Choices: Privacy & Surveillance in a Once & Future Internet

Susan Landau

Abstract: The Internet's original design provided a modicum of privacy for users; it was not always possible to determine where a device was or who was using it. But a combination of changes, including "free" Internet services, increasing use of mobile devices to access the network, and the coming "Internet of Things" (sensors everywhere) make surveillance much easier to achieve and privacy more difficult to protect. Yet there are also technologies that enable communications privacy, including address anonymizers and encryption. Use of such technologies complicate law-enforcement and national-security communications surveillance, but do not completely block it. Privacy versus surveillance in Internet communications can be viewed as a complex set of economic tradeoffs – for example, obtaining free services in exchange for a loss of privacy; and protecting communications in exchange for a more expensive, and thus less frequently used, set of government investigative techniques – and choices abound.

SUSAN LANDAU is Professor of Cybersecurity Policy in the Department of Social Science and Policy Studies at Worcester Polytechnic Institute. Previously, she served as Senior Staff Privacy Analyst for Google and as Distinguished Engineer at Sun Microsystems. She is the author of *Surveillance or Security? The Risks Posed by New Wiretapping Technologies* (2011) and *Privacy on the Line: The Politics of Wiretapping and Encryption* (with Whitfield Diffie, 1998; rev. ed. 2007), and contributed to the National Research Council's *Bulk Collection of Signals Intelligence: Technical Options* (2015).

Electronic communications create challenges. In enabling citizens to connect at a distance, they would appear to loosen governmental control. But signals can be eavesdropped on and recorded, and communications surveillance gives tremendous power. Even if the communication itself is encrypted, communications metadata – the who, when, where of a message – are not. Anyone who can collect metadata has vast opportunity to know who is where and connecting with whom.

The ability to remotely eavesdrop has existed for at least as long as electronic communications. Because a signal can be plucked from the air, rather than visibly tapped into, radio is easier to eavesdrop on than are wired communications. In many ways, the Internet has made such surveillance easier still. The tremendous flexibility afforded by the network – the medium supports applications as diverse as search engines, maps, online social networks (OSNs), Twitter, YouTube, Netflix, Uber, and MOOCs (Massive Open Online Courses) – makes the Internet indispensable to citizens and nations alike, and its signals

doi:10.1162/DAED_a_00365

provide rich content for anyone listening in. This can include governments, suppliers of the services, and eavesdroppers.

Yet it does not have to be that way. Communications can be encrypted, and, like speech, they can be ephemeral. The record of a communication's path through a network can be essentially undiscoverable. There are many ways to provide Internet communications, some of which do not impinge on privacy. This article is about those choices.

It starts, as it often does, with the underlying technology. The Internet was developed as a medium for sharing data. Its basic architectural principles – to break data into a numbered set of small packets and transmit the packets as efficiently as possible – reflects that underlying premise. Each packet is transmitted using the Internet protocol (IP). Packets typically have three parts. The *header* says where the packet is from (the sender's IP address), where the packet is going (the receiver's IP address), the type of communications protocol (email, Web page, video, voice, and so on), and its position (packet number) in that particular transmission. The *payload* – the actual content – follows. Finally a *trailer* marks the end of a packet. Applications – an http connection to a Web page, an email connection, a Voice over IP call (VoIP) – are broken into packets and then reassembled at the receiver's end.

Mobility of devices means that the user's IP address at the café at 10 a.m. is different from that in the seminar room at 11 a.m. Each time the user connects back to the network, her IP address is transmitted to her service provider. That is how Facebook communications and your email reach her even when she has moved locations and her IP address has changed.

IP location provides partial identification. While an IP address delimits a location from which and to which packets are transmitted, that address is, for a number of reasons, not necessarily useful in identification. The IP address may be one used temporarily, and without strong identification, such as at an Internet café or an airport. Without ancillary information, such an IP address may provide minimal identifying information. Another reason that an IP address may not provide definitive identification is that few routers along the transmission check a sender's address; so spoofing an IP address is easy.[1]

Even if the IP address is correct, it may not provide an investigator with information to determine who is responsible for a particular action.[2] That is because in such instances, the connecting machine may be just a way station. Consider, for example, *DDoS attacks* (Distributed Denial of Service attacks), in which hundreds of thousands of computers simultaneously send messages to an online service, overwhelming it and taking it offline. The machines sending these messages are simply intermediaries that have been compromised themselves. DDoS is an example of a *multistage attack*, in which a perpetrator infiltrates a series of machines to launch an attack. *Cyberexploits* – theft of information from networked systems – are also typically multistage attacks. The first machine to receive the exfiltrated data is often itself compromised, and the stolen data will be quickly moved from that machine to another and another – a lengthy chain of compromised machines – before the data end up in the attacker's hands. An investigation may lead to the initial machine that was used in the scam, but is unlikely to lead all the way to the real attacker.[3]

The fact that IP addresses do not provide precise identity matters very little in certain cases. Spoofing does not concern the Recording Industry Association of America, which uses an IP address as a jumping-off point for copyright infringement suits.[4]

IP addresses have also served law enforcement as a starting point for investigations.[5] They can also be useful in investigations in which the participants' addresses are related; for example, if they all work at Enron.

Since an IP address is typically not enough to identify an individual, a user browsing generic sites such as *The New York Times*' without logging in achieves some anonymity. Actions the user takes, however, can alter that. In particular, a series of sufficiently personal searches that can be linked to each other may suffice to identify an individual.[6]

The point is that IP addresses have a fungibility, at least when it comes to identity. They provide a starting point for linking a person with an action, but they are also insufficient to be definitive. Thus, IP addresses can provide surveillance capabilities and privacy; the specific circumstances determine which it might be.

Investigators often seek identity, though not necessarily at the level of an IP address. Following users across the Internet became important with the arrival of free services such as Facebook, Google search, and Yik-Yak. Such services are supported through advertising. In this instance, identity does not mean identifying a user in the sense of "Alison Clark is visiting honda.com," but rather that the browser currently viewing nytimes.com is the same that ran an Internet search for compact cars earlier in the day. This enables the search provider, for example, to serve a Honda ad on *The New York Times* website that the user is browsing. Identification is derived through cookies in the browser, not an IP address.

There are times when identity on the Internet at the level of a person matters. A bank does not particularly care what a user's IP address is, but if there's a transaction occurring, the bank seeks assurance that the person is who she claims to be and wants her to authenticate her identity at the bank's site. For many situations, including transactions with high value, authentication conducted within an application is sufficient proof of identity.

Increasingly, identity is required for accessing services. *Federated identity management* – facilitating access to different sites once a user has been authenticated to an "identity provider" – is one way to do this. For example, a corporate login could permit seamless electronic access to outsourced services such as HR or travel booking; a university login could allow access to electronic resources at a federated institution.

Some approaches to identity management carefully protect privacy. One example is Shibboleth, which is used for sharing secured Web resources and services among a consortium of universities, research labs, publishers, and medical libraries. To access a resource, the user must establish her right to it, such as by being a faculty member or a student. The user's ID is shared only if access to the resource requires it.[7] Another case of privacy-protective identity management comes courtesy of the U.S. government, which employs private-sector identity providers for accessing government websites, but requires that the users' information be employed only for authentication, audit, and complying with the law – and not for ads or sharing with third parties.[8] So if a user is looking at Veterans Affairs benefits and then at information about sexually transmitted diseases, that information should neither be tracked nor stored by the identity provider.

Other systems take a very different approach, using user data to entice services to work with them. Thus, for example, when the Facebook login is used to authenticate a user to an app, Facebook shares with the app the user's name and gender, and provides a list of the user's friends who also use the application. This makes the Facebook login valuable to the app, but not to a user seeking privacy.

The existing model of advertising and tracking in exchange for services is not the only possible model for the Internet. One alternative would be to charge for services: a tenth of a cent for a search, a monthly charge for email support, and so on. And there is no reason the two systems could not coexist: charges for users seeking privacy-protective services, and an advertise-and-track model for those who are indifferent to the privacy issue or unable to pay.

By making the network indispensable to daily life, the Internet drove the development of smartphones. Most Internet accesses now occur through mobile devices, a fact with profound implications for privacy and surveillance. While a laptop can be "on" but not connected to the network – functioning as a computer, not a communications device – if a smartphone is on (and not in "airplane mode"), it will be connected to the telephone network whenever the provider's system is within range. Thus, a phone's location, which is broadcast several times an hour to announce "I am here," is a relatively public piece of information. The phone's connection is through the nearest base station: the cell tower closest to the user. As the user moves to new locations, the phone connection is "handed off" to the next base station. That is information that the phone network – or an interceptor – will learn.

Where an individual is calling from, or whom they called, may be much more interesting than what they actually said; communications metadata, for example, can reveal the structure of an organization. One striking example of this type of analysis concerns the case of former Lebanese Prime Minister Rafik Hariri, who was assassinated in Beirut in 2005, when a truck bomb exploded near his motorcade.[9] The planning behind the assassination was well-hidden, but analysis of cell phone traffic in Beirut and other locations exposed a pattern of communications that revealed who did it – and how.[10]

Desktops, laptops, and tablets are, to some extent, multiuser machines; but smartphones are more strictly associated with individuals. Thus, just tracking the phone's location provides an extremely accurate way of determining a phone's user.[11] Know the recipients of a person's calls, and you can infer who she is and what is happening in her life: whether she has just lost her job, her mother is ill, or her son has just gone off to college. Because people carry personal transmitters and receivers, government investigators no longer need to tail individuals and monitor phone booths to capture conversations and movements; they simply track mobile phones. Because communications patterns are so revealing, if a government can fully surveil a nation's communications network, it can even track "burner" phones (anonymous prepaid phones) through correlations in location and use.

Governments are not the only ones following users' locations; in fact they may be collecting far less information than many companies. To provide the Internet with services for which smartphones are valued – finding a local restaurant and making dinner reservations and then determining the best route there – the phone must provide location information to the app. This is done through GPS, which typically operates on a resolution within ten meters.[12] So the network provider knows where the phone is and with which service the user is communicating, while the app provider learns phone location and what information is delivered through the app.[13]

This is an interesting design choice in location data tracking: Apple's iOS8 does not allow apps to collect location information when the app is not in use, but there is no such restriction for Android phones (of course, if location collection is shut off, then Android apps cannot collect it). The

latter situation might change: in February 2015, the U.S. Federal Trade Commission told app developers, "If you access users' locations when they're not using your app, it's a good idea to clearly disclose what you're doing and provide users with choices."[14]

The real gold in the Internet advertising world is "conversion tracking," learning what customers do after clicking on an ad: whether they bought the product or followed up in another way (such as visiting a product's website). When the Web access, user location, and payment are all on the same device, it becomes even easier for an Internet service and an advertiser to determine an ad's effectiveness. The phone might not announce, "This is Alison Clark" at the Honda dealership, but if her phone shows an identifier from the search she conducted, that provides the relevant information. For this reason, companies are at least as eager as governments to use smartphones to track users.

With such interest in following the user and such capabilities for surveillance, it becomes difficult to imagine that any privacy is possible. Yet there are many technical solutions for protecting privacy. It is particularly striking that there are even technical solutions for obscuring with whom you are communicating. In the mid-1990s, the Naval Research Laboratory began work on a system that makes it difficult to determine who is connecting with whom on the Internet.[15]

The onion routing network, commonly known as Tor, protects against traffic analysis through deployment of a "Tor network," a collection of servers with encryption and decryption software. A path is determined for each communication, which is then routed through a network of Tor nodes (servers) that strip off the encryption "one layer at a time." Encryption keys are based on the nodes and route.

Anyone who is eavesdropping on Internet traffic can determine that one Tor node is communicating with another. More specifically, if there is surveillance of connections to a website – such as who in Iran is reading about international sanctions – the interceptor will see a connection from the Tor network to the forbidden website. But the eavesdropper will not see the IP address that initiated the Internet connection unless the eavesdropper can view the entire network at once and thus correlate times and sizes for all network transmissions. In such cases, interceptors can deanonymize Tor communications, but otherwise Tor makes such identification extremely difficult.[16] Browsers and instant messaging apps can be used on the Tor network, enabling truly anonymous communication through which it is infeasible for the receiver to determine the original sender's IP address.

It might be surprising that a U.S. government agency supports anonymous Internet accesses. But there are good reasons for the government to seek such capabilities. A military unit in a safe house in the Middle East would not want to let the local Internet service provider (ISP) know that it is communicating with the Naval Academy in Annapolis, while an FBI agent investigating a child pornography chat room does not want to use an IP address that resolves to "fbi.com." So a system that makes it appear that the Web connection is from somewhere else provides useful investigative capability. Tor is widely available and popular with journalists, human-rights workers, and others seeking privacy of communication. And it provides cover for military personnel and law-enforcement investigators, whose identities as U.S. government employees are masked by the system's broader set of users.

In many ways, confidentiality of communications is simpler to achieve than privacy. Encryption – encoding messages so that only the sender and receiver can read

it – accomplishes this. But simple answers belie simple understanding. By now it should be clear that nothing about protecting communications is entirely simple.

For a quarter-century, from the 1970s to the late 1990s, the U.S. government battled academia and industry over encryption used to support confidentiality. This fight came to a head during the "Crypto Wars" two decades ago, at the dawn of the Internet era. In 1999, the European Union loosened its controls on the export of products with strong encryption; a few months later the United States did the same.[17] This change made it much simpler to deploy cryptography in commercial products.

While use of encryption for confidentiality had been controlled, its use for authentication – assurance that a person or site is who they say they are – had not. Https, the secure version of the http linking protocol, is used to authenticate a website (for example, confirming that the site is bankofamerica.com and not an imposter that is like bankofamerica.com) and encrypt communications between a user and the site. This protocol was essential for electronic commerce, and was already deployed by the mid- to late 1990s. Given that https was widely deployed quite early for ecommerce, it is surprising and somewhat striking that Web mail, the service that provides email through a browser, was not similarly protected. Let us examine how such services work.

Suppose a user with the email account boris@yahoo.com is communicating with another user, natasha@gmail.com. When Boris sends an email to Natasha, he logs on to his Yahoo! Mail account, writes his message, hits send, and the mail travels to Natasha's Gmail account. She will read his message once she logs onto the Gmail server (many users, including those with Android phones, are always logged on).

From the beginning, Web mail providers used the https protocol when authenticating users to their accounts; this encrypted the user password from the user's browser to the site. But for many years, the large Web mail providers – Gmail, Hotmail, Yahoo! Mail – did not encrypt the connection between the user and her mail account; that is, *the emails themselves traveled in the clear* between the user's machine and the provider. Anyone eavesdropping on the Internet connection between Boris and his Yahoo! Mail account, relatively easy to do, could read Boris's incoming and outgoing mail. In response to the Snowden revelations, Google changed their connection to a secure one, and other providers are following suit. Mail traveling between Natasha's computer and her Gmail account are on an encrypted channel; interceptors cannot read it.

But this change does not fully encrypt the mail from Boris to Natasha. Although systems began securing communications between the user and the mail provider, the communications themselves still are not encrypted "end-to-end" from sender to receiver. If Boris and Natasha both happen to be using Gmail, then their communication will be encrypted between Boris and the Gmail server and between the Gmail server and Natasha. Contents on the Gmail servers are encrypted, but there will be a time when Boris's mail to Natasha is in the clear at Google. That is because Google uses the mail to serve ads and to provide personalized services. For example, a plane reservation in an email account will trigger a notification in the Google Now app to inform the user about traffic on her preferred route to the airport.

There are other models for email, some of which provide greater confidentiality. One such service was Lavabit. Mail on Lavabit servers arrived encrypted and stayed that way; they were decrypted only when a user was reading the communication. Users received keys through a secure https connection.

Lavabit was shut down by owner Ladar Levison after the U.S. government requested the encryption key securing the https connection between Lavabit and its users. Although government investigators appeared to be interested only in a single user's communications, giving up that key would have allowed access to all https connections, thus potentially exposing all customer passwords. Levison felt that would violate his privacy commitment to his customers. Instead of doing so, he closed the service.[18]

Another example of alternative privacy protection is Off-the-Record (OTR) chat. Google's OTR chat does not store chat histories in users' accounts, or in the accounts of the people with whom they are chatting. But Google policy does not preclude storing the communications elsewhere.[19] A more protective version would be not to store the communications whatsoever. Even more protective would be not storing and providing encryption for the chat. Most protective would be to encrypt using a technique called *forward secrecy*, so that even if the encryption key is compromised at some point, no previously intercepted messages can be decrypted.[20] There are OTR systems that provide this level of security.[21]

Alternatives in designing applications lead to varying degrees of privacy. Such safeguards do not come for free. They cost extra development time and can decrease efficiency by preventing reuse of data in other applications. And, as Lavabit's owner discovered, sometimes privacy-protective systems lead to conflicts with the government.[22]

Encryption's knotty issue is that legal access to decrypted content may be granted to an investigator, but technology prevents such access. And although electronic communications now provide much richer investigative information than ever before – consider the Hariri case – sometimes content provides information that these other tools cannot. There is, however, a way around this problem.

As the Snowden disclosures confirmed, national-security agencies may exploit vulnerabilities in communications devices to exfiltrate data from targets.[23] Such capabilities are used not only by intelligence agencies, but by law enforcement as well.[24] As encryption becomes increasingly common, such "lawful hacking" will increasingly be used when communications content cannot be retrieved in other ways.[25] It is no silver bullet; a vulnerabilities approach is more complex legally and technically, and more expensive than if unencrypted communications can be made available.

The privacy situation is about to grow far more difficult. While Internet transactional information is remarkably revelatory, the information from sensors on toothbrushes, watches, clothes, heart monitors, phones – and everything else – will be many times more so. Cheap sensors communicating with the Internet will soon be everywhere: sensors to measure tire pressure and bridge structural health; sensors to report on the freshness of food in the fridge, the dampness in the soil, and the movement of an elderly person at night; sensors to determine whether the car driver or passenger is making a call. The number of devices from the Internet of Things (IoT) will dwarf the current number of devices connected to the Internet.

A user has some control over whether information on her smartphone is shared with the app; she can always shut the application off or completely remove it from her phone. While in some cases – as with smart toothbrushes – the user might have the same capability, she is unlikely to be provided with such control on many other applications (such as tire sensors).

For security's sake, one approach might be enclaves: creating domains with extreme-

ly limited ability to communicate outside a narrow realm. Consider the type of connectivity a smart refrigerator should have. Fully connecting to the Internet creates an unnecessary security risk. The fridge needs to be able to communicate with the manufacturer for updates and with the owner for the you-need-milk notifications. A smartphone app that puts milk on the shopping list does not need notifications directly from the fridge; it could do so instead by accessing owner updates. Limiting information flows from sensors and controlling where those data initially go provides a measure of privacy and security.

Enclaves are likely to be for systems of similar purposes (medical devices, infrastructure sensors). A patient-sensor network in a hospital intensive care unit should not be accessible outside that area, while a sensor network for medical research might span wide geographies. In some cases, data can be aggregated before reaching a larger network – such as combining data from sensors on soil conditions within a region – providing privacy to individuals. Flows of information – which data are shared with whom – will be determined by enclaves.

Determining appropriate enclaves for sensor networks – should the enclave for medical research networks be strictly separate from that used for patient networks? – is complex, but provides only a partial solution for privacy. This is partly because keeping enclaves truly separate remains a difficult technical problem. "Car-hack" attacks, such as when in 2010 researchers remotely took control of a car's brakes and engine,[26] were possible because enclaves lacked clear separation. In addition, putting tight legal and policy controls on the data's use will be crucial for privacy.[27]

Communications between people at a distance have never been entirely private. Delivery is variable, seals can be broken, messages decrypted.[28] Communications that were once ephemeral now have a trail, and being anonymous in modern society is no longer plausible. It not only means eschewing the use of smartphones (and credit cards, transit passes, and so forth), but also requiring companions to do the same. You cannot hide from network detection if your known companions' phones broadcast their whereabouts.

In the wake of the Snowden disclosures, privacy-enhancing technologies such as Tor, and Google and Apple's encrypted phones, in which decryption is only possible with the user key (though, of course, much of the data may also be stored elsewhere), have drawn much interest. Privacy-enhancing technologies enable different levels of ability to conceal identity and increase the cost for monitors to determine information about an individual, but data collection is so vast that these tools are unlikely to be sufficient for people with specific needs to protect themselves, including journalists and human-rights workers, as well as criminals, terrorists, and spies. Indeed, serious efforts to defend against electronic traces may only draw increased attention from intelligence agencies or other eavesdroppers.[29]

Privacy has always been about economics. How much does it cost to use Lavabit's encrypted email services instead of free Gmail services? Or how much more does it cost to use cash at the bookstore instead of ordering over Amazon? On the flip side, how many resources must be devoted to investigations if communications are protected through privacy-protective technologies?

The Internet changed the equation in various ways. In the initial development of Internet applications, we tipped in one direction, allowing collection and release of massive amounts of information about ourselves. Application design, however, provides a plethora of possibilities. As

long as "free" is the model for Internet services, it is unlikely that the tracking industry, developed to support Internet advertising, will disappear. The information amassed by private industry, including the vast collection of data afforded through the Internet of Things, will also be accessed by governments.

Our current Internet design is a world in which applications sometimes provide privacy-protective solutions for those who want them. But these give only a modicum of privacy. Changing the ease with which surveillance can be performed, making it more difficult to track user preferences and activity, is largely a matter of choices. (Of course, under some governments, there are no such choices. But in the United States, private industry is not required to know who users are in order to provide them a service.) Choices for more privacy-protective solutions can come from government regulations, and they can come from customer demand. But such alternatives in application design do exist.

Humans are a highly communicative species, and the Internet fed this aspect of our nature. That the Internet grew spectacularly alongside the terrorist attacks of September 11th and their aftermath meant that privacy, always on a societal pendulum, largely suffered over the last decade and a half. Now choices abound; we may be reaching a time when the pendulum swings back. But the market will only provide effective privacy-protective solutions if enough users demand them.

ENDNOTES

1 Robert Beverly, Ryan Coga, and kc claffy, "Initial Longitudinal Analysis of IP Source Spoofing Capability on the Internet," July 25, 2013, http://calhoun.nps.edu/bitstream/handle/10945/36775/Beverly_Initial_Longitudinal_2013.pdf.

2 This discussion on the value of IP addresses for attribution is based on David Clark and Susan Landau, "Untangling Attribution," *Harvard National Security Journal* 2 (2) (2011): 25 – 40.

3 In the case of DDoS attacks, the machine sending connection requests to the service has itself been infected with malware. By examining ISP logs, it will usually be possible to determine from where it is receiving instructions to attack (flood the service with connection requests). But determining which machine, or set of machines, are sending these instructions may be a multistep process, since instructions may be laundered through several machines. A similar situation exists for data exfiltration, with the wrinkle that the data can be followed only until they reach a jurisdiction in which the trail is made opaque. For further discussion, see ibid.

4 This technique was successfully used by the RIAA for a number of years, although determining an infringing user from an IP address is not straightforward; see *Recording Industry Association of America, Inc. v. Verizon Internet Service*, 351 F. 2nd 1229, D.C. Circuit (2003). Since 2012, a number of rulings have gone against the RIAA.

5 See Clark and Landau, "Untangling Attribution," fn 3.

6 AOL released information about users' searches over a three-month period; identifying a particular user was not hard to do. See Michael Barbaro and Tom Zeller Jr., "A Face is Exposed for AOL Searcher No. 4417749," *The New York Times*, August 9, 2006. The ability to link a set of searches to a user requires first being able to link the user to her searches.

7 R. L. Morgan, Scott Cantor, Steven T. Carmody, Walter Hoehn, and Kenneth J. Klingenstein, "Federated Security: The Shibboleth Approach," *EDUCAUSE Quarterly* 27 (4) (2004): 12 – 17.

8 Georgia Tech Research Institute, "GTRI NSTIC Trustmark Pilot" (October 7, 2014), https://trustmark.gtri.gatech.edu/operational-pilot/trustmark-definitions/ficam-privacy-activity-tracking-requirements-for-csps-and-bae-responders/1.0/.

9 The case is currently being decided in the Hague by the UN's Special Tribunal for Lebanon.

10 The cell phone traffic showed several groups coordinating actions while tracking Hariri through Beirut, and possibly even conducting a dry run of the attack. See Ronen Bergman, "The Hezbollah Connection," *New York Times Magazine*, February 10, 2015.

11 Phillippe Golle and Kurt Partridge, "On the Anonymity of Home/Work Location Pairs," in *Pervasive Computing: 7th International Conference, Pervasive 2009, Nara, Japan, May 11 – 14, 2009, Proceedings*, ed. Hideyuki Tokuda, Michael Beigl, Adrian Friday, A. J. Bernheim Brush, and Yoshito Tobe (Berlin: Springer Berlin Hedelberg, 2009), 390 – 397.

12 Matt Blaze, Testimony to the House Committee on the Judiciary Subcommittee on the Constitution, Civil Rights, and Civil Liberties, *Hearing on ECPA Reform and the Revolution in Location Based Technologies and Services*, June 24, 2010.

13 If the user is signed in, as are most Android users, then Google will learn which apps are being used and how frequently, though not what information is being communicated (unless the apps are Google apps).

14 See Amanda Koulousias, "Location, Location, Location," Federal Trade Commission, February 11, 2015, https://www.ftc.gov/news-events/blogs/business-blog/2015/02/location-location-location.

15 See https://www.torproject.org.

16 James Ball, Glenn Greenwald, and Bruce Schneier, "NSA and GCHQ Target Tor Network that Protects Anonymity of Web Users," *The Guardian*, October 4, 2014.

17 Export controls had effectively prevented the deployment of cryptography in domestic products. While the change in regulations did not permit export of cryptography in all products, it worked well enough to support the needs of the expanding Internet ecosystem. See Whitfield Diffie and Susan Landau, "The Export of Cryptography in the 20th Century and the 21st," in *The History of Information Security: A Comprehensive Handbook*, ed. Karl De Leeuw and Jan Bergstra (Amsterdam: Elsevier, 2007), 725 – 736.

18 See Nicole Perlroth and Scott Shane, "As F.B.I. Pursued Snowden, an E-Mail Service Stood Firm," *The New York Times*, October 2, 2013.

19 Google Support, Chat Help, "Chatting Off the Record," https://support.google.com/chat/answer/29291?hl=en (accessed March 29, 2015).

20 Whitfield Diffie, Paul van Oorschot, and Michael Wiener, "Authentication and Authenticated Key Exchanges," *Designs, Codes, and Cryptography* 2 (2) (June 1992): 107 – 125.

21 Surveillance Self-Defense, "How to: Use OTR for Mac," and "How to: Use OTR for Windows," Electronic Frontier Foundation, https://ssd.eff.org/en/index (accessed March 30, 2015).

22 Levison had previously complied with court orders for targeted access. His objection to the FBI request was that the agency sought the encryption key for his SSL certificate, which would have compromised the privacy of all Lavabit users. See Perlroth and Shane, "As F.B.I. Pursued Snowden, an E-Mail Service Stood Firm."

23 Spiegel Staff, "Inside TAO: Documents Reveal NSA Top Hacking Unit," *Spiegel Online International*, December 29, 2013.

24 In this case, to determine IP addresses; but the method can also be used to exfiltrate data, including encryption keys. See Steve Bellovin, Matt Blaze, Sandy Clark, and Susan Landau, "Going Bright: Wiretapping without Weakening Communications Infrastructure," *IEEE Security and Privacy* 11 (1) (January/February 2013): 62 – 72; and Steve Bellovin, Matt Blaze, Sandy Clark, and Susan Landau, "Lawful Hacking: Using Existing Vulnerabilities for Wiretapping on the Internet," *Northwestern Journal of Technology and Intellectual Property* 12 (1) (2014).

25 See Bellovin et al., "Going Bright."

26 Karl Koscher, Alexei Czeskis, Franziska Roesner, et al., "Experimental Security Analysis of a Modern Automobile," in *Proceedings of IEEE Symposium on Security and Privacy (Oakland) 2010*,

ed. David Evans and Giovanni Vigna (Washington, D.C.: IEEE Computer Society, 2010), 447–462.

27 See President's Council of Advisors on Science and Technology, *Big Data and Privacy: A Technological Perspective* (Washington, D.C.: Executive Office of the President, May 2014), https://www.whitehouse.gov/sites/default/files/microsites/ostp/PCAST/pcast_big_data_and_privacy_-_may_2014.pdf.

28 Message decryption led to the downfall of Mary, Queen of Scots, in the sixteenth century; see The National Archives of the United Kingdom, Codes and Ciphers, "Mary's Ciphers," http://www.nationalarchives.gov.uk/spies/ciphers/mary/ma1.htm.

29 The lack of connection to communication networks was one hint that bin Laden was in the villa in Abbottabad; see Mark Mazzeti, Helen Cooper, and Peter Baker, "Behind the Hunt for Bin Laden," *The New York Times*, May 2, 2011.