **administrative control** shifts from one entity to another (this might even be at a link between a departmental LAN and the campus backbone, on campus)
 - the point where traffic from one **network address block** leaves that address block and enters a link whose other end has an address controlled by another entity
 - the point at which the **autonomous system number** (see <http://www.uoregon.edu/~joe/one-pager-asn.pdf>) associated with traffic changes from one entity to another (but some ASNs, like AS701, represent phenomenally large aggregations of disparate customers!)
- The law should be amended to clarify this key point.

Overcollection and Minimization

CALEA: What Was Ordered, And ONLY What Was Ordered

- Another important provision of CALEA is that it requires delivery of what was approved for interception, and ONLY what was approved for interception.
- For example, a court may issue an order for either "**pen-register/trap-and-trace**" data or for "**full communication contents**" (what's often referred to as a "Title III" order in the U.S.). The geek way to think about the difference between the two is to think about the difference between doing Netflow, and doing a full packet capture.
- If the court orders the production of **just** flow level data, you **cannot** be lazy and give LE full packet captures instead (if you tried to do so, LE should and hopefully *will* refuse to accept it). You must provide **only** what was ordered.

Minimization

- Another example: if you receive an order specifying the interception of traffic associated with a particular IP address, you cannot respond to that order by simply providing a copy of everyone's traffic – you must minimize what's delivered to be **ONLY** the traffic for the entity specified in the order.
- ATIS-1000013.2007 says a **subject** may be identified by:
 - their IP address (or set of IP addresses)
 - an account session ID assigned to the subject at loginIn others cases, the interception order may specify particular **equipment**, rather than a particular person, via:
 - a MAC address
 - an IP address (or set of IP addresses)
 - a circuit ID or ATM or Frame Relay PVC
- Sometimes it may be hard to suitably minimize traffic, and avoid overcollection using those identifiers. For example...

Firewalls with NAT

- Hardware firewalls are a common feature in many network architectures, and are intended to shelter interior devices from external scans and from attack traffic.
- Some hardware firewalls, in addition to deflecting unwanted external traffic, also do network address translation ("NAT"). Linksys Cable/DSL "routers" are one example of a popular consumer hardware "firewall" device which does NAT.
- When doing NAT, all traffic from a NAT box can be made to share a single public IP address, making it extremely difficult to determine if a publicly observable flow is coming from user A, user B, user C or ...
- Attempting to attribute traffic to a particular user typically requires access to the NAT box's log files (which may not even exist, particularly in consumer environments), accurate time stamps, and the cooperation of the NAT administrator.

Some NAT'ing Firewalls Aren't Sitting In Front of "Just a Few" Folks

- Over time, particularly as worries about breaches involving personally identifiable information have increased, there's been a growing tendency at some universities to NAT entire campuses, potentially putting thousands of individuals behind one (or just a handful) of IP addresses.
- While one might think that this strategy enhances the security of the campus ("Hey! Everything's behind a firewall, we must be safe(r), right?"), a campus-wide NAT actually creates a sort of fate-sharing. All users inherit the reputation of the worst user working from behind the NAT device, and a court order asking for "all network traffic" associated with a single public IP address (which could be the public IP address of the NAT box) might potentially include traffic associated with **thousands** of users.

Post-Hoc Minimization

- When there's absolutely no way to minimize intercepted traffic in advance, it may be possible for post-hoc minimization to be done.
- Presumably that could be done for NAT'd traffic just as post-hoc minimization is used to deal with other tricky mixed traffic scenarios, but the process of teasing out one user's NAT'd traffic from the traffic of thousands of other users would potentially be quite daunting, and if done improperly, could jeopardize the privacy of a large number of innocent users who are also behind that NAT.

Web Caches

- Sometimes an illegal behavior, such as making a threat of violence against an elected official via a web email account or visiting a child porn site, takes place via the web. When that occurs, backtracking normally begin with the IP address (and time stamp & time zone) of the requesting browser.
- Sometimes, however, the IP address that's identified turns out to be the address of a web cache.
- For those who may not be familiar with web cache boxes, they are a type of server which is designed to insure that popular pages don't need to be repeatedly re-retrieved over wide area links – the first time a page gets retrieved it gets saved to memory or disk in the web cache, and then subsequent requests for that page can be serviced from the local copy, eliminating the need to repeatedly request a popular (but unchanged) page from a remote web site.

Backtracking Beyond the Web Cache

- When a web cache is used (just as when NAT is used), it will usually be necessary to have the cooperation of the local administrator to attribute access via that web cache to a particular individual users.
- Most (but by no means all) web cache servers log the use of the web cache so that with a destination of interest and fairly precise time stamps, a web cache administrator can potentially tie a visit made through her web cache to an IP address (and ultimately to a user) by reviewing her logs.
- Log files rapidly grow large, though, so many administrators may truncate or rotate their logs on a periodic basis, which means that an investigator's ability to identify a web site which was visited through a web cache server may go away forever after just a period of days (e.g., when the admin rotates or removes the web cache server's logs).

"X-Forwarded-For" (XFF) Header

- Some web cache products add an X-Forwarded-For header, nominally adding the IP addresses of the client for which a given web request has been made. For example:
X-Forwarded-For: client1, proxy1, proxy2
(see <http://wiki.squid-cache.org/SquidFaq/ConfiguringSquid> and <http://en.wikipedia.org/wiki/X-Forwarded-For>)
- Some things to note about X-Forwarded-For:
 - web caches may be hierarchical, so be aware that you may see multiple IPs addresses on the header...
 - in some cases you may not see any IP address for a given client (e.g., it is possible for a proxy to suppress relevant IP addresses, acting as an anonymizer)
 - the contents of the X-Forwarded-For: header, while often accurate, can be trivially forged (e.g., you may see intentionally misleading data provided in an XFF header)₃

Web Caching and "Hidden" User Access

- Another phenomena associated with web caching is that while one retrieval of a web page may be externally visible, and potentially even associated with a particular local username, there may be **additional** local users who retrieve a page of interest from the local cache, and who thus also gain access to that material (from the local cache) even if they're completely unknown to the ultimate provider of that content (that party won't know they exist since they were serviced from the locally cached copy of the content).
- This is another example of a problem associated with an "Internet Gateway compliance only" model – traffic for a popular page may be delivered from the web cache without ever touching the Internet Gateway. Thus, LE might hypothetically catch one local visitor to an Internet child porn web site while missing another local user who hit those same unlawful pages (but only via the local web cache). ⁷⁴

Shell Accounts on Shared Systems

- A shell account is username/password pair which allow someone to use ssh or telnet to login and get a % or \$ prompt. At that point the user can type in Unix commands. Shell accounts are another example of a traffic source which may be hard to minimize as required by law.
- For example, some shell servers may have **thirty or forty thousand accounts**, and it may be common to see thousands of those users connected at any one time.
- Attributing realtime traffic emitted from a large host of that sort to a particular user can be hard, although some large hosts support identd (see RFC1413), which may provide some (but not much!) assistance when it comes to tying traffic to a particular user. Application accounting logs (if kept), and application content may also help to identify the source of traffic, although one must be wary of spoofing. 75

Dynamic Addressing and Timestamps

- Another complication is dynamic addressing. In general, when a desktop or laptop system connects to a university network, it uses DHCP (dynamic host configuration protocol) to get an IP address, to learn its broadcast address/netmask/default route, the right name servers to use, etc.
- When dealing with dynamic addresses, as when dealing with traffic that's flowed through a NAT or web cache, having accurate time stamps (with time zone information!) can be absolutely key to correctly identifying a party of interest.
- Things can get very complicated if time synchronization is poor, connect times are brief, and IP utilization is high with little idle time between sessions.
- Sometimes users of dynamic addresses authenticate, which can be a big help, but other times users of dynamic addresses may not do so.

Dynamic Addresses Without Auth

- While authentication will usually be required for hosts connecting via dialup, or for hosts connecting via wireless, hosts connecting via a hardwired 10/100/1000 Mbps ethernet connection typically do NOT require authentication.
- The rationale behind not requiring authentication for all dynamic addresses is that:
 - we know where a given ethernet jack is physically located, so...
 - if you know a physical location you should usually be able to identify the user.
- For example, if a connection is coming from Jack #11 on Switch #34, your network wiring documentation might tell you that connection goes to Williams Hall Room 178, and once we know that location, we could then check our files determine that that office is being used by Bill Smith.

Network Jacks and Attributability

- The assumption that if we know a jack's location ==> we know who's using that jack breaks down when:
 - it is applied to shared public spaces, such as classrooms, libraries, etc., where anyone can plug a host in without auth
 - per-port documentation is wrong/out-of-date/non-existent
 - wiring closets are insecure, and someone can add a new connection, or move existing connections without authorization;
 - when wiring runs are insecure, and someone can "split" an existing connection with a hub, switch, optical splitter, etc.
- True anecdote: student is observed engaged in behavior inconsistent with the AUP; student's port get's turned off. Student moves to roommate's port. That port goes off, too. Student comes back on a third time, seemingly on next door neighbor's connection. ??? Visit site, see an ethernet cable running *out the window* from the room of interest to next door

CALEA and Advanced Protocols

CALEA Is NOT Supposed to Discourage Advanced Protocols

- 47 USC 1002(b) makes it clear that

"This subchapter does not authorize any law enforcement agency or officer

[...]

(B) to prohibit the adoption of any equipment, facility, service, or feature by any provider of a wire or electronic communication service, any manufacturer of telecommunications equipment, or any provider of telecommunications support services."

- Everyone recognizes that Internet innovation is critical to American competitiveness....

That's Great, But...

- Does anyone have an out-of-the-box CALEA solution which handles even the following half dozen areas?
 - 1) IP multicast
 - 2) IPv6
 - 3) jumbo frame support
 - 4) non-layer 3 wide area traffic
 - 5) non-local authentication and authorization (e.g., Shibboleth-based access control and authorization), and
 - 6) extreme data rates?
- Let's look briefly at these advanced protocols and architectures, and why they're relevant to lawful intercept.

1) IP Multicast

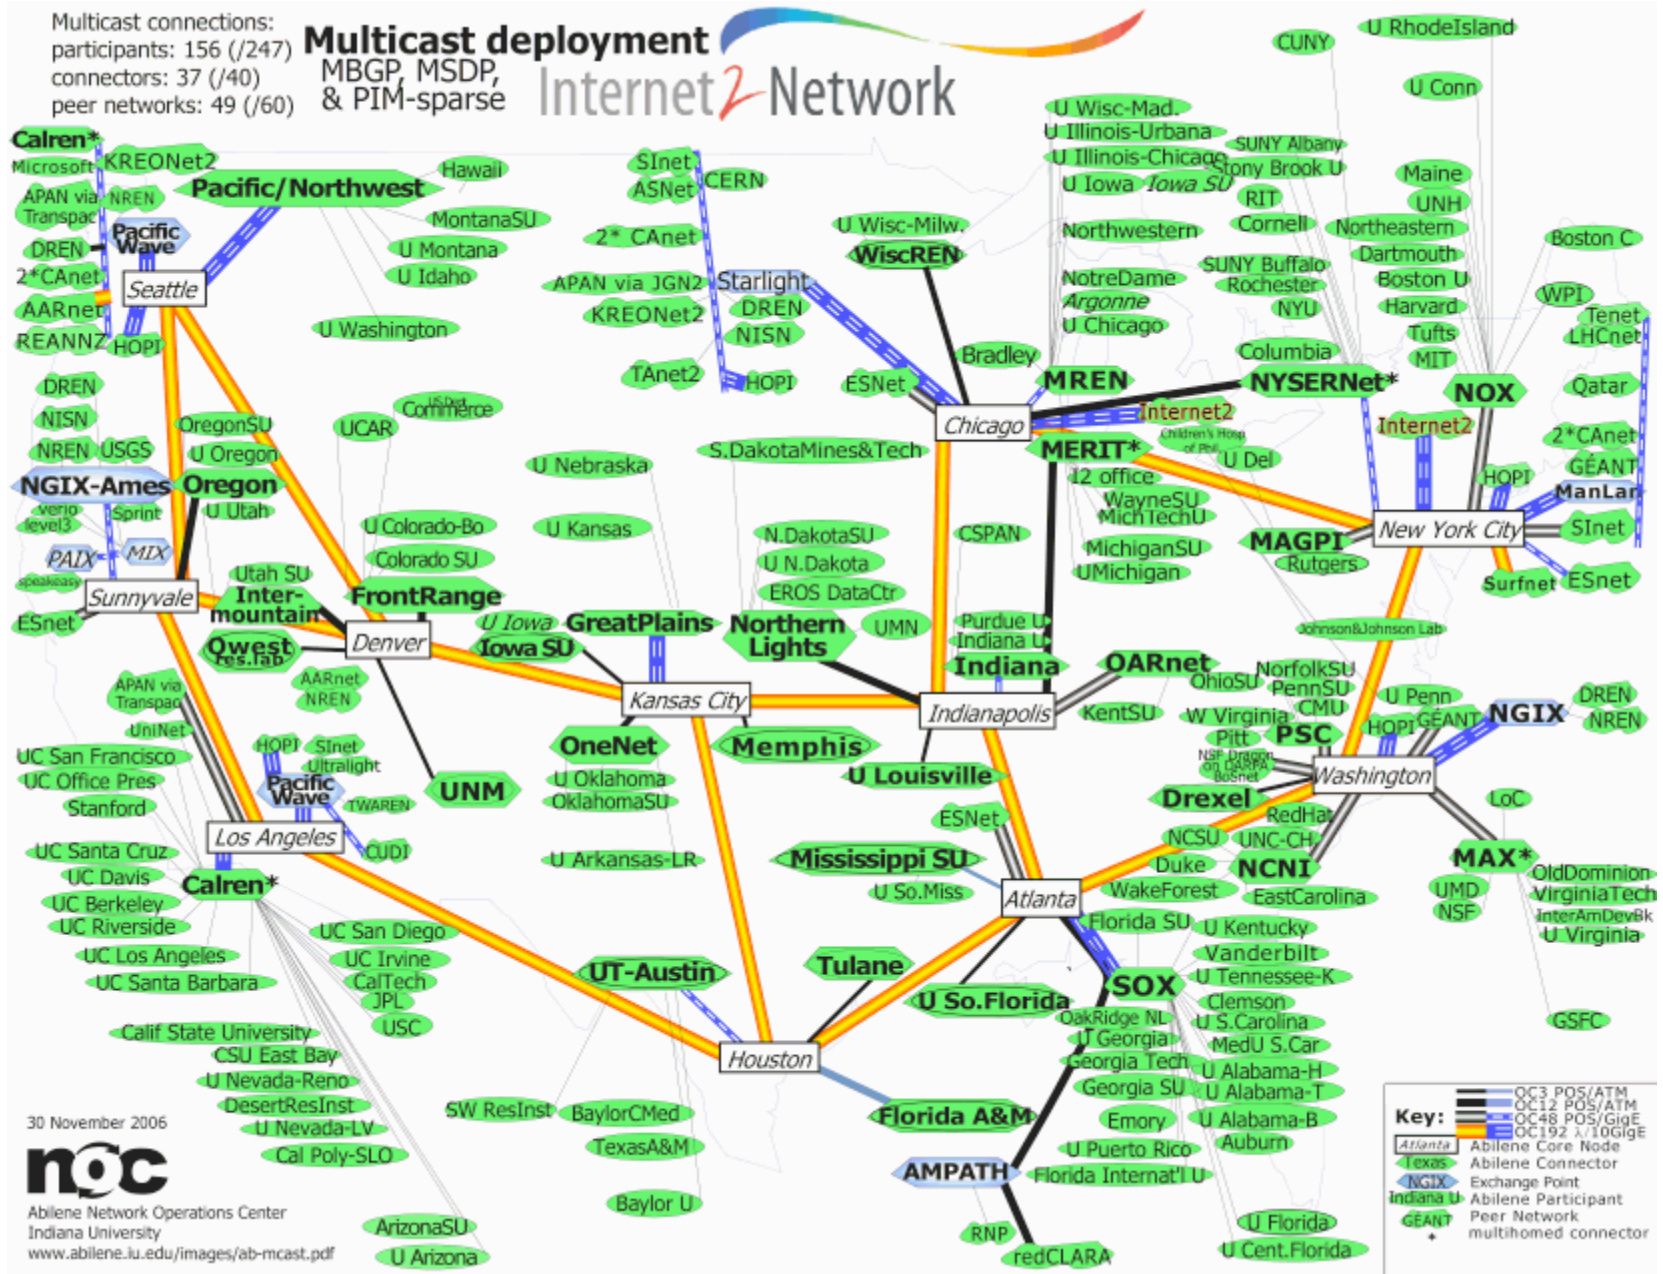
- Normally, a network session is "unicast" between one source and one destination. For example, you might request a streaming video broadcast from a news web site, and that web site would deliver that video on demand from their server to your workstation. If another person wanted to see that same video, the news web site would then create a second independent video stream, iterating for each viewer interested in a particular video.
- IP multicast is an alternative approach which allows a server to distribute a **shared** network stream which can serve one user, or a dozen, or a thousand, or a million. Because IP multicast is so efficient, a content originator (like an online news site) can afford to distribute TV quality video (MPEG1) instead of postage-stamp-sized herky-jerky video.
- For a technical overview of multicast, see www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/mcst_sol/mcst_ovr.pdf⁸²

IP Multicast and CALEA

- Why mention IP multicast today? Well, IP multicast is rare in the commercial ISP space but **quite common** in higher ed (see the map on the next slide), and I don't think anyone has given much thought to how IP multicast traffic would be handled under CALEA. Relevant issues might include:
 - IP multicast content is typically delivered via a network tree which gets built to a local router (rather than directly to the interested party), so it may be hard for gateway-based lawful intercept software to recognize that particular IP multicast content is associated with a user of interest
 - multiple participants (some who may be the subject of an interception order, and others who may not be the subject of an interception order) may be accessing or contributing content to the same IP multicast group – what can/should/needs to be done then to protect 3rd party users' privacy?

Multicast connections:
 participants: 156 (/247)
 connectors: 37 (/40)
 peer networks: 49 (/60)

Multicast deployment MBGP, MSDP, & PIM-sparse Internet2 Network



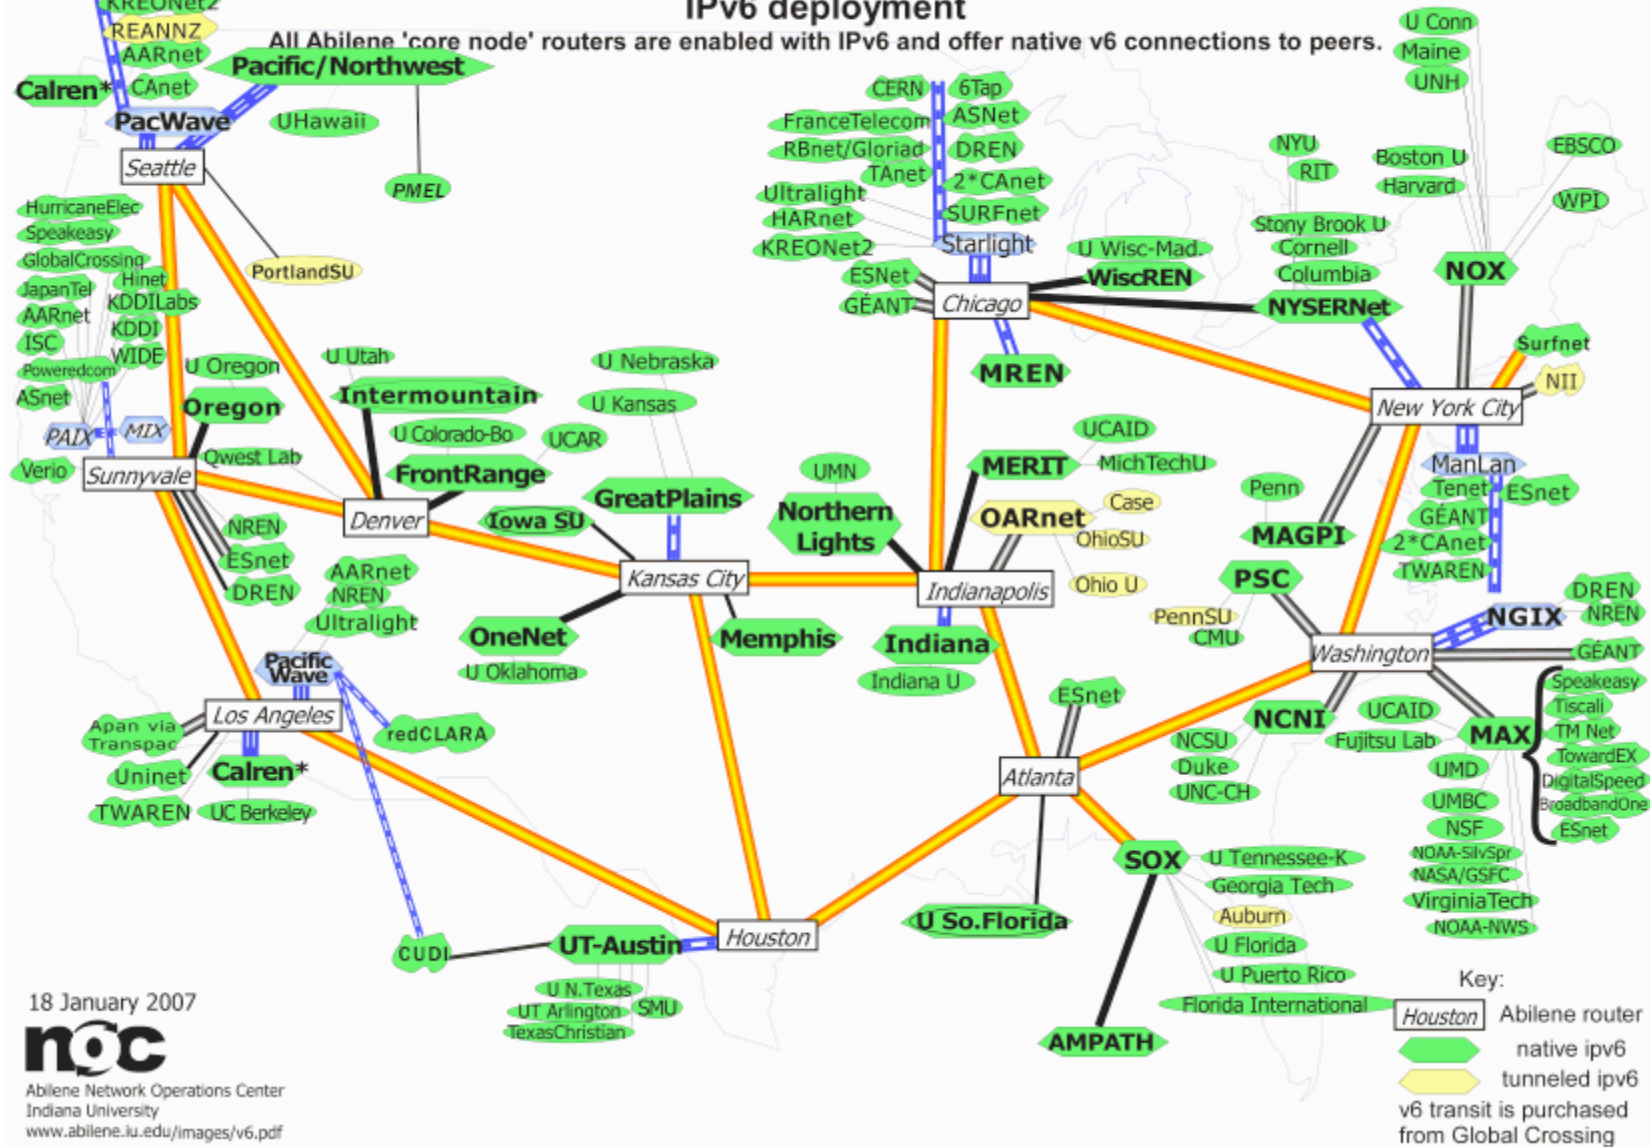
30 November 2006
n9c
 Abilene Network Operations Center
 Indiana University
 www.abilene.iu.edu/images/ab-mcast.pdf

2) IPv6: Hey, It's Real Too, Folks!

- Most network traffic on the Internet today uses IPv4, but it is projected that within 4 to 5 years we will exhaust available IPv4 address space (see: <http://bgp.potaroo.net/ipv4/>)
- IPv6 modifies the traditional IPv4 packet format in numerous ways, the most important of which is that with IPv6 network addresses go from 32 bits to 128 bits, thereby dramatically increasing the number of addresses available for allocation.
- Numerous operating systems are IPv6 aware today, including Microsoft Windows Vista, Apple Mac OS X, Linux, Solaris and others. Numerous networks carry native IPv6 traffic in the United States and overseas, including higher education research and education networks such as Internet2 (see the map on the next page), and all .gov core networks must be ready to pass IPv6 traffic by 6/08 (see http://www.cio.gov/documents/IPv6_Transition_Guidance.doc).

Internet² Network IPv6 deployment

All Abilene 'core node' routers are enabled with IPv6 and offer native v6 connections to peers.



18 January 2007
noc
 Abilene Network Operations Center
 Indiana University
www.abilene.iu.edu/images/v6.pdf

<http://www.abilene.iu.edu/media/i2network/documents/v6.gif>

ATIS-1000013.2007 and IPv6

- Checking ATIS-1000013.2007, "Lawfully Authorized Electronic Surveillance (LAES) For Internet Access and Services," issued April 2nd, 2007, the string "IPv6" appears 6 times in 75 pages, with the only substantive reference appearing on PDF page 16 where it is mentioned that "The Subject Domain, Access Network, Intermediate Network and ISP Network may be using IPv4, IPv6, or any combination of IPv4 and IPv6 involving translation or tunneling."
- That implies an ATIS-1000013.2007-compliant CALEA implementation should be able to deal with (for example), both native IPv6 frames, and IPv6 traffic tunneled via protocols such as Teredo (see <http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/teredo.msp>).
- Looking at the marketplace, I'm not seeing any commercial CALEA product which fully supports IPv6 intercepts, nor does OpenCALEA as of v0.5.

3) Jumbo Frames

- Normal ethernet frames are 1500 bytes long, but that's far too short for optimum performance on long distance, high bandwidth networks characteristic of advanced networks in higher education. Currently both Internet2 and the Federal Joint Engineering Team (JET) recommend use of a 9K MTU:

www.nitrd.gov/subcommittee/lsn/jet/9000_mtu_statement.pdf

noc.net.internet2.edu/i2network/documentation/policy-statements/rrsum-almes-mtu.html

- Will CALEA interception devices and CALEA delivery links be engineered to accommodate jumbo frames? Lawful interception gear may or may not be prepared to even see/intercept 9K (or larger!) frames, and simple delivery of 9K MTU traffic may also pose issues.

ATIS-1000013.2007 and MTU Issues

- ATIS-1000013.2007 addresses fragmentation in Appendix D.4 "IC-APDU Fragmentation and Optimization (Informative)" (which is not officially part of the standard)
- Because jumbo frames are frequently associated with high throughput connections, practical issues associated with the fragmentation of high bandwidth jumbo frame traffic may be non-trivial to resolve.
- Heck, given encapsulation overhead, even transferring 1500 byte packets will require fragmentation and reassembly.
- It would be great if a normative (rather than just informative) statement on the fragmentation and delivery of jumbo frames could be made part of a future version of this standard.

4) Non-Layer 3 Wide Area Traffic

- While there's generally an assumption that wide area Internet traffic will be routed at layer three, under some circumstances that may not be true. This is relevant to LI because standards typically focus on the packet transfer function (PT-F) at layer 3. For example, ATIS-1000013.2007 specifies that *"Only those network elements that recognize the Layer 3 packet structure (i.e., IP header fields) and handle the packets can perform the packet transfer function. A Layer 2 network (e.g., an Asynchronous Transfer Mode (ATM) network) or a Layer 1 network (e.g., a Time Division Multiplexing Network) may be used to transport the Layer 3 packets from the customer premises to the ISP network. Network elements in the Layer 1 and Layer 2 networks cannot recognize the Layer 3 packet structure and do not use the Layer 3 packet header information. These Layer 1 or Layer 2 networks do not perform PT-F."*

Could You Have A Layer 2 "Internet Gateway?"

- So what happens, then, if traffic exits a university campus "Internet Gateway" at "just" L1 or L2?
- For example, consider National Lambda Rail's FrameNet service, which has been described as "a large ethernet switch fabric" (where "large" in this case refers to a nationwide network interconnecting institutions from one side of the country to the other) – that's clearly a wide area L2 network.
- Where would the "Internet Gateway" be on that sort of network? (Or would this just be another example of a "private network," albeit an extraordinarily large one, by definition?)

How About Layer 1?

- Going even further down the OSI model, what of wide area lambda (optical wave-based) connections, perhaps connecting just one host (or a small subnet of hosts) at an institution with a host (or a small subnet of hosts) at another institution? Is this another example of a sort of nationwide "private network?"
- Does it matter if network traffic can be shifted "on the fly" from a packet mode connection (where it might be subject to lawful interception) onto a dynamic wave (where it might not be)? Dynamic traffic shifting **IS** being done today, by the way -- see slide 17 in FermiLab's "Lambda Station" Joint Techs talk, <http://events.internet2.edu/2006/jt-albuquerque/sessionDetails.cfm?session=2440&event=243>
- Soon a dynamic lambda-based connection will be part of the connection package most Internet2 connectors buy!

4) Federated Authentication

- We're all awash in site specific usernames & passwords, so if you trust my school (or other institution), why not agree to:
 - let them authenticate me via a username & password,
 - let them share (just) relevant attributes about me with you
 - then, as may be appropriate, give me access to resourcesSee, for example: <http://shibboleth.internet2.edu/> and <http://www.incommonfederation.org/>
- Some examples of federated authentication:
 - access to a proprietary online database (may be limited by contract to just faculty in a particular department)
 - access to the now-legal Napster music service (perhaps limited to just undergraduate students at a university)
 - wireless access to the Internet while visiting another site
- It is the last type of case, which is tricky for CALEA – you may know **what** I am (e.g., I'm a faculty member from Alpha University), but not have any idea **who** I am.

Example of What A User Might See...

Select an Identity Provider

The Service you are trying to access requires that you identify yourself. Please select an identity provider from the list below.

NOTE: If you need to sign up for a new account, select **ProtectNetwork** from the InCommon list and you can register for a personal account there using a valid e-mail address.

Choose from the list:

Federation

US Higher Education
UK Federation
MAMS Testbed Federation
All Sites

Institution

Internet2
New York University
Ohio State University
Ohio University Main Campus
OhioLink
ProtectNetwork
Stanford University
Stony Brook University
The Pennsylvania State University
The State University of New York at Buffalo

Isn't Shib-Mediated Auth Just An Example of a CALEA Access Network?

- No. ATIS 1000013.2007 describes an Internet Access and Service Model which includes "Access Networks" potentially performing a registration function, REG-F, and a resource function, RES-F, e.g., see figure 1 in that spec. The Packet Transfer Function (PT-F), however always takes place via the ISP network. That model runs into trouble when **authentication** occurs via one's "home institution" but **network traffic** only flows via the facilities of the institution which someone happens to be visiting.
- Yes, retrospectively the home institution could identify a person using the network at the visiting institution, but at that point the identification would be retrospective, and too late -- a party subject to a LI order might have had access to the Internet without his or her traffic being ID'd & monitored.

6) Extreme Data Rates

- Unlike dialup/POTS/cell links, broadband connections are (by definition) associated with higher than normal data rates. Typical **consumer** broadband connections might range from ISDN speeds of ~128Kbps to 15-20 Mbps (for things like cable modem connections and fiber-to-the-home services).
- In the case of **higher education**, however, it is routine to see **far** faster connections, including:
 - 100/1000Mbps on the local area network and
 - speeds up to 10Gbps to Internet2 or to the regular Internet
- Because these speeds exceed the speeds which are commonly/affordably available to LE via the commodity Internet, there exists the possibility that a subject connecting via a high speed connection may be able to generate network traffic in excess of what LE can cost effectively transport from a local interception point to an offsite LE wireroom.

Two Excerpts from DOJ Report 06-13

- "During our site visits, many law enforcement officials noted that CALEA addresses what carriers need to provide to law enforcement agencies without addressing how **data is delivered**. For example, CALEA does not address whether carriers can use digital or audio phone lines to deliver the audio portions of intercepts. As a result, the delivery method of intercepted data varies by carrier. Due to the various delivery methods, law enforcement agencies must purchase additional equipment to receive the intercepted data from a carrier. The four delivery methods are **dial-out, VPN, frame relay, and T-1 lines**." [emphasis added]
- Discussion of non-VPN delivery methods may indicate a failure to consider the realities associated with higher education's DS3 (45Mbps), 100Mbps, OC3 (155Mbps), OC12 (622Mbps), gigabit, and 10gigabit-class connections.

Heck, LE Doesn't Want to Even Buy T1's

- "A law enforcement official in California stated that his office was informed by two in-state wireline carriers that they are CALEA-compliant but law enforcement would need to build a T-1 line to each of the carriers' switches. The law enforcement official explained that this concept is unreasonable considering his agency's jurisdiction has about 95 switches from one carrier and about 130 switches from the other. Therefore, it would cost his agency about \$292,500 to install T-1 lines to each of the switches. This scenario would not be cost beneficial to his agency because a T-1 line is only used for wireline intercepts, and approximately 70 percent of this agency's wiretaps are performed on wireless phones." [from DOJ Report 06-13]
- $\$292,500 / (130 + 95) = \$1,300/\text{line}$ (e.g., that's just installation); monthly reoccurring costs would add another \$575-1,800/line

It's Not Just The Pipes, It's Also The Devices On The End of The Pipes

- Just as LE might balk at the cost of purchasing connectivity sufficient to transport sub-gigabit, gigabit, or multi-gigabit class traffic, LE may also find the cost of purchasing **routers, storage systems and analysis engines** sufficient to absorb and process sustained high speed real time traffic flows to be cost prohibitive.
- In particular, **10gig class routers and router interfaces have a non-trivial cost**, and even assuming you need to store "just" traffic averaging a gigabit per second per day, that implies you'd need $(10000000000 \text{ bps}) / (8 \text{ bits/byte}) * 86400 \text{ sec/day} \Rightarrow 10.8 \text{ TB/day}$ or 324 TB/month
- Thus, I believe there is, or should be, a real question when it comes to LE's ability to sustainably deliver and archive CALEA intercept traffic for high bandwidth intercepts.

"But Are Sustained Gigabit Class Flows Actually Seen???"

- Excellent question! We know from the Internet2 daily netflow report (see <http://netflow.internet2.edu>) that a typical distribution of bulk TCP flows looks like:

Percentile	Throughput (b/s)
5	1.456M
10	1.552M
50	2.825M
90	9.823M
95	15.58M <== 95% of all bulk flows
99	50.21M are ~15 Mbps or less
99.9	992.4M
99.99	1.065G
99.999	2.147G

On the Other Hand, Don't Forget About Multiple Parallel/Concurrent Flows...

- Because a single user (including potentially a subject of LE interest) might generate multiple parallel/concurrent flows, it would be an error to just look at the average single stream throughput when attempting to provision the capacity potentially required to transport intercepted traffic.
- For instance, assume a user of LE interest was transporting communications via encrypted binaries over a private NNTP network, or exchanging steganographic messages hidden within BitTorrent downloads (just to mention a couple of possible examples) – in those cases **an agency might genuinely need the ability to transport and store sustained aggregate flows at the gigabit level or beyond for a single subject of interest at a single site.**

Could LI Traffic Be Minimized Through Local Policy?

- Confronted with sustained high volume traffic, one LE alternative to building out "budget-busting" wide area, high capacity circuits might be encouraging "traffic minimization" via local policy. How might this be done? Two possibilities:
 - DMCA notices (assuming distribution of unauthorized copyrighted content is involved), or
 - provider traffic management via traffic shapers or provider charges for "excess" bandwidth consumption.
- Obviously strategies of this sort could not be used just on a target of interest, they'd need to be deployed site-wide.
- If a site's traffic **could** be throttled down from gigabit or even 100 Mbps down to some lower bandwidth level, that translates to that much less traffic which LE would need to potentially transport & process (but would LE ever ask this?)

Buffering Flows

- Another potential approach, at least for short duration "bursty" high bandwidth flows, would be to buffer those flows, and to then do "near real time delivery" of the flows in the background over a slower speed circuit with LE's consent
- For example, if a subject of interest averaged 15Mbps worth of traffic over the course of an hour, one option might be to locally buffer that traffic and then deliver the traffic at a lower rate, say, 1.5Mbps (T1 speed), over a ten hour period.
- This strategy fails if:
 - genuine real time (unbuffered) access is required,
 - there are orders of magnitude disparities between the capacity of the subject's connectivity and LE's connectivity
 - sustained average traffic volumes exceed the delivery capacity of the LE circuit (even with compression, etc.), or peaking loads exceed the capacity of deployed buffers

On-Site/Local Wire Rooms

- Another possibility which LE might consider would be the build out of local wire rooms to handle processing of high bandwidth flows without the need to purchase wide area, high-capacity, circuits. Once traffic had been transported to the secure on-site wire room, it could then be summarized or otherwise processed, including potentially being written to portable media for offline transport to analysis resources.
- This strategy begins to break down as the number of high bandwidth sites (and thus the number of local wire rooms required) gets large.
- This approach also assumes either a persistent interest in a particular site (justifying permanent facilities), or perhaps the ability to easily deliver a portable wireroom (conceptually imagine something like a preconfigured semi-trailer-based mobile wire room, which could be driven up and plugged in).

Dynamic Circuits

- Conceptually, LE or a university might also consider using a dynamic circuit to deliver high volume CALEA intercept traffic to a LE-designated location.
- What do I mean by a dynamic circuit? Well, that might be a MPLS VPN or other tunnel running over the institution's existing wide area connection (assuming sufficient unused capacity exists), or the use of emerging dynamic lambda facilities (assuming those facilities are available and not already committed).
- Those sort of strategies might allow intercept traffic to be backhauled to one of a small number of regional analysis centers, thereby eliminating the need for LE to either build out local wirerooms at each high bandwidth site, or the need for LE to purchase and maintain dedicated high bandwidth physical circuits to each site.

Conclusion and Recommendations

CALEA Deserves Statutory Cleanup

- CALEA is showing its age, and trying to make a 1994 law that was aimed at traditional telephone services fit broadband Internet providers hasn't been, and never will be, very successful. Federal agencies, assuming they want CALEA to unambiguously cover broadband Internet access, including broadband in higher ed, should do the right thing and pass the necessary amendments to the existing law.
- Be sure to start with the basics – what's the Internet? What's a private network? What precisely is covered? What's not?
- And please, admit reality: just as encryption was identified early on as something where providers simply couldn't do much, recognize that there may be other corner cases as well, and be realistic about what a provider can and can't do in good faith. Advanced protocols should be exempted unless/until they begin to be exploited.

Help Sites To Do The Right Thing

- Governments could do a far better job of showing what they need and expect by building a collection of case studies showing how various sites met CALEA's requirements to the government's satisfaction. Show **real life examples** (with the consent of the profiled sites) illustrating how typical small colleges, large universities, statewide networks, etc., were all able to successfully comply with the law, and what it cost for them to do so.
- Negotiate arrangements with standards bodies so that privately developed CALEA-related industry standards can be freely downloaded and do not need to be purchased one-off by every site that's trying to comply with the law.
- Negotiate standardized discounted pricing for CALEA lawful intercept products and services so that RFPs aren't needed.

Speaking of Pricing: Financial Support Is Necessary

- Supporting lawful intercept isn't cheap. If lawful intercept capabilities are important, the government needs to step up and financially support that requirement – don't leave that burden on the shoulders of colleges, students and families.
- The original statute recognized the importance of financial support, and provided funding to underwrite the work which was required, but that funding dried up while CALEAifying the telephone system, and now funding isn't available for Internet providers (and colleges and universities) who must become CALEA compliant now. That **MUST** be corrected.
- The costs which law enforcement face also need to be recognized. It does no good to require providers and universities to build out lawful intercept capabilities if law enforcement literally can't afford to take advantage of the facilities we're collectively being made to install.

Give Broadband Providers More Time

- Recognizing that it took the better part of ten years to get most of the wireless and wireline telephone infrastructure CALEA-ready, it is unrealistic to expect broadband service providers and universities to be able to become CALEA-compliant at the drop of a hat.
- For context, note that the extension of CALEA to broadband providers was only contemplated in July, 2003 (see <http://news.com.com/2100-1028-5056424.html>); ACE vs. FCC wasn't decided until June, 2006; and the broadband CALEA technical specification, ATIS-1000013.2007, wasn't approved until April 2, 2007 -- yet all facilities based broadband providers were to be compliant by May 14, 2007!
- If you look at the implementation of CALEA as stretching back to the "dark ages" of 1994, that seems like an awfully long time to be working on building CALEA out, but in the case of broadband, we've only had a matter of **months!** ¹¹⁰

Leverage Lawful Intercept to Help Attack Cybercrime

- We're currently losing the war on cyber crime, and an important part of changing that dynamic may be beginning to more broadly use lawful intercept against cyber criminals.
- Currently lawful intercept is such a politically charged "third rail" (e.g., touch it and die), and so lumbering, many law enforcement officers won't even consider employing it, and thus cyber criminals are able to work online with impunity.
- Lawful intercept has long been associated with foreign intelligence and combating drug lords, but it needs to also be applied appropriately to the king pin cyber criminals who are destroying the usability and stability of the Internet, too. Existing laws (see 18 USC 2516) should be amended to allow the judicious use of lawful interception when investigating major DDoS and botnet-related cases, too. 111

Thanks!

- I'd like to conclude by acknowledging the thoughtful and detailed comments I received from a number of people who reviewed a draft version of this talk, including Jack Bates, Steven Bellovin, and Neil Schwartzman, as well as a number of additional individuals who either provided comments explicitly on a not-for-attribution basis, or who provided comments but didn't confirm whether they'd like to be publicly acknowledged or not. In any event, the content of this talk remain solely my responsibility.
- And with that, thanks for the chance to talk! Does anyone have any questions?
- Reminder: an expanded version of this talk is available at: <http://www.uoregon.edu/~joe/calea-requirements/>