Security Implications of Applying the Communications Assistance to Law Enforcement Act to Voice over IP

Steven Bellovin, Columbia University
Matt Blaze, University of Pennsylvania
Ernest Brickell, Intel Corporation
Clinton Brooks, NSA (retired)
Vinton Cerf, Google
Whitfield Diffie, Sun Microsystems
Susan Landau, Sun Microsystems
Jon Peterson, NeuStar
John Treichler, Applied Signal Technology

June 13, 2006

1 Introduction

For many people, Voice over Internet Protocol (VoIP) looks like a nimble way of using a computer to make phone calls. Download the software, pick an identifier and then wherever there is an Internet connection, you can make a phone call. From this perspective, it makes perfect sense that anything that can be done with the telephone system — such as E911[1] and the graceful accommodation of wiretapping — should be able to be done readily with VoIP as well.

[1]Enabling E911 for VOIP is complex and will involve new protocol development or enhancement. In the long run, VOIP implementation of E911 could be superior to its manifestation in conventional mobile and wireline telephone systems.

This simplified view of VoIP misses the point of the new technology. The network architectures of the Internet and the Public Switched Telephone Network (PSTN) are substantially different. Lack of understanding of the implications of the differences has led to some difficult — and potentially dangerous — policy decisions. One of these is the recent FBI request to apply the Communications Assistance for Law Enforcement Act (CALEA) to VoIP. The FCC has issued an order for all "interconnected" and all broadband access VoIP services to comply with CALEA (without issuing specific regulations on what that would mean). The FBI has suggested that CALEA should apply to all forms of VoIP, regardless of the technology involved in its implementation[17].

Some cases — intercept against a VoIP call made from a fixed location with a fixed Internet address[2] connecting directly to a big Internet provider's access router — are the equivalent to a normal phone call, and such interceptions are relatively easy to do. But if *any* of these conditions is not met, then the problem of assuring interception is enormously harder.

In order to extend authorized interception much beyond the easy scenario outlined above, it is necessary either to eliminate the flexibility that Internet communications allow — thus making VoIP essentially a copy of the PSTN — or else introduce serious security risks to domestic VoIP implementations. The former would have significant negative effects on U.S. ability to innovate, while the latter is simply dangerous. The current FBI and FCC direction on CALEA applied to VoIP carries great risks. In this paper, we amplify and expand upon these issues.

2 Briefly: What is VoIP?

It is useful to begin with a brief explanation of VoIP, which is not one service but rather a multitude of possible services.

VoIP is an application conveying real-time audio information such as human voice, in a manner emulating traditional telephone service. VoIP relies on the fundamental principle of Internet architecture that any computer with an IP address can send whatever data it is instructed to to any other

[2]Internet address, usually called IP, or Internet Protocol, address, is a unique number that devices use to communicate across a computer network. All urls, for example, translate into IP addresses; www.nsa.gov is 12.110.110.204, while www.pm.gov.uk is 194.201.189.210.

computer with an IP address. Many users are familiar with the client/server network architecture which has the client (a PC, a PDA, a cellphone) sending requests to another computer on the network called a server. Although VoIP can work in this mode, VoIP is quite flexible, and it does not need the client/server model. Instead VoIP traffic is commonly sent peer-to-peer — that is, from one endpoint computer and its user to another. VoIP only requires an Internet connection and a program on the endpoint computer capable of encoding and transmitting speech.

Much of the significance of the Internet is the way in which it supports a mobile lifestyle, and that leads to a slight complication for VoIP: Internet users do not necessarily know the IP address of the person they seek to contact for a VoIP conversation. Given the nature of the Internet, which enables, and, indeed, encourages mobility, IP addresses are, more often than not, allocated dynamically (that is, each time the computer is connected to the Internet); users may migrate between multiple environments (the office, the cafe, the train station, the hotel lobby). As such, almost all VoIP systems have an associated rendezvous service, whose purpose is to take a familiar identifier, a telephone number, a screen name, or an email address, and transform it into the specific IP address of the computer where the designated user can currently be reached.

Once the IP address has been established, the data connection — the conversation — can travel peer to peer. Consider the VoIP network shown in Figure 1. Alice and Bob are both currently connected via the ISP C using router R1 and ISP D using router R2, respectively. Alice, however, uses VoIP Provider 1, a customer of ISP A, while Bob gets his service from VoIP Provider 2, a customer of ISP B. Both Alice and Bob travel and thus are in varying locations; they connect via different ISPs without changing their VoIP providers.

Two of the best-known VoIP service providers demonstrate the variety of models available. Skype, which builds on the technology of the Kazaa peer-to-peer file-sharing application, allows computers to connect with one another free of charge using Skype-registered screennames. In contrast, Vonage, a Session Initiation Protocol (SIP)-based service, is primarily a PSTN interworking application, permitting computers to dial out to the Public Switched Telephone Networks (PSTN) using traditional telephone numbers; it does so at a cost that is competitive with existing local and long-distance service. In addition, AOL has integrated VoIP with its popular AIM instant messaging system. Further deployments of VoIP will be integrated into a

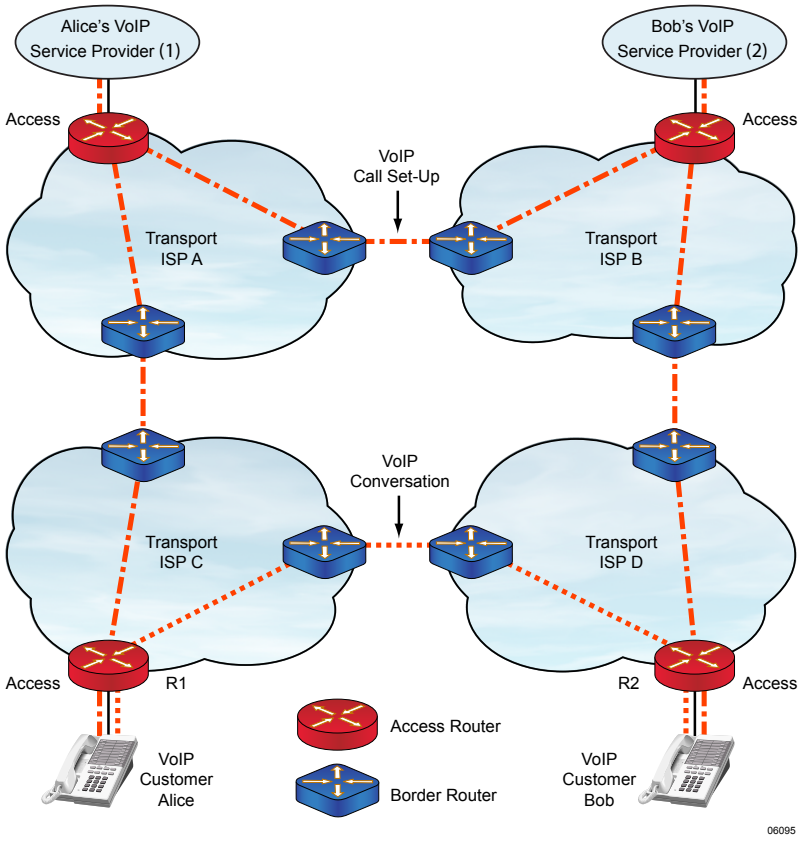

Fig. 1: Alice and Bob using VoIP

host of Internet-enabled services, including instant messaging, Internet gaming, commercial web pages (for customer support), and enterprise private branch exchange (PBX) replacements for IP-enabled offices. In the future, it is very likely that VoIP will become more prevalent than now.

3 How Wiretapping Works

Pre-CALEA wiretapping was conducted at a point along the local loop, the pair of wires running from the local telephone switch to the subscriber's phone. Early wiretaps were literally that, connections made to wires on telephone poles. Later they migrated onto the premises of the switch, which is typically housed in the telephone company's *central office,* also known as the *local exchange.* The technique matured into one employing two elements: the loop extender and the friendly circuit (the law-enforcement wiretapping line). The loop extender was the actual tap creating a logical fork in the subscriber's local loop. Rather than install earphones and a tape recorder at that point, however, the signal was routed to another phone line (the "friendly" line) and transported to a more convenient and secure location for monitoring.

A local-loop wiretap has the capability of receiving all and only the information that passes over the local loop. If the subscriber has call forwarding, the call will be forwarded directly from within the switch, never reaching the local loop, and thus will not be available to the wiretap. If the subscriber has caller id, the numbers of calling parties will be available to the wiretap; if the subscriber does not have caller id, they will not, even though the numbers are available to the switch. The FBI named this collection of shortcomings of local-loop wiretaps the "digital telephony problem." The Bureau insistently lobbied from the late 1980s until its success in 1994 for a law requiring the telephone companies to build wiretapping into their central office switches.

CALEA wiretaps

CALEA requires that the communications infrastructure be made wiretap ready. In essence, the conference calling capabilities of switches are adapted to turning wiretapped calls into conference calls with an unacknowledged silent listener. All of the information available to the switch — call forwarding information, speed call lists, true caller identities (beyond those offered as caller ID) — is placed at the disposal of the wiretap. Much of this information

either never appears on the local loop or, like a speed call association, need not appear during a call to which it applies.

By requiring that digitally-switched networks be built in accordance with federal specifications for wiretapping, CALEA changed the design process. Disagreements between the telephone companies and the FBI resulted in lawsuits and a delayed implementation of the law. Controversy continues, and delays are not the only cost. A recent report by the Office of Inspector General in the Department of Justice observed that CALEA taps are expensive, "[W]hile we found that fees vary widely, a wiretap with CALEA features costs law enforcement approximately $2,200, according to law enforcement officials and carrier representatives,"[18, p. xiii]. Yet applying CALEA to the centralized architecture of the PSTN is a piece of cake compared to applying the law to the decentralized architecture of the Internet. One instance of the difficulties is that all central office switches must conform to a single technical standard, which makes the implementation of CALEA relatively straightforward. There are no such general standards for VoIP, which can be implemented in a variety of ways. The difficulties that have arisen in the PSTN environment are a harbinger of the problems likely to arise in applying CALEA to VoIP.

Roving Wiretaps

Traditional wiretaps name both the subject and the particular phone number to be tapped. In the case of a target deliberately eluding wiretaps through switching telephones, the Communications Privacy Act of 1986 allows "roving" wiretap orders: orders in which the telephone number does not have to be specified on the wiretap warrant. This enables wiretapping on such things as banks of payphones. From an implementation standpoint, we note that nonetheless wiretap law requires *minimization* of communications not subject to the wiretap order [15, pp. 325-6]. In effect, this means that the payphone wiretap would be activated only when the wiretap target was actually using that particular payphone.

Let us consider Alice and Bob in more detail. When Alice calls Bob, her VoIP phone sends a message across the Internet to her VoIP provider, which contacts Bob's VoIP provider, which in turn notifies Bob. (The flow of the call setup messages is shown via dashed lines.) The actual data flow of the phone conversation though, goes directly between the two (dotted line).

Suppose we are trying to wiretap Alice's calls to Bob. The obvious points to do the tapping are access routers R1 and R2 (to the extent that there is

an Internet analogy of the local exchange, these would be the access routers). However, neither router knows who Alice or Bob are; instead, it is the two VoIP providers who do. For the tap to succeed, R1 or R2 would have to receive a "start recording" instruction from one of the two VoIP providers. But these providers can be located at arbitrary places on the Internet, and they need have no business or technical relationship to any ISP other than their own. In fact, they could easily be located in and owned by foreign (and even hostile) countries. How can Alice's ISP trust such a wiretap request?

If Alice's VoIP Provider is owned by her ISP (that is, ISP A and ISP C are one and the same), the issue is simpler. Indeed, many broadband ISPs have their own VoIP operations. This, however, is not required nor even expected to be the norm. Skype, for example, is a non-U.S. company, and is not associated with any ISP. The disassociation of the VoIP provider from the ISP combined with the mobility of the VoIP user makes CALEA applied to VoIP exceedingly complex. As things stand, investigations against people who are constantly on the move are likely either to fail or to violate the privacy of innocent bystanders.

4 How the PSTN and the Internet are the Same and Different

The PSTN provides communications that are reliable, reasonably secure, and moderately expensive. Telecommunications have served as the foundation and infrastructure for a vast range of business services from telemarketing to travel agencies, but there were various limitations. Because of the technology available, all of these services were operated by people. Furthermore, the relationship between phone numbers and locations frequently remained constant for years. The PSTN architecture concentrates investment in the telephone companies, providing a system that is smart in the center and dumb at the edges. In addition, calls have a high setup cost. Thus, although the PSTN can carry data traffic, it is ill suited to services that require patterns of short messages among multiple locations.

While built upon the same semiconductor and fiber technology as the PSTN, the Internet is different in almost every characteristic. Its basic offering is the unreliable transmission of a small packet of data at very low cost. The Internet concentrates investment (and particularly "smarts") at

the edges. The center is a computationally powerful but fundamentally dumb collection of routers and transmission channels.

Until low-cost computing became available in the mid 1980s, there was no commercially feasible way to build a network with the smarts at the edges. The genius of the PSTN was to use intelligence selectively inside the network to offer an extraordinary range of voiceband services using dumb and inexpensive terminals — telephones. Customers owned personal computers years before the arrival of the Internet made the devices so much more valuable. The genius of the Internet Service Providers has been to take advantage of the fact that users were accustomed to providing their own expensive terminals.

Transmission costs dominated the cost of telephony from the invention of the telegraph in 1844 until the late 1990s. Although switches and the local loop cost money, through the mid 1980s these costs were small in comparison with the cost of long-distance transmission. Since 1900 the cost of carrying a single voice circuit for a mile over a long-distance transmission system has fallen by a factor of one million (in inflation-adjusted dollars) [5, p. 779]. Fiber, which became far more widespread in the last decade, dropped costs further. It costs less than a hundredth of cent for a three-minute conventionally switched phone call now, which is less than it costs to print and mail the item on the bill. Fiber enabled the high-bandwidth low-cost internet[3].

The differences in the way the two networks operate is what makes the application of CALEA to VoIP so fraught with difficulties. In a circuit-switched network such as the PSTN, when two parties create a call, they establish a direct path between themselves. For the duration of the call, only these two parties use this path; it is a temporary, but dedicated, connection.

The Internet is, instead, a "packet-routed" network. Rather than fixed circuits, the data that are sent are broken into small packets and each packet travels its own route over the Internet. The packets are reassembled when they are received at the other end. In this respect, internet communications are resource light. That enables such applications as Instant Messaging which, for the PSTN, would require keeping a channel open (in fact, multiple channels) for a long time. Packet routing also enables great flexibility, such as web redirects, that would be much too expensive to accomplish in a circuit-switched environment.

[3]Fiber also drove the long-distance telephone companies out of business since there was no longer any cost basis for the charge dependency on distance. Of course, increased competition played a role in this as well.

The PSTN and the Internet now exist side-by-side with some interaction and some overlap but provide basically different services. It is important to understand the similarities and differences of the two networks. We start with the similarities:

- Both use the same types of transmission facilities (e.g., DSL over twisted pair locally and fiber optics to span long distances). In fact, the two services usually share the same transmission cables.
- Both use electronic routing/switching devices at central nodes to efficiently move bits from one user to another through the network.
- Both use transmission links and switching/routing equipment parsimoniously to serve the largest number of customers with the smallest amount of equipment and transmission capacity.
- Many facilities-based companies operate networks, and they must work together to deliver one user's traffic to another if the two users "belong" to separate networks. (A carrier is facilities based if it provides the switches and transmission between the end user and the ISP.)
- Both the PSTN and the Internet began with the "all-you-can-eat" model for local access pricing, owing, in both cases (in 1876 and more than a hundred years later respectively), to the lack of technology to meter individual usage. In both cases the technology has improved sufficiently to do usage-based pricing, but the culture (and regulation for the PSTN) has not followed suit.
- Both use digital transmission and some form of time-division multiplexing.

In some fundamental ways the two networks are quite different:

- The PSTN has historically used expensive switches to provide end-to-end service with guaranteed quality. In contrast, the Internet and its predecessor, the ARPAnet, have historically used relatively inexpensive routers to minimize the cost of data transfer in trade for only "best-effort delivery"[4]. The Internet is migrating toward switch-based tech-

[4]Users may not be aware that the Internet Protocol makes only a "best-effort" to deliver, but provides no guarantee of data delivery.

niques to achieve the guarantees on quality of service that industrial-grade users demand.

- To minimize cost, the Internet eschews intelligence in the network, in the sense that its inner workings do not discriminate based upon the application type. The PSTN introduced network-based intelligence so as to be able to add new services using dumb terminals, thus permitting the continued use of legacy telephones.

One of the Internet's great virtues arose accidentally: transmission of small quantities of data is inherently cheap, so originally no billing capabilities were built in to measure it. The absence of billing removed a source of overhead and took control of costs out of the hands of the carriers who were left with flat rate billing. (The situation has now changed and it is possible to bill based on usage). Moreover, the natural monopoly of the local loop does not propagate upward into the communication system as it did in the PSTN. Most Internet communications by private parties and small businesses consist of a local phone call to an Internet Service Provider. The cost of entry to the ISP business is low, and competition abounds, holding costs down.

The inexpensive transmission of data in small packets through an extensible switching fabric that need not be reliable creates a supportive medium for complex services run in host computers. What has characterized the Internet's development is the steady appearance of unexpected services from unexpected places. The premier example of such a service is the World Wide Web which emerged from the European Organization for Nuclear Research and has become the backbone of a large segment of worldwide commerce and culture. By contrast, the source of innovation in the Public Switched Telephone Network (PSTN) has been largely limited to the telephone companies themselves.

The differences in the network architectures arose from a combination of policy decisions and available technology. With digitization and fiber, the two networks are becoming more alike. Yet there remain some fundamental differences between the networks that arise out of their distinct architectures; these are subtle but go to the heart of the issues raised in this paper. Any attempt to apply CALEA to the Internet would have to fully accommodate the genuine and fundamental differences between the PSTN and the Internet.

5 Security Issues if CALEA is Applied to VoIP

The PSTN works in a hierarchical manner. Callers using a fixed phone always connect through the same local exchange. For wired telephones and for cellular phones operating inside their home region (see box), this switch is where the wiretap is placed. Designing wiretapping into the communication system raises a fundamental security issue: can the capability be controlled so that only authorized parties can employ it? In the case of the circuit switched telephone system the answer appears to be 'yes.' The wiretapping capability is located primarily within the software of the switch.

Wiretapping Cellular Calls

Since cellular calls superficially appear to share the characteristics of roaming locations of VoIP calls, we discuss wiretapping cellular communications in order to understand how the situations differ.

If a cellphone is operating from within a cell connected to its "home" switch, from the intercept viewpoint it might as well be a wireline phone. Wiretap software running in the switch will be able to identify, copy, and route calls going to and from the cellphone.

When the cellphone is roaming — meaning it is being used outside its normal service area — the problem is quite different. When the roaming phone is initially turned on, and maybe every fifteen minutes after that, a signaling message is sent to the home switch. (Actually the signaling message is sent to the home location register, a database containing the identity of the subscriber and her service profile.) Note that at this point no call content has been transferred to the home network, only signalling information has. If the roaming cellphone is called, the cellphone's home system is consulted during call setup.

Once the phone is "registered" with the home switch, if a call is made locally by the cell phone, there is *no* immediate notification about that call to the home switch (or billing system) and the call is *not* routed through the home switch unless that is the call's destination. In other words, when roaming, the cellphone effectively joins the local network in which it is roaming for the purpose of making outgoing calls. This prevents wiretapping outgoing calls from roaming cellphones by their home switches. By artificially routing the call to the target's home system and back again one could wiretap, but such routing might well be detectable by the target as a result of changes in timing, voice quality, or billing.

Although the switching software is trade-secret and thus its security cannot readily be assessed by outside parties, the switch premises, hardware and software are all owned by the telephone company and, at least in the United States, are reasonably well guarded. Once it is in operation, the wiretap will convey the intercepted material to a remote location via telephone, but the procedure for enabling the wiretap is local to the phone company. The law enforcement agency contacts the communications carrier at an administrative level, and the wiretap is enabled by the carrier's own employees. There have been incidents in which systems of this type have been corrupted — the recent wiretapping of Greek government ministers[5] appears to be such an instance — but in general this form of administration controls eavesdropping capability sufficiently so as to assure that wiretapping is done only under authorized circumstances. The in-switch wiretapping is effective because the wiretap is targeted at a phone number served by the switch. Any call to or from that phone number must pass through the switch. Even in the case of an incoming call forwarded to another number, the call must reach the local switch before being forwarded and therefore comes within the domain of the wiretap.

The centralized nature of the telephone network makes secure wiretapping of a known and fixed phone number a relatively simple prospect. VoIP presents the problem that the switch is not owned by the carrier. It presents an additional problem in that a VoIP call is inherently one that is not tied to a fixed location. In some instances, the computer's Internet address is fixed. In most, however, whether it be the wireless hotel lobby, the Internet cafe, the airport lounge, most home networks, and even the average office computer, the IP address changes with each connection. As society increasingly uses mobile communication devices, there will be an accompanying shift to dynamic IP addressing.

A VoIP provider under a wiretap order might be able to guide the targeted caller to a law-enforcement-controlled rendezvous point at which the tap could be installed (note that in the wireless case law enforcement might even be able to arrange connectivity so that the target is redirected through a law-enforcement access point). The paradigm of VoIP intercept difficulty is a call between two road warriors who constantly change locations and who,

[5]Vodafone used CALEA-like software provided by Ericsson, a telecommunications supplier. The software included "locked" eavesdropping capabilities. An insider at Vodafone — who remains unknown at this writing — activated the eavesdropping capabilities and had the targeted communications delivered to prepaid, untraceable mobile telephones.

for example, may call from a cafe in Boston to a hotel room in Paris and an hour later from an office in Cambridge to a giftshop at the Louvre.

Building a comprehensive, unavoidable, VoIP intercept capability into the Internet would appear to require the cooperation of a very large portion of the routing infrastructure. The fact that packets are carrying voice is largely irrelevant at the level at which tapping is conducted (which is largely the Internet Protocol, or addressing, layer). Most of the provisions of the wiretap law do not distinguish among different types of electronic communications. While currently the FBI is focused on applying CALEA's design mandates to VoIP, there is nothing in wiretapping law that would argue against the extension of intercept design mandates to all types of Internet communications. Indeed, the changes necessary to meet CALEA requirements for VoIP would likely have to be implemented in a way that covered all forms of Internet communication.

There is a danger that intercept design features adopted for the benefit of legitimate law enforcement agencies could be used by others, rendering the entire Internet's application space more vulnerable than it already is. This is very dangerous (and has more than privacy implications). In 2000, the Internet Engineering Task Force[6] Network Working Group examined the issue and declined to consider wiretapping requirements as part of the standards process [9] — because of the potential security problems involved. Various attacks, including man-in-the-middle alteration of data (done by attacker interposed between the communication endpoints), capture of identity information and passwords, and many other pernicious behaviors could well be enabled by CALEA-like accommodations. Furthermore, because these accommodations would apply only to U.S.-based applications, there is the potential to drive traffic to locations unaffected by the U.S. government requirements. Indeed, tunneling and end-to-end cryptographic methods might make it possible for users to "escape" intercept mechanisms in place in the U.S., instead taking advantage of services offered outside U.S. borders. This would not only be bad for American business, it would destroy certain advantages currently enjoyed by U.S. intelligence.

Tricks like creating controlled rendezvous points may work in some cases, but the only certain way to catch the communications between Alice in her cafe and Bob in his hotel room is to create an intercept process in real time at one or both of the routers local to Alice and Bob. This would

[6]The IETF, http://www.ietf.org, develops Internet standards.

be both technically and legally challenging. At a minimum, the routers in question would need to be under the authority of the jurisdiction that had authorized the wiretap. The switch operators would have to receive real time (authenticated) messages ordering them to start the (probably short duration) tapping process. They would have to feel legally comfortable in complying with these orders. There are just under fifteen hundred ISPs that have fewer than one thousand employees in the United States, the vast majority of which have fewer than one hundred employees[11]. Would these ISPs have the resources to properly configure and maintain the complex support that real-time wiretapping of VoIP communications would entail? Or might the wiretapping requirements drive the small ISPs out of business?

Nor would large service providers be immune from problems. VoIP is also *identity agile*, much as though you could select a phone number at will and begin making calls with it immediately. Even if the entities against which wiretap warrants were issued were individuals, recognizing and tracking the multiple identities that are so natural to the Internet lifestyle would be taxing. If you are logging the traffic coming into a location watching for a pattern of calls from some targeted person, you are dependent on being able to recognize when calls are from the same person. If the target has lots of VoIP accounts, then what the pen register (which records all outgoing numbers) lists will be insufficient to recognize the actual identity making the calls, although access to the calls themselves would probably yield this information.

Thus the single biggest problem for VoIP call interception is VoIP mobility, followed closely by VoIP identity agility. But there are other issues as well. We summarize the security problems in building CALEA capabilities into the VoIP environment:

- Physical security of the switching/routing equipment into which wiretap instructions are inserted. This is made particularly difficult because the switching and routing equipment for the VoIP call cannot be predicted in advance (and in this, VoIP differs from both all wired calls and at least all incoming calls on cellular telephones). Compounding the problem is the possibility that the initial ISP used may be one of the thirteen hundred domestic ISPs with fewer than one hundred employees (and thus less likely to have the expertise to secure the switching and routing equipment).

- Physical control of the mechanism for inserting the wiretap instructions. Unlike the PSTN, which is made up of large corporations with

attendant security, VoIP providers run the gamut in size. On average, the physical security of the systems is much weaker.

- Ease of creating new identities on the Internet. As the New Yorker cartoon put it in a different context, "On the Internet, no one knows you're a dog." It is vastly simpler to change an Internet identity than it is to change a phone number. This greatly complicates obtaining all the VoIP communications of the target.

- Secure transport of the selected signals to the law enforcement facility. By opening up the communications to an unacknowledged third party, wiretapping is an architected security breach; the combination of wiretapping with remote delivery elevates the risk that communications security can be violated *with minimal risk of discovery.*

- Increases the risk that the target discovers a wiretap is in place. The smart edges/dumb networks architecture increases the risk of discovery of surveillance by the target. This risk is considerably higher than in the dumb edges/smart network world of the PSTN.

- Ensuring proper "minimization" in the wiretapping process. U.S. law requires minimization — only the target of a court authorization, and only those communications pertaining to the court authorization may be tapped. Due to mobility and identity agility issues, the difficulty of isolating the VoIP communication raises concerns about proper minimization. Widescale wiretapping of non-targeted individuals would diminish respect for the law and lose public support for such type of investigations.

- Increases risk of introducing a vulnerability into the communications system, either through the installation of a general wiretap capability or a specific wiretap. This is the concern raised by the Internet Engineering Task Force.

People call people, not IP addresses. Exactly what makes VoIP so valuable as a communications mechanism — beyond its low cost —- is its ability to enable communication in a highly mobile society. VoIP simplifies communications from people who call from constantly varying places. As we noted earlier, interception against a VoIP call made from a fixed location with a

fixed IP address directly to a big internet provider's access router is equivalent to wiretapping a normal phone call and is easy to accomplish. But if *any* of the conditions listed is not met, then the problem of assuring a high probability of intercept is enormously much harder.

Ways of not meeting these conditions include, but are not limited to, using DHCP at your company or ISP to connect to the Internet (DHCP dynamically configures your Internet connection, which means that may have a different IP address each time you connect even though you have the same physical location), using NAT[7], which makes the IP address invisible to the wiretapper, using different media (dialup versus DSL versus cable modem), moving from place to place, having different URLs for home and work, using freeware VoIP software (Skype), using a non-facilities-based provider (Vonage). Calls of this type are shortly likely to become the norm for VoIP communications.

While it would indeed be technically feasible to build a network with intercept facilities and adequate security — there are defense communications networks that do this — it is unlikely to be politically or socially possible to do so now. Fifteen years ago the Internet was more of a U.S. phenomenon, and international cooperation was not an issue. That is no longer the case. In considering the application of CALEA to VoIP, the lesson of Clipper [4, pp. 212-216], in which foreign governments were simply not interested in a program in which the U.S. government held the encryption keys, speaks loudly. What is theoretically possible is not practically so.

6 Innovation Concerns if CALEA is Applied to VoIP

A major advantage of VoIP is cost savings. CALEA is expensive. The recent report by the Inspector General of the Department of Justice observed, "A VoIP provider contracted to pay approximately $100,000 to a trusted third party (TTP) to develop its CALEA solution. In addition, the TTP will charge a monthly fee of $14,000 to $15,000 and $2,000 for each intercept. These amounts do not include the cost of labor for writing code into

[7]NATs, or network address translation boxes, rewrite source and/or destination of IP addresses as they pass through routers or firewalls and are generally used to support multiple devices on a single public IP address (these are very common in home networks, for example).

the software to accommodate the CALEA solution ... [Telephone company] officials were concerned that the government would mandate that every new feature would have to be CALEA-compliant prior to being offered to the public. Such a restriction would cost the company revenue and place them at a disadvantage in comparison to non-U.S. based providers, who do not have to comply with CALEA."[18, pp. 54-55]. CALEA applied to VoIP may be much more expensive, especially if "opportunity costs" — the costs of not investing in new services — are included.

Voice over IP is the immediate target of the FBI's CALEA efforts. The Internet architecture is rich and flexible, and VoIP is not the only real-time communication in which Internet users indulge. Current real-time applications include Instant Messaging, massively multi-player online role-playing games (MMORPGs) — even music "jamming" sessions. IM and MMORPGs represent huge markets. These communication types fall under the wiretap laws, even if neither the FBI nor the FCC has currently sought to include them in the CALEA requirements.

MMORPGs would probably be completely stifled under such a regime. The Inspector General report comments that the CALEA " standards development process is slow and contentious," [18, p. 30]. Were the U.S. government to adopt a CALEA-type regimen for VoIP (or other real-time Internet communications), the time delays caused by the standards development process would create serious problems for U.S.-based innovation in an industry where an Internet year is a matter of a few months. There is no reason to believe that Japan and Korea, which have very high numbers of MMORPG players, would shoot themselves in the foot by applying CALEA to real-time Internet communications. Opportunity costs could be high indeed.

7 Summing Up

VOIP implementations vary substantially across the Internet making it impossible to implement CALEA uniformly. It appears that CALEA may be effectively applied to those VoIP services that look most like conventional telephony. Intercept against a VoIP call made from a fixed location with a fixed IP address directly to a big internet provider's access router is equivalent to wiretapping a normal phone call, and classical PSTN-style CALEA concepts could be applied directly. In fact, they could be exactly the same if the ISP properly secured its infrastructure and wiretap control process

as the PSTN's central offices are assumed to do. On the other hand, the feasibility of applying CALEA to more decentralized VoIP services seems quite problematic. Neither the manageability of such a wiretapping regime nor whether it can be made secure against subversion seem clear. Rather it seems fairly clear that a CALEA-type regimen is likely to introduce serious vulnerabilities through its "architected security breach."

The fundamental difficulty of applying CALEA to VoIP lies in law-enforcement's desire to achieve 100% compliance with an authorized wiretap order. If law enforcement were to adopt the practice of the intelligence agencies and settle for the best intelligence at a reasonable cost, it might do quite well.

Beyond VoIP lie internet applications such as multi-player games that are not modeled on existing communications and computing services. Just as eBay has become the platform on which many new businesses rest, these may be the basis for future social and business structures that will give the societies that adopt them a major competitive advantage. Although for those who are less than net savvy, it may appear that the Internet is not much more than a place for teens to blog and eBay to offer used Mustangs for sale, the Internet is not a toy. More bits are now carried by the Internet in the United States than our phone companies use to carry conventional phone calls. In slightly over a decade the Internet has become an inherent deeply embedded part of U.S. communications. Regulatory tinkering to enable law-enforcement wishes will impose enormous costs on an extensive established infrastructure.

The real cost of a poorly conceived "packet CALEA" requirement would be the destruction of American leadership in the world of telecommunications and the services built on them. This would cause enormous and very serious national-security implications. Blindly applying CALEA to VoIP and real-time Internet communications is simply not worth this risk.

References

[1] Berson, Tom, *Skype Security Evaluation,* 18 October 2005.

[2] Biondi, Philippe and Fabrice Desclaux, "Silver Needles in the Skype," BlackHat Europe, 2-3 March 2006

[3] Communications Assistance for Law Enforcement Act, Pub. Law No. 103-414, 18 Stat. 4279 (1994).

[4] Diffie, Whitfield and Susan Landau, *Privacy on the Line: The Politics of Wiretapping and Encryption,* MIT Press, 1998.

[5] O'Neill, E. F., *A History of Engineering and Science in the Bell System (1925-1975),* AT& T Bell Laboratories, 1985.

[6] Federal Communications Commission, *Policy Statement FCC 05-151,* 5 August 2005.

[7] Federal Bureau of Investigation, *Electronic Surveillance Needs for Carrier-Grade Voice over Packet (CGVoP) Service,* 29 June 2003.

[8] Foreign Intelligence Surveillance Act, 50 U.S.C. §1801 (2006) et. seq.

[9] Internet Engineering Task Force, *NWG, RFC2804 — IETF Policy on Wiretapping,* May 2000.

[10] Landau, Susan, "National Security on the Line," *Journal of Telecommunications and High Technology Law,* to appear.

[11] OneSource, *High-Technology Product Code: Internet infrastructure services (U.S. only),* run 27 April 2006.

[12] RfC3261 SIP- Session Initiation Protocol, June 2002.

[13] RfC3924 Cisco Architecture for Lawful Intercept in IP Networks, October 2004.

[14] Sherr, Micah, Eric Cronin, Sandy Clark and Matt Blaze, "Signaling Vulnerabilities in Wiretapping Systems," *IEEE Security and Privacy,* November/December 2005, pp. 13-25.

[15] Solove, Daniel and Marc Rotenberg, *Information Privacy Law,* Aspen Publishers, 2003.

[16] Omnibus Crime Control and Safe Streets Act, 18 U.S.C. §§2510-2521 (1968).

[17] U.S. Department of Justice, *In the Matter of Communications Assistance for Law Enforcement Act and Broadband Access and Services,* ET Docket No. 04-295, RM 10865, 14 November 2005.

[18] U.S. Department of Justice, Office of Inspector General, Audit Division, *Implementation of the Communications Assistance for Law Enforcement Act,* Audit Report 06-13, March 2006.

Authors

- Steven M. Bellovin, a professor of computer science at Columbia University, is a member of the National Academy of Engineering and a former Security Area director for the Internet Engineering Task Force.
- Matt Blaze is an associate professor of computer science at the University of Pennsylvania whose research focuses on the design of secure systems. In 1994, Blaze discovered a serious flaw in the US Government's *Clipper* encryption system.
- Ernie Brickell designs security and privacy architectures as a senior professional engineer at Intel Corporation and is the founding editor-in-chief of the *Journal of Cryptology.*
- Clinton Brooks, Ph.D., retired as a senior executive at the National Security Agency, where his career involved a number of assignments relevant to the considerations in this report.
- Vinton Cerf, Chief Internet Evangelist for Google, is one of the founding fathers of the Internet.
- Whitfield Diffie, Chief Security Officer for Sun Microsystems, is one of the fathers of internet security.
- Susan Landau is a Distinguished Engineer at Sun Microsystems, where she works at the intersection of security, cryptography, and policy.
- Jon Peterson is a Fellow at NeuStar Inc, serves as an Area Director of the Real-time Applications and Infrastructure (RAI) Area of the Internet Engineering Task Force (IETF), and has authored or co-authored numerous IETF standards related to the Session Initiation Protocol (SIP).
- John Treichler is a founder, director, and the Chief Technical Officer of Applied Signal Technology, Inc.

Appendix

Wiretapping Law and CALEA

Who practices wiretapping? The most visible practitioners are the police, using wiretapping to collect evidence for use in prosecution. The fact that police must generally introduce their evidence in court where it is subject to examination by the defense makes police wiretapping reasonably tractable to regulation. The second body of wiretappers are intelligence agencies. Their operations are far less visible than those of police and are rarely examined in public proceedings. Wiretapping is by and large illegal for all other parties, though the laws regulating radio reception vary substantially from one jurisdiction to another. Most non-state practitioners of wiretapping are therefore, essentially by definition, criminals.

In the United States, wiretapping is fundamentally governed by two laws: the 1968 Omnibus Crime Prevention and Safe Streets Act, Title III of which pertains to wiretapping, for criminal investigations, and the 1978 Foreign Intelligence Surveillance Act, or FISA, which governs wiretapping for intelligence purposes.

These laws requires telecommunications providers to cooperate with law enforcement, providing them access to facilities and assistance in carrying out their tasks. They do not, however, require the telecommunications companies to make wiretapping convenient or inexpensive. They only require them to make reasonable efforts to accommodate the police in installing wiretaps in the existing telecommunications system, not to make alterations designed to make wiretapping easier.

Law enforcement's capabilities in wiretapping took a giant leap forward in 1994 with the passage of the Communications Assistance for Law Enforcement Act (CALEA), which require the telecommunication providers to design their systems to law-enforcement standards for wiretapping. Failure to deliver content for which there is legal authorization was made punishable by draconian fines *regardless of the cause*.