Keys Under Doormats:

MANDATING INSECURITY BY REQUIRING GOVERNMENT ACCESS TO ALL DATA AND COMMUNICATIONS

Harold Abelson, Ross Anderson, Steven M. Bellovin, Josh Benaloh, Matthew Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Peter G. Neumann, Susan Landau, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael Specter, Daniel J. Weitzner

Abstract

Twenty years ago, law enforcement organizations lobbied to require data and communication services to engineer their products to guarantee law enforcement access to all data. After lengthy debate and vigorous predictions of enforcement channels "going dark," these attempts to regulate the emerging Internet were abandoned. In the intervening years, innovation on the Internet flourished, and law enforcement agencies found new and more effective means of accessing vastly larger quantities of data. Today we are again hearing calls for regulation to mandate the provision of exceptional access mechanisms. In this report, a group of computer scientists and security experts, many of whom participated in a 1997 study of these same topics, has convened to explore the likely effects of imposing extraordinary access mandates.

We have found that the damage that could be caused by law enforcement exceptional access requirements would be even greater today than it would have been 20 years ago. In the wake of the growing economic and social cost of the fundamental insecurity of today's Internet environment, any proposals that alter the security dynamics online should be approached with caution. Exceptional access would force Internet system developers to reverse "forward secrecy" design practices that seek to minimize the impact on user privacy when systems are breached. The complexity of today's Internet environment, with millions of apps and globally connected services, means that new law enforcement requirements are likely to introduce unanticipated, hard to detect security flaws. Beyond these and other technical vulnerabilities, the prospect of globally deployed exceptional access systems raises difficult problems about how such an environment would be governed and how to ensure that such systems would respect human rights and the rule of law.

Executive Summary

Political and law enforcement leaders in the United States and the United Kingdom have called for Internet systems to be redesigned to ensure government access to information — even encrypted information. They argue that the growing use of encryption will neutralize their investigative capabilities. They propose that data storage and communications systems must be designed for *exceptional access* by law enforcement agencies. These proposals are unworkable in practice, raise enormous legal and ethical questions, and would undo progress on security at a time when Internet vulnerabilities are causing extreme economic harm.

As computer scientists with extensive security and systems experience, we believe that law enforcement has failed to account for the risks inherent in exceptional access systems. Based on our considerable expertise in real-world applications, we know that such risks lurk in the technical details. In this report we examine whether it is technically and operationally feasible to meet law enforcement's call for exceptional access without causing large-scale security vulnerabilities. We take no issue here with law enforcement's desire to execute lawful surveillance orders when they meet the requirements of human rights and the rule of law. Our strong recommendation is that anyone proposing regulations should first present concrete technical requirements, which industry, academics, and the public can analyze for technical weaknesses and for hidden costs.

Many of us worked together in 1997 in response to a similar but narrower and better-defined proposal called the Clipper Chip [1]. The Clipper proposal sought to have all strong encryption systems retain a copy of keys necessary to decrypt information with a trusted third party who would turn over keys to law enforcement upon proper legal authorization. We found at that time that it was beyond the technical state of the art to build key escrow systems at scale. Governments kept pressing for key escrow, but Internet firms successfully resisted on the grounds of the enormous expense, the governance issues, and the risk. The Clipper Chip was eventually abandoned. A much more narrow set of law enforcement access requirements have been imposed, but only on regulated telecommunications systems. Still, in a small but troubling number of cases, weakness related to these requirements have emerged and been exploited by state actors and others. Those problems would have been worse had key escrow been widely deployed. And if all information applications had had to be designed and certified for exceptional access, it is doubtful that companies like Facebook and Twitter would even exist. Another important lesson from the 1990's is that the decline in surveillance capacity predicted by law enforcement 20 years ago did not happen. Indeed, in 1992, the FBI's Advanced Telephony Unit warned that within three years Title III wiretaps would be useless: no

more than 40% would be intelligible and that in the worst case all might be rendered useless [2]. The world did not "go dark." On the contrary, law enforcement has much better and more effective surveillance capabilities now than it did then.

The goal of this report is to similarly analyze the newly proposed requirement of exceptional access to communications in today's more complex, global information infrastructure. We find that it would pose far more grave security risks, imperil innovation, and raise thorny issues for human rights and international relations.

There are three general problems. First, providing exceptional access to communications would force a U-turn from the best practices now being deployed to make the Internet more secure. These practices include *forward secrecy* — where decryption keys are deleted immediately after use, so that stealing the encryption key used by a communications server would not compromise earlier or later communications. A related technique, *authenticated encryption*, uses the same temporary key to guarantee confidentiality and to verify that the message has not been forged or tampered with.

Second, building in exceptional access would substantially increase system complexity. Security researchers inside and outside government agree that complexity is the enemy of security — every new feature can interact with others to create vulnerabilities. To achieve widespread exceptional access, new technology features would have to be deployed and tested with literally hundreds of thousands of developers all around the world. This is a far more complex environment than the electronic surveillance now deployed in telecommunications and Internet access services, which tend to use similar technologies and are more likely to have the resources to manage vulnerabilities that may arise from new features. Features to permit law enforcement exceptional access across a wide range of Internet and mobile computing applications could be particularly problematic because their typical use would be surreptitious — making security testing difficult and less effective.

Third, exceptional access would create concentrated targets that could attract bad actors. Security credentials that unlock the data would have to be retained by the platform provider, law enforcement agencies, or some other trusted third party. If law enforcement's keys guaranteed access to everything, an attacker who gained access to these keys would enjoy the same privilege. Moreover, law enforcement's stated need for rapid access to data would make it impractical to store keys offline or split keys among multiple keyholders, as security engineers would normally do with extremely high-value credentials. Recent attacks on the United States Government Office of Personnel Management (OPM) show how much harm can arise when many organizations rely on a single institution that itself has security vulnerabilities. In the case of OPM, numerous federal agencies lost sensitive data because OPM had insecure infrastructure. If service providers implement exceptional

access requirements incorrectly, the security of all of their users will be at risk.

Our analysis applies not just to systems providing access to encrypted data but also to systems providing access directly to plaintext. For example, law enforcement has called for social networks to allow automated, rapid access to their data. A law enforcement backdoor into a social network is also a vulnerability open to attack and abuse. Indeed, Google's database of surveillance targets was surveilled by Chinese agents who hacked into its systems, presumably for counterintelligence purposes [3].

The greatest impediment to exceptional access may be jurisdiction. Building in exceptional access would be risky enough even if only one law enforcement agency in the world had it. But this is not only a US issue. The UK government promises legislation this fall to compel communications service providers, including US-based corporations, to grant access to UK law enforcement agencies, and other countries would certainly follow suit. China has already intimated that it may require exceptional access. If a British-based developer deploys a messaging application used by citizens of China, must it provide exceptional access to Chinese law enforcement? Which countries have sufficient respect for the rule of law to participate in an international exceptional access framework? How would such determinations be made? How would timely approvals be given for the millions of new products with communications capabilities? And how would this new surveillance ecosystem be funded and supervised? The US and UK governments have fought long and hard to keep the governance of the Internet open, in the face of demands from authoritarian countries that it be brought under state control. Does not the push for exceptional access represent a breathtaking policy reversal?

The need to grapple with these legal and policy concerns could move the Internet overnight from its current open and entrepreneurial model to becoming a highly regulated industry. Tackling these questions requires more than our technical expertise as computer scientists, but they must be answered before anyone can embark on the technical design of an exceptional access system.

In the body of this report, we seek to set the basis for the needed debate by presenting the historical background to exceptional access, summarizing law enforcement demands as we understand them, and then discussing them in the context of the two most popular and rapidly growing types of platform: a messaging service and a personal electronic device such as a smartphone or tablet. Finally, we set out in detail the questions for which policymakers should require answers if the demand for exceptional access is to be taken seriously. Absent a concrete technical proposal, and without adequate answers to the questions raised in this report, legislators should reject out of hand any proposal to return to the failed cryptography control policy of the 1990s.

Contents

1 Background of today's debate on exceptional access **5**
1.1 Summary of the current debate 5
1.2 Findings from the 1997 analysis of key escrow systems 6
1.3 What has changed and what remains the same since 1990s? 7

2 Scenarios **11**
2.1 Scenario 1: Providing exceptional access to globally distributed, encrypted messaging applications 11
2.2 Scenario 2: Exceptional access to plaintext on encrypted devices such as smartphones 14
2.3 Summary of risks from the two scenarios 15

3 Security impact of common law enforcement requirements with exceptional access **18**
3.1 Access to communications content 18
3.2 Access to communications data 19
3.3 Access to data at rest 20

4 Principles at stake and unanswered questions **20**
4.1 Scope, limitations, and freedoms 21
4.2 Planning and design 22
4.3 Deployment and operation 23
4.4 Evaluation, assessment, and evolution 24

5 Conclusion **24**

6 Author Biographies **30**

7 Acknowledgments **31**

1 Background of today's debate on exceptional access

The encryption debate has been reopened in the last year with both FBI Director James Comey and UK Prime Minister David Cameron warning, as in the early 1990s, that encryption threatens law enforcement capabilities, and advocating that the providers of services that use encryption be compelled by law to provide access to keys or to plaintext in response to duly authorized warrants. We have therefore reconvened our expert group to re-examine the impact of mandatory exceptional access in today's Internet environment.[1]

In the 1990s, the governments of United States and a number of other industrialized countries advocated weakening encryption. Claiming that widespread encryption would be disastrous for law enforcement, the US government proposed the use of the *Clipper Chip*, an encryption device that contained a government master key to give the government access to encrypted communications. Other governments followed suit with proposals for encryption licensing that would require copies of keys to be held in escrow by *trusted third parties* — companies that would be trusted to hand over keys in response to warrants. The debate engaged industry, NGOs, academia, and others. Most of the authors of the present paper wrote a report on the issues raised by key escrow or trusted-third-party encryption that analyzed the technical difficulties, the added risks, and the likely costs of such an escrow system[1]. That push for key escrow was abandoned in 2000 because of pressure from industry during the dotcom boom and because of political resistance from the European Union, among others.

1.1 Summary of the current debate

The current public policy debate is hampered by the fact that law enforcement has not provided a sufficiently complete statement of their requirements for technical experts or lawmakers to analyze. The following exhortation from United States FBI Director James Comey is as close as we come:

> "We aren't seeking a back-door approach. We want to use the front door, with clarity and transparency, and with clear guidance provided by law. We are completely comfortable with court orders and legal process — front doors that provide the evidence and information we need to investigate crime and

[1] We follow the 1996 National Academies CRISIS report in using the phrase "exceptional access" to "stress that the situation is not one that was included within the intended bounds of the original transaction." [4, p. 80]

prevent terrorist attacks."

> "Cyber adversaries will exploit any vulnerability they find. But it makes more sense to address any security risks by developing intercept solutions during the design phase, rather than resorting to a patchwork solution when law enforcement comes knocking after the fact. And with sophisticated encryption, there might be no solution, leaving the government at a dead end — all in the name of privacy and network security." [5]

Prime Minister David Cameron simply wants the police to have access to everything. Speaking in the wake of the Charlie Hebdo murders in Paris, he said:

> "In our country, do we want to allow a means of communication between people which, even in extremis, with a signed warrant from the home secretary personally, that we cannot read? ...The question remains: are we going to allow a means of communications where it simply is not possible to do that? My answer to that question is: no, we must not." [6]

So, we must ask, is it possible to build in such exceptional access without creating unacceptable risk? In order to understand the technical and operational issues, we first review the results of our 1997 report and consider what has changed since then. We next try to clarify ideal law enforcement requirements and understand the kinds of risks that are likely to arise if these generic requirements are imposed broadly in the global Internet environment. Then, we present two technology scenarios typical of the landscape facing modern electronic surveillance. Combining what is publicly known about surveillance practices today, along with common legal requirements, we are able to present scenarios that illustrate many of the key risks that exceptional access will entail.

We do not suggest that our own interpretation of Comey's stated requirements serve as a basis for regulation but merely as a starting point for discussion. If officials in the UK or US disagree with our interpretation, we urge them to state their requirements clearly. Only then can a rigorous technical analysis be conducted in an open, transparent manner. Such analysis is crucial in a world that is so completely reliant on secure communications for every aspect of daily lives, from nations' critical infrastructure, to government, to personal privacy in daily life, to all matters of business from the trivial to the global.

1.2 Findings from the 1997 analysis of key escrow systems

We begin by reviewing the findings on the risks of key recovery/key escrow systems from a paper that many of us wrote almost 20 years ago[1]. Many of us came together then to

examine the security risks of ensuring law enforcement access to encrypted information. We found that any key escrow system had basic requirements that placed substantial costs on end users, and that these costs would have been too difficult and expensive to implement. For law enforcement to have quick and reliable access to plaintext, every key escrow system required the existence of highly sensitive yet perennially available secret keys. This requirement alone inevitably leads to an increased risk of exposure, inflated software complexity, and high economic costs.

The first downside is increased risk of a security incident. An organization that holds an escrow key could have a malicious insider that abuses its power or leaks that organization's key. Even assuming an honest agency, there is an issue of competence: cyberattacks on keyholders could easily result in catastrophic loss.

The additional complexity of a key escrow system compounds these risks. At the time, all openly proposed key escrow solutions had major flaws that could be exploited; even normal encryption was difficult to implement well, and key escrow made things much harder. Another source of complexity was the scale of a universal key recovery system — the number of agents, products, and users involved would be immense, requiring an escrow system well beyond the technology of the time. Further, key escrow threatened to increase operational complexity: a very large number of institutions would have to securely and safely negotiate targeting, authentication, validity, and information transfer for lawful information access.

All of the above factors raise costs. Risks of exposure, for instance, change the threat landscape for organizations, which must then worry about mistaken or fraudulent disclosures. The government would have increased bureaucracy to test and approve key recovery systems. Software vendors would have to bear the burden of increased engineering costs. In 1997, we found that systems enabling exceptional access to keys would be inherently less secure, more expensive, and much more complex than those without. This result helped policymakers decide against mandated exceptional access.

1.3 What has changed and what remains the same since 1990s?

It is impossible to operate the commercial Internet or other widely deployed global communications network with even modest security without the use of encryption. An extensive debate in the 1980s and 1990s about the role of encryption came to this conclusion once before. Today, the fundamental technical importance of strong cryptography and the difficulties inherent in limiting its use to meet law enforcement purposes remain the same. What has changed is that the scale and scope of systems dependent on strong encryption are far greater, and our society is far more reliant on far-flung digital networks that are

under daily attack.

In the early 1990s, the commercialization of the Internet was being thwarted by US government controls on encryption — controls that were in many ways counterproductive to long-term commercial and national security interests. A 1996 United States National Academy of Science study concluded that, "On balance, the advantages of more widespread use of cryptography outweigh the disadvantages" [4, p. 6]. Four years later, partly in response to pressures from industry, partly in response to the loosening of cryptographic export controls by the European Union, partly because crypto export controls were declared unconstitutional by US Circuit Courts, and partly because of increasing reliance on electronic communications and commerce, the US relaxed export controls on encryption [7].

The Crypto Wars actually began in the 1970s, with conflicts over whether computer companies such as IBM and Digital Equipment Corporation could export hardware and software with strong encryption, and over whether academics could publish cryptographic research freely. They continued through the 1980s over whether the NSA or the National Institute of Standards and Technology (NIST) would control the development of cryptographic standards for the non-national security side of the government (NIST was given the authority under the 1987 Computer Security Act). They came to full force during the 1990s, when the US government, largely through the use of export controls, sought to prevent companies such as Microsoft and Netscape from using strong cryptography in web browsers and other software that was at the heart of the growing Internet. The end of the wars — or the apparent end — came because of the Internet boom.

In many ways, the arguments are the same as two decades ago. US government cryptographic standards — the Data Encryption Standard then, the Advanced Encryption Standard now — are widely used both domestically and abroad. We know more now about how to build strong cryptosystems, though periodically we are surprised by a break. However, the real security challenge is not the mathematics of cryptosystems; it is engineering, specifically the design and implementation of complex software systems. Two large government efforts, healthcare.gov and the FBI Trilogy program, demonstrate the difficulties that scale and system integration pose in building large software systems. Healthcare.gov, the website implementing the president's signature healthcare program, failed badly in its initial days, unable to serve more than a tiny percentage of users [8]. A decade earlier, five years of effort spent building an electronic case file system for the FBI — an effort that cost $170 million — was abandoned as unworkable [9].

At one level, the worst has not come to pass — the power grid, the financial system, critical infrastructure in general, and many other systems all function reliably using com-

plex software. On another level, the worst is occurring daily. Recent breaches for financial gain include: T.J. Maxx, theft of 45 million credit card records [10]; Heartland Payment Systems, compromise of 100 million credit cards [11]; Target, compromise of 40 million credit cards; Anthem, collection of names, addresses, birthdates, employment and income information, and Social Security numbers of 80 million people that could result in identity theft [12].

Attacks on government agencies are also increasing. A set of 2003 intrusions targeting US military sites collected such sensitive data as specifications for Army helicopter mission planning systems, Army and Air Force flight-planning software, and schematics for the Mars Orbiter Lander [13]. Such theft has not only been from the defense industrial base, but has included the pharmaceuticals, Internet, biotechnology and energy industries. In 2010, then Deputy Secretary of Defense William Lynn concluded, "Although the threat to intellectual property is less dramatic than the threat to critical national infrastructure, it may be the most significant cyberthreat that the United States will face over the long term" [14].

The December 2014 North Korean cyberattacks against Sony, the first such by a nation-state, resulted in large headlines. But the 2011 theft from RSA/EMC of the seed keys — initial keys used to generate other keys — in hardware tokens used to provide two-factor authentication [15], and the recent theft of personnel records from the US Office of Personnel Management are far more serious issues. The former undermined the technical infrastructure for secure systems, while the latter, by providing outsiders with personal information of government users, creates leverage for many years to come for potential insider attacks, undermining the social infrastructure needed to support secure governmental systems — including any future system for exceptional access. And while attacks against critical infrastructure have not been significant, the potential to do so has been demonstrated in test cases [16] and in an actual attack on German steel mill that caused significant damage to a blast furnace [17].

As exceptional access puts the security of Internet infrastructure at risk, the effects will be felt every bit as much by government agencies as by the private sector. Because of cost and Silicon Valley's speed of innovation, beginning in the mid-1990s, the US government moved to a commercial off the shelf (COTS) strategy for information technology equipment, including communications devices. In 2002, Information Assurance Technical Director Richard George told a Black Hat audience that "NSA has a COTS strategy, which is: when COTS products exist with the needed capabilities, we will encourage their use whenever and wherever appropriate . . . "[18]. Such a COTS solution makes sense, of course, only if the private sector technologies the government uses are secure.

Communications technologies designed to comply with government requirements for backdoors for legal access have turned out to be insecure. For ten months in 2004 and 2005, 100 senior members of the Greek government (including the Prime Minister, the head of the Ministry of National Defense and the head of the Ministry of Justice) were wiretapped by unknown parties through lawful access built into a telephone switch owned by Vodafone Greece [19]. In 2010 an IBM researcher observed that a Cisco architecture for enabling lawful interception in IP networks was insecure.[2] This architecture had been public for several years, and insecure versions had been implemented by several carriers in Europe [20]. And when the NSA examined telephone switches built to comply with government-mandated access for wiretapping, it discovered security problems with *all* the switches submitted for testing[21]. Embedding exceptional access requirements into communications technology will ensure even more such problems, putting not only private-sector systems, but government ones, at risk.

Speaking on the topic of law enforcement access and systems security, Vice Chairman of the Joint Chiefs of Staff Admiral James A. Winnefeld recently remarked, "But I think we would all win if our networks are more secure. And I think I would rather live on the side of secure networks and a harder problem for Mike [NSA Director Mike Rogers] on the intelligence side than very vulnerable networks and an easy problem for Mike and part of that, it's not only is the right thing to do, but part of that goes to the fact that we are more vulnerable than any other country in the world, on our dependence on cyber. I'm also very confident that Mike has some very clever people working for him, who might actually still be able to get some good work done."

While the debate over mandated law enforcement access is not new, it does take on added urgency in today's world. Given our growing dependence on the Internet, and the urgent need to make this and other digital infrastructures more secure, any move in the direction of decreased security should be looked upon with extreme skepticism. Once before, when considering this issue, governments around the world came to the conclusion that designing in exceptional access provisions to vital systems would increase security risk and thwart innovation. As the remainder of this paper will show, such measures are even riskier today.

[2] It is worth noting that the router's design was based on standards put forth by the European Telecommunications Standards Institute.

2 Scenarios

Law enforcement authorities have stated a very broad requirement for exceptional access. Yet there are many details lacking including the range of systems to which such requirements would apply, the extraterritorial application, whether anonymous communications would be allowed, and many other variables. To analyze the range of security risks that may arise in commonly used applications and services, we examine two popular scenarios: encrypted real-time messaging services and devices such as smartphones that use strong encryption to lock access to the device.

2.1 Scenario 1: Providing exceptional access to globally distributed, encrypted messaging applications

Imagine a massively distributed global messaging application on the Internet currently using end-to-end encryption. Many examples of such systems actually exist, including Signal, which is available on iPhone and Android, Off-the-Record (OTR), a cryptography-enabling plug-in for many popular computer chat programs, and the often cited TextSecure and WhatsApp. Could one provide a secure application while meeting law enforcement exceptional access requirements?

To provide law enforcement access to encrypted data, one natural approach is to provide law enforcement direct access to keys that can be used to decrypt the data, and there is a frequently suggested and seemingly quite attractive mechanism for escrowing decryption keys. Data is typically encrypted — either for storage or transmission — with a symmetric key,[3] and many data transmission protocols (e.g., the Transport Layer Security (TLS) protocol) can operate in a mode where the data to be sent is encrypted with a symmetric key that is in turn encrypted with a public key[4] associated with the intended recipient. This encrypted symmetric key then travels with the encrypted data, and the recipient accesses the data by first using its private key to decrypt the symmetric key and then using the symmetric key to decrypt the data.

A common suggestion is to augment this approach by encrypting the symmetric key a second time — this time with a special escrowing public key. If the data is then transmitted, two encryptions of the symmetric key accompany the data — one with the public key of the intended recipient and one with a public key associated with an escrow agent. If the data has been encrypted with a symmetric key for storage rather than

[3] A symmetric key is one that is used for both encryption and decryption.

[4] A public key is used to encrypt data that can then be decrypted only by an entity in possession of an associated private key.

transmission, the symmetric key might be encrypted with the public key of an escrow agent and this escrowed key could remain with the encrypted data. If a law enforcement entity obtains this encrypted data either during transmission or from storage the escrow agent could be enlisted to decrypt the symmetric key, which could then be used to decrypt the data.

There are, however, three principal impediments to using this approach for third-party escrow. Two are technical and the third is procedural.

The first technical obstacle is that although the mode of encrypting a symmetric key with a public key is in common use, companies are aggressively moving away from it because of a significant practical vulnerability: *if an entity's private key is ever breached, all data ever secured with this public key is immediately compromised.* Because it is unwise to assume a network will never be breached, a single failure should never compromise all data that was ever encrypted.

Thus, companies are moving towards *forward secrecy*, an approach that greatly reduces the exposure of an entity that has been compromised. With forward secrecy, a new key is negotiated with each transaction, and long-term keys are used only for authentication. These transaction (or *session*) keys are discarded after each transaction — leaving much less for an attacker to work with. When a system with forward secrecy is used, an attacker who breaches a network and gains access to keys can only decrypt data from the time of the breach until the breach is discovered and rectified; historic data remains safe. In addition, since session keys are destroyed immediately after the completion of each transaction, an attacker must interject itself into the process of each transaction in real time to obtain the keys and compromise the data.[5]

The security benefits make clear why companies are rapidly switching to systems that provide forward secrecy.[6] However, the requirement of key escrow creates a long-term vulnerability: if *any* of the private escrowing keys are *ever* compromised, then *all* data that *ever* made use of the compromised key is permanently compromised. That is, in order to accommodate the need for surreptitious, third-party access by law enforcement agencies, messages will have to be left open to attack by anyone who can obtain a copy of one of the many copies of the law enforcement keys. *Thus all known methods of achieving third-party escrow are incompatible with forward secrecy.*

Innovations providing better forward secrecy also support a broad social trend: users are moving en masse to more ephemeral communications. Reasons for moving to ephemeral communications range from practical decisions by corporations to protect proprietary in-

[5] Lack of forward secrecy was identified in the 1997 paper [1] as a weakness of key escrow systems then. Since that time, the need for forward secrecy has grown substantially.

[6] See [22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32].

formation from industrial espionage to individuals seeking to protect their ability to communicate anonymously and avoid attack by repressive governments. Many corporations delete email after 90 days, while individuals are moving from email to chat and using services like Snapchat where messages vanish after reading. Leading companies such as Twitter, Microsoft, and Facebook are supporting the move to transient messaging, and using modern security mechanisms to support it. This social and technical development is not compatible with retaining the means to provide exceptional access.

The second technical obstacle is that current best practice is often to use *authenticated encryption*, which provides *authentication* (ensuring that the entity at the other end of the communication is who you expect, and that the message has not been modified since being sent) as well as *confidentiality* (protecting the privacy of communications, including financial, medical, and other personal data). However, disclosure of the key for authenticated encryption to a third party means the message recipient is no longer provided with technical assurance of the communication's integrity; disclosure of the key allows the third party not only to *read* the encrypted traffic but also to *forge* traffic to the recipient and make it look as if it is coming from the original sender. Thus disclosing the key to a third party creates a new security vulnerability. Going back to the encryption methods of the 1990s, with separate keys for encryption and authentication, would not only double the computational effort required, but introduce many opportunities for design and implementation errors that would cause vulnerabilities.

The third principal obstacle to third-party key escrow is procedural and comes down to a simple question: who would control the escrowed keys? Within the US, one could postulate that the FBI or some other designated federal entity would hold the private key necessary to obtain access to data and that judicial mechanisms would be constructed to enable its use by the plethora of federal, state, and local law enforcement entities. However, this leaves unanswered the question of what happens outside a nation's borders. Would German and French public- and private-sector organizations be willing to use systems that gave the US government access to their data — especially when they could instead use locally built systems that do not? What about Russia? Would encrypted data transmitted between the US and China need to have keys escrowed by both governments? Could a single escrow agent be found that would be acceptable to both governments? If so, would access be granted to just one of the two governments or would both need to agree to a request?

These difficult questions must be answered before any system of exceptional access can be implemented. Such an architecture would require global agreements on how escrow would be structured, often against the best interests of certain countries' domestic goals,

together with mandates in virtually all nations to only sell and use compliant systems.

2.2 Scenario 2: Exceptional access to plaintext on encrypted devices such as smartphones

Imagine a smartphone platform vendor that seeks to accommodate law enforcement exceptional demands. When law enforcement comes into possession of a device, perhaps at a crime scene, and then obtains the necessary legal authorization (in the US this would be a warrant as a result of *Riley v. California*), the agent collects a unique identifying number from the device through some service mechanism, and then sends a request to the platform vendor to unlock the device remotely or provide the keys necessary for law enforcement to unlock the device locally.

At first glance, providing access to plaintext on devices — laptop hard drives, smartphones, tablets — is straightforward. Indeed, many corporations already escrow device encryption keys. However, and as is frequently the case, scaling up a corporate mechanism to a global one is hard.

When encrypting device storage, the user-entered passphrase is generally not used directly as an encryption key. There are many reasons for this; from a usability perspective, the most important one is to make it easier for the user to change the passphrase. If the key were used directly, it would be a time-consuming process to decrypt and re-encrypt the entire device when the passphrase is changed. Instead, a random key is used for bulk encryption; the user-supplied key (called the Key-Encrypting Key, or KEK) is used to encrypt the random key.

To protect against brute-force attacks against the user's passphrase, the device vendor may go a step further and combine it with a device-specific unique identifier to produce the KEK. In the iPhone, the KEK is stored in a special tamper-resistant processor that limits the guess rate to once every 80 milliseconds. This protects device owners against, for example, sophisticated thieves who might try to gain access to things like banking passwords. But regardless of how the KEK is generated, obtaining access to the plaintext requires that the device-encrypting key be encrypted under some additional key or keys. These could be manufacturer-owned keys or keys belonging to one or more law enforcement agencies. Either choice is problematic[33].

If a vendor-supplied key is used, some sort of network protocol to decrypt the device key is necessary. This request must be authenticated, but how? How can the vendor have secure credentials for all of the thousands of law enforcement agencies around the world? How can the result be strongly bound to the device, to prevent unscrupulous agencies from requesting keys to devices not in their lawful possession? These are not

easy requirements to meet, especially for devices that will not even boot without a valid key. They are likely to require changes to security hardware or to the software that drives it; both are difficult to do properly. Fixing glitches — especially security glitches — in deployed hardware is expensive and often infeasible.

Providing devices with law enforcement keys is equally difficult. Again, how can the vendor know who supplied the keys? How are these keys to be changed? [7] How many keys can be installed without causing unacceptable slowdowns? Another alternative is to require that law enforcement ship devices back to the vendor for exceptional access decryption. However, it will still be necessary to store over long periods of time keys that can decrypt all of the sensitive data on devices. This only shifts the risks of protecting these keys to the device manufacturers.

Some would argue that per-country keys could be a sales requirement. That is, all devices sold within the US would be required to have, say, a preinstalled FBI-supplied key. That, however, does not suffice for devices brought in by travelers — and those are the devices likely to be of interest in terrorism investigations. A requirement that keys be installed at the border is also problematic. There are no standard input ports or key-loading mechanisms; furthermore, it would expose American travelers to malware installed by border guards in other countries [34, 35].

2.3 Summary of risks from the two scenarios

Designing exceptional access into today's information services and applications will give rise to a range of critical security risks. First, major efforts that the industry is making to improve security will be undermined and reversed. Providing access over any period of time to thousands of law enforcement agencies will necessarily increase the risk that intruders will hijack the exceptional access mechanisms. If law enforcement needs to look backwards at encrypted data for one year, then one year's worth of data will be put at risk. If law enforcement wants to assure itself real time access to communications streams, then intruders will have an easier time getting access in real time, too. This is a trade-off space in which law enforcement cannot be guaranteed access without creating serious risk that criminal intruders will gain the same access.

Second, the challenge of guaranteeing access to multiple law enforcement agencies in multiple countries is enormously complex. It is likely to be prohibitively expensive and also an intractable foreign affairs problem.

Simple requirements can yield simple solutions (e.g. a door lock). But the requirements

[7] We note that some pieces of malware, such as Stuxnet and Duqu 2, have relied on code-signing keys issued to legitimate companies. When a key is compromised, it must be replaced.

of law enforcement access to encrypted data are inherently complex and, as we have already shown, nearly contradictory. Complex or nearly contradictory requirements yield brittle, often-insecure solutions. As NSA's former head of research testified in 2013:

> "When it comes to security, complexity is not your friend. Indeed it has been said that complexity is the enemy of security. This is a point that has been made often about cybersecurity in a variety of contexts including, technology, coding and policy. The basic idea is simple: as software systems grow more complex, they will contain more flaws and these flaws will be exploited by cyber adversaries." [36]

We have a very real illustration of the problem of complexity in a recent analysis of one of the most important security systems on the Internet: SSL/TLS. Transport Layer Security (TLS) and its predecessor Secure Socket Layer (SSL) are the mechanisms by which the majority of the web encrypts its traffic — every time a user logs into a bank account, makes an electronic purchase, or communicates over a social network, that user is trusting SSL/TLS to function properly. All a user needs to know of all of this complexity is that the lock or key icon shows up in the browser window. This indicates that the communication between the user and the remote website is secure from interception.

Unfortunately, writing code that correctly implements such cryptographic protocols has proven difficult; weakened protections makes it harder still. For instance, OpenSSL, the software used by about two-thirds of websites to do TLS encryption, has been plagued with systems-level bugs resulting in catastrophic vulnerabilities. The now-infamous Heartbleed bug was caused by a missing bounds check, an elementary programming error that lurked in the code for two years, leaving 17% of *all* websites vulnerable to data theft. More recent vulnerabilities, however, were caused by legacy restrictions on the exportation of cryptographic algorithms, dating back to the Crypto Wars. The fact that there are so many different implementations of TLS, all of which have to interoperate to make the Web secure, has proven to be a real source of security risk [37]. Website operators are reluctant to switch to more secure protocols if this will lose them even a few percent of prospective customers who are still using old software, so vulnerabilities introduced deliberately during the Crypto Wars have persisted to this day. Introducing complex new exceptional access requirements will similarly add more security bugs that will lurk in our software infrastructure for decades to come.

Third, there are broader risks for poorly deployed surveillance technology. Exceptional access mechanisms designed for law enforcement use have been exploited by hostile actors in the past. Between 1996 and 2006, it appears that insiders at Telecom Italia enabled the

wiretapping of 6,000 people, including business, financial, and political leaders, judges, and journalists [38]. In a country of 60 million, this means that no major business or political deal was truly private. The motivation here appeared to be money, including the possibility of blackmail. As we mentioned earlier, from 2004 to 2005, the cell phones of 100 senior members of the Greek government, including the Prime Minister, the head of the Ministry of National Defense, the head of the Ministry of Justice, and others. Vodafone Greece had purchased a telephone switch from Ericsson. The Greek phone company had not purchased wiretapping capabilities, but these were added during a switch upgrade in 2003. Because Vodafone Greece had not arranged for interception capabilities, the company did not have the ability to access related features, such as auditing. Nevertheless, someone acting without legal authorization was able to activate the intercept features and keep them running for ten months without being detected. The surveillance was uncovered only when some text messages went awry. Although the techniques of how it was done are understood, who was behind the surveillance remains unknown[19].

Next, there are the broader costs to the economy. Economic growth comes largely from innovation in science, technology, and business processes. At present, technological progress is largely about embedding intelligence — software and communications — everywhere. Products and services that used to be standalone now come with a mobile phone app, an online web service, and business models that involve either ads or a subscription. Increasingly these are also "social", so you can chat to your friends and draw them into the vendor's marketing web. Countries that require these new apps and web services to have their user-to-user communications functions authorized by the government will be at a significant disadvantage. At present, the world largely uses US apps and services, rather than the government-approved ones from Russia and China. This provides enormous leverage to US businesses.

Finally, this market advantage gives real benefits not just economically but in terms of soft power and moral leadership. The open Internet has long been a foreign policy goal of the US and its allies for a lot of good reasons. The West's credibility on this issue was damaged by the Snowden revelations, but can and must recover. Lawmakers should not risk the real economic, geopolitical, and strategic benefits of an open and secure Internet for law enforcement gains that are at best minor and tactical.

3 Security impact of common law enforcement requirements with exceptional access

Since there is no specific statement of law enforcement requirements for exceptional access, we consider what we understand to be a very general set of electronic surveillance needs applicable in multiple jurisdictions around the world. Our goal here is to understand the general nature of security risks associated with the application of exceptional access requirements in the context of traditional categories of electronic surveillance. Law enforcement agencies in different countries have presented different requirements at different times, which we will treat under four headings: access to communications content, access to communications data, access to content at rest, and covert endpoint access. All types of access must be controlled and capable of being audited according to local legal requirements; for example, under the requirements of US law, one must respect the security and privacy of non-targeted communications.[8]

3.1 Access to communications content

Most police forces are permitted to access suspect data. In countries with respect for the rule of law, such access is carefully regulated by statute and supervised by an independent judiciary, though most of the world's population do not enjoy such legal protections. Law enforcement access might be to a central database of unencrypted messages where this exists at a central provider. Where there is no central database, such as for a telephone or video call, the police must tap the communication as it happens. How might an exceptional access requirement be implemented to enable for access to communications content? If the data is encrypted, the most obvious mechanism to allow for police access would require that traffic between Alice in country X and Bob in country Y would have its session key also encrypted under the public keys of the police forces in both X and Y, or of third parties trusted by them. This, however, raises serious issues.

First, any escrow requirement will restrict other important security functionality such as forward secrecy, the use of transient identities, and strong location privacy. As illustrated in the scenario analysis above, an exceptional access requirement overlaid on the traditional content surveillance will put the security of the content at risk. To the extent that capabilities exist to provide law enforcement exceptional access, they can be abused by others.

Second, the global nature of Internet services makes compliance with exceptional access

[8] In the USA, 47 USC 1002(a)(4)

rules both hard to define and hard to enforce. If software sold in country X will copy all keys to that country's government, criminals might simply buy their software from countries that don't cooperate; thus, US crooks might buy their software from Russia. And if software automatically chooses which governments to copy using a technique such as IP geolocation, how does one prevent attacks based on location spoofing? While it is possible to design mobile phone systems so that the host jurisdictions have access to the traffic (so long as the users do not resort to VoIP), this is a much harder task for general-purpose messaging applications.

Third, one might have to detect or deter firms that do not provide exceptional access, leading to issues around certification and enforcement. For example, if the US or the UK were to forbid the use of messaging apps that are not certified under a new escrow law, will such apps be blocked at the national firewall? Will Tor then be blocked, as in China? Or will it simply become a crime to use such software? And what is the effect on innovation if every new communications product must go through government-supervised evaluation against some new key escrow protection profile?

3.2 Access to communications data

Communications data traditionally meant call detail records and (since mobile phones became common) caller location history; it was obtained by subpoena from phone companies, and is used in the investigation of most serious violent crimes such as murder, rape, and robbery. Communications data remains widely available as service providers keep it for some time for internal purposes. However, police forces outside the US complain that the move to globalized messaging services makes a lot of data harder to obtain. For example, emails are now typically encrypted using TLS; that is, the message is encrypted between the user's computer and the service provider (e.g., Google for Gmail, Microsoft for Hotmail, etc.). Thus, to acquire the communications in plaintext, law enforcement must serve the email provider with a court order. A new UK surveillance law may require message service firms like Apple, Google, and Microsoft to honor such requests expeditiously and directly as a condition of doing business in the UK. So will there be uniform provisions for access to communications data subject to provisions for warrants or subpoenas, transparency, and jurisdiction?

As already noted, determining location is not trivial, and cheating (using foreign software, VPNs, and other proxies) could be easy. Criminals would turn to noncompliant messaging apps, raising issues of enforcement; aggressive enforcement might impose real costs on innovation and on industry generally.

3.3 Access to data at rest

Communications data are one instance of the general problem of access to data at rest. Almost all countries allow their police forces access to data. Where basic rule of law is in place, access is under the authority of a legal instrument such as a warrant or subpoena, subject to certain limits. Many corporations already insist on escrowing keys used to protect corporate data at rest (such as BitLocker on corporate laptops). So this is one field with an already deployed escrow "solution": a fraud investigator wanting access to a London rogue trader's laptop can simply get a law enforcement officer to serve a decryption notice on the bank's CEO. But still, many of the same problems arise. Suspects may use encryption software that does not have escrow capability, or may fail to escrow the key properly, or may claim they have forgotten the password, or may actually have forgotten it. The escrow authority may be in another jurisdiction, or may be a counterparty in litigation. In other words, what works tolerably well for corporate purposes or in a reasonably well-regulated industry in a single jurisdiction simply does not scale to a global ecosystem of highly diverse technologies, services, and legal systems.

Another thorny case of access to data at rest arises when the data is only present on, or accessible via, a suspect's personal laptop, tablet, or mobile phone. At present, police officers who want to catch a suspect using Tor services may have to arrest him while his laptop is open and a session is live. Law enforcement agencies in some countries can get a warrant to install malware on a suspect's computer. Such agencies would prefer antivirus companies not to detect their malware; some might even want the vendors to help them, perhaps via a warrant to install an upgrade with a remote monitoring tool on a device with a specific serial number. The same issues arise with this kind of exceptional access, along with the issues familiar from covert police access to a suspect's home to conduct a surreptitious search or plant a listening device. Such exceptional access would gravely undermine trust and would be resisted vigorously by vendors.

4 Principles at stake and unanswered questions

With people's lives and liberties increasingly online, the question of whether to support law enforcement demands for guaranteed access to private information has a special urgency, and must be evaluated with clarity. From a public policy perspective, there is an argument for giving law enforcement the best possible tools to investigate crime, subject to due process and the rule of law. But a careful scientific analysis of the likely impact of such demands must distinguish what might be desirable from what is technically possible. In this regard, a proposal to regulate encryption and guarantee law enforcement access

centrally feels rather like a proposal to require that all airplanes can be controlled from the ground. While this might be desirable in the case of a hijacking or a suicidal pilot, a clear-eyed assessment of how one could design such a capability reveals enormous technical and operational complexity, international scope, large costs, and massive risks — so much so that such proposals, though occasionally made, are not really taken seriously.

We have shown that current law enforcement demands for exceptional access would likely entail very substantial security risks, engineering costs, and collateral damage. If policy-makers believe it is still necessary to consider exceptional access mandates, there are technical, operational, and legal questions that must be answered in detail before legislation is drafted. From our analysis of the two scenarios and general law enforcement access requirements presented earlier in the paper, we offer this set of questions.

4.1 Scope, limitations, and freedoms

The first set of questions that an exceptional access proposal must address concerns the scope of applicability of the exceptional access requirement, any limitations on the mandate, and what user freedoms would remain protected under such proposals. Questions such as these arise in this category:

1. Are all systems that use encryption covered, or just some? Which ones?

2. Do all online communications and information platforms have to provide access to plain text, or merely provide keys to agencies that had already collected ciphertext using technical means?

3. Would individuals, corporations, nonprofit institutions, or governments be allowed to deploy additional encryption services on top of those systems with exceptional access? Would those user-installed systems also have to meet exceptional access requirements?

4. Would machine-to-machine systems be covered? What about Internet of Things and industrial control (SCADA) systems? Much information exchange is from one machine to another, such as communicating personal health data from a sensor to a smartphone, field-based agricultural sensing devices to tractors, or load balancing controls in electric power, gas, oil and water distribution systems.

5. How would cross-border regulatory differences be resolved? Would technology developers have to meet different exceptional access requirements in each jurisdiction where their systems are used? Or would there be a globally harmonized set of regulatory requirements?

6. How can the technical design of an exceptional access system prevent mass surveillance that would covertly violate the rights of entire populations, while still allowing covert targeted surveillance of small numbers of suspects as an actual "exception" to a general rule of citizen privacy?

7. Would there be an exception for research and teaching?

8. Could companies refuse to comply with exceptional access rules based on a fear of violating human rights?

9. Would anonymous communications, widely recognized as vital to democratic societies, be allowed?

4.2 Planning and design

Designing the technology and planning the administrative procedures that would be needed to implement a comprehensive exceptional access system raises many questions:

1. What are the target cost and benefit estimates for such a program? No system is cost-free and this one could be very expensive, especially if it has to accommodate a large number of providers, such as today's millions of app developers.

2. What security and reliability measures would be established for the design? How would system prototypes be tested? How long would companies have to comply with exceptional access rules?

3. How would existing services and products be treated if they do not comply with exceptional access rules? Would providers have to redesign their systems? What if those systems cannot accommodate exceptional access requirements?

4. Who would be involved in the design of the systems and procedures — just the US government, or would other governments be invited to participate? Could foreign technology providers such as Huawei participate in the design discussions?

5. Would the technical details of the program be made public and open for technical review? What level of assurance would be provided for the design?

6. We note that it generally takes many years after a cryptographic protocol is published before it is deemed secure enough for actual use. For example, the Needham-Schroeder public-key protocol, first published in 1978 [39], was discovered to have security flaw only in 1995 by Gavin Lowe (17 later!) [40].

4.3 Deployment and operation

Once regulations are established and technical design parameters set, there would remain questions about how systems would be deployed, who would supervise and regulate compliance, and how the design of the system would evolve to address inevitable technical and operational bugs that emerge. We know of no system that is designed perfectly the first time, and it is well understood that maintenance, support, and evolution of existing systems constitutes a major expense.

1. Who would supervise compliance? Would an existing regulatory agency such as the FCC be given jurisdiction over the entire process? How would other countries regulate US domestic and foreign services? Would there be a global harmonization of rules regulation and enforcement? Would the International Telecommunications Union have a role in setting and enforcing requirements?

2. Would global technical standards be required? How would these be developed and enforced? How would be such standards be changed/improved/patched? Would traditional standards bodies such as the UN International Telecommunications Union T-sector or ISO set standards, or would the world look to Internet standards bodies such as the IETF and the World Wide Web Consortium? How would the world converge on one set of standards?

3. Would the US government provide reference software libraries implementing the desired functionality?

4. Would programs and apps need to be certified before they were allowed to be sold? Who would test or certify that programs produced operate as intended?

5. Who would be liable if the plaintext-disclosure mechanisms were buggy (either in design or in implementation), causing the disclosure of all citizens' information? More generally, what would happen when (not if) critical secret information was revealed, such as the private keys that allow encrypted data to be read by anyone, that destroyed the privileged position of law enforcement?

6. How many companies would withdraw all but local sales staff from markets where exceptional access was mandated in ways that clashed with their business strategies or the rights of users in other countries, as Google already has done from China and Russia?

4.4 Evaluation, assessment, and evolution

Large systems exist because successful systems evolve and grow. Typically, this evolution happens through interaction guided by the institution (software company, government agency, or open-source community) responsible for the system. A system that evolves subject to a set of constraints, such as medical systems that need to maintain a safety case or flight control systems that need to maintain not just a safety case but also need to meet real-time performance requirements, evolve less quickly and at more cost. If all systems that communicate must in future evolve subject to an exceptional access constraint, there will be real costs, which are hard to quantify, since the question of who exactly would be responsible for establishing and policing the exceptional access constraint is not clear. However that question is answered, the following further issues will arise.

1. What oversight program would be required to monitor the effectiveness, cost, benefits, and abuse of exceptional access?

2. What sunset provisions would be build into legislation for such a program? What conditions would be in place for its termination (e.g., for lack of sufficient benefit, for excessive cost, or for excessive abuse)?

3. One unintended consequence of such a program may be a much-reduced use of crypto altogether. This would further weaken our already fragile and insecure information infrastructure, so how do we incentivize companies to continue encrypting sensitive user communications?

4. A further unintended consequence of such a program might be to make the US and other participating countries less welcoming to technological innovation; diminishing or displacing innovation may have consequences for economic growth and national security. How will these economic impacts be assessed before an exceptional access program is mandated? Further, what economic effect would be considered too impactful for exceptional access to be considered worthwhile?

5 Conclusion

Even as citizens need law enforcement to protect themselves in the digital world, all policy-makers, companies, researchers, individuals, and law enforcement have an obligation to work to make our global information infrastructure more secure, trustworthy, and resilient. This report's analysis of law enforcement demands for exceptional access to private communications and data shows that such access will open doors through which

criminals and malicious nation-states can attack the very individuals law enforcement seeks to defend. The costs would be substantial, the damage to innovation severe, and the consequences to economic growth difficult to predict. The costs to developed countries' soft power and to our moral authority would also be considerable. Policy-makers need to be clear-eyed in evaluating the likely costs and benefits. It is no surprise that this report has ended with more questions than answers, as the requirements for exceptional access are still vague. If law enforcement wishes to prioritize exceptional access, we suggest that they need to provide evidence to document their requirements and then develop genuine, detailed specifications for what they expect exceptional access mechanisms to do. As computer scientists and security experts, we are committed to remaining engaged in the dialogue with all parts of our governments, to help discern the best path through these complex questions.

References

[1] H. Abelson, R. N. Anderson, S. M. Bellovin, J. Benaloh, M. Blaze, W. Diffie, J. Gilmore, P. G. Neumann, R. L. Rivest, J. I. Schiller, and others, "The risks of key recovery, key escrow, and trusted third-party encryption," 1997. [Online]. Available: http://academiccommons.columbia.edu/catalog/ac:127127

[2] Advanced Telephony Unit, Federal Bureau of Investigation, "Telecommunications Overview, slide on Encryption Equipment," 1992. [Online]. Available: https://www.cs.columbia.edu/~smb/Telecommunications_Overview_1992.pdf

[3] E. Nakashima, "Chinese hackers who breached Google gained access to sensitive data, U.S. officials say," *The Washington Post*, May 2013. [Online]. Available: https://www.washingtonpost.com/world/national-security/chinese-hackers-who-breached-google-gained-access-to-sensitive-data-us-officials-say/2013/05/20/51330428-be34-11e2-89c9-3be8095fe767_story.html

[4] K. W. Dam, H. S. Lin, and others, *Cryptography's role in securing the information society*. National Academies Press, 1996.

[5] James B. Comey, "Going Dark: Are Technology, Privacy, and Public Safety on a Collision Course?" Oct. 2014, speech at the Brookings Institution. [Online]. Available: https://www.fbi.gov/news/speeches/going-dark-are-technology-privacy-and-public-safety-on-a-collision-course

[6] David Cameron, "PM: spy agencies need more powers to protect Britain," Jan. 2015. [Online]. Available: https://embed.theguardian.com/embed/video/uk-news/video/2015/jan/12/david-cameron-spy-agencies-britain-video

[7] W. Diffie and S. Landau, *Privacy on the Line: The Politics of Wiretapping and Encryption*. Cambridge, Mass: The MIT Press, Jan. 1998.

[8] Paul Ford, "The Obamacare Website Didn't Have to Fail. How to Do Better Next Time," Oct. 2013. [Online]. Available: http://www.bloomberg.com/bw/articles/2013-10-16/open-source-everything-the-moral-of-the-healthcare-dot-gov-debacle

[9] D. Eggen and G. Witte, "The FBI's Upgrade That Wasn't," *The Washington Post*, Aug. 2006. [Online]. Available: http://www.washingtonpost.com/wp-dyn/content/article/2006/08/17/AR2006081701485.html

[10] Jaikumar Vijayan, "TJX data breach: At 45.6m card numbers, it's the biggest ever," Mar. 2007. [Online]. Available: http://www.computerworld.com/article/2544306/security0/tjx-data-breach--at-45-6m-card-numbers--it-s-the-biggest-ever.html

[11] Brian Krebs, "Security fix - payment processor breach may be largest ever," Jan. 2009. [Online]. Available: http://voices.washingtonpost.com/securityfix/2009/01/payment_processor_breach_may_b.html

[12] R. Abelson and M. Goldstein, "Anthem Hacking Points to Security Vulnerability of Health Care Industry," *The New York Times*, Feb. 2015. [Online]. Available: http://www.nytimes.com/2015/02/06/business/experts-suspect-lax-security-left-anthem-vulnerable-to-hackers.html

[13] N. Thornburgh, "The Invasion of the Chinese Cyberspies," *Time*, Aug. 2005. [Online]. Available: http://content.time.com/time/magazine/article/0,9171,1098961,00.html

[14] William J. Lynn III, "Defending a New Domain," *Foreign Affairs*, Oct. 2010. [Online]. Available: https://www.foreignaffairs.com/articles/united-states/2010-09-01/defending-new-domain

[15] Arthur Coviello, "Open Letter from Arthur Coviello, Executive Chairman, RSA, Security Division of EMC, to RSA customers," Mar. 2011.

[16] Jeanne Meserve, "Sources: Staged cyber attack reveals vulnerability in power grid - CNN.com," *CNN*, Sep. 2007. [Online]. Available: http://www.cnn.com/2007/US/09/26/power.at.risk/index.html?iref=topnews

[17] K. Zetter, "A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever," Jan. 2015. [Online]. Available: http://www.wired.com/2015/01/german-steel-mill-hack-destruction/

[18] R. George, "Views on the future direction of information assurance," Jul. 2002, remarks by Richard George at Blackhat Las Vegas. [Online]. Available: https://www.blackhat.com/presentations/bh-usa-02/bh-us-02-george-keynote.doc

[19] V. Prevelakis and D. Spinellis, "The athens affair," *Spectrum, IEEE*, vol. 44, no. 7, pp. 26–33, 2007. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4263124

[20] Tom Cross, "Exploring Lawful Intercept to Wiretap the Internet," Washington, DC, USA, 2010. [Online]. Available: https://www.blackhat.com/presentations/bh-dc-10/Cross_Tom/BlackHat-DC-2010-Cross-Attacking-LawfulI-Intercept-slides.pdf

[21] Richard George, "Private communication between Richard George, Former Technical Director, Information Assurance Directorate, NSA and Susan Landau," Dec. 2011.

[22] Nicole Perlroth and Vindu Goel, "Twitter Toughening Its Security to Thwart Government Snoops," Nov. 2013. [Online]. Available: http://bits.blogs.nytimes.com/2013/11/22/twitter-toughening-its-security-to-thwart-government-snoops/

[23] Larry Seltzer, "Google moves forward towards a more perfect SSL," Nov. 2013. [Online]. Available: http://www.zdnet.com/article/google-moves-forward-towards-a-more-perfect-ssl/

[24] D. Gupta, "Google Enables 'Forward Secrecy (PFS)' by 'Default' for HTTPS Services," Nov. 2011. [Online]. Available: http://www.ditii.com/2011/11/23/google-enables-forward-secrecy-pfs-by-default-for-https-services/

[25] Selena Larson, "After Heartbleed, "Forward Secrecy" Is More Important Than Ever," Apr. 2014. [Online]. Available: http://readwrite.com/2014/04/15/heartbleed-perfect-forward-secrecy-security-encryption

[26] Adam Langley, "Protecting data for the long term with forward secrecy," Nov. 2011. [Online]. Available: http://googleonlinesecurity.blogspot.com/2011/11/protecting-data-for-long-term-with.html

[27] J. Kiss, "Twitter adds more security to thwart predators and government agencies," Nov. 2013. [Online]. Available: http://www.theguardian.com/technology/2013/nov/23/twitter-security-google-facebook-data-nsa

[28] Parker Higgins, "Pushing for Perfect Forward Secrecy, an Important Web Privacy Protection," Aug. 2013. [Online]. Available: https://www.eff.org/deeplinks/2013/08/pushing-perfect-forward-secrecy-important-web-privacy-protection

[29] Michael Mimoso, "Microsoft Expands TLS, Forward Secrecy Support | Threatpost | The first stop for security news," Jul. 2014. [Online]. Available: https://threatpost.com/microsoft-expands-tls-forward-secrecy-support/106965

[30] ——, "Microsoft Brings Perfect Forward Secrecy to Windows | Threatpost | The first stop for security news," May 2015. [Online]. Available: https://threatpost.com/new-crypto-suites-bring-perfect-forward-secrecy-to-windows/112783

[31] P. Bright, "Microsoft expands the use of encryption on Outlook, OneDrive," Jul. 2014. [Online]. Available: http://arstechnica.com/security/2014/07/microsoft-expands-the-use-of-encryption-on-outlook-onedrive/

[32] Liam Tung, "Yahoo finally enables HTTPS encryption for email by default," Jan. 2014. [Online]. Available: http://www.zdnet.com/article/yahoo-finally-enables-https-encryption-for-email-by-default/

[33] Apple, "iOS Security on iOS 8.3 or Later," Tech. Rep., Apr. 2015. [Online]. Available: https://www.apple.com/business/docs/iOS_Security_Guide.pdf

[34] N. Perlroth, "Electronic Security a Worry in an Age of Digital Espionage," *The New York Times*, Feb. 2012. [Online]. Available: http://www.nytimes.com/2012/02/11/technology/electronic-security-a-worry-in-an-age-of-digital-espionage.html

[35] Ben Thompson, "UAE Blackberry update was spyware," *BBC*, Jul. 2009. [Online]. Available: http://news.bbc.co.uk/2/hi/8161190.stm

[36] Frederick R. Chang, "Is Your Data on the Healthcare.gov Website Secure?" Written Testimony, U.S. House of Representatives, Nov. 2013. [Online]. Available: http://docs.house.gov/meetings/SY/SY00/20131119/101533/HHRG-113-SY00-Wstate-ChangF-20131119.pdf

[37] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and J. K. Zinzindohoue, "A messy state of the union: Taming the composite state machines of TLS," in *IEEE Symposium on Security and Privacy*, 2015. [Online]. Available: https://www.smacktls.com/smack.pdf

[38] Piero Colaprico, “”Da Telecom dossier sui Ds” Mancini parla dei politici - cronaca - Repubblica.it,” Jan. 2007. [Online]. Available: http://www.repubblica.it/2006/12/sezioni/cronaca/sismi-mancini-8/dossier-ds/dossier-ds.html

[39] R. M. Needham and M. D. Schroeder, “Using encryption for authentication in large networks of computers,” *Communications of the ACM*, vol. 21, no. 12, pp. 993–999, 1978. [Online]. Available: http://dl.acm.org/citation.cfm?id=359659

[40] G. Lowe, “An Attack on the Needham-Schroeder Public-key Authentication Protocol,” *Information Processing Letters*, vol. 56, no. 3, pp. 131–133, Nov. 1995. [Online]. Available: http://dx.doi.org/10.1016/0020-0190(95)00144-2

6 Author Biographies

Harold "Hal" Abelson is a Professor of Electrical Engineering and Computer Science at MIT, a fellow of the IEEE, and a founding director of both Creative Commons and the Free Software Foundation.

Ross Anderson is Professor of Security Engineering at the University of Cambridge.

Steven M. Bellovin is the Percy K. and Vida L.W. Hudson Professor of Computer Science at Columbia University.

Josh Benaloh is Senior Cryptographer at Microsoft Research where his research focuses on verifiable election protocols and related technologies.

Matt Blaze is Associate Professor of Computer and Information Science at the University of Pennsylvania where he directs the Distributed Systems Lab.

Whitfield "Whit" Diffie is an American cryptographer whose 1975 discovery of the concept of public-key cryptography opened up the possibility of secure, Internet-scale communications.

John Gilmore is an entrepreneur and civil libertarian. He was an early employee of Sun Microsystems, and co-founded Cygnus Solutions, the Electronic Frontier Foundation, the Cypherpunks, and the Internet's *alt* newsgroups.

Matthew Green is a Research Professor at the Johns Hopkins University Information Security Institute. His research focus is on cryptographic techniques for maintaining users' privacy, and on new techniques for deploying secure messaging protocols.

Peter G. Neumann, Senior Principal Scientist at the SRI International Computer Science Lab, and moderator of the ACM Risks Forum for thirty years.

Susan Landau is a professor of cybersecurity policy at Worcester Polytechnic Institute. She is the author of *Surveillance or Security? The Risks Posed by New Wiretapping Technologies* (MIT Press, 2011) and co-author, with Whitfield Diffie, of *Privacy on the Line: The Politics of Wiretapping and Encryption* (MIT Press, 1998).

Ronald L. Rivest is an MIT Institute Professor, and well known for his co-invention of the RSA public-key cryptosystem, as well for founding RSA Security and Verisign.

Jeffrey I. Schiller was the Internet Engineering Steering Group Area Director for Security (1994–2003).

Bruce Schneier is a security technologist, author, Fellow at the Berkman Center for Internet and Society at Harvard Law School, and the CTO of Resilient Systems, Inc. He has written a number of books, including *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World* (Norton, 2015).

Michael A. Specter is a security researcher and PhD candidate in Computer Science at MIT's Computer Science and Artificial Intelligence Laboratory.

Daniel J. Weitzner is Principal Research Scientist at the MIT Computer Science and Artificial Intelligence Lab and Founding Director, MIT Cybersecurity and Internet Policy Research Initiative. From 2011–2012, he was United States Deputy Chief Technology Officer in the White House.

7 Acknowledgments

The authors thank several individuals who were extremely helpful in the production of this report. Alan Davidson was instrumental in the early discussions that led to this report while he was Vice President and Director of the Open Technology Institute at the New America Foundation. Beth Friedman, Technical Communicator at Resilient Systems, provided invaluable editing support. The MIT Cybersecurity and Internet Policy Research Initiative helped with convening the authors and producing the final version of the report.