

# Can Smartphones and Privacy Coexist?

## Assessing Technologies and Regulations Protecting Personal Data on Android and iOS Devices

*Arkady Yerukhimovich, Rebecca Balebako, Anne Boustead, Robert K. Cunningham, William Welser IV, Richard Housley, Richard Shay, Chad Spensky, Karlyn D. Stanley, Jeffrey Stewart, Ari Trachtenberg, Zev Winkelman*

### Key findings

- Google Android and Apple iOS ecosystems differ fundamentally, largely because of their parent companies' different business models. However, the platforms' tools and protections appear to be converging: Android is adopting run-time permissions requests, which iOS has used for years; both are incorporating stronger encryption.
- On Android, some apps requesting no permissions can capture spoken conversations, can access several types of sensitive data that may allow fingerprinting a phone, and have the ability to cause low-level system faults.
- While most banks use encryption properly, a few still exhibit significant faults in their implementations, and several banks transmit more information than appears necessary.
- We propose a tool that is based on the data life cycle and Fair Information Practice Principles and that allows policymakers to analyze gaps and strengths in smartphone privacy protections during each phase in the life cycle of smartphone data.

**SUMMARY** ■ As smartphones become more ubiquitous around the globe, policymakers inevitably have to grapple with issues related to the security and privacy of these devices. Studies show that smartphone users want and expect privacy (Balebako, Jung, et al., 2013; Boyles, Smith, and Madden, 2012; Felt, Egelman, and Wagner, 2012; and Muslukhov et al., 2012); however, these desires and expectations depend on policymakers gaining greater insights into technological, social, and governmental forces that shape today's evolving smartphone environment.

To aid in this understanding, in 2015, the Defense Advanced Research Projects Agency (DARPA) commissioned a team of researchers from the Massachusetts Institute of Technology (MIT) Lincoln Laboratory and the RAND Corporation to assess smartphone users' privacy from both technical and regulatory perspectives. This report documents the team's approach and findings. On the technical side, it describes a literature review and experiments performed by MIT Lincoln Laboratory investigating the state of privacy of the two major smartphone platforms in 2015: Google's Android and Apple's iOS (the operating system in Apple's mobile devices, such as iPhone and iPad). On the regulatory side, this report describes a review conducted by RAND of major federal regula-

tory mechanisms for protecting privacy in the United States and provides a tool to understand both privacy regulation and technology.

We found that although privacy-preserving technology is improving, users' privacy concerns have not been fully addressed by the technology itself. Appropriate regulatory protections also play

a role in protecting smartphone users' privacy. Currently, many gaps exist between regulation and technology: The two are not adequately paired to provide the desired protections. We believe that many of these gaps can be identified using a tool that the project team developed for policymakers. By combining technical and regulatory components associated with smartphone privacy, this matrix-based tool will help policymakers guide directions for future research and assess the impact of technical and regulatory solutions that have been or will be implemented.

---

## INTRODUCTION

Several well-known technology observers have proclaimed that privacy is dead. Scott McNealy, cofounder of Sun Microsystems, is famous for making this statement as far back as 1999: "You have zero privacy. . . . Get over it." Recently, others, including Margo Seltzer of the World Economic Forum, have also made pronouncements about the death of privacy: "Privacy as we knew it in the past is no longer feasible. . . . How we conventionally think of privacy is dead." While these statements and others like them have gathered publicity, many question whether they are accurate. Certainly, these statements prompt the question of whether the death of privacy is desirable.

Americans clearly desire privacy, as demonstrated by several recent polls conducted by the Pew Research Center (Boyles, Smith, and Madden, 2012; and Madden and Rainie, 2015). A majority of Americans surveyed have avoided using one or more smartphone apps because of privacy concerns. Many people say that it is important to be in control of access to their private data and do not want people watching them without permission (Boyles, Smith, and Madden, 2012; and Madden and Rainie, 2015).

In these polls, *privacy* is either left undefined or described as the smartphone owner wishing to maintain control over how their personal information is shared. The challenge with defining privacy is that often what is considered *personal* or *private* may depend on context and on the individual. Americans may also have differing opinions on who gets the data; privacy concerns may vary depending on whether information is shared with corporations or the government.

Such companies as Apple and Google have recognized the privacy concerns and have responded by using privacy as

You have zero privacy. . . .  
Get over it.

—Scott McNealy,  
Sun Microsystems, 1999

a marker of competitive advantage in the marketplace. Apple regularly makes public statements about its privacy posture to entice users to use its products. Apple and Google tout their privacy-preserving technology offerings to demonstrate a competitive advantage over each other.

Privacy is complicated wherever it is pursued (Westin, 1968), but especially so in the smartphone ecosystem. Smartphone technologies are new compared with computers and laptops, and the protections they offer are still evolving. There are many different players involved in building the different components of smartphone platforms, each with different sets of incentives. Additionally, smartphones tend to actively collect information because they are always on, always with the user, and equipped with a wide variety of sensors that enable the gathering of vast amounts of potentially private information about the user. Furthermore, the small screen size of such devices can limit what can be conveyed. This can be particularly relevant for complicated ideas, such as security and privacy (Harris, Goodman, and Traynor, 2012).

This report considered the state of smartphone privacy in the face of these difficulties and complications. Our goal is to understand the state of privacy in the smartphone world, including identifying gaps and opportunities.

More concretely, this report looked at the privacy offerings on devices running Apple's iOS and Google's Android operating system. Both systems manage Internet-connected mobile phones, often called *smartphones*, that are widely deployed and used by consumers. For these two platforms, we aimed to look broadly at the privacy offered today, considering both the technical and regulatory protections available. We used our findings to identify future directions along which smartphone privacy is headed and to identify gaps between technology and regulation that can be filled by future research or technological innovation.

Any discussion about smartphone privacy needs to begin with an understanding of the components of the problem space. First, there is the device itself, with all of its technological capa-

bilities and various sensors for collecting information. Next, we have to consider the functionality or utility that users wish to get from this phone. Specifically, does the user intend to use the phone for surfing the web, conducting business, or playing games? Each type of use may result in different preferences for the level of privacy desired by the user, as well as different private information being collected and used. Finally, we have to consider the business models of the providers of the different services offered on the phone. Such providers include the phone manufacturer, operating system manufacturer, phone carrier, app developers, and many others. The way that these different providers earn money from the services they provide differs greatly and has implications for the users' privacy. For example, Google is a technology company that derives the vast majority of its revenue from advertising, which relies on data gathered about users. Instead of building its own hardware, Google tailors its operating system to work on hardware developed by a variety of third-party vendors, enabling wider deployment of Google-provided services and advertising. Apple is primarily a hardware manufacturer and makes most of its money from the sale of physical devices (of which iOS platforms dominate). Additionally, many app developers make their money through incorporating advertising libraries in their apps, requiring access to private data that may not be necessary for the core functionality of the app.

These considerations must be kept in mind when looking at the tools used to ensure privacy in the smartphone space. First are various privacy-preserving technologies, such as encryption, sandboxing technologies that isolate code, and permissions models that aim to control what private data may be collected (and protect that data after it has been collected). Second are various privacy-promoting regulations, such as legal, contractual, and policy mechanisms, that aim to protect individuals' privacy by enforcing the rules about the type of private data that may be collected and how the data can be used.

The intersection of these two mechanisms for ensuring privacy is important. Specifically, technology developers need

to understand the existing regulatory protections to ensure that they adhere to policies and standards. On the other side, regulatory mechanisms need to take into account what is realizable and enforceable by technological means.

---

## DATA LIFE CYCLE


The *data life cycle* is a useful abstraction that allows us to reason about how private data are protected during smartphone operations. RAND is not the first to recognize the role of understanding the data life cycle in privacy. For example, "Privacy by Design" (Cavoukian, 2009), a process of engineering privacy into systems by design, refers to providing privacy and security protection throughout a data life cycle, but does not define what that is. In this report, RAND broke down the data life cycle into the phases through which information may need to be protected. Figure 1 depicts those phases.

There are benefits to defining the data life cycle's phases. Defining the phases allows a system designer to consider whether privacy is protected in each phase. Additionally, identifying the life-cycle phases gives stakeholders from different disciplines a common vocabulary. It allows people who work on different elements of an ecosystem to identify at which phases of the life cycle their respective technology works. While the data life cycle allows for the identification of the phases relevant to a particular system and provides a terminology for different phases, the data life cycle is flexible. It is neither a rigid requirement nor a formal definition. It is possible that multiple stakeholders make decisions at each life-cycle phase; for example, hardware and software decisions are made at the conception stage.

We do not formally define data or information. We consider data to be something that users would like to protect. Examples may include users' location information or health records.

Privacy as we knew it in the past is no longer feasible. . . .  
How we conventionally think of privacy is dead."  
—Margo Seltzer, World Economic Forum, 2015

Figure 1. Data Life Cycle



RAND RR1393-1

## Phases

The data life cycle begins with the *conception phase*, during which decisions are being made about the user's privacy well before the data are created. In this phase, rules are made about what data can be accessed and how they are protected. There can be a number of decisions made during this phase that strongly emphasize system design and architecture, much of which is determined prior to data creation. Examples include when manufacturers are deciding what sensors the phone will have, when operating system developers are deciding how apps will be sandboxed or protected in the operating system, or even when app developers are making decisions about how an app will run and what the business model is. The conception phase could be divided into subphases that generally do not overlap. The subphases include decisions made during platform creation, app store creation, phone manufacturing, and app development. Notably, these all occur before the data are actually created.

The *creation phase* describes decisions made at the time data are first generated and/or captured by the smartphone. In some current smartphone systems, users may be able to make decisions during the creation phase if they are given a “run-time” notice about data collection. The run-time notices are sometimes called *just-in-time notices*. The creation phase can be distinguished from the conception phase because there will be more contexts about the data being created and collected, as well as more perspective about the prospective value of the data.

The *transmission phase* occurs when data are moving. Typically, data are not stored solely on the phone and can be transmitted to storage or shared with other users or companies. Decisions about transmission may include the type of communication that is used, such as transmission over a telecommunications wireless network or a Wi-Fi network. The transmission

phase may involve encryption to prevent man-in-the-middle attacks, in which attackers will secretly relay and possibly change the communication between two parties without their knowledge.

In the *storage phase*, data are at rest in a repository. In the smartphone ecosystem, data are often stored on cloud-storage networks by cloud-storage providers. Cloud-storage providers may make decisions about how the data are organized and maintained, such as where the data will be stored, when the data will expire, and who can access the data.

Data may also be stored on the device, either in addition to or instead of being stored on the cloud. In the case of storage on the device only, when data are stored in the same location as where they were created, the transmission phase may have been bypassed. As previously stated, the data life cycle is flexible enough to allow for this.

Data may then enter the *exploitation phase*. Exploitation in this context means that the data are analyzed or used—sometimes in a way that was originally intended and other times in new ways following fusion with other data. In this context, exploitation phase does *not* refer to malware exploitation, which implies malicious activity. Examples of the exploitation phase include app developers analyzing the data to understand how their apps are being used or third-party advertisers using single or multiple sets of data to decide which ads will be provided to that phone.

Figure 1 shows a loop from exploitation to creation. This demonstrates that data in the exploitation phase may be reborn as new information to be created, such as new inferences or profiles of the users. Data rarely cease to exist; they continue to be copied, retransmitted, and exploited.<sup>1</sup> Data do not have a prescribed lifetime and may exist in various forms for indeterminate periods of time. At the exploitation phase, the data

## Research Methodology

This work focused on the technology and regulation used to protect privacy on smartphones and tried to understand how they intersect to provide protections to the users. This two-pronged approach was necessary because system developers, as well as regulatory bodies, must consider both aspects in order for any privacy protective system to make the transition from research to wider deployment and market adoption. To understand these aspects of smartphone privacy better, we introduced our interpretation of the data life cycle to capture the privacy concerns, decisions, and decisionmakers involved in each state of smartphone operations.

**Technology.** We evaluated the privacy-preserving technology available on the two dominant smartphone platforms: Google’s Android and Apple’s iOS. We began with a review of the app-store models, permissions models, sandboxing, and data encryption. We highlight the key differences and similarities between the protections offered on these platforms and the effects on user privacy.

We also engaged in several experiments to evaluate how well the available privacy-preserving technology meets its stated goals. In particular, we investigated how much private data can be accessed by apps requesting no permissions—either by collecting sensitive data that are given to the application by the system or by fuzzing the system Application Programmer Interface (API) to extract sensitive information.

Additionally, we investigated whether encryption is used appropriately to protect private data in banking applications and used the data life cycle to illustrate gaps and strengths in protections.

**Regulation.** To provide background on smartphone privacy regulation, we reviewed the regulatory enforcement mechanisms available for safeguarding consumer privacy. We began by describing a number of frameworks that can be used to understand the myriad of federal statutes and regulatory mechanisms for protecting privacy in the United States. This included looking at the Fair Information Practice Principles (FIPPs), which are principles for protecting privacy that have been used by regulators for decades.

Which framework was the most useful for identifying gaps in privacy protection in regulation and in technology? We looked at them using several criteria—completeness, applicability to technology, and applicability to regulation—but none met all of those criteria. To address that, we took two steps. First, we delineated and described the stages of the data life cycle for smartphones. The data life cycle provides a continuous framework for identifying privacy-protective technologies and decisionmakers. Second, we combined the data life cycle with the FIPPs, creating a new tool that we refer to as the Data Life cycle–Fair Information Practice Principles (DL-FIPPs) tool. The DL-FIPPs Tool allows one to analyze the gaps and strengths in protections in each cell of the matrix.

may be copied and shared. Data are often shared beyond the app with third parties; they can be shared with advertising companies, analytics companies, or the government. Users are concerned with all these areas (Balebako, Jung, et al., 2013; and Urban, Hoofnagle, and Li, 2012). The technology aspects of this report specifically focus on the data on the device. The regulatory aspects and the two-dimensional framework discussed toward the end of this report leave room to consider third-party sharing and reuse.

Various groups may be making decisions at each phase of the life cycle. This is important for understanding regulation throughout the life cycle, as each decisionmaker may be subject to different laws or oversight by different regulatory bodies. The

decisionmakers’ skills and resources may affect what technical protections are put in place. Therefore, identifying the decisionmakers leads to a greater understanding of the privacy protections likely to be employed. For example, app developers without security expertise may fail to correctly implement privacy-preserving technologies (Balebako, Marsh, et al., 2014).

## Decisionmakers

Various decisionmakers are involved throughout the life cycle. This section discusses decisionmakers as illustrative points, as the full roster may depend on the specific app and platform. At the conception phase, *platform developers* make decisions

about the hardware, including what sensors might exist on a smartphone. This impacts the type of data that can be collected by the phone. Platform developers also make decisions about the operating system architecture that can impact privacy and security. Examples include whether encryption is system-wide or whether apps are sandboxed from one another. The *app store* controls what apps are available for download, and the security reviews and privacy reviews (or lack of privacy reviews) made by the app store influence whether privacy-invasive apps or malware can be downloaded (Apple, Inc., 2016; and Lockheimer, 2012). The *app developer*, when coding the app, makes many decisions about what functionality triggers data collection and what data are collected.

At the creation phase, users may be able to choose, through run-time permissions, whether to allow data collection. App developers and advertisers may request data that already have been approved by the user; advertisers and app developers may control the frequency with which data are collected, unbeknownst to users who have already approved collection.

At the transmission phase, the decisionmakers may change depending on the method of transmission. For example, a *telecommunications provider* may use different technologies and may be regulated differently from an Internet service provider. Therefore, the privacy implications of transferring data over the Internet (via Wi-Fi or Ethernet cable) may be different from the implications of transferring data using a cellular network.

At the storage phase, if data are stored on the cloud, the *cloud provider* may make decisions that impact the privacy of the data. The technical decisions may include who gets access control, which standards they follow, and whether encryption or deletion is available. Other decisions may impact how they are regulated. For example, the location of cloud servers may have an impact, and different countries and states may have different regulations on data privacy. Other stakeholders may be involved if the data are not stored only on the cloud or are only stored locally on the phone.

The exploitation phase is particularly interesting, as different stakeholders may wish to access the data and may decide to use the data in different ways. Many smartphone users expect that the app developer will access the data collected, and they expect it to be used to improve the app (Balebako, Jung, et al., 2013). However, smartphone users may not expect app developers to resell the data to other parties. Advertisers may profile the user during the exploitation phase, and these profiles may impact the ads to which users are exposed. The *government* may

also request access to information, which may cause privacy concerns for users.

Overall, the data life cycle in smartphone data is a way to consider how privacy is protected in the smartphone ecosystem. It provides an opportunity for technologists and regulators to consider whether there is protection throughout the life cycle or to define how their technology or regulation fits into a particular phase of the life cycle.

---

## PRIVACY-PRESERVING TECHNOLOGY

A number of privacy-preserving technologies are available on both platforms. These include such tools as the app market model and permissions models for regulating the capabilities of the apps installed on the phone; sandboxing, full-disk encryption, and in-transit encryption for protecting the data that are collected, stored, and transmitted by the phone; and secure boot and trusted-execution environment for protecting the device itself from attacks that may lead to private data being exposed. In comparing the way these protections are implemented on iOS and Android devices, we focused on implications for the privacy of users' data. Overall, it appears that the privacy techniques used by these two platforms are converging, but fundamental differences remain due to the foundational differences between the iOS and Android ecosystems.

## Ecosystems

The biggest difference between the Android and iOS ecosystems is that Apple is the sole manufacturer of the iPhone. Google, on the other hand, allows different hardware manufacturers to build devices running the Android operating system. This difference leads to several important differences in privacy protection.

Google's ecosystem that allows multiple vendors to create hardware makes it difficult to patch existing systems, as there are many versions to be patched and many parties responsible for applying the patches. This results in a slow patching process. Apple, on the other hand, can, at any time, apply a global patch to all iOS devices. This affects user privacy, as known potential vulnerabilities exposing private data are likely to remain open for a longer time in the Google ecosystem (Vidas, Votipka, and Christin, 2011).

Additionally, Google's distributed ecosystem means that the different carriers can add their own apps to the base

Android distribution and force users to install them on their phone to use their services. Such apps, commonly called *bloatware* (McDaniel, 2012), often allow the providers to collect large amounts of private user data to use for their own purposes.

Apple, on the other hand, is the sole code provider for all iOS devices, minimizing the variety of pre-installed apps (considered bloatware by some) and their sources. However, Apple's centralized ecosystem does not come without a cost. Specifically, the homogeneous attack surface of all iOS devices means that a single exploit may be able to compromise private data on all such devices. The many different versions of Android potentially avoid such a shared vulnerability.

The bottom line is that iOS's centralized model better controls patching and privacy-invasive bloatware but is more susceptible to universal exploits.

## App Store

This difference in ecosystems also translates onto the app store models of iOS and Android, which control what apps are approved for installation and operation on the devices. Both iOS and Android attempt to vet all apps that are added to their respective app stores. This vetting consists of static, dynamic, and manual analysis to try to ensure that the app operates as advertised, thus providing some guarantees to the user that operating system-specific apps will not be malicious. However, as with the overall ecosystem, Google does not exert full control over the spectrum of Android-compatible apps and still allows users to install apps from unknown sources even if they have not been approved for the app store ("Google Report: Android Security 2014 Year in Review," 2014). This makes it possible to install apps that, among other things, may not properly preserve and protect private data. Given this, while both iOS and Android are moving toward a "distrust and verify" model for apps, the fundamental difference in the ecosystem resides in the level of control that they have.

## Permissions Models

We also examined the permissions models on iOS and Android devices that control what information an app is allowed to collect. Prior to Nougat (the most recent operating system for Android), the system required that all permissions be approved when an app is first installed, whereas iOS required that permissions be approved at run-time (such as with just-in-time

requests). This meant that Android users could see all the permissions that an app would ever need at install time, whereas iOS users only would see a subset of the necessary permissions when choosing which app to install.

In Marshmallow (which preceded Nougat), Android moved closer to the iOS model: Certain groups of permissions are additionally requested as needed, rather than all approvals being available from the beginning. These run-time approvals differ somewhat between iOS and Android, with Android only requiring user approval for permissions that it deems "dangerous" (Taylor, 2015) because they involve the user's private information. Notably, these do not include such permissions as setting the alarm clock or accessing the Internet. iOS also uses run-time approval for entitlements that allow for specific non-default capabilities, such as permitting the use of iCloud storage (Atienza et al., 2015).

Finally, iOS and Android impose limitations on their permissions models in the name of usability. Android puts permissions into permissions groups, such that approving any permission in the group actually grants access to all permissions in that group. iOS instead has some default permissions that grant some access to the app without requiring user approval. Both of these limit the amount of control that the user of the phone has over private data.

## Sandboxing and Encryption

Both Android and iOS use sandboxing and encryption to protect data that are collected by the phone. Sandboxing tries to segment apps to prevent one app from accessing potentially sensitive data that are collected by other apps. Although iOS and Android use different technical tools, they provide similar levels of sandboxing protection.

iOS and Android additionally provide disk encryption to protect data collected by apps while that data are stored on the phone. iOS turns on full-disk encryption by default, ensuring that all data are always encrypted. Android currently allows full-disk encryption to be optional, thus potentially leaving data vulnerable. However, Google has plans to make encryption the default on Android as well. iOS and Android use hardware-backed key signing to prevent offline attacks on the encryption keys.

---

## EXPERIMENTS TO EVALUATE EFFECTIVENESS OF PRIVACY-PRESERVING TECHNOLOGY

Project team members from MIT held experiments to understand how well these privacy-preserving technologies protect user privacy in the real world. Specifically, the researchers ran two different experiments to evaluate the effectiveness of the permissions model by looking at what private data are accessible by someone bypassing the permissions model either through private application programmer interface (API) calls or by exploiting the existing system APIs. To understand what private information may be accessible to a network eavesdropper, the team performed several experiments to evaluate whether encryption is used correctly by various classes of apps and what private information such apps transmit.

### What Private Information Is Accessible by Android Apps That Have Not Requested Permissions?

Our first experiment aimed to identify what private information would be accessible by an Android app requesting no permissions. To measure this, we wrote a scraper script that scanned the public Android API references to identify API calls that require no permissions. The resulting list, including roughly 36,000 unique API calls, was then sifted manually to identify potential privacy leaks. We found a number of such privacy leaks. For example, we discovered that apps could identify lists of packages installed on the device. This could be used for fingerprinting the phone, identifying vulnerable apps, and identifying location of private data for exploitation. Fingerprinting may be a privacy concern, as it allows for the unique identification of the phone in a way that the platform developer may not have intended. Fingerprinting may allow a device or user to be uniquely identified across sources, enabling the aggregation of data from different sources. This allows for new inferences to be made that were outside of the expectation of the user (Turnow, 2012).

Listing actual accounts on the device requires special permissions. However, listing account authenticators does not. Account authenticators often tell us what accounts the phone user has, which may be sensitive and private, although some services might share an account authenticator, limiting the damage from this information.

Additionally, it is possible to access the per-app network throughput. This includes the granular transmission and reception of data. For example, we were able to see how many bytes of data the Skype process received over a period of time. This information is particularly sensitive. Other research has shown that the throughput of Voice over Internet Protocol (VoIP) applications can be used to identify the language of speech (Wright, Ballard, Monroe, et al., 2007) and even to detect specific spoken phrases (Wright, Ballard, Coull, et al., 2010), thus allowing a no-permission app to eavesdrop on private phone conversations. This side channel may also allow an app to learn coarse-grain location information by monitoring the throughput of mapping apps as they load map segments as the user moves around.

### What Private Information Is Accessible by Exploiting System API?

We ran a second experiment to evaluate what private information is accessible by exploiting the system API. Specifically, we created a tool we call the *reflector app*, which uses reflection (Microsoft, 2016). Reflection is a method to observe program calls at execution, and we used it to identify and call all possible methods of any objects created during phone operations. This allowed us to find additional ways for an app to co-opt users' private data using methods not described in the API documentation. This method is closely related to fuzzing, a technique commonly used in software vulnerability testing (Miller, 2007) in which large amounts of random data are input into the code to see if the program will crash. This is the first use of this technique to explore privacy that we are aware of.

This experiment led to several interesting findings about Android. First, there is some publicly accessible information, such as shared libraries and other system features, that can allow the fingerprinting of a phone. Perhaps more worrying is the ability to learn the keyboard layout on the phone, which may reveal a user's preferred language. Second, we uncovered several potentially dangerous methods that caused low-level faults on the device and were not caught by the virtual machine. Low-level faults may be more of a security concern than a privacy concern but potentially indicate the presence of a low-level vulnerability that could be exploited by an attacker. However, as a positive, we were unable to access any extremely private data, such as wireless data or any unique forms of identification, that could be used to identify the device imme-

diately. Thus, while there are some concerns, it appears that reflection does not find any critically private information.

## To What Extent Do Apps Use Cryptography or Transmit Nonessential Private Data?

In a third experiment, we focused on financial and banking apps, as these should be relatively mature and would most likely demonstrate best practices in managing cryptography and private data. Specifically, we analyzed 50 banking apps on both iOS and Android to understand how well these apps use cryptography.

First, we looked at whether these apps properly validate the server certificates that they use to authenticate the server and protect their communications. Failure to do so may open the door for man-in-the-middle attacks, which could steal private data by impersonating a banking server. On both iOS and Android, the majority of banking apps properly validated the certificates. However, a small number of banking apps did not correctly implement certificate validation. This implies that certificate validation is still not a straightforward process and less-mature apps are likely to get this wrong, potentially allowing private data to be intercepted in transit.

Additionally, while properly validating certificates is the current industry standard, this leaves users vulnerable to attackers who can subvert the certificate authority (CA) process. Examples of attackers who can subvert the CA process include state actors and attackers who can compromise a CA (e.g., DigiNotar, Comodo, VeriSign). We note that there have been prior examples of CA compromise to generate fake certificates (Whitney, 2011). Pinning certificates is currently considered a way to provide more assurance that the end point to which the app is talking is actually the expected end point. We note that few banking apps actually choose to use this stronger approach to give extra protection to their users.

Next, we intercepted and analyzed the types of private data that are transmitted by these banking apps during operations. We noticed that a fair share of Android apps seem to transmit permanent identifiers, which could be used to fingerprint the device. Additionally, on Android and iOS, several of the apps seemed to transmit location data when it did not seem necessary, as the location-data count excludes map-related requests. In particular, on iOS, two of the seven apps that send location data wait until the user has granted the location data for mapping purposes and then transmit the data while not using the map. Essentially, they are taking advantage of the fact that

the user has granted access to location data for one purpose, but then they are using the data for another. Furthermore, in addition to location information sent explicitly, location can be inferred from other information that is sent. Several of the apps also transmit the Basic Service Set Identification, which is the media access control address for the wireless access point that the device is using, letting the app locate the user. Transmission of such identifying or unnecessary data indicates that these banking apps are likely using users' private data for purposes other than their stated functionality, thus potentially compromising the users' privacy.


---

## THE STATE OF PRIVACY-PRESERVING TECHNOLOGY

Building on the data life cycle depicted in Figure 1, we examined how the available privacy-preserving technologies protect the data throughout their life cycles and identify some gaps in existing protections. Our findings are shown in Figure 2, with protections described in green and gaps described in red.

In the conception phase, the app market-vetting strategy protects users by disallowing certain malicious apps from entering the app marketplace. However, this protection is not perfect because the vetting can be bypassed, and vetted apps can dynamically load malicious applications during operations. In the creation phase, permissions models and sandboxing protect user privacy by limiting what type of private data can be collected and who can access those data. However, even these technologies have their shortcomings, with many permissions defaulting to giving all apps access to some user data and with many apps requesting—and receiving—more permissions than necessary for their operations. In the transmission phase, encryption does offer the users a set of tools to protect their data. However, this is nontrivial to use and is often used incorrectly, which creates possible vulnerabilities. In the storage phase, full-disk encryption can be used to protect data stored on the smartphone, but once the data are taken off the phone, phone users are not guaranteed protection of their private data. Finally, the exploitation stage is not well covered by any existing privacy-preserving technology, which leaves essentially no constraints on what can be computed based on the collected private data. This analysis makes it clear that, while technology tries to protect user privacy, many gaps remain to really protect privacy through the whole data life cycle.

**Figure 2. Privacy Protections and Gaps in Each Data Life-Cycle Phase**



RAND RR1393-2

## Malware Effects

Additionally, we looked briefly at how malware affects the privacy environment. Specifically, we reviewed existing literature to understand the types of private data that malware authors pursue and how these differ from the type of private data typically captured by nonmalicious apps. For the case of Android, there were extensive data on what malware does. Specifically, a study in “Andrubis-1,000,000 Apps Later: A View on Current Android Malware Behaviors” (Lindorfer et al., 2014) used dynamic and static analysis to analyze roughly 400,000 malware apps taken from various malware corpora. We found much less information on iOS malware and were thus forced to focus on a small list of known, existing iOS malware and to examine how the data that these pursue lined up with the more-detailed statistics available on Android.

The interesting takeaway from this analysis is that malware in both operating systems focuses on permanent identifiers such as the International Mobile Station Equipment Identity and credentials for malware owners to break into and exploit user accounts. Malware seems less focused on stealing private information, such as location data. Surprisingly, nonmalicious apps are more likely to use and leak this data for such purposes as advertising.

## OVERVIEW OF REGULATORY ENFORCEMENT

As described at the beginning of this report, it is important to understand the regulatory environment for consumer privacy. Without understanding regulation, privacy technologies or systems may fail to meet legal and social requirements. This may lead to a failure to transfer from research to the consumer market. Therefore, we provide a brief overview of some regulatory levers that may be used to protect privacy in the smartphone ecosystem. We do not focus on protections offered by the U.S. Constitution for privacy from governmental intrusions; instead, we focus on protection offered to consumers who may feel privacy invasions when apps, analytics companies, and platform manufacturers collect data about them. This discussion focuses on federal law for consumer privacy in the United States and does not address state or international laws. However, these can also be influential. For example, the California Attorney General has been particularly active in the smartphone privacy space (State of California Department of Justice, Office of the Attorney General, 2012a, 2012b).

The U.S. Federal Trade Commission (FTC) enforcement actions have been particularly important to consumer privacy protections in the United States, and the FTC is the federal agency most actively involved in consumer privacy enforcement. The FTC uses its authority to protect consumers by enforcing companies’ privacy policies. According to the FTC, it uses law enforcement, policy initiatives, and consumer and

business education to protect consumers' personal information and ensure that they have the confidence to take advantage of the many benefits of the ever-changing marketplace (U.S. Federal Trade Commission, 2000).

## Enforcement of Privacy Policies

Here, we start with smartphone users, the app developer who creates apps for the smartphone, and the app stores that allow users to purchase these apps. There are agreements in place between all of these different actors, but we focus on one type of agreement in particular: the privacy policies that mobile app developers create to tell users what information about them will be collected and shared.

However, these privacy policies may not be an easy source of information about the data collected by mobile apps. They may fail to describe how data are shared with third parties, such as advertisers, and give users little information about the exploitation phase of data. MIT Lincoln Laboratory examined 223 Android policies and 126 iOS policies and found that the average reading level required to understand the text was that of a college sophomore and that just more than 20 percent could be understood by someone with less than a high school diploma. In other words, the privacy policies for apps are not doing a particularly good job at providing notice and choice to users.

Although privacy policies are agreements between users and app developers, courts may hold that they do not meet the legal requirements to be contracts. However, the FTC can use consumer-protection statutes to ensure that app developers abide by their privacy policies. Therefore, despite the limitations of privacy policies, they have become an important part of FTC enforcement of privacy. In addition, there are more-specific consumer-protection statutes that the FTC can use to protect certain classes of information or users. For example, the Children's Online Privacy Protection Act (COPPA) is intended to prevent the sharing of information about children without parental permission. If an app developer collects information about children under 13 years of age without parental consent, then the FTC can pursue an enforcement action against the app developer (United States Code, 1998). This line of protection and responsibility could be drawn from app developers to app users provided by COPPA, but it could also be the app's privacy policy that offers protection. If the FTC determines that COPPA or the app's privacy policy has not been respected, it can bring an enforcement mechanism against the app.

In one example of how privacy enforcement can work in the United States, the FTC recently pursued a similar enforcement action against Yelp for improperly collecting children's information and settled out of court with the FTC for about \$450,000 (U.S. Federal Trade Commission, undated).

---

## FRAMEWORKS FOR EXAMINING PRIVACY REGULATION

Our goal was to identify gaps and strengths in protection in regulation and technology. The data life cycle offers one way of examining protection. But other frameworks also exist. One could, for example, look at the types of data that are protected. This would align nicely with some existing statutes that, for example, protect health information (e.g., Health Insurance Portability and Accountability Act [HIPAA]) or consumer credit information (e.g., Fair Credit Reporting Act). Alternatively, one could consider what class of people are involved, such as who the data subject or data controller is. This method has also been used in some statutes, such as COPPA, and the Family Educational Rights and Privacy Act, which protects students' information.

However, these frameworks are domain specific. They are also in silos, in that they protect people or data in that domain but do not offer protections outside of that domain. This naturally creates gaps in protections.

## Fair Information Practice Principles

One framework that has long been used to understand and frame privacy regulation is the FTC's Fair Information Practice Principles (FIPPs). The FIPPs, originally developed by a committee chaired by the privacy pioneer Willis Ware in the 1970s (Ware, 2008), have advanced the idea that there are several aspects that must be addressed to protect consumer privacy. Many variations of the FIPPs exist, and they have been incorporated into many regulations, both in the United States and internationally.

The FTC, for example, relies on FIPPs and uses a particular definition, which is reproduced in the text box on the following page (U.S. Federal Trade Commission, 2000).

Note that the FTC defines *access* differently than the security community does. Access is not something to be prevented but something required for users, allowing them to view and correct data about themselves. One familiar example in the

### FTC Description of Fair Information Practice Principles

1. Notice: Data collectors must disclose their information practices before collecting personal information from consumers.
2. Choice: Consumers must be given options with respect to whether and how personal information collected from them may be used for purposes beyond those for which the information was provided.
3. Access: Consumers should be able to view and contest the accuracy and completeness of data collected about them.
4. Security: Data collectors must take reasonable steps to assure that information collected from consumers is accurate and secure from unauthorized use.
5. Enforcement: the use of a reliable mechanism to impose sanctions for noncompliance with these fair information practices.

United States is the Fair Credit Reporting Act (United States Code, 2012), which requires that users have the ability to access and correct their credit reports.

### New Matrix Tool to Identify Privacy Gaps and Strengths

Our analysis of frameworks found that neither FIPPs nor the data life cycle alone is appropriate for understanding gaps in regulation and technology. As a result, the RAND team, in collaboration with the MIT Lincoln Laboratory team, combined the data life-cycle phases and the FIPPs elements into a two-dimensional matrix tool for identifying privacy gaps and strengths.

The matrix tool is shown in Table 1. It combines phases of the life cycle (represented as columns) with the FIPPs (shown as rows). We refer to this as the DL-FIPPs Tool, to recognize the role of the data life cycle and FIPPs.

By stepping through each row in the matrix, a system designer can consider whether that type of protection exists for each phase of the data life cycle. For example, a designer can consider whether the notices provided to the user cover each phase of the life cycle. Similarly, a designer can consider whether the system settings provide an informed choice for each of the phases of the life cycle.

Access is another interesting area. In general, apps have not focused on providing access to data about the user. Currently, when access is provided, it is typically only in the storage phase. One example from the web space is the ads-preferences tool offered by Google that tells users about the profile that has been built about them, including Google's guesses about their age and gender. Less attention has been given to how users can access or correct data in transmission, but this may become more relevant as users' information is continuously collected and transmitted through sensors and as decisions and feedback are provided in real-time.

Security technologies certainly can be evaluated in light of the data life cycle. Enforcement of any of these aspects is an open question, as some technologies and aspects of compliance are more open to auditing than others. One question to consider is how a privacy-protective system can verify compliance with regulation at all phases of the data life cycle and how compliance can be audited.

The inclusion of the exploitation phase of the data life cycle emphasizes the importance of considering third-party access to data and how (or whether) the FIPPs are respected. Therefore, this column allows us to consider, for example, what happens when an app provides data to parties such as ad or analytics companies. When data are resold or copied, are users given appropriate notice and choice? Will users have access to the information that is transferred? And is the transfer done with best security practices? More work is needed on how protections at the exploitation phase can be enforced.

### Using the Tool

To demonstrate how the matrix tool can be used, we pursued two examples that allowed us to identify either gaps or strengths between privacy technology and regulation. The first, which shows the intersection of creation and choice, is depicted in the beige cell in Table 2. In this case, we observed a gap between technology and regulation. On the one hand, technology has provided run-time permissions that allow users to choose what data are collected about them. On the other

**Table 1. Data Life Cycle–Fair Information Practice Principles Matrix (DL-FIPPs Tool)**

Life-Cycle Phases / FIPPs	Conception	Creation	Transmission	Storage	Exploitation
Notice					
Choice					
Access					
Security					
Enforcement					
Total					

hand, regulation in the form of COPPA requires that parents have a choice about the data collected about children under the age of 13. Thus, while run-time permissions provide choice, they do not *require* parental approval. A child could accept the permissions request. Although technology and regulation are addressing this space, there is nevertheless a gap in how the two intersect in this square.

In the second example, shown in the purple cell, we looked at the intersection of transmission and security. In technology, we see that efforts in in-transit encryption are making progress in preventing access to data or modifications by outsiders. At the same time, regulation, such as the HIPAA statute, states that covered entities are required to implement protections for transmitted information. In this case, we see that technology is working to offer protections required by regulation. While we do not claim strong protection in a provable sense, we observed that this area has stronger privacy protection and work than other areas.

Therefore, these examples show us how the tool can be used to identify areas of strength and weakness in technology and regulation in the smartphone ecosystem. We believe this tool will be particularly effective at identifying gaps.

---

## WHERE IS SMARTPHONE PRIVACY HEADED?

On the technological front, we found that the two smartphone platforms, while fundamentally different, are converging. The differing platform ecosystems have led to fundamental differ-

ences between privacy protections and guarantees on iOS and Android devices. However, the permissions models controlling what data can be collected by their apps are converging in significant ways. Further, both are adopting increased encryption to secure the data that are collected.

With respect to regulation, policymakers have several options to protect privacy. Some, however, put the onus on the user to recognize and prove that harm has occurred and to identify the perpetrator. This can be particularly difficult in privacy encroachments in the digital ecosystem, where harm—such as feeling exposed—may be intangible or where it can be difficult to identify who is responsible for the privacy invasion.

A comprehensive policy overhaul relating to privacy is unlikely to occur in the United States in the short term. The FTC is the main federal enforcer of consumer privacy, and it is likely to continue in that role for the time being.

While privacy-preserving technology continues to improve overall, there are certain problems that cannot be solved by technology alone, and it is necessary to understand how the technology pairs with existing regulation. Currently, many gaps remain between regulation and technology, leading to weaknesses in protections and enforceability. The DL-FIPPs Tool that we presented may help policymakers gain insight into smartphone technology and regulation issues and to identify gaps and strengths in current and future policy approaches.

**Table 2. Using the DL-FIPPs Tool: Two Examples**

Life-Cycle Phases		Creation		Transmission		Storage		Exploitation	
FIPPs	Conception	Creation		Transmission		Storage		Exploitation	
Notice									
Choice		<b>Technology run-time permissions:</b> User has choice about what personal data are collected.	<b>Regulation COPPA:</b> Requires that parents have choice about data collected on their children.						
		<b>Gap:</b> Run-time permissions do not require parental approval.							
Access									
Security				<b>Technology in-transit encryption:</b> Prevent access to data or modification by outsiders.	<b>Regulation HIPPA:</b> Covered entities are required to implement protections for transmitted information.				
				<b>Strength:</b> Technology provides protections required by regulation.					
Enforcement									
Total									

**NOTES**

<sup>1</sup> There are dimensions of data sharing that users care about but are not captured by our description of the data life cycle. These dimensions include the frequency of data sharing, with whom data are being shared, and what the data are used for. We recommend keeping these dimensions in mind at each stage of the life cycle.

<sup>2</sup> We also attempted to perform the same analysis on iOS devices, but the phones kept crashing and thus prevented us from collecting any interesting information. We plan to investigate this further in future work.

<sup>3</sup> There are several other legal mechanisms (e.g., tort, criminal law, contract law) that may be appropriate for protecting some aspects of privacy against invasion by some parties; whether they will apply in a particular instance will depend on the facts presented. One difficulty with relying on those mechanisms is that they put the burden on the user (whose privacy has been violated) not only to recognize the violation but also to identify who is responsible for the violation and be prepared to take the case to court. FTC enforcement does not put the burden on the user. In addition, there are other sources of regulatory mechanisms (e.g., international law or state law). International law includes the EU US Safe Harbor Privacy Principles (see, for example, “U.S.-EU Safe Harbor List,” 2015). While these mechanisms do provide additional privacy protections, our research maintained focus

on U.S. federal regulatory mechanisms.

<sup>4</sup>We use the FTC version instead of, for example, the Organisation for Economic Co-operation and Development (OECD) version of the FIPPs (OECD Publishing, 2002) or other versions, as we are focused on U.S. federal regulations. While this is consistent with our focus, we note that the OECD version includes some principles regarding limitations on data collection and use that also may be useful to consider for full privacy protection.

<sup>5</sup>We intend this tool to be used as a thought exercise for anyone developing or evaluating a privacy-protective system.

<sup>6</sup>See Google, “Control Your Google Ads,” undated, for the link to Google’s ads-preferences tool.

---

## REFERENCES

Apple, Inc., “iOS Security—Whitepaper,” May 2016. As of July 7, 2016:

[https://www.apple.com/business/docs/iOS\\_Security\\_Guide.pdf](https://www.apple.com/business/docs/iOS_Security_Guide.pdf)

Atienza, Audie A., Christina Zarcadoolas, Wendy Vaughn, Penelope Hughes, Vaishali Patel, Wen-Ying Sylvia Chou, and Joy Pritts, “Consumer Attitudes and Perceptions on mHealth Privacy and Security: Findings from a Mixed-Methods Study,” *Journal of Health Communications*, Vol. 20, No. 6, 2015, pp. 673–679.

Balebako, Rebecca, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen, “‘Little Brothers Watching You’: Raising Awareness of Data Leaks on Smartphones,” *ACM Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, Newcastle, United Kingdom, July 24–26, 2013.

Balebako, Rebecca, Abigail Marsh, Jialiu Lin, Jason I. Hong, and Lorrie Faith Cranor, “The Privacy and Security Behaviors of Smartphone App Developers,” *Proceedings of Internet Society’s NDSS Workshop on Usable Security (USEC)*, 2014.

Boyles, Jan Lauren, Aaron Smith, and Mary Madden, *Privacy and Data Management on Mobile Devices*, Pew Internet and American Life Project, September 5, 2012.

Cavoukian, Ann, “Privacy by Design: 7 Foundational Principles,” Toronto, Canada: Information and Privacy Commissioner of Ontario, August 2009 (revised January 2011). As of October 20, 2015: <http://www.privacybydesign.ca/index.php/about-pbd/7-foundational-principles/>

Felt, Andrew Porter, Serge Egelman, and David Wagner, “I’ve Got 99 Problems, but Vibration Ain’t One: A Survey of Smartphone Users’ Concerns,” *Proceedings of the Second Annual ACM Conference on Computer and Communications Security (CCS) Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)*, Raleigh, N.C., October 19, 2012.

Google, “Control Your Google Ads,” undated. As of July 8, 2016: <https://www.google.com/settings/u/0/ads/authenticated>

“Google Report: Android Security 2014 Year in Review,” 2014. As of October 2015: [https://static.googleusercontent.com/media/source.android.com/en//devices/tech/security/reports/Google\\_Android\\_Security\\_2014\\_Report\\_Final.pdf](https://static.googleusercontent.com/media/source.android.com/en//devices/tech/security/reports/Google_Android_Security_2014_Report_Final.pdf)

Harris, Andrew, Seymour Goodman, and Patrick Traynor, “Privacy and Security Concerns Associated with Mobile Money Applications in Africa,” *Washington Journal of Law, Technology, and Arts*, Vol. 8, No. 3, 2012, p. 245.

Lindorfer, Martina, Matthias Neugschwandtner, Lukas Weichselbaum, Yanick Fratantonio, Victor Van Der Veen, and Christian Platzer, “Andrubis-1,000,000 Apps Later: A View on Current Android Malware Behaviors,” *Proceedings of the the Third International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, September 11, 2014.

Lockheimer, Hiroshi, “Android and Security,” *Google Mobile Blog*, February 2, 2012. As of July 7, 2016: <http://googlemobile.blogspot.com/2012/02/android-and-security.html>

Madden, Mary, and Lee Rainie, *Americans’ Views About Data Collection and Security*, Pew Research Center, May 20, 2015. As of November 17, 2015: <http://www.pewinternet.org/2015/05/20/americans-views-about-data-collection-and-security/>

McDaniel, Patrick, “Bloatware Comes to the Smartphone,” *IEEE Security and Privacy*, Vol. 10, No. 4, July–August 2012, pp. 85–87.

Microsoft, “Security Considerations for Reflection,” *Microsoft Developer Network*, 2016. As of February 17, 2016: [https://msdn.microsoft.com/en-us/library/stfy7tfc\(v=vs.110\).aspx](https://msdn.microsoft.com/en-us/library/stfy7tfc(v=vs.110).aspx)

Miller, Charlie, “Real World Fuzzing,” *Independent Security Evaluators*, October 19, 2007. As of May 24, 2016: <https://crypto.stanford.edu/cs155/papers/fuzzing.pdf>

Muslukhov, Ildar, Yazan Boshmaf, Cynthia Kuo, Jonathan Lester, and Konstantin Beznosov, “Understanding Users’ Requirements for Data Protection in Smartphones,” *2012 IEEE 28th Technical Conference on Data Engineering Workshops*, April 1–5, 2012, pp. 228–235. As of July 7, 2016: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6313685>

OECD Publishing, “OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,” Organisation for Economic Co-operation and Development, 2002. As of July 7, 2016: <http://www.oecd.org/sti/ieconomy/oecdguidelinesonthe protection ofprivacyandtransborderflowsofpersonaldata.htm>

State of California Department of Justice, Office of the Attorney General, "Press Release: Attorney General Kamala D. Harris Announces Privacy Enforcement and Protection Unit," Sacramento, Calif., July 19, 2012a. As of June 2, 2016:

<http://oag.ca.gov/news/press-releases/attorney-general-kamala-d-harris-announces-privacy-enforcement-and-protection>

———, "Press Release: Attorney General Kamala D. Harris Notifies Mobile App Developers of Non-Compliance with California Privacy Law," Sacramento, Calif., October 30, 2012b. As of June 2, 2016: <http://oag.ca.gov/news/press-releases/attorney-general-kamala-d-harris-notifies-mobile-app-developers-non-compliance>

Taylor, Linnet, "No Place to Hide? The Ethics and Analytics of Tracking Mobility Using Mobile Phone Data," *Environment and Planning D: Society and Space*, Vol. 34, No. 2, 2015, pp. 319-336, DOI: 0263775815608851.

Turnow, Joseph, *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth*, New Haven, Conn.: Yale University Press, 2012.

United States Code, Title 15, Sections 6501–6506, Children's Online Privacy Protection, 1998.

———, Title 15, Section 1681, Fair Credit Reporting Act, 2012.

Urban, Jennifer M., Chris Jay Hoofnagle, and Su Li, "Mobile Phones and Privacy," *UC Berkeley Public Law Research Paper Series*, No. 2103405, July 10, 2012.

"U.S.-EU Safe Harbor List," export.gov, 2015. As of July 14, 2016: <https://safeharbor.export.gov/list.aspx>

U.S. Federal Trade Commission, "Protecting Consumer Privacy," undated. As of October 20, 2015: <https://www.ftc.gov/news-events/media-resources/protecting-consumer-privacy>

———, "Privacy Online: Fair Information Practices in the Electronic Marketplace, a Report to Congress," May 2000. As of November 19, 2015: <http://www.ftc.gov/reports/privacy2000/privacy2000.pdf>

Vidas, Timothy, Daniel Votipka, and Nicolas Christin, "All Your Droid Are Belong to Us: A Survey of Current Android Attacks," *WOOT '11 Proceedings of 5th USENIX Workshop on Offensive Technologies*, San Francisco, Calif., August 8–12, 2011, pp. 81–90.

Ware, Willis H., *RAND and the Information Evolution: A History in Essays and Vignettes*, Santa Monica, Calif.: RAND Corporation, CP-537-RC, 2008. As of July 7, 2016: [http://www.rand.org/pubs/corporate\\_pubs/CP537.html](http://www.rand.org/pubs/corporate_pubs/CP537.html)

Westin, Alan F., "Privacy and Freedom," *Washington and Lee Law Review*, Vol. 25, No. 1, March 1, 1968, p. 166.

Whitney, Lance, "Comodohacker Returns in Diginotar Incident," *CNET*, September 6, 2011. As of July 7, 2016: <http://www.cnet.com/news/comodohacker-returns-in-diginotar-incident/>

Wright, Charles V., Lucas Ballard, Scott E. Coull, Fabian Monrose, and Gerald M. Masson, "Uncovering Spoken Phrases in Encrypted Voice over IP Conversations," *ACM Transactions on Information and System Security (TISSEC)*, Vol. 13, No. 4, December 2010, p. 35.

Wright, Charles V., Lucas Ballard, Fabian Monrose, and Gerald M. Masson, "Language Identification of Encrypted VoIP Traffic: Alejandra y Roberto or Alice and Bob?" *The 16th USENIX Security Symposium Proceedings*, 2007, p. 3.

---

## About the Authors

Researchers from the RAND Corporation and MIT Lincoln Laboratory combined their expertise in technology, policy, and law to get a broader understanding of smartphone privacy.

### **RAND Team Members**

Rebecca Balebako is a computer scientist whose research focuses on digital privacy and human-computer interaction.

Anne Boustead, a doctoral candidate in policy analysis and a lawyer, studies issues related to privacy, technology, and interstate legal variation.

Karlyn Stanley, a lawyer specializing in emerging technologies, has extensive experience representing mobile communications carriers, wireless infrastructure (cell tower), and other telecom clients, including AT&T.

William (Bill) Welser IV is the director of RAND's Engineering and Applied Sciences Department. His research expertise includes applied cryptography and technology policy.

Zev Winkelman is a computer engineer with a Ph.D. in public policy and a research focus on privacy and security issues.

### **MIT Lincoln Laboratory Team Members**

Arkady Yerukhimovich is a technical staff member at MIT Lincoln Laboratory. His research expertise is in applied and theoretical cryptography. His recent research has focused on leveraging theoretical cryptography tools for protecting privacy and security in real-world applications.

Rob Cunningham is the leader of the Secure, Resilient Systems and Technology Group at MIT Lincoln Laboratory. His research expertise is broad, covering machine learning, image processing, and computer security. He is interested in a range of privacy-preserving technologies.

Rick Housley is a computer engineering student pursuing his bachelor's and master's degrees at Stevens Institute of Technology. His work largely focuses on embedded-device prototyping and embedded device security.

Richard Shay is a technical staff member at MIT Lincoln Laboratory. His research expertise is usable privacy and security.

Chad Spensky is a computer researcher and a member of the Seclab at the University of California, Santa Barbara, and the Cyber System Assessments group at MIT Lincoln Laboratory. His current research interests include low-level security protocols on mobile platforms, embedded-systems security, smart card security, usable authentication, and novel introspection techniques.

Jeffrey Stewart is an associate staff member of the Cyber System Assessments Group at MIT Lincoln Laboratory. His expertise includes embedded systems and low-level system security.

Ari Trachtenberg is a professor of electrical and computer engineering at Boston University and participated in this project while on sabbatical at MIT Lincoln Laboratory.

---

## About This Report

This report documents research findings resulting from collaboration between Massachusetts Institute of Technology (MIT) Lincoln Laboratory and the RAND Corporation. The RAND project is titled “Overview of Regulation Protecting the Privacy of Mobile Device Users Data,” and the MIT Lincoln Laboratory project is titled “User Privacy on iOS and Android Devices.” The Brandeis Program within the Information Innovation Office of the Defense Advanced Research Projects Agency (DARPA) sponsored this work.

This report provides an assessment of smartphone users’ privacy from technical and regulatory perspectives. From a technical perspective, it describes a review of literature and experiments performed by MIT Lincoln Laboratory investigating the state of privacy on the current major smartphone platforms in 2015: Google’s Android and Apple’s iPhone operating system (iOS). From a regulatory perspective, this report describes the major federal regulatory mechanisms for protecting privacy in the United States. Additionally, it introduces a framework for identifying gaps and strengths in privacy protection from both the technical and regulatory perspectives.

The RAND portion of this research was sponsored by DARPA and conducted within the Acquisition and Technology Policy Center of the RAND National Defense Research Institute, a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, the Unified Combatant Commands, the Department of the Navy, the Marine Corps, the defense agencies, and the defense Intelligence Community.

For more information on the RAND Acquisition and Technology Policy Center, see [www.rand.org/nsrd/ndri/centers/atp](http://www.rand.org/nsrd/ndri/centers/atp) or contact the director (contact information is provided on the web page).

The MIT Lincoln Laboratory portion of this research is based upon work supported by the Defense Advanced Research Projects Agency under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency.

## Limited Print and Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents for commercial use. For information on reprint and linking permissions, please visit [www.rand.org/pubs/permissions.html](http://www.rand.org/pubs/permissions.html).

For more information on this publication, visit [www.rand.org/t/rr1393](http://www.rand.org/t/rr1393).

© Copyright 2016 RAND Corporation



MIT Lincoln Laboratory is a Department of Defense (DoD) federally funded research and development center working on problems critical to national security.



The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

RAND’s publications do not necessarily reflect the opinions of its research clients and sponsors. **RAND®** is a registered trademark.

[www.rand.org](http://www.rand.org)