

# The Android Platform Security Model\*

RENÉ MAYRHOFER, Google and Johannes Kepler University Linz

JEFFREY VANDER STOEP, Google

CHAD BRUBAKER, Google

NICK KRALEVICH, Google

Android is the most widely deployed end-user focused operating system. With its growing set of use cases encompassing communication, navigation, media consumption, entertainment, finance, health, and access to sensors, actuators, cameras, or microphones, its underlying security model needs to address a host of practical threats in a wide variety of scenarios while being useful to non-security experts. The model needs to strike a difficult balance between security, privacy, and usability for end users, assurances for app developers, and system performance under tight hardware constraints. While many of the underlying design principles have implicitly informed the overall system architecture, access control mechanisms, and mitigation techniques, the Android security model has previously not been formally published. This paper aims to both document the abstract model and discuss its implications. Based on a definition of the threat model and Android ecosystem context in which it operates, we analyze how the different security measures in past and current Android implementations work together to mitigate these threats. There are some special cases in applying the security model, and we discuss such deliberate deviations from the abstract model.

CCS Concepts: • **Security and privacy** → **Software and application security**; **Domain-specific security and privacy architectures**; **Operating systems security**; • **Human-centered computing** → **Ubiquitous and mobile devices**.

Additional Key Words and Phrases: Android, security, operating system, informal model

## ACM Reference Format:

René Mayrhofer, Jeffrey Vander Stoep, Chad Brubaker, and Nick Kralevich. 2020. The Android Platform Security Model. *ACM Trans. Priv. Sec.* 1, 1 (May 2020), 32 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Android is, at the time of this writing, the most widely deployed end-user operating system. With more than 2 billion monthly active devices [8] and a general trend towards mobile use of Internet services, Android is now the most common interface for global users to interact with digital services. Across different form factors (including e.g. phones, tablets, wearables, TV, Internet-of-Things, automobiles, and more special-use categories) there is a vast – and still growing – range of use cases from communication, media consumption, and entertainment to finance, health, and physical sensors/actuators. Many of these applications are increasingly security and privacy critical,

---

\*Last updated in March 2020 based on Android 10 as released and some pre-release public sources of Android 11.

---

Authors' addresses: René Mayrhofer, Google and Johannes Kepler University Linz, [rmayrhofer@google.com](mailto:rmayrhofer@google.com); Jeffrey Vander Stoep, Google, [jeffv@google.com](mailto:jeffv@google.com); Chad Brubaker, Google, [cbrubaker@google.com](mailto:cbrubaker@google.com); Nick Kralevich, Google, [nnk@google.com](mailto:nnk@google.com).

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2019 Association for Computing Machinery.

2471-2566/2020/5-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

50 and Android as an OS needs to provide sufficient and appropriate assurances to users as well as  
51 developers.

52 To balance the different (and sometimes conflicting) needs and wishes of users, application  
53 developers, content producers, service providers, and employers, Android is fundamentally based  
54 on a multi-party consent<sup>1</sup> model: *an action should only happen if all involved parties consent to it*.  
55 If any party does not consent, the safe-by-default choice is for that action to be blocked. This is  
56 different to the security models that more traditional operating systems implement, which are  
57 focused on user access control and do not explicitly consider other stakeholders.

58 While the multi-party model has implicitly informed architecture and design of the Android  
59 platform from the beginning, it has been refined and extended based on experience gathered from  
60 past releases. This paper aims to both document the Android security model and determine its  
61 implications in the context of ecosystem constraints and historical developments. Specifically, we  
62 make the following contributions:

- 63 (1) We motivate and for the first time define the Android security model based on security  
64 principles and the wider context in which Android operates. Note that the core multi-party  
65 consent model described and analyzed in this paper has been implicitly informing Android  
66 security mechanisms since the earliest versions, and we therefore systematize knowledge  
67 that has, in parts, existed before, but that was not formally published so far.
- 68 (2) We define the threat model and how the security model addresses it and discuss implications  
69 as well as necessary special case handling.
- 70 (3) We explain how AOSP (Android Open Source Project, the reference implementation of  
71 the Android platform) enforces the security model based on multiple interacting security  
72 measures on different layers.
- 73 (4) We identify currently open gaps and potential for future improvement of this implementation.  
74

75 *Android as a platform.* This paper focuses on security and privacy measures in the Android  
76 platform itself, i.e. code running on user devices that is part of AOSP. Within the scope of this  
77 paper, we define the *platform* as the set of AOSP components that together form an Android  
78 system passing the Compatibility Test Suite (CTS). While some parts of the platform may be  
79 customized or proprietary for different vendors, AOSP provides reference implementations for  
80 nearly all components, including the e.g. Linux kernel<sup>2</sup>, Trusty as an ARM TEE<sup>3</sup>, or libavb for boot  
81 loader side verified boot<sup>4</sup> that are sufficient to run a fully functional Android system on reference  
82 development hardware<sup>5</sup>. Note that Google Mobile Services (GMS), including Google Play Services  
83 (also referred to as GmsCore), Google Play Store, Google Search, Chrome, and other standard apps  
84 are sometimes considered part of the platform, as they provide dependencies for common services  
85 such as location estimation or cloud push messaging. Android devices that are certified to support  
86 GMS are publicly listed<sup>6</sup>. While replacements for these components exist (including an independent,  
87 minimal open source version called microG<sup>7</sup>), they may not be complete or behave differently.  
88 Concerning the security model described in this paper, we do not consider GMS to be part of the  
89 platform, as they are also subject to the security policy defined and enforced by AOSP components.  
90

91 <sup>1</sup>Throughout the paper, the term ‘consent’ is used to refer to various technical methods of declaring or enforcing a party’s  
92 intent, rather than the legal requirement or standard found in many privacy legal regimes around the world.

93 <sup>2</sup><https://android.googlesource.com/kernel/common/>

94 <sup>3</sup><https://android.googlesource.com/trusty/vendor/google/aosp/>

95 <sup>4</sup><https://android.googlesource.com/platform/external/avb/>

96 <sup>5</sup><https://source.android.com/setup/build/devices>

97 <sup>6</sup>[https://storage.googleapis.com/play\\_public/supported\\_devices.html](https://storage.googleapis.com/play_public/supported_devices.html)

98 <sup>7</sup>[https://github.com/microg/android\\_packages\\_apps\\_GmsCore/wiki](https://github.com/microg/android_packages_apps_GmsCore/wiki)

99 In terms of higher-level security measures, there are services complementary to those imple-  
100 mented in AOSP in the form of Google Play Protect (GPP) scanning applications submitted to  
101 Google Play and on-device (Verify Apps or Safe Browsing as opt-in services) as well as Google Play  
102 policy and other legal frameworks. These are out of scope of the current paper, but are covered by  
103 related work [17, 43, 67, 115]. However, we explicitly point out one policy change in Google Play  
104 with potentially significant positive effects for security: Play now requires that new apps and app  
105 updates target a recent Android API level, which will allow Android to deprecate and remove APIs  
106 known to be abused or that have had security issues in the past [54].

107  
108 *Structure.* In the following, we will first introduce Android security principles, and the ecosystem  
109 context and threat analysis that are the basis of the Android security model (Section 2). Then, we  
110 define the central security model (Section 3) and its implementation in the form of OS architecture  
111 and enforcement mechanisms on different OS layers (Section 4). Note that all implementation  
112 specific sections refer to Android 10 at the time of its initial release unless mentioned otherwise  
113 (cf. [39] for changes in Android 10 and [101] for changes in Android 9). We will refer to earlier  
114 Android version numbers instead of their code names: 4.1–4.3 (Jelly Bean), 4.4 (KitKat), 5.x (Lollipop),  
115 6.x (Marshmallow), 7.x (Nougat), 8.x (Oreo), and 9.x (Pie). All tables are based on an analysis of  
116 security relevant changes to the whole AOSP code base between Android releases 4.x and 10  
117 (inclusive), spanning about 9 years of code evolution. Finally, we discuss special cases (Section 5)  
118 and related work in terms of other security models (Section 6).

## 119 2 ANDROID BACKGROUND

120  
121 Before introducing the security model, we explain the context in which it needs to operate, both in  
122 terms of ecosystem requirements and platform security principles.

### 123 2.1 Ecosystem context

124  
125 Some of the design decisions need to be put in context of the larger *ecosystem*, which does not  
126 exist in isolation. A successful ecosystem is one where all parties benefit when it grows, but also  
127 requires a minimum level of mutual trust. This implies that a platform must create safe-by-default  
128 environments where the main parties (end user, application developer, operating system) can define  
129 mutually beneficial terms of engagement. If these parties cannot come to an agreement, then the  
130 most trust building operation is to disallow the action (default-deny). The Android platform security  
131 model introduced below is based on this notion.

132 This section is not comprehensive, but briefly summarizes those aspects of the Android ecosystem  
133 that have direct implications to the security model:

134 *Android is an end user focused operating system.* Although Android strives for flexibility, the main  
135 focus is on typical users. The obvious implication is that, as a consumer OS, it must be useful to  
136 users and attractive to developers.

137 The end user focus implies that user interfaces and workflows need to be safe by default and  
138 require explicit intent for any actions that could compromise security or privacy. This also means  
139 that the OS must not offload technically detailed security or privacy decisions to non-expert users  
140 who are not sufficiently skilled or experienced to make them [16].

141  
142 *The Android ecosystem is immense.* Different statistics show that in the last few years, the majority  
143 of a global, intensely diverse user base already used mobile devices to access Internet resources (i.e.  
144 63% in the US [1], 56% globally [2], with over 68% in Asia and over 80% in India). Additionally, there  
145 are hundreds of different OEMs (Original Equipment Manufacturers, i.e. device manufacturers)  
146 making tens of thousands of Android devices in different form factors [102] (including, but not  
147

148 limited to, standard smartphones and tablets, watches, glasses, cameras and many other Internet of  
149 things device types, handheld scanners/displays and other special-purpose worker devices, TVs,  
150 cars, etc.). Some of these OEMs do not have detailed technical expertise, but rely on ODMs (Original  
151 Device Manufacturers) for developing hardware and firmware and then re-package or simply  
152 re-label devices with their own brand. Only devices shipping with Google services integration need  
153 to get their firmware certified, but devices simply based off AOSP can be made without permission  
154 or registration. Therefore, there is no single register listing all OEMs, and the list is constantly  
155 changing with new hardware concepts being continuously developed. One implication is that  
156 changing APIs and other interfaces can lead to large changes in the device ecosystem and take time  
157 to reach most of these use cases.

158 However, devices using Android as a trademarked name to advertise their compatibility with  
159 Android apps need to pass the Compatibility Test Suite (CTS). Developers rely on this compatibility  
160 when writing apps for this wide variety of different devices. In contrast to some other platforms,  
161 Android explicitly supports installation of apps from arbitrary sources, which led to the development  
162 of different app stores and the existence of apps outside of Google Play. Consequently, there is a  
163 long tail of apps with a very specific purpose, being installed on only few devices, and/or targeting  
164 old Android API releases. Definition of and changes to APIs need to be considerate of the huge  
165 number of applications that are part of the Android ecosystem.

166  
167 *Apps can be written in any language.* As long as apps interface with the Android framework using  
168 the well-defined Java language APIs for process workflow, they can be written in any programming  
169 language, with or without runtime support, compiled or interpreted. Android does not currently  
170 support non-Java language APIs for the basic process lifecycle control, because they would have to  
171 be supported in parallel, making the framework more complex and therefore more error-prone.  
172 Note that this restriction is not directly limiting, but apps need to have at least a small Java language  
173 wrapper to start their initial process and interface with fundamental OS services. The important  
174 implication of this flexibility for security mechanisms is that they cannot rely on compile-time  
175 checks or any other assumptions on the build environment. Therefore, Android security needs to  
176 be based on runtime protections around the app boundary.

## 178 2.2 Android security principles

179 From the start, Android has assumed a few basic security and privacy principles that can be seen  
180 as an implicit contract between many parties in this open ecosystem:

181  
182 *Actors control access to the data they create.* Any actor that creates a data item is implicitly granted  
183 control over this particular instance of data representation. Note that this refers to the technical  
184 act of protecting data, either on the filesystem or in memory – but does not automatically imply  
185 ownership over data in the legal sense. While this model has long been the default for filesystem  
186 access control (DAC, cf. section 4.3.1 below), we consider it a wider design principle. One intention  
187 of making this principle explicit is that exceptions such as device backup (cf. section 5) can be  
188 argued about within the scope of the security model.

189  
190 *Consent is informed and meaningful.* Actors consenting to any action must be empowered to base  
191 their decision on information about the action and its implications and must have meaningful ways  
192 to grant or deny this consent. This applies to both users and developers, although very different  
193 technical means of enforcing (lack of) consent apply. Consent is not only required from the actor  
194 that created a data item, but from all involved actors. Consent decisions should be enforced and  
195 not self-policed.

196

197 *Safe by design/default.* Components should be safe by design. That is, the default use of an  
198 operating system component or service should always protect security and privacy assumptions,  
199 potentially at the cost of blocking some use cases. This principle applies to modules, APIs, com-  
200 munication channels, and generally to interfaces of all kinds. When variants of such interfaces  
201 are offered for more flexibility (e.g. a second interface method with more parameters to override  
202 default behavior), these should be hard to abuse, either unintentionally or intentionally. Note that  
203 this architectural principle targets developers, which includes device manufacturers, but implicitly  
204 includes users in how security is designed and presented in user interfaces. Android targets a wide  
205 range of developers and intentionally keeps barriers to entry low for app development. Making it  
206 hard to abuse APIs not only guards against malicious adversaries, but also mitigates genuine errors  
207 resulting e.g. from incomplete knowledge of an interface definition or caused by developers lacking  
208 experience in secure system design. As in the defense in depth approach, there is no single solution  
209 to making a system safe by design. Instead, this is considered a guiding principle for defining new  
210 interfaces and refining – or, when necessary, deprecating and removing – existing ones.

211  
212 *Defense in depth.* A robust security system is not sufficient if the acceptable behavior of the  
213 operating system allows an attacker to accomplish all of their goals without bypassing the security  
214 model (e.g. ransomware encrypting all files it has access to under the access control model).  
215 Specifically, violating any of the above principles should require such bypassing of controls on-  
216 device (in contrast to relying on off-device verification e.g. at build time).

217 Therefore, the primary goal of any security system is to enforce its model. For Android operating  
218 in a multitude of environments (see below for the threat model), this implies an approach that  
219 does not immediately fail when a single assumption is violated or a single implementation bug is  
220 found, even if the device is not up to date. Defense in depth is characterized by rendering individual  
221 vulnerabilities more difficult or impossible to exploit, and increasing the number of vulnerabilities  
222 required for an attacker to achieve their goals. We primarily adopt four common security strategies to  
223 prevent adversaries from bypassing the security model: *isolation and containment*, *exploit mitigation*,  
224 *integrity*, and *patching/updates*. Their implementation will be discussed in more detail in section 4.

### 225 226 2.3 Threat model

227 Threat models for mobile devices are different from those commonly used for desktop or server  
228 operating systems for two major reasons: by definition, mobile devices are easily lost or stolen, and  
229 they connect to untrusted networks as part of their expected usage. At the same time, by being  
230 close to users at most times, they are also exposed to even more privacy sensitive data than many  
231 other categories of devices. Recent work [95] previously introduced a layered threat model for  
232 mobile devices which we adopt for discussing the Android security model within the scope of this  
233 paper:

234  
235 *Adversaries can get physical access to Android devices.* For all mobile and wearable devices, we  
236 have to assume that they will potentially fall under physical control of adversaries at some point.  
237 The same is true for other Android form factors such as things, cars, TVs, etc. Therefore, we assume  
238 Android devices to be either directly accessible to adversaries or to be in physical proximity to  
239 adversaries as an explicit part of the threat model. This includes loss or theft, but also multiple  
240 (benign but potentially curious) users sharing a device (such as a TV or tablet). We derive specific  
241 threats due to physical or proximal access:

- 242  
243 **T1** Powered-off devices under complete physical control of an adversary (with potentially high  
244 sophistication up to nation state level attackers), e.g. border control or customs checks.

**T2** Screen locked devices under complete physical control of an adversary, e.g. thieves trying to exfiltrate data for additional identity theft.

**T3** Screen unlocked (shared) devices under control of an authorized but different user, e.g. intimate partner abuse, voluntary submission to a border control, or customs check.

**T4** (Screen locked or unlocked) devices in physical proximity to an adversary (with the assumed capability to control all available radio communication channels, including cellular, WiFi, Bluetooth, GPS, NFC, and FM), e.g. direct attacks through Bluetooth [5, 52]. Although NFC could be considered to be a separate category to other proximal radio attacks because of the scale of distance, we still include it in the threat class of proximity instead of physical control.

*Network communication is untrusted.* The standard assumption of network communication under complete control of an adversary certainly also holds for Android devices. This includes the first hop of network communication (e.g. captive WiFi portals breaking TLS connections and malicious fake access points) as well as other points of control (e.g. mobile network operators or national firewalls), summarized in the usual Dolev-Yao model [57] with additional relay threats for short-range radios (e.g. NFC or BLE wormhole attacks [105]). For practical purposes, we mainly consider two network-level threats:

**T5** Passive eavesdropping and traffic analysis, including tracking devices within or across networks, e.g. based on MAC address or other device network identifiers.

**T6** Active manipulation of network traffic, e.g. MITM on TLS connections or relaying.

These two threats are different from [T4] (proximal radio attacks) in terms of scalability of attacks. Controlling a single choke point in a major network can be used to attack a large number of devices, while proximal (last hop) radio attacks require physical proximity to target devices.

*Untrusted code is executed on the device.* One fundamental difference to other mobile operating systems is that Android intentionally allows (with explicit consent by end users) installation of application code from arbitrary sources, and does not enforce vetting of apps by a central instance. This implies attack vectors on multiple levels (cf. [95]):

**T7** Abusing APIs supported by the OS with malicious intent, e.g. spyware.

**T8** Exploiting bugs in the OS, e.g. kernel, drivers, or system services [6, 9, 10, 12].

**T9** Abusing APIs supported by other apps installed on the device [11].

**T10** Untrusted code from the web (i.e. JavaScript) is executed without explicit consent.

**T11** Mimicking system or other app user interfaces to confuse users (based on the knowledge that standard in-band security indicators are not effective [56, 103]), e.g. to input PIN/password into a malicious app [66].

**T12** Reading content from system or other app user interfaces, e.g. to screen-scrape confidential data from another app [78, 84].

**T13** Injecting input events into system or other app user interfaces [69].

*Untrusted content is processed by the device.* In addition to directly executing untrusted code, devices process a wide variety of untrusted data, including rich (in the sense of complex structure) media. This directly leads to threats concerning processing of data and metadata:

**T14** Exploiting code that processes untrusted content in the OS or apps, e.g. in media libraries [4]. This can be both a local as well as a remote attack surface, depending on where input data is taken from.

**T15** Abusing unique identifiers for targeted attacks (which can happen even on trusted networks), e.g. using a phone number or email address for spamming or correlation with other data sets, including locations.



### 3 THE ANDROID PLATFORM SECURITY MODEL

The basic security model described in this section has informed the design of Android, and has been refined but not fundamentally changed. Given the ecosystem context and general Android principles explained above, the Android security model balances security and privacy requirements of users with security requirements of applications and the platform itself. The threat model described above includes threats to all stakeholders, and the security model and its enforcement by the Android platform aims to address all of them. The Android platform security model is informally defined by 5 rules:

① *Multi-party consent*. No action should be executed unless all main parties agree – in the standard case, these are *user*, *platform*, and *developer* (implicitly representing stakeholders such as content producers and service providers). Any one party can veto the action. This multi-party consent spans the traditional two dimensions of subjects (users and application processes) vs. objects (files, network sockets and IPC interfaces, memory regions, virtual data providers, etc.) that underlie most security models (e.g. [113]). Focusing on (regular and pseudo) files as the main category of objects to protect, the default control over these files depends on their location and which party created them:

- Data in shared storage is controlled by users.
- Data in private app directories and app virtual address space is controlled by apps.
- Data in special system locations is controlled by the platform (e.g. list of granted permissions).

However, it is important to point out that, under multi-party consent, even if one party primarily controls a data item, it may only act on it if the other involved parties consent. Control over data also does not imply ownership (which is a legal concept rather than a technical one and therefore outside the scope of an OS security model).

There are corner cases in which only a subset of all parties may need to consent (for actions in which the user only uses platform/OS services without involvement of additional apps) or an additional party may be introduced (e.g. on devices or profiles controlled by a mobile device management, this policy is also considered as a party for consenting to an action).

Public information and resources are out of scope of this access control and available to all parties; particularly all static code and data contained in the AOSP system image and apps (mostly in the Android Package (APK) format) is considered to be public (cf. Kerckhoff's principle). However, it is generally accepted that such public code and data is read-only to all parties and its integrity needs to be protected, which is explicitly in scope of the security measures.

② *Open ecosystem access*. Both users and developers are part of an open ecosystem that is not limited to a single application store. Central vetting of developers or registration of users is not required. This aspect has an important implication for the security model: generic app-to-app interaction is explicitly supported. Instead of creating specific platform APIs for every conceivable workflow, app developers are free to define their own APIs they offer to other apps.

③ *Security is a compatibility requirement*. The security model is part of the Android specification, which is defined in the Compatibility Definition Document (CDD) [20] and enforced by the Compatibility (CTS), Vendor (VTS), and other test suites. Devices that do not conform to CDD and do not pass CTS are not Android. Within the scope of this paper, we define *rooting* as modifying the system to allow starting processes that are not subject to sandboxing and isolation. Such rooting, both intentional and malicious, is a specific example of a non-compliant change which violates CDD. As such, only CDD-compliant devices are considered. While many devices support unlocking their

bootloader and flashing modified firmware<sup>8</sup>, such modifications may be considered incompatible under CDD if security assurances do not hold. Verified boot and hardware key attestation can be used to validate if currently running firmware is in a known-good state, and in turn may influence consent decisions by users and developers.

④ *Factory reset restores the device to a safe state.* In the event of security model bypass leading to a persistent compromise, a factory reset, which wipes/reformats the writable data partitions, returns a device to a state that depends only on integrity protected partitions. In other words, system software does not need to be re-installed, but wiping the data partition(s) will return a device to its default state. Note that the general expectation is that the read-only device software may have been updated since originally taking it out of the box, which is intentionally not downgraded by factory reset. Therefore, more specifically, factory reset returns an Android device to a state that only depends on system code that is covered by Verified Boot, but does not depend on writable data partitions.

⑤ *Applications are security principals.* The main difference to traditional operating systems that run apps in the context of the logged-in user account is that Android apps are not considered to be fully authorized agents for user actions. In the traditional model typically implemented by server and desktop OS, there is often no need to even exploit the security boundary because running malicious code with the full permissions of the main user is sufficient for abuse. Examples are many, including file encrypting ransomware [80, 107] (which does not violate the OS security model if it simply re-writes all the files the current user account has access to) and private data leakage (e.g. browser login tokens [92], history or other tracking data, cryptocurrency wallet keys, etc.).

*Summary.* Even though, at first glance, the Android security model grants less power to users compared to traditional operating systems that do not impose a multi-party consent model, there is an immediate benefit to end users: if one app cannot act with full user privileges, the user cannot be tricked into letting it access data controlled by other apps. In other words, requiring application developer consent – enforced by the platform – helps avoid user confusion attacks and therefore better protects private data.

## 4 IMPLEMENTATION

Android's security measures implement the security model and are designed to address the threats outlined above. In this section we describe security measures and indicate which threats they mitigate, taking into account the architectural security principles of 'defense in depth' and 'safe by design'.

### 4.1 Consent

Methods of giving meaningful consent vary greatly between actors, as well as potential issues and constraints.

#### 4.1.1 Developer(s)

Unlike traditional desktop operating systems, Android ensures that the developer consents to actions on their app or their app's data. This prevents large classes of abusive behavior where unrelated apps inject code into or access/leak data from other applications on a user's device.

<sup>8</sup>Google Nexus and Pixel devices as well as many others support the standard `fastboot oem unlock` command to allow flashing any firmware images to actively support developers and power users. However, executing this unlocking workflow will forcibly factory reset the device (wiping all data) to make sure that security guarantees are not retroactively violated for data on the device.



393 Consent for developers, unlike the user, is given via the code they sign and the system executes,  
394 uploading the app to an app store and agreeing to the associated terms of service, and obeying  
395 other relevant policies (such as CDD for code by an OEM in the system image). For example, an  
396 app can consent to the user sharing its data by providing a respective mechanism, e.g. based on OS  
397 sharing methods such as built-in implicit Intent resolution chooser dialogs [21]. Another example  
398 is debugging: as assigned virtual memory content is controlled by the app, debugging from an  
399 external process is only allowed if an app consents to it (specifically through the debuggable  
400 flag in the app manifest). By uploading an app to the relevant app store, developers also provide  
401 the consent for this app to be installed on devices that fetch from that store under appropriate  
402 pre-conditions (e.g. after successful payment).

403 Meaningful consent then is ensuring that APIs and their behaviors are clear and the developer  
404 understands how their application is interacting with or providing data to other components. Addi-  
405 tionally, we assume that developers of varying skill levels may not have a complete understanding  
406 of security nuances, and as a result APIs must also be safe by default and difficult to incorrectly use  
407 in order to avoid accidental security regressions.

408 In order to ensure that it is the app developer and not another party that is consenting, applications  
409 are signed by the developer (or when using key rotation functionality, a key that was previously  
410 granted this ability by the app). This prevents third parties – including the app store – from  
411 replacing or removing code or resources in order to change the app’s intended behavior. However,  
412 the app signing key is trusted implicitly upon installation, so replacing or modifying apps in transit  
413 (e.g. when side-loading apps) is currently out of scope of the platform security model and may  
414 violate developer consent.

#### 415 4.1.2 The Platform

416 While the platform, like the developer, consents via code signing, the goals are quite different: the  
417 platform acts to ensure that the system functions as intended. This includes enforcing regulatory  
418 or contractual requirements (e.g. communication in cell-based networks) as well as taking an  
419 opinionated stance on what kinds of behaviors are acceptable (e.g. mitigating apps from applying  
420 deceptive behavior towards users). Platform consent is enforced via Verified Boot (see below  
421 for details) protecting the system images from modification, internal compartmentalization and  
422 isolation between components, as well as platform applications using the platform signing key and  
423 associated permissions, much like applications.

424  
425 *Note on the platform as a party:* Depending on how the involved stakeholders (parties for  
426 consent) and enforcing mechanisms are designated, either an inherent or an apparent asymmetry  
427 of power to consent may arise:

428  
429 (a) If the Android “platform” is seen as a single entity (composed of hardware, firmware, OS  
430 kernel, system services, libraries, and app runtime), then it may be considered omniscient in the  
431 sense of having access to and effectively controlling all data and processes on the system. Under  
432 this point of view, the conflict of interest between being one party of consent and simultaneously  
433 being the enforcing agent gives that entity – the platform – overreaching power over all other  
434 parties.

435 (b) If Android as a platform is considered in depth, it consists of many different components.  
436 These can be considered individual representatives of the platform for a particular interaction  
437 involving multi-party consent, while other components act as enforcing mechanism for that  
438 consent. In other words, the Android platform is structured in such a way as to minimize trust in  
439 itself and contain multiple mechanisms of isolating components from each other to enforce each  
440 other’s limitations (cf. section 4.3). One example is playing media files: even when called by an  
441

app, a media codec cannot directly access the underlying resources if the user has not granted this through the media server, because MAC policies in the Linux kernel do not allow such bypass (cf. section 4.3.3). Another example is storage of cryptographic keys, which is isolated even from the Linux kernel itself and enforced through hardware separation (cf. section 4.3.5). While this idealized model of platform parties requiring consent for their actions is the abstract goal of the security model we describe, in practice there still are individual components that sustain the asymmetry between the parties. Each new version of Android continues to further strengthen the boundaries of platform components among each other, as described in more detail below.

Within the scope of this paper, we take the second perspective when it comes to notions of consent involving the platform itself, i.e. considering the platform to be multiple parties whose consent is being enforced by independent mechanisms (mostly the Linux kernel isolating platform components from each other). However, when talking about the whole system implementing our Android security model, in favor of simpler expression we will generally refer to the platform as the combination of all (AOSP) components that together act as an enforcing mechanism for other parties, as defined in the introduction.

#### 4.1.3 *User(s)*

Achieving meaningful user consent is by far the most difficult and nuanced challenge in determining meaningful consent. Some of the guiding principles have always been core to Android, while others were refined based on experiences during the 10 years of development so far:

- **Avoid over-prompting.** Over-prompting the user leads to prompt fatigue and blindness (cf. [18]). Prompting the user with a yes/no prompt for every action does not lead to meaningful consent as users become blind to the prompts due to their regularity.
- **Prompt in a way that is understandable.** Users are assumed not to be experts or understand nuanced security questions (cf. [65]). Prompts and disclosures must be phrased in a way that a non-technical user can understand the effects of their decision.
- **Prefer pickers and transactional consent over wide granularity.** When possible, we limit access to specific items instead of the entire set. For example, the Contacts Picker allows the user to select a specific contact to share with the application instead of using the Contacts permission. These both limit the data exposed as well as present the choice to the user in a clear and intuitive way.
- **The OS must not offload a difficult problem onto the user.** Android regularly takes an opinionated stance on what behaviors are too risky to be allowed and may avoid adding functionality that may be useful to a power user but dangerous to an average user.
- **Provide users a way to undo previously made decisions.** Users can make mistakes. Even the most security and privacy-savvy users may simply press the wrong button from time to time, which is even more likely when they are being tired or distracted. To mitigate against such mistakes or the user simply changing their mind, it should be easy for the user to undo a previous decision whenever possible. This may vary from denying previously granted permissions to removing an app from the device entirely.

Additionally, it is critical to ensure that the user who is consenting is the legitimate user of the device and not another person with physical access to the device ([T1]-[T3]), which directly relies on the next component in the form of the Android lock screen. Implementing model rule ① is cross-cutting on all system layers.

We use two examples to better describe the consent parties:

- Sharing data from one app to another requires:

- user consent through the user selecting a target app in the share dialog;
- developer consent of the source app by initiating the share with the data (e.g. image) they want to allow out of their app;
- developer consent of the target app by accepting the shared data; and
- platform consent by arbitrating the data access between different components and ensuring that the target app cannot access any other data than the explicitly shared item through the same link, which forms a temporary trust relationship between two apps.
- Changing mobile network operator (MNO) configuration option requires:
  - user consent by selecting the options in a settings dialog;
  - (MNO app) developer consent by implementing options to change these configuration items, potentially querying policy on backend systems; and
  - platform consent by verifying e.g. policies based on country regulations and ensuring that settings do not impact platform or network stability.

## 4.2 Authentication

Authentication is a gatekeeper function for ensuring that a system interacts with its owner or legitimate user. On mobile devices the primary means of authentication is via the lockscreen. Note that a lockscreen is an obvious trade-off between security and usability: On the one hand, users unlock phones for short (10-250 seconds) interactions about 50 times per day on average and even up to 200 times in exceptional cases [62, 75], and the lockscreen is obviously an immediate hindrance to frictionless interaction with a device [73, 74]. On the other hand, devices without a lockscreen are immediately open to being abused by unauthorized users ([T1]-[T3]), and the OS cannot reliably enforce user consent without authentication.

In their current form, lockscreens on mobile devices largely enforce a binary model – either the whole phone is accessible, or the majority of functions (especially all security or privacy sensitive ones) are locked. Neither long, semi-random alphanumeric passwords (which would be highly secure but not usable for mobile devices) nor swipe-only lockscreens (usable, but not offering any security) are advisable. Therefore, it is critically important for the lockscreen to strike a reasonable balance between security and usability.

Towards this end, recent Android releases use a tiered authentication model where a secure knowledge-factor based authentication mechanism can be backed by convenience modalities that are functionally constrained based on the level of security they provide. The added convenience afforded by such a model helps drive lockscreen adoption and allows more users to benefit both from the immediate security benefits of a lockscreen and from features such as file-based encryption that rely on the presence of an underlying user-supplied credential. As an example of how this helps drive lockscreen adoption, starting with Android 7.x we see that 77% of devices with fingerprint sensors have a secure lockscreen enabled, while only 50% of devices without fingerprints have a secure lockscreen<sup>9</sup>.

As of Android 10, the tiered authentication model splits modalities into three tiers.

- *Primary Authentication* modalities are restricted to knowledge-factors and by default include PIN, pattern, and password. Primary authentication provides access to all functions on the phone.
- *Secondary Authentication* modalities are biometrics – which offer easier, but potentially less secure (than Primary Authentication), access into a user’s device. Secondary modalities are themselves split into sub-tiers based on how secure they are, as measured along two axes:

<sup>9</sup>These numbers are from internal analysis that has not yet been formally published.

- (1) *Spoofability* as measured by the Spoof Acceptance Rate (SAR) of the modality [100]. Accounting for an explicit attacker in the threat model on the level of [T1-T2] helps reduce the potential for insecure unlock methods [97].
- (2) *Security of the biometric pipeline*, where a biometric pipeline is considered secure if neither platform or kernel compromise confer the ability to read raw biometric data or inject data into the biometric pipeline to influence an authentication decision.

These axes are used to categorize secondary authentication modalities into three sub-tiers, where each sub-tier has constraints applied in proportion to the level of security they provide. Secondary modalities are also prevented from performing some actions — for example, they do not decrypt file-based or full-disk encrypted user data partitions (such as on first boot) and are required to fallback to primary authentication once every 72 hours. If a weak biometric does not meet either of the criteria (spoofability and pipelines security), then they cannot unlock Keymaster auth-bound keys and have a shorter fallback period. Android 10 introduced support for implicit biometric modalities in BiometricPrompt for modalities that do not require explicit interaction, for example face recognition.

- *Tertiary Authentication* modalities are alternate modalities such as unlocking when paired with a trusted Bluetooth device, or unlocking at trusted locations. Tertiary modalities are subject to all the constraints of secondary modalities. Additionally, like the weaker secondary modalities, tertiary modalities are also restricted from granting access to Keymaster auth-bound keys (such as those required for payments) and also require a fallback to primary authentication after any 4-hour idle period.

The Android lockscreen is currently implemented by Android system components above the kernel, specifically Keyguard and the respective unlock methods (some of which may be OEM specific). User knowledge factors of secure lockscreens are passed on to Gatekeeper/Weaver (explained below) both for matching them with stored templates and deriving keys for storage encryption. One implication is that a kernel compromise could lead to bypassing the lockscreen — but only after the user has logged in for the first time after reboot.

As of April 2019, lockscreen authentication on Android 7+ can now be used for FIDO2/WebAuthn [13, 126] authentication to web pages, additionally making Android phones second authentication factors for desktop browsers through implementing the Client to Authenticator Protocol (CTAP) [111]. While this support is currently implemented in Google Play Services [68], the intention is to include support directly in AOSP in the future when standards have sufficiently settled down to become stable for the release cycle of multiple Android releases.

#### 4.2.1 Identity Credential

While the lockscreen is the primary means for user-to-device (U2D) authentication and various methods support device-to-device (D2D) authentication (both between clients and client/server authentication such as through WebAuthn), identifying the device owner to other parties has not been in focus so far. Through the release of a JetPack library<sup>10</sup>, apps can make use of a new “Identity Credential” subsystem to support privacy-first identification [77] (and, to a certain degree, authentication). One example are upcoming third-party apps to support mobile driving licenses (mDL) according to the ISO 18013-5 standard [14]. The first version of this subsystem targets in-person presentation of credentials, and identification to automated verification systems is subject to future work.

<sup>10</sup>Note to reviewers: The library is expected to be released in a production version before the final version of this paper is submitted. A preview draft is available as open source at <https://android.googlesource.com/platform/hardware/interfaces/+refs/heads/master/identity/aid/android/hardware/identity/IIdentityCredentialStore.aidl>.

Android 11 will include the Identity Credential subsystem in the form of a new HAL, a new system daemon, and API support in AOSP [22]. If the hardware supports direct connections between the NFC controller and tamper-resistant dedicated hardware, credentials will be able to be marked for “Direct Access”<sup>11</sup> to be available even when the main application processor is no longer powered (e.g. in a low-battery case).

### 4.3 Isolation and Containment

One of the most important parts of enforcing the security model is to enforce it at runtime against potentially malicious code already running on the device. The Linux kernel provides much of the foundation and structure upon which Android’s security model is based. Process isolation provides the fundamental security primitive for sandboxing. With very few exceptions, the process boundary is where security decisions are made and enforced – Android intentionally does not rely on in-process compartmentalization such as the Java security model. The security boundary of a process is comprised of the process boundary and its entry points and implements rule ②: an app does not have to be vetted or pre-processed to run within the sandbox. Strengthening this boundary can be achieved by a number of means such as:

- Access control: adding permission checks, increasing the granularity of permission checks, or switching to safer defaults (e.g. default deny) to address the full range of threats [T7-T15].
- Attack surface reduction: reducing the number of entry points, particularly [T7-T9], i.e. the principle of least privilege.
- Containment: isolating and de-privileging components, particularly ones that handle untrusted content as in [T10] and [T14].
- Architectural decomposition: breaking privileged processes into less privileged components and applying attack surface reduction for [T8-T14].
- Separation of concerns: avoiding duplication of functionality.

In this section we describe the various sandboxing and access control mechanisms used on Android on different layers and how they improve the overall security posture.

#### 4.3.1 Permissions

Android uses three distinct permission mechanisms to perform access control:

- **Discretionary Access Control (DAC):** Processes may grant or deny access to resources that they own by modifying permissions on the object (e.g. granting world read access) or by passing a handle to the object over IPC. On Android this is implemented using UNIX-style permissions that are enforced by the kernel and URI permission grants. Processes running as the root user often have broad authority to override UNIX permissions (subject to MAC permissions – see below). URI permission grants provide the core mechanism for app-to-app interaction allowing an app to grant selective access to pieces of data it controls.
- **Mandatory Access Control (MAC):** The system has a security policy that dictates what actions are allowed. Only actions explicitly granted by policy are allowed. On Android this is implemented using SELinux [110] and primarily enforced by the kernel. Android makes extensive use of SELinux to protect system components and assert security model requirements during compatibility testing.
- **Android permissions** gate access to sensitive data and services. Enforcement is primarily done in userspace by the data/service provider (with notable exceptions such as INTERNET).

<sup>11</sup>See the HAL definition at <https://android-review.googlesource.com/c/platform/hardware/interfaces/+/1151485/30/identity/1.0/IIdentityCredentialStore.hal>.

638 Permissions are defined statically in an app's `AndroidManifest.xml` [23], though not all  
639 permissions requested may be granted.

640 Android 6.0 brought a major change by no longer guaranteeing that all requested permissions  
641 are granted when an application is installed. This was a direct result of the realization that  
642 users were not sufficiently equipped to make such a decision at installation time (cf. [64, 65,  
643 104, 128]).

644 The second major change in Android permissions was introduced with Android 10 in the  
645 form of non-binary, context dependent permissions: in addition to *Allow* and *Deny*, some  
646 permissions (particularly location, and potentially starting with Android 11 others like camera  
647 and microphone) can now be set to *Allow only while using the app*. This third state only grants  
648 the permission when an app is in the foreground, i.e. when it either has a visible activity or  
649 runs a foreground service with permanent notification [50].

650 At a high level Android permissions fall into one of five classes in increasing order of severity:

- 651 (1) Audit-only Permissions: These are install time permissions with the 'normal' protection  
652 level.
- 653 (2) Runtime Permissions: These are permissions that the user must approve as part of a runtime  
654 prompt dialog. These permissions are guarding commonly used sensitive user data, and  
655 depending on how critical they are for the current functioning of an application, different  
656 strategies for requesting them are recommended [38].
- 657 (3) Special Access Permissions: For permissions that expose more or are higher risk than  
658 runtime permissions there exists a special class of permissions with much higher granting  
659 friction that the application cannot show a runtime prompt for. In order for a user to allow  
660 an application to use a special access permission the user must go to settings and manually  
661 grant the permission to the application.
- 662 (4) Privileged Permissions: These permissions are for pre-installed privileged applications only  
663 and allow privileged actions such as carrier billing.
- 664 (5) Signature Permissions: These permissions are only available to components signed with  
665 the same key as the component which declares the permission e.g. the platform signing  
666 key. They are intended to guard internal or highly privileged actions, e.g. configuring the  
667 network interfaces.

668 Permission availability is defined by their `protectionLevel` attribute [24] with two parts  
669 (the level itself and a number of optional flags) which may broaden which applications may  
670 be granted a permission as well as how they may request it. The protection levels are:

- 671 – normal: Normal permissions are those that do not pose much privacy or security risk  
672 and are granted automatically at install time. These permissions are primarily used for  
673 auditability of app behavior.
- 674 – dangerous: Permissions with this `protectionLevel` are runtime permissions, and apps  
675 must both declare them in their manifest as well as request users grant them during use.  
676 These permissions, which are fairly fine-grained to support auditing and enforcement, are  
677 grouped into logical permissions using the `permissionGroup` attribute. When requesting  
678 runtime permissions, the group appears as a single permission to avoid over-prompting.
- 679 – signature: Applications can only be granted such permission if they are signed with the  
680 same key as the application that defines the permission, which is the platform signing key  
681 for platform permission. These permissions are granted at install time if the application is  
682 allowed to use them.

683 Additionally, there are a number of protection flags that modify the grantability of per-  
684 missions. For example, the `BLUETOOTH_PRIVILEGED` permission has a `protectionLevel`

685

686



of signature|privileged, with the privileged flag allowing privileged applications to be granted the permission (even if they are not signed with the platform key).

Each of the three permission mechanisms roughly aligns with how one of the three parties of the multi-party grant consent (rule ①). The platform utilizes MAC, apps use DAC, and users consent by granting Android permissions. Note that permissions are not intended to be a mechanism for obtaining consent in the legal sense but a technical measure to enforce auditability and control. It is up to the app developer processing personal user data to meet applicable legal requirements.

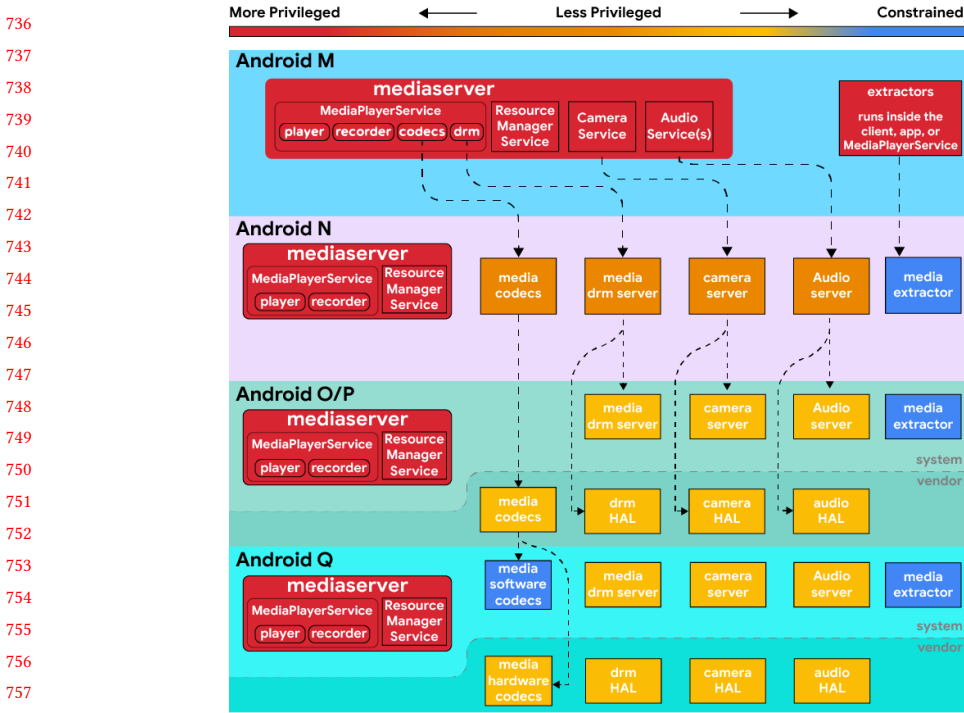
#### 4.3.2 Application sandbox

Android's original DAC application sandbox separated apps from each other and the system by providing each application with a unique UNIX user ID (UID) and a directory owned by the app. This approach was quite different from the traditional desktop approach of running applications using the UID of the physical user. The unique per-app UID simplifies permission checking and eliminates per-process ID (PID) checks, which are often prone to race conditions. Permissions granted to an app are stored in a centralized location (`/data/system/packages.xml`), to be queried by other services. For example, when an app requests location from the location service, the location service queries the permissions service to see if the requesting UID has been granted the location permission.

The UID sandbox had a number of shortcomings. Processes running as root were essentially unsandboxed and possessed extensive power to manipulate the system, apps, and private app data. Likewise, processes running as the system UID were exempt from Android permission checks and permitted to perform many privileged operations. Use of DAC meant that apps and system processes could override safe defaults and were more susceptible to dangerous behavior, such as symlink following or leaking files/data across security boundaries via IPC or `fork/exec`. Additionally, DAC mechanisms can only apply to files on file systems that support access controls lists (respectively simple UNIX access bits). The main implication is that the FAT family of file systems, which is still commonly used on extended storage such as (micro-) SD cards or media connected through USB, does not directly support applying DAC. On Android, each app has a well-known directory on external storage devices, where the package name of the app is included into the path (e.g. `/sdcard/Android/data/com.example`). Since the OS already maintains a mapping from package name to UID, it can assign UID ownership to all files in these well-known directories, effectively creating a DAC on a filesystem that doesn't natively support it. From Android 4.4 to Android 7.x, this mapping was implemented through FUSE, while Android 8.0 and later implement an in-kernel `sdcardfs` for better performance. Both are equivalent in maintaining the mapping of app UIDs to implement effective DAC. Android 10 introduced *scoped storage*, which limits app access to its own external directory path as well as media files that itself created in the shared media store.

Despite its deficiencies, the UID sandbox laid the groundwork and is still the primary enforcement mechanism that separates apps from each other. It has proven to be a solid foundation upon which to add additional sandbox restrictions. These shortcomings have been mitigated in a number of ways over subsequent releases, partially through the addition of MAC policies but also including many other mechanisms such as runtime permissions and attack surface reduction (cf. Table 1).

Rooting, as defined above, has the main aim of enabling certain apps and their processes to break out of this application sandbox in the sense of granting "root" user privileges [76], which override the DAC rules (but not automatically MAC policies, which led to extended rooting schemes with processes intentionally exempt from MAC restrictions). Malware may try to apply these rooting approaches through temporary or permanent exploits and therefore bypass the application sandbox.



759 Fig. 1. Changes to mediaserver and codec sandboxing from Android 6 to Android 10

760

761

762

### 763 4.3.3 Sandboxing system processes

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

### 785 4.3.4 Sandboxing the kernel

786 Security hardening efforts in Android userspace have increasingly made the kernel a more attractive target for privilege escalation attacks [123]. Hardware drivers provided by System on a Chip (SoC) vendors account for the vast majority of kernel vulnerabilities on Android [125]. Reducing app/system access to these drivers was described above, but sandboxing code inside the kernel itself also improved significantly over the various releases (cf. Table 3).

### 787 4.3.5 Sandboxing below the kernel

788 In addition to the kernel, the trusted computing base (TCB) on Android devices starts with the boot loader (which is typically split into multiple stages) and implicitly includes other components below the kernel, such as the trusted execution environment (TEE), hardware drivers, and userspace components `init`, `ueventd`, and `voId` [40]. It is clear that the sum of all these creates sufficient complexity that, given current state of the art, we have to assume bugs in some of them. For highly

sensitive use cases, even the mitigations against kernel and system process bugs described above may not provide sufficient assurance against potential vulnerabilities.

Therefore, we explicitly consider the possibility of a kernel compromise (e.g. through directly attacking some kernel interfaces based on physical access in [T2]-[T4] or chaining together multiple bugs from user space code to reach kernel surfaces in [T8]), misconfiguration (e.g. with incorrect or overly permissive SELinux policies [51]), or bypass (e.g. by modifying the boot chain to boot a different kernel with deactivated security policies) as part of the threat model for some select scenarios. To be clear, with a compromised kernel, Android no longer meets the compatibility requirements and many of the security and privacy assurances for users and apps no longer hold. However, we can still defend against some threats even under this assumption:

- **Keymaster** implements the Android key store in TEE to guard cryptographic key storage and use in the case of a run-time kernel compromise [26]. That is, even with a fully compromised kernel, an attacker cannot read key material stored in Keymaster<sup>12</sup>. Apps can explicitly request keys to be stored in Keymaster, i.e. to be hardware-bound, to be only accessible after user authentication (which is tied to Gatekeeper/Weaver), and/or request attestation certificates to verify these key properties [27], allowing verification of compatibility in terms of rule ③.
  - **Strongbox**, specified starting with Android 9.0, implements the Android keystore in separate tamper resistant hardware (TRH) for even better isolation. This mitigates [T1] and [T2] against strong adversaries, e.g. against cold boot memory attacks [72] or hardware bugs such as Spectre/Meltdown [82, 91], Rowhammer [48, 121], or Clkscrew [114] that allow privilege escalation even from kernel to TEE. From a hardware perspective, the main application processor (AP) will always have a significantly larger attack surface than dedicated secure hardware. Adding a separate TRH affords another sandboxing layer of defense in depth.
- Note that only storing and using keys in TEE or TRH does not completely solve the problem of making them unusable under the assumption of a kernel compromise: if an attacker gains access to the low-level interfaces for communicating directly with Keymaster or Strongbox, they can use it as an oracle for cryptographic operations that require the private key. This is the reason why keys can be authentication bound and/or require user presence verification, e.g. by pushing a hardware button that is detectable by the TRH to assure that keys are not used in the background without user consent.
- **Gatekeeper** implements verification of user lock screen factors (PIN/password/pattern) in TEE and, upon successful authentication, communicates this to Keymaster for releasing access to authentication bound keys [28]. **Weaver** implements the same functionality in TRH and communicates with Strongbox. Specified for Android 9.0 and initially implemented on the Google Pixel 2 and Pixel 3 phones, we also add a property called ‘Insider Attack Resistance’ (IAR): without knowledge of the user’s lock screen factor, an upgrade to the Weaver/Strongbox code running in TRH will wipe the secrets used for on-device encryption [96, 129]. That is, even with access to internal code signing keys, existing data cannot be exfiltrated without the user’s cooperation.
  - **Protected Confirmation**, also introduced with Android 9.0 [29], partially mitigates [T11] and [T13]. In its current scope, apps can tie usage of a key stored in Keymaster or Strongbox to the user confirming (by pushing a physical button) that they have seen a message displayed on the screen. Upon confirmation, the app receives a hash of the displayed message, which can

<sup>12</sup>Note: This assumes that hardware itself is still trustworthy. Side-channel attacks such as [86] are currently out of scope of this (software) platform security model, but influence some design decisions on the system level, e.g. to favor dedicated TRH over on-chip security partitioning.

834 be used to remotely verify that a user has confirmed the message. By controlling the screen  
835 output through TEE when protected confirmation is requested by an app, even a full kernel  
836 compromise (without user cooperation) cannot lead to creating these signed confirmations.  
837

#### 838 4.4 Encryption of data at rest

839 A second element of enforcing the security model, particularly rules ① and ③, is required when  
840 the main system kernel is not running or is bypassed (e.g. by reading directly from non-volatile  
841 storage).

842 Full Disk Encryption (FDE) uses a credential protected key to encrypt the entire user data  
843 partition. FDE was introduced in Android 5.0, and while effective against [T1], it had a number  
844 of shortcomings. Core device functionality (such as emergency dialer, accessibility services, and  
845 alarms) were inaccessible until password entry. Multi-user support introduced in Android 6.0 still  
846 required the password of the primary user before disk access.

847 These shortcomings were mitigated by File Based Encryption (FBE) introduced in Android 7.0.  
848 On devices with TEE or TRH, all keys are derived within these secure environments, entangling the  
849 user knowledge factor with hardware-bound random numbers that are inaccessible to the Android  
850 kernel and components above.<sup>13</sup> FBE allows individual files to be tied to the credentials of different  
851 users, cryptographically protecting per-user data on shared devices [T3]. Devices with FBE also  
852 support a feature called *Direct Boot* which enables access to emergency dialer, accessibility services,  
853 alarms, and receiving calls all before the user inputs their credentials.

854 Android 10 introduced support for Adiantum [53], a new wide-block cipher mode based on  
855 AES, ChaCha, and Poly1305 to enable full device encryption without hardware AES acceleration  
856 support. While this does not change encryption of data at rest for devices with existing AES support,  
857 lower-end processors can now also encrypt all data without prohibitive performance impact. The  
858 significant implication is that all devices shipping originally with Android 10 are required to encrypt  
859 all data by default without any further exemptions, homogenizing the Android ecosystem in that  
860 aspect.

861 Note that encryption of data at rest helps significantly with enforcing rule ④, as effectively  
862 wiping user data only requires to delete the master key material, which is much quicker and not  
863 subject to the complexities of e.g. flash translation layer interactions.  
864

#### 865 4.5 Encryption of data in transit

866 Android assumes that all networks are hostile and could be injecting attacks or spying on traffic.  
867 In order to ensure that network level adversaries do not bypass app data protections, Android  
868 takes the stance that *all* network traffic should be end-to-end encrypted. Link level encryption is  
869 insufficient. This primarily protects against [T5] and [T6].

870 In addition to ensuring that connections use encryption, Android focuses heavily on ensuring that  
871 the encryption is used correctly. While TLS options are secure by default, we have seen that it is easy  
872 for developers to incorrectly customize TLS in a way that leaves their traffic vulnerable to MITM  
873 attacks [60, 61, 70]. Table 4 lists recent improvements in terms of making network connections safe  
874 by default.  
875

#### 876 4.6 Exploit mitigation

877 A robust security system should assume that software vulnerabilities exist and actively defend  
878 against them. Historically, about 85% of security vulnerabilities on Android result from unsafe  
879

880 <sup>13</sup>A detailed specification and analysis of key entanglement is subject to related work and currently in progress. A reference  
881 to this detail will be added to a later version of this paper.  
882

883 memory access (cf. [83, slide 54]). While this section primarily describes mitigations against memory  
884 unsafety, we note that the best defense is the memory safety offered by languages such as Java or  
885 Kotlin. Much of the Android framework is written in Java, effectively defending large swathes of  
886 the OS from entire categories of security bugs.

887 Android mandates the use of a number of mitigations including ASLR [44, 109], RWX memory  
888 restrictions (e.g.  $W \oplus X$ , cf. [108]), and buffer overflow protections (such as stack-protector for  
889 the stack and allocator protections for the heap). Similar protections are mandated for Android  
890 kernels [118].

891 In addition to the mitigations listed above, Android is selectively enabling new mitigations,  
892 focusing first on code areas which are remotely reachable (e.g. the media frameworks [41]) or have  
893 a history of high severity security vulnerabilities (e.g. the kernel). Android has pioneered the use  
894 of LLVM undefined behavior sanitizer (UBSAN) in production devices to protect against integer  
895 overflow vulnerabilities in the media frameworks and other security sensitive components. Android  
896 is also rolling out Control Flow Integrity (CFI) [119] in the kernel and security sensitive userspace  
897 components including media, Bluetooth, WiFi, NFC, and parsers [93] in a fine-grained variant as  
898 implemented by current LLVM [117] that improves upon previous, coarse-grained approaches  
899 that have been shown to be ineffective [55]. Starting with Android 10, the common Android  
900 kernel as well as parts of the Bluetooth stack can additionally be protected against backwards-  
901 edge exploitation through the use of Shadow Call Stack (SCS), again as implemented by current  
902 LLVM [112] as the best trade-off between performance overhead and effectiveness [47].

903 These mitigation methods work in tandem with isolation and containment mechanisms to form  
904 many layers of defense; even if one layer fails, other mechanisms aim to prevent a successful  
905 exploitation chain. Mitigation mechanisms also help to uphold rules ② and ③ without placing  
906 additional assumptions on which languages apps are written in.

#### 907 4.7 System integrity

909 Finally, system (sometimes also referred to as device) integrity is an important defense against  
910 attackers gaining a persistent foothold. AOSP has supported *Verified Boot* using the Linux kernel  
911 `dm-verity` support since Android KitKat, providing strong integrity enforcement for the Trusted  
912 Computing Base (TCB) and system components to implement rule ④. Verified Boot [31] has been  
913 mandated since Android Nougat (with an exemption granted to devices which cannot perform AES  
914 crypto above 50MiB/sec. up to Android 8, but no exemptions starting with Android 9.0) and makes  
915 modifications to the boot chain detectable by verifying the boot, TEE, and additional vendor/OEM  
916 partitions, as well as performing on-access verification of blocks on the system partition [32]. That  
917 is, attackers cannot permanently modify the TCB even after all previous layers of defense have  
918 failed, leading to a successful kernel compromise. Note that this assumes the primary boot loader  
919 as root of trust to still be intact. As this is typically implemented in a ROM mask in sufficiently  
920 simple code, critical bugs at that stage are less likely.

921 Additionally, rollback protection with hardware support (counters stored in tamper-proof persist-  
922 ent storage, e.g. a separate TRH as used for Strongbox or enforced through RPMB as implemented  
923 in a combination of TEE and eMMC controller [19]) prevents attacks from flashing a properly  
924 signed but outdated system image that has known vulnerabilities and could be exploited. Finally,  
925 the Verified Boot state is included in key attestation certificates (provided by Keymaster/Strongbox)  
926 in the `deviceLocked` and `verifiedBootState` fields, which can be verified by apps as well as  
927 passed onto backend services to remotely verify boot integrity [33].

928 Starting with Android 10 on some devices supporting the latest Android Verified Boot (AVB,  
929 the recommended default implementation for verifying the integrity of read-only partitions [34])  
930 version 2, the `VMBeta` struct `digest` (a top-level hash over all parts) is included in these key  
931



932 attestation certificates to support firmware transparency by verifying that digest match released  
 933 firmware images [34, 96]. In combination with server side validation, this can be used as a form of  
 934 remote system integrity attestation akin to PCR verification with trusted platform modules (TPMs).  
 935 Integrity of firmware for other CPUs (including, but not limited to, the various radio chipsets,  
 936 the GPU, touch screen controllers, etc.) is out of scope of AVB at the time of this writing, and is  
 937 typically handled by OEM-specific boot loaders.

#### 938 4.7.1 Verification key hierarchy and updating

939 While the details for early boot stages are highly dependent on the respective chipset hardware  
 940 and low-level boot loaders, Android devices generally use at least the following keys for verifying  
 941 system integrity:

- 942 (1) The first (and potentially multiple intermediate) boot loader(s) is/are signed by a key  $K_A$  held  
 943 by the hardware manufacturer and verified through a public key embedded in the chipset  
 944 ROM mask. This key cannot be changed.
- 945 (2) The (final) bootloader responsible for loading the Android Linux kernel is verified through  
 946 a key  $K_B$  embedded in a previous bootloader. Updating this signing key is chipset specific,  
 947 but may be possible in the field by updating a previous, intermediate bootloader block.  
 948 Android 10 strongly recommends that this bootloader use the reference implementation  
 949 of Android Verified Boot [34] and VBMeta structs for verifying all read-only (e.g. system,  
 950 vendor, etc.) partitions.
- 951 (3) A VBMeta signing key  $K_C$  is either directly embedded in the final bootloader or retrieved  
 952 from a separate TRH to verify flash partitions before loading the kernel. AVB implementations  
 953 may also allow a user-defined VBMeta signing key  $K'_C$  to be set (typically in a TEE or TRH) –  
 954 in this case, the Verified Boot state will be set to YELLOW to indicate that non-manufacturer  
 955 keys were used to sign the partitions, but that verification with the user-defined keys has still  
 956 been performed correctly (see Figure 2).

957 Updating this key  $K_C$  used to sign any partitions protected through AVB is supported through  
 958 the use of chained partitions in the VBMeta struct (resulting in partition-specific signing keys  
 959  $K_D^i$  for partition  $i$  that are in turn signed by  $K_C/K'_C$ ), by updating the key used to sign the  
 960 VBMeta struct itself (through flashing a new version of the final bootloader in an over-the-air  
 961 update), or – in the case of user-defined keys – using direct physical access<sup>15</sup>.

- 962 (4) The digest(s) embedded in VBMeta struct(s) are used by the Android Linux kernel to verify  
 963 blocks within persistent, read-only partitions on-access using `dm-verity` (or for small parti-  
 964 tions, direct verification before loading them atomically into memory). Inside the system  
 965 partition, multiple public signing keys are used for different purposes, e.g. the platform  
 966 signing key mentioned in section 4.3.1 or keys used to verify the download of over-the-air  
 967 (OTA) update packages before applying them. Updating those keys is trivial through simply  
 968 flashing a new system partition.
- 969 (5) All APKs are individually signed by the respective developer key  $K_E^j$  for APK  $j$  (some may  
 970 be signed by the platform signing key to be granted signature permissions for those com-  
 971 ponents), which in turn are stored on the system or data partition. Integrity of updateable  
 972 (system or user installed) apps is enforced via APK signing [35] and is checked by Android's  
 973 PackageManager during installation and update. Every app is signed and an update can  
 974 only be installed if the new APK is signed with the same identity or by an identity that was  
 975 delegated by the original signer.

976  
 977  
 978 <sup>15</sup>E.g. Pixel devices support this through `fastboot flash avb_custom_key` as documented online at <https://source.android.com/security/verifiedboot/device-state>.



981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
1026  
1027  
1028  
1029




Fig. 2. Verified Boot flow and different states: (YELLOW): warning screen for LOCKED devices with custom root of trust set; (ORANGE): warning screen for UNLOCKED devices; (RED): warning screen for dm-verity corruption or no valid OS found [36].

For run-time updateable apps, the APK Signature Scheme version 3 was introduced with Android 9.0 to support rotation of these individual signing keys [35].

### 4.8 Patching

Orthogonal to all the previous defense mechanisms, vulnerable code should be fixed to close discovered holes in any of the layers. Regular patching can be seen as another layer of defense. However, shipping updated code to the huge and diverse Android ecosystem is a challenge [116] (which is one of the reasons for applying the defense in depth strategy).

Starting in August 2015, Android has publicly released a monthly security bulletin and patches for security vulnerabilities reported to Google. To address ecosystem diversity, project Treble launched with Android 8.0, with a goal of reducing the time/cost of updating Android devices [94, 98].

In 2018, the Android Enterprise Recommended program as well as general agreements with OEMs added the requirement of 90-day guaranteed security updates [37].

Starting with Android 10, some core system components can be updated through Google Play Store as standard APK files or – if required early in the boot process or involving native system libraries/services – as an APEX loopback filesystems in turn protected through dm-verity [71].

## 5 SPECIAL CASES

There are some special cases that require intentional deviations from the abstract security model to balance specific needs of various parties. This section describes some of these but is not intended to be a comprehensive list. One goal of defining the Android security model publicly is to enable researchers to discover potential additional gaps by comparing the implementation in AOSP with the model we describe, and to engage in conversation on those special cases.

- **Listing packages:** The ability for one app to discover what other apps are installed on the device can be considered a potential information leak and violation of user consent (rule ①). However, app discovery is necessary for some direct app-to-app interaction which

is derived from the open ecosystem principle (rule ②). As querying the list of all installed apps is potentially privacy sensitive and has been abused by malware, Android 11 supports more specific app-to-app interaction using platform components and limits general package visibility for apps targeting this API version<sup>16</sup>. While this special case is still supported at the time of this writing, it will require the new `QUERY_ALL_PACKAGES` and may be limited further in the future.

- **VPN apps may monitor/block network traffic for other apps:** This is generally a deviation from the application sandbox model since one app may see and impact traffic from another app (*developer* consent). VPN apps are granted an exemption because of the value they offer users, such as improved privacy and data usage controls, and because *user* consent is clear. For applications which use end-to-end encryption, clear-text traffic is not available to the VPN application, partially restoring the confidentiality of the application sandbox.
- **Backup:** Data from the private app directory is backed up by default. Android 9 added support for end-to-end encryption of backups to the Google cloud by entangling backup session keys with the user lockscreen knowledge factor (LSKF) [79]. Apps may opt out by setting fields in their manifest.
- **Enterprise:** Android allows so-called Device Owner (DO) or Profile Owner (PO) policies to be enforced by a Device Policy Controller (DPC) app. A DO is installed on the primary/main user account, while a PO is installed on a secondary user that acts as a work profile. Work profiles allow separation of personal from enterprise data on a single device and are based on Android multi-user support. This separation is enforced by the same isolation and containment methods that protect apps from each other but implement a significantly stricter divide between the profiles [7].  
A DPC introduces a fourth party to the consent model: only if the policy allows an action (e.g. within the work profile controlled by a PO) in addition to consent by all other parties can it be executed. The distinction of personal and work profile is enhanced by the recent support of different user knowledge factors (handled by the lockscreen as explained above in subsection 4.2), which lead to different encryption keys for FBE. Note that on devices with a work profile managed by PO but no full-device control (i.e. no DO), privacy guarantees for the personal profile still need to hold under this security model.
- **Factory Reset Protection (FRP):** is an exception to not storing any persistent data across factory reset (rule ④), but is a deliberate deviation from this part of the model to mitigate the threat of theft and factory reset ([T1][T2]).

## 6 RELATED WORK

Classical operating system security models are primarily concerned with defining access control (read/write/execute or more finely granular) by subjects (but most often single users, groups, or roles) to objects (typically files and other resources controlled by the OS, in combination with permissions sometimes also called protection domains [113]). The most common data structures for efficiently implementing these relations (which, conceptually, are sparse matrices) are Access Control Lists (ACLs) [106] and capability lists (e.g. [127]). One of the first well-known and well-defined models was the Bell-LaPadula multi-level security model [42], which defined properties for assigning permissions and can be considered the abstract basis for Mandatory Access Control and Type Enforcement schemes like SELinux. Consequently, the Android platform security model implicitly builds upon these general models and their principle of least privilege.

<sup>16</sup>Preview release changes are described at <https://developer.android.com/preview/privacy/package-visibility> and are expected to become final (potentially in a slightly changed form) before the final version of this paper is published.

1079 One fundamental difference is that, while classical models assume processes started by a user to  
1080 be a proxy for their actions and therefore execute directly with user privileges, more contemporary  
1081 models explicitly acknowledge the threat of malware started by a user and therefore aim to  
1082 compartmentalize their actions. Many mobile OS (including Symbian as an earlier example) assign  
1083 permissions to processes (i.e. applications) instead of users, and Android uses a comparable approach.  
1084 A more detailed comparison to other mobile OS is out of scope in this paper, and we refer to other  
1085 surveys [58, 85, 99] as well as previous analysis of Android security mechanisms and how malware  
1086 exploited weaknesses [15, 59, 63, 88–90, 130].  
1087

## 1088 7 CONCLUSION

1089 In this paper, we described the Android platform security model and the complex threat model and  
1090 ecosystem it needs to operate in. One of the abstract rules is a multi-party consent model that is  
1091 different to most standard OS security models in the sense that it implicitly considers applications  
1092 to have equal veto rights over actions in the same sense that the platform implementation and,  
1093 obviously, users have. While this may seem restricting from a user point of view, it effectively  
1094 limits the potential abuse a malicious app can do on data controlled by other apps; by avoiding  
1095 an all-powerful user account with unfiltered access to all data (as is the default with most current  
1096 desktop/server OS), whole classes of threats such as file encrypting ransomware or direct data  
1097 exfiltration become impractical.

1098 AOSP implements the Android platform security model as well as the general security principles  
1099 of ‘defense in depth’ and ‘safe by default’. Different security mechanisms combine as multiple  
1100 layers of defense, and an important aspect is that even if security relevant bugs exist, they should  
1101 not necessarily lead to exploits reachable from standard user space code. While the current model  
1102 and its implementation already cover most of the threat model that is currently in scope of Android  
1103 security and privacy considerations, there are some deliberate special cases to the conceptually  
1104 simple security model, and there is room for future work:  
1105

- 1106 • Keystore already supports API flags/methods to request hardware- or authentication-bound  
1107 keys. However, apps need to use these methods explicitly to benefit from improvements like  
1108 Strongbox. Making encryption of app files or directories more transparent by supporting  
1109 declarative use similar to network security config for TLS connections would make it easier  
1110 for app developers to securely use these features.
- 1111 • It is common for malware to dynamically load its second stage depending on the respective  
1112 device it is being installed on, to both try to exploit specific detected vulnerabilities and  
1113 hide its payload from scanning in the app store. One potential mitigation is to require all  
1114 executable code to: a) be signed by a key that is trusted by the respective Android instance  
1115 (e.g. with public keys that are pre-shipped in the firmware and/or can be added by end-users)  
1116 or b) have a special permission to dynamically load/create code during runtime that is not  
1117 contained in the application bundle itself (the APK file). This could give better control over  
1118 code integrity, but would still not limit languages or platforms used to create these apps. It  
1119 is recognized that this mitigation is limited to executable code. Interpreted code or server  
1120 based configuration would bypass this mitigation.
- 1121 • Advanced attackers may gain access to OEM or vendor code signing keys. Even under such  
1122 circumstance, it is beneficial to still retain some security and privacy assurances to users. One  
1123 recent example is the specification and implementation of ‘Insider Attack Resistance’ (IAR)  
1124 for updateable code in TRH [129], and extending similar defenses to higher-level software is  
1125 desirable [96]. Potential approaches could be reproducible firmware builds or logs of released  
1126 firmware hashes comparable to e.g. Certificate Transparency [87].  
1127

- Hardware level attacks are becoming more popular, and therefore additional (software and hardware) defense against e.g. RAM related attacks would add another layer of defense, although, most probably with a trade-off in performance overhead.

However, all such future work needs to be done considering its impact on the wider ecosystem and should be kept in line with fundamental Android security principles.

## ACKNOWLEDGMENTS

We would like to thank Dianne Hackborn for her influential work over a large part of the Android platform security history and insightful remarks on earlier drafts of this paper. Additionally, we thank Joel Galenson, Ivan Lozano, Paul Crowley, Shawn Willden, Jeff Sharkey, Haining Chen, and Xiaowen Xin for input on various parts, and particularly Vishwath Mohan for direct contributions to the Authentication section. We also thank the enormous number of security researchers (<https://source.android.com/security/overview/acknowledgements>) who have improved Android over the years and anonymous reviewers who have contributed highly helpful feedback to earlier drafts of this paper.

## REFERENCES

- [1] [n. d.]. <https://www.stonetemple.com/mobile-vs-desktop-usage-study/>
- [2] [n. d.]. <http://gs.statcounter.com/platform-market-share/desktop-mobile-tablet>
- [3] [n. d.]. [https://kernsec.org/wiki/index.php/Exploit\\_Methods/Userspace\\_execution](https://kernsec.org/wiki/index.php/Exploit_Methods/Userspace_execution)
- [4] 2015. Stagefright Vulnerability Report. <https://www.kb.cert.org/vuls/id/924951>
- [5] 2017. BlueBorne. [https://go.armis.com/hubfs/BlueBorne%20-%20Android%20Exploit%20\(20171130\).pdf?t=1529364695784](https://go.armis.com/hubfs/BlueBorne%20-%20Android%20Exploit%20(20171130).pdf?t=1529364695784)
- [6] 2017. CVE-2017-13177. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-13177>
- [7] 2018. Android Enterprise Security White Paper. [https://source.android.com/security/reports/Google\\_Android\\_Enterprise\\_Security\\_Whitepaper\\_2018.pdf](https://source.android.com/security/reports/Google_Android_Enterprise_Security_Whitepaper_2018.pdf)
- [8] 2018. Android Security 2017 Year In Review. [https://source.android.com/security/reports/Google\\_Android\\_Security\\_2017\\_Report\\_Final.pdf](https://source.android.com/security/reports/Google_Android_Security_2017_Report_Final.pdf)
- [9] 2018. CVE-2017-17558: Remote code execution in media frameworks. <https://source.android.com/security/bulletin/2018-06-01#kernel-components>
- [10] 2018. CVE-2018-9341: Remote code execution in media frameworks. <https://source.android.com/security/bulletin/2018-06-01#media-framework>
- [11] 2018. SVE-2018-11599: Theft of arbitrary files leading to emails and email accounts takeover. <https://security.samsungmobile.com/securityUpdate.smsb>
- [12] 2018. SVE-2018-11633: Buffer Overflow in Trustlet. <https://security.samsungmobile.com/securityUpdate.smsb>
- [13] 2019. Android Now FIDO2 Certified. <https://fidoalliance.org/android-now-fido2-certified-accelerating-global-migration-beyond-pas>
- [14] 2020. Personal identification – ISO-compliant driving licence – Part 5: Mobile driving licence (mDL) application. Draft International Standard: ISO/IEC DIS 18013-5.
- [15] Y. Acar, M. Backes, S. Bugiel, S. Fahl, P. McDaniel, and M. Smith. 2016. SoK: Lessons Learned from Android Security Research for Appified Software Platforms. In *2016 IEEE Symposium on Security and Privacy (SP)*. 433–451. <https://doi.org/10.1109/SP.2016.33>
- [16] Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (Dec. 1999), 40–46. <https://doi.org/10.1145/322796.322806>
- [17] Andrew Ahn. 2018. How we fought bad apps and malicious developers in 2017. <https://android-developers.googleblog.com/2018/01/how-we-fought-bad-apps-and-malicious.html>
- [18] Bonnie Brinton Anderson, Anthony Vance, C. Brock Kirwan, Jeffrey L. Jenkins, and David Eargle. 2016. From Warning to Wallpaper: Why the Brain Habituates to Security Warnings and What Can Be Done About It. *Journal of Management Information Systems* 33, 3 (2016), 713–743. <https://doi.org/10.1080/07421222.2016.1243947>
- [19] Anil Kumar Reddy, P. Paramasivam, and Prakash Babu Vemula. 2015. Mobile secure data protection using eMMC RPMB partition. In *2015 International Conference on Computing and Network Communications (CoCoNet)*. 946–950. <https://doi.org/10.1109/CoCoNet.2015.7411305>
- [20] AOSP. [n. d.]. <https://source.android.com/compatibility/cdd>
- [21] AOSP. [n. d.]. <https://developer.android.com/guide/components/intents-filters>
- [22] AOSP. [n. d.]. <https://developer.android.com/reference/android/security/identity/package-summary>

- 1177 [23] AOSP. [n. d.]. <https://developer.android.com/guide/topics/manifest/manifest-intro>
- 1178 [24] AOSP. [n. d.]. <https://developer.android.com/guide/topics/manifest/permission-element>
- 1179 [25] AOSP. [n. d.]. <https://developer.android.com/guide/components/activities/background-starts>
- 1180 [26] AOSP. [n. d.]. <https://source.android.com/security/keystore/>
- 1181 [27] AOSP. [n. d.]. <https://developer.android.com/reference/android/security/keystore/KeyGenParameterSpec>
- 1182 [28] AOSP. [n. d.]. <https://source.android.com/security/authentication/gatekeeper>
- 1183 [29] AOSP. [n. d.]. <https://developer.android.com/preview/features/security#android-protected-confirmation>
- 1184 [30] AOSP. [n. d.]. <https://developer.android.com/training/articles/security-config>
- 1185 [31] AOSP. [n. d.]. <https://source.android.com/security/verifiedboot/verified-boot>
- 1186 [32] AOSP. [n. d.]. <https://android.googlesource.com/platform/external/avb/+/-/pie-release/README.md>
- 1187 [33] AOSP. [n. d.]. <https://developer.android.com/training/articles/security-key-attestation>
- 1188 [34] AOSP. [n. d.]. <https://android.googlesource.com/platform/external/avb/+/-/master/README.md>
- 1189 [35] AOSP. [n. d.]. <https://source.android.com/security/apksigning/v3>
- 1190 [36] AOSP. [n. d.]. <https://source.android.com/security/verifiedboot/boot-flow>
- 1191 [37] AOSP. [n. d.]. <https://www.android.com/enterprise/recommended/requirements/>
- 1192 [38] AOSP. [n. d.]. Android platform permissions requesting guidance. [https://material.io/design/platform-guidance/  
android-permissions.html#request-types](https://material.io/design/platform-guidance/android-permissions.html#request-types)
- 1193 [39] AOSP. [n. d.]. [https://source.android.com/  
security/enhancements/enhancements10](https://source.android.com/security/enhancements/enhancements10)
- 1194 [40] AOSP. [n. d.]. Security Updates and Resources - Process Types. [https://source.android.com/security/overview/  
updates-resources#process\\_types](https://source.android.com/security/overview/<br/>updates-resources#process_types)
- 1195 [41] Dan Austin and Jeff Vander Stoep. 2016. Hardening the media stack. [https://android-developers.googleblog.com/  
2016/05/hardening-media-stack.html](https://android-developers.googleblog.com/<br/>2016/05/hardening-media-stack.html)
- 1196 [42] D. Bell and L. LaPadula. 1975. *Secure Computer System Unified Exposition and Multics Interpretation*. Technical Report MTR-2997. MITRE Corp., Bedford, MA.
- 1197 [43] James Bender. 2018. Google Play security metadata and offline app distribution. [https://android-developers.  
googleblog.com/2018/06/google-play-security-metadata-and.html](https://android-developers.<br/>googleblog.com/2018/06/google-play-security-metadata-and.html)
- 1198 [44] Sandeep Bhatkar, Daniel C. DuVarney, and R. Sekar. 2003. Address Obfuscation: An Efficient Approach to Combat a Board Range of Memory Error Exploits. In *Proc. USENIX Security Symposium - Volume 12*. USENIX Association, Berkeley, CA, USA, 8–8. <http://dl.acm.org/citation.cfm?id=1251353.1251361>
- 1200 [45] Chad Brubaker. 2014. Introducing nogotofail – a network traffic security testing tool. [https://security.googleblog.  
com/2014/11/introducing-nogotofail-a-network-traffic.html](https://security.googleblog.<br/>com/2014/11/introducing-nogotofail-a-network-traffic.html)
- 1201 [46] Chad Brubaker. 2018. Protecting users with TLS by default in Android P. [https://android-developers.googleblog.  
com/2018/04/protecting-users-with-tls-by-default-in.html](https://android-developers.googleblog.<br/>com/2018/04/protecting-users-with-tls-by-default-in.html)
- 1202 [47] N. Burrow, X. Zhang, and M. Payer. 2019. SoK: Shining Light on Shadow Stacks. In *2019 IEEE Symposium on Security  
and Privacy (SP)*. 985–999. <https://doi.org/10.1109/SP.2019.00076>
- 1203 [48] Pierre Carru. 2017. Attack TrustZone with Rowhammer. [http://www.eshard.com/wp-content/plugins/  
email-before-download/download.php?dl=9465aa084ff070a3acedb566bc34f5](http://www.eshard.com/wp-content/plugins/<br/>email-before-download/download.php?dl=9465aa084ff070a3acedb566bc34f5)
- 1204 [49] Dan Cashman. 2017. SELinux in Android O: Separating Policy to Allow for Independent Updates. [https://events.  
static.linuxfound.org/sites/events/files/slides/LSS%20-%20Treble%20%27n%27%20SELinux.pdf](https://events.<br/>static.linuxfound.org/sites/events/files/slides/LSS%20-%20Treble%20%27n%27%20SELinux.pdf) Linux Security Summit.
- 1205 [50] Jen Chai. 2019. Giving users more control over their location data. [https://android-developers.googleblog.com/2019/  
03/giving-users-more-control-over-their.html](https://android-developers.googleblog.com/2019/<br/>03/giving-users-more-control-over-their.html)
- 1206 [51] Haining Chen, Ninghui Li, William Enck, Yousra Aafer, and Xiangyu Zhang. 2017. Analysis of SEAndroid Policies: Combining MAC and DAC in Android. In *Proceedings of the 33rd Annual Computer Security Applications Conference (ACSAC 2017)*. ACM, New York, NY, USA, 553–565. <https://doi.org/10.1145/3134600.3134638>
- 1207 [52] Jiska Classen and Matthias Hollick. 2019. Inside job: diagnosing bluetooth lower layers using off-the-shelf devices. In *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2019, Miami, Florida, USA, May 15-17, 2019*. ACM, 186–191. <https://doi.org/10.1145/3317549.3319727>
- 1208 [53] Paul Crowley and Eric Biggers. 2018. Adiantum: length-preserving encryption for entry-level processors. *IACR Transactions on Symmetric Cryptology* 2018, 4 (Dec. 2018), 39–61. <https://doi.org/10.13154/tosc.v2018.i4.39-61>
- 1209 [54] Edward Cunningham. 2017. Improving app security and performance on Google Play for years to come. [https://  
android-developers.googleblog.com/2017/12/improving-app-security-and-performance.html](https://android-developers.googleblog.com/2017/12/improving-app-security-and-performance.html)
- 1210 [55] Lucas Davi, Ahmad-Reza Sadeghi, Daniel Lehmann, and Fabian Monrose. 2014. Stitching the Gadgets: On the Ineffectiveness of Coarse-Grained Control-Flow Integrity Protection. In *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX Association, San Diego, CA, 401–416. [https://www.usenix.org/conference/usenixsecurity14/  
technical-sessions/presentation/davi](https://www.usenix.org/conference/usenixsecurity14/<br/>technical-sessions/presentation/davi)
- 1211
- 1212
- 1213
- 1214
- 1215
- 1216
- 1217
- 1218
- 1219
- 1220
- 1221
- 1222
- 1223
- 1224
- 1225



- 1226 [56] Rachna Dhamija, J. D. Tygar, and Marti Hearst. 2006. Why Phishing Works. In *Proceedings of the SIGCHI Conference*  
 1227 *on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 581–590. [https://doi.org/10.1145/](https://doi.org/10.1145/1124772.1124861)  
 1228 [1124772.1124861](https://doi.org/10.1145/1124772.1124861)
- 1229 [57] Danny Dolev and Andrew Chi chih Yao. 1983. On the security of public key protocols. *IEEE Transactions on Information*  
 1230 *Theory* 29 (1983), 198–208. Issue 2. <https://doi.org/10.1109/TIT.1983.1056650>
- 1231 [58] Andre Egners, Björn Marschollek, and Ulrike Meyer. 2012. *Hackers in Your Pocket: A Survey of Smartphone Security*  
 1232 *Across Platforms*. Technical Report 2012,7. RWTH Aachen University. [https://itsec.rwth-aachen.de/publications/ae\\_](https://itsec.rwth-aachen.de/publications/ae_hacker_in_your_pocket.pdf)  
 1233 [hacker\\_in\\_your\\_pocket.pdf](https://itsec.rwth-aachen.de/publications/ae_hacker_in_your_pocket.pdf)
- 1234 [59] W. Enck, M. Ongtang, and P. McDaniel. 2009. Understanding Android Security. *IEEE Security Privacy* 7, 1 (Jan 2009),  
 1235 50–57. <https://doi.org/10.1109/MSP.2009.26>
- 1236 [60] Sascha Fahl, Marian Harbach, Thomas Muders, Lars Baumgärtner, Bernd Freisleben, and Matthew Smith. 2012. Why  
 1237 Eve and Mallory Love Android: An Analysis of Android SSL (in)Security. In *Proceedings of the 2012 ACM Conference*  
 1238 *on Computer and Communications Security (CCS '12)*. ACM, New York, NY, USA, 50–61. [https://doi.org/10.1145/](https://doi.org/10.1145/2382196.2382205)  
 1239 [2382196.2382205](https://doi.org/10.1145/2382196.2382205)
- 1240 [61] Sascha Fahl, Marian Harbach, Henning Perl, Markus Koetter, and Matthew Smith. 2013. Rethinking SSL Development  
 1241 in an Appified World. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer &#38; Communications Security*  
 1242 *(CCS '13)*. ACM, New York, NY, USA, 49–60. <https://doi.org/10.1145/2508859.2516655>
- 1243 [62] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin.  
 1244 2010. Diversity in Smartphone Usage. In *Proc. 8th International Conference on Mobile Systems, Applications, and*  
 1245 *Services (MobiSys '10)*. ACM, New York, NY, USA, 179–194. <https://doi.org/10.1145/1814433.1814453>
- 1246 [63] P. Faruki, A. Bharmal, V. Laxmi, V. Ganmoor, M. S. Gaur, M. Conti, and M. Rajarajan. 2015. Android Security: A  
 1247 Survey of Issues, Malware Penetration, and Defenses. *IEEE Communications Surveys Tutorials* 17, 2 (2015), 998–1022.  
 1248 <https://doi.org/10.1109/COMST.2014.2386139>
- 1249 [64] Adrienne Porter Felt, Serge Egelman, Matthew Finifter, Devdatta Akhawe, and David A. Wagner. 2012. How to Ask  
 1250 for Permission. In *HotSec*.
- 1251 [65] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android  
 1252 Permissions: User Attention, Comprehension, and Behavior. In *Proceedings of the Eighth Symposium on Usable Privacy*  
 1253 *and Security (SOUPS '12)*. ACM, New York, NY, USA, Article 3, 14 pages. <https://doi.org/10.1145/2335356.2335360>
- 1254 [66] Earlene Fernandes, Qi Alfred Chen, Justin Paupore, Georg Essl, J. Alex Halderman, Z. Morley Mao, and Atul Prakash.  
 1255 2016. Android UI Deception Revisited: Attacks and Defenses. In *Financial Cryptography and Data Security (Lecture*  
 1256 *Notes in Computer Science)*. Springer, Berlin, Heidelberg, 41–59. [https://doi.org/10.1007/978-3-662-54970-4\\_3](https://doi.org/10.1007/978-3-662-54970-4_3)
- 1257 [67] Nate Fischer. 2018. Protecting WebView with Safe Browsing. [https://android-developers.googleblog.com/2018/04/](https://android-developers.googleblog.com/2018/04/protecting-webview-with-safe-browsing.html)  
 1258 [protecting-webview-with-safe-browsing.html](https://android-developers.googleblog.com/2018/04/protecting-webview-with-safe-browsing.html)
- 1259 [68] Google APIs for Android. [n. d.]. <https://developers.google.com/android/reference/com/google/android/gms/fido/Fido>
- 1260 [69] Yanick Fratantonio, Chenxiong Qian, Simon Chung, and Wenke Lee. 2017. Cloak and Dagger: From Two Permissions  
 1261 to Complete Control of the UI Feedback Loop. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)*.  
 1262 San Jose, CA.
- 1263 [70] Martin Georgiev, Subodh Iyengar, Suman Jana, Rishita Anubhai, Dan Boneh, and Vitaly Shmatikov. 2012. The most  
 1264 dangerous code in the world: validating SSL certificates in non-browser software. In *ACM Conference on Computer*  
 1265 *and Communications Security*. 38–49.
- 1266 [71] Anwar Ghuloum. 2019. Fresher OS with Projects Treble and Mainline. [https://android-developers.googleblog.com/](https://android-developers.googleblog.com/2019/05/fresher-os-with-projects-treble-and-mainline.html)  
 1267 [2019/05/fresher-os-with-projects-treble-and-mainline.html](https://android-developers.googleblog.com/2019/05/fresher-os-with-projects-treble-and-mainline.html)
- 1268 [72] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J.  
 1269 Feldman, Jacob Appelbaum, and Edward W. Felten. 2009. Lest We Remember: Cold-boot Attacks on Encryption Keys.  
 1270 *Commun. ACM* 52, 5 (May 2009), 91–98. <https://doi.org/10.1145/1506409.1506429>
- 1271 [73] Daniel Hintze, Rainhard D. Findling, Muhammad Muaaz, Sebastian Scholz, and René Mayrhofer. 2014. Diversity in  
 1272 Locked and Unlocked Mobile Device Usage. In *Proceedings of the 2014 ACM International Joint Conference on*  
 1273 *Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp 2014)*. ACM Press, 379–384. [https://doi.org/10.](https://doi.org/10.1145/2638728.2641697)  
 1274 [1145/2638728.2641697](https://doi.org/10.1145/2638728.2641697)
- 1275 [74] Daniel Hintze, Rainhard D. Findling, Sebastian Scholz, and René Mayrhofer. 2014. Mobile Device Usage Characteristics:  
 1276 The Effect of Context and Form Factor on Locked and Unlocked Usage. In *Proc. MoMM 2014: 12th International*  
 1277 *Conference on Advances in Mobile Computing and Multimedia*. ACM Press, New York, NY, USA, 105–114. <https://doi.org/10.1145/2684103.2684156>
- 1278 [75] Daniel Hintze, Philipp Hintze, Rainhard Dieter Findling, and René Mayrhofer. 2017. A Large-Scale, Long-Term  
 1279 Analysis of Mobile Device Usage Characteristics. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 2, Article  
 1280 13 (June 2017), 21 pages. <https://doi.org/10.1145/3090078>



- 1275 [76] Sebastian Höbarth and René Mayrhofer. 2011. A framework for on-device privilege escalation exploit execution on  
1276 Android. In *Proc. IWSSI/SPMU 2011: 3rd International Workshop on Security and Privacy in Spontaneous Interaction and*  
1277 *Mobile Phone Use, colocated with Pervasive 2011*.
- 1278 [77] Michael Hölzl, Michael Roland, and René Mayrhofer. 2017. Real-world Identification for an Extensible and Privacy-  
1279 preserving Mobile eID. In *Privacy and Identity Management. The Smart Revolution. Privacy and Identity 2017*. IFIP  
1280 AICT, Vol. 526/2018. Springer, Ispra, Italy, 354–370. [https://doi.org/10.1007/978-3-319-92925-5\\_24](https://doi.org/10.1007/978-3-319-92925-5_24)
- 1281 [78] Yeongjin Jang, Chengyu Song, Simon P. Chung, Tielei Wang, and Wenke Lee. 2014. A11Y Attacks: Exploiting Accessi-  
1282 bility in Operating Systems. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications*  
1283 *Security (CCS '14)*. ACM, New York, NY, USA, 103–115. <https://doi.org/10.1145/2660267.2660295>
- 1284 [79] Troy Kensing. 2018. Google and Android have your back by protecting your backups. <https://security.googleblog.com/2018/10/google-and-android-have-your-back-by.html>
- 1285 [80] Amin Kharraz, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. 2015. Cutting the Gordian Knot:  
1286 A Look Under the Hood of Ransomware Attacks. In *Detection of Intrusions and Malware, and Vulnerability Assessment*,  
1287 Magnus Almgren, Vincenzo Gulisano, and Federico Maggi (Eds.). Springer International Publishing, Cham, 3–24.
- 1288 [81] Erik Kline and Ben Schwartz. 2018. DNS over TLS support in Android P Developer Preview. <https://android-developers.googleblog.com/2018/04/dns-over-tls-support-in-android-p.html>
- 1289 [82] Paul Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas  
1290 Prescher, Michael Schwarz, and Yuval Yarom. 2018. Spectre Attacks: Exploiting Speculative Execution. *arXiv:1801.01203 [cs]* (2018). arXiv:1801.01203 <http://arxiv.org/abs/1801.01203>
- 1291 [83] Nick Kralevich. 2016. The Art of Defense: How vulnerabilities help shape security features and mitigations in Android. [https://www.blackhat.com/docs/us-16/materials/us-16-Kralevich-The-Art-Of-Defense-How-Vulnerabilities-Help-Shape-](https://www.blackhat.com/docs/us-16/materials/us-16-Kralevich-The-Art-Of-Defense-How-Vulnerabilities-Help-Shape-Security-Features-And-Mitigations-In-Android.pdf)  
1292 [Security-Features-And-Mitigations-In-Android.pdf](https://www.blackhat.com/docs/us-16/materials/us-16-Kralevich-The-Art-Of-Defense-How-Vulnerabilities-Help-Shape-Security-Features-And-Mitigations-In-Android.pdf) BlackHat.
- 1293 [84] Joshua Kraunelis, Yinjie Chen, Zhen Ling, Xinwen Fu, and Wei Zhao. 2014. On Malware Leveraging the Android  
1294 Accessibility Framework. In *Mobile and Ubiquitous Systems: Computing, Networking, and Services*, Ivan Stojmenovic,  
1295 Zixue Cheng, and Song Guo (Eds.). Springer International Publishing, Cham, 512–523.
- 1296 [85] Mariantonietta La Polla, Fabio Martinelli, and Daniele Sgandurra. 2013. A Survey on Security for Mobile Devices. 15  
1297 (01 2013), 446–471.
- 1298 [86] Ben Lapid and Avishai Wool. 2019. Cache-Attacks on the ARM TrustZone Implementations of AES-256 and AES-256-  
1299 GCM via GPU-Based Analysis. In *Selected Areas in Cryptography – SAC 2018*, Carlos Cid and Michael J. Jacobson Jr.  
1300 (Eds.). Springer International Publishing, Cham, 235–256.
- 1301 [87] B. Laurie, A. Langley, and E. Kasper. 2013. Certificate Transparency. <https://www.rfc-editor.org/info/rfc6962>
- 1302 [88] Li Li, Alexandre Bartel, Jacques Klein, Yves Le Traon, Steven Arzt, Siegfried Rasthofer, Eric Bodden, Damien Oceau,  
1303 and Patrick McDaniel. 2014. I know what leaked in your pocket: uncovering privacy leaks on Android Apps with  
1304 Static Taint Analysis. *arXiv:1404.7431 [cs]* (April 2014). <http://arxiv.org/abs/1404.7431>
- 1305 [89] Li Li, Tegawendé F. Bissyandé, Mike Papadakis, Siegfried Rasthofer, Alexandre Bartel, Damien Oceau, Jacques  
1306 Klein, and Le Traon. 2017. Static analysis of Android apps: A systematic literature review. *Information and Software*  
1307 *Technology* 88 (2017), 67 – 95. <https://doi.org/10.1016/j.infsof.2017.04.001>
- 1308 [90] M. Lindorfer, M. Neugschwandtner, L. Weichselbaum, Y. Fratantonio, V. v. d. Veen, and C. Platzer. 2014. ANDRUBIS  
1309 – 1,000,000 Apps Later: A View on Current Android Malware Behaviors. In *2014 Third International Workshop on*  
1310 *Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*. 3–17. <https://doi.org/10.1109/BADGERS.2014.7>
- 1311 [91] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Stefan Mangard, Paul Kocher, Daniel  
1312 Genkin, Yuval Yarom, and Mike Hamburg. 2018. Meltdown. *arXiv:1801.01207 [cs]* (2018). arXiv:1801.01207 <http://arxiv.org/abs/1801.01207>
- 1313 [92] T. Lodderstedt, M. McGloin, and P. Hunt. 2013. OAuth 2.0 Threat Model and Security Considerations. <https://www.rfc-editor.org/info/rfc6819>
- 1314 [93] Ivan Lozano. 2018. Compiler-based security mitigations in Android P. <https://android-developers.googleblog.com/2018/06/compiler-based-security-mitigations-in.html>
- 1315 [94] Iliyan Malchev. 2017. Here comes Treble: A modular base for Android. <https://android-developers.googleblog.com/2017/05/here-comes-treble-modular-base-for.html>
- 1316 [95] René Mayrhofer. 2014. An Architecture for Secure Mobile Devices. *Security and Communication Networks* (2014).  
1317 <https://doi.org/10.1002/sec.1028>
- 1318 [96] René Mayrhofer. 2019. Insider Attack Resistance in the Android Ecosystem. In *Enigma 2019*. USENIX Association,  
1319 Burlingame, CA.
- 1320 [97] René Mayrhofer, Stephan Sigg, and Vishwath Mohan. [n. d.]. Adversary Models for Mobile Device Authentication.  
1321 ([n. d.]). submitted for review.
- 1322
- 1323

- 1324 [98] T. McDonnell, B. Ray, and M. Kim. 2013. An Empirical Study of API Stability and Adoption in the Android Ecosystem.  
 1325 In *2013 IEEE International Conference on Software Maintenance*. 70–79. <https://doi.org/10.1109/ICSM.2013.18>
- 1326 [99] I. Mohamed and D. Patel. 2015. Android vs iOS Security: A Comparative Study. In *2015 12th International Conference*  
 1327 *on Information Technology - New Generations*. 725–730. <https://doi.org/10.1109/ITNG.2015.123>
- 1328 [100] Vishwath Mohan. 2018. Better Biometrics in Android P. [https://android-developers.googleblog.com/2018/06/  
 1329 better-biometrics-in-android-p.html](https://android-developers.googleblog.com/2018/06/better-biometrics-in-android-p.html)
- 1330 [101] Vikrant Nanda and René Mayrhofer. 2018. Android Pie à la mode: Security & Privacy. [https://android-developers.  
 1331 googleblog.com/2018/12/android-pie-la-mode-security-privacy.html](https://android-developers.googleblog.com/2018/12/android-pie-la-mode-security-privacy.html)
- 1332 [102] Sundar Pichai. 2018. Android has created more choice, not less. [https://blog.google/around-the-globe/google-europe/  
 1333 android-has-created-more-choice-not-less/](https://blog.google/around-the-globe/google-europe/android-has-created-more-choice-not-less/)
- 1334 [103] Peter Riedl, Rene Mayrhofer, Andreas Möller, Matthias Kranz, Florian Lettner, Clemens Holzmann, and Marion Koelle.  
 1335 2015. Only play in your comfort zone: interaction methods for improving security awareness on mobile devices. *Personal and Ubiquitous Computing* (27 March 2015), 1–14. <https://doi.org/10.1007/s00779-015-0840-5>
- 1336 [104] Franziska Roesner, Tadayoshi Kohno, Er Moshchuk, Bryan Parno, Helen J. Wang, and Crispin Cowan. 2012. User-  
 1337 driven access control: Rethinking permission granting in modern operating systems. In *Proceedings of the 2012 IEEE  
 1338 Symposium on Security and Privacy, ser. SP '12*. 224–238. <https://doi.org/10.1109/SP.2012.24>
- 1339 [105] Michael Roland, Josef Langer, and Josef Scharinger. 2013. Applying Relay Attacks to Google Wallet. In *Proceedings  
 1340 of the Fifth International Workshop on Near Field Communication (NFC 2013)*. IEEE, Zurich, Switzerland. <https://doi.org/10.1109/NFC.2013.6482441>
- 1341 [106] R. S. Sandhu and P. Samarati. 1994. Access control: principle and practice. *IEEE Communications Magazine* 32, 9 (Sept  
 1342 1994), 40–48. <https://doi.org/10.1109/35.312842>
- 1343 [107] N. Scaife, H. Carter, P. Traynor, and K. R. B. Butler. 2016. CryptoLock (and Drop It): Stopping Ransomware Attacks  
 1344 on User Data. In *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*. 303–312. <https://doi.org/10.1109/ICDCS.2016.46>
- 1345 [108] Arvind Seshadri, Mark Luk, Ning Qu, and Adrian Perrig. 2007. SecVisor: A Tiny Hypervisor to Provide Lifetime  
 1346 Kernel Code Integrity for Commodity OSES. In *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating  
 1347 Systems Principles (SOSP '07)*. ACM, New York, NY, USA, 335–350. <https://doi.org/10.1145/1294261.1294294>
- 1348 [109] Hovav Shacham, Matthew Page, Ben Pfaff, Eu-Jin Goh, Nagendra Modadugu, and Dan Boneh. 2004. On the Effectiveness  
 1349 of Address-space Randomization. In *Proceedings of the 11th ACM Conference on Computer and Communications  
 1350 Security (CCS '04)*. ACM, New York, NY, USA, 298–307. <https://doi.org/10.1145/1030083.1030124>
- 1351 [110] Stephen Smalley and Robert Craig. 2013. Security Enhanced (SE) Android: Bringing Flexible MAC to Android. In  
 1352 *Proc. of NDSS 2013*. 18.
- 1353 [111] Sampath Srinivas and Karthik Lakshminarayanan. 2019. Simplifying identity and access management  
 1354 of your employees, partners, and customers. [https://cloud.google.com/blog/products/identity-security/  
 1355 simplifying-identity-and-access-management-of-your-employees-partners-and-customers](https://cloud.google.com/blog/products/identity-security/simplifying-identity-and-access-management-of-your-employees-partners-and-customers)
- 1356 [112] Jeff Vander Stoep and Chong Zhang. 2019. Queue the Hardening Enhancements. [https://android-developers.  
 1357 googleblog.com/2019/05/queue-hardening-enhancements.html](https://android-developers.googleblog.com/2019/05/queue-hardening-enhancements.html)
- 1358 [113] Andrew S. Tanenbaum and Herbert Bos. 2014. *Modern Operating Systems* (4th ed.). Prentice Hall Press, Upper Saddle  
 1359 River, NJ, USA.
- 1360 [114] Adrian Tang, Simha Sethumadhavan, and Salvatore Stolfo. 2017. CLKSCREW: Exposing the Perils of Security-Oblivious  
 1361 Energy Management. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association, Vancouver, BC,  
 1362 1057–1074. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/tang>
- 1363 [115] Sai Deep Tetali. 2018. Keeping 2 Billion Android devices safe with machine learning. [https://android-developers.  
 1364 googleblog.com/2018/05/keeping-2-billion-android-devices-safe.html](https://android-developers.googleblog.com/2018/05/keeping-2-billion-android-devices-safe.html)
- 1365 [116] Daniel R. Thomas, Alastair R. Beresford, and Andrew Rice. 2015. Security Metrics for the Android Ecosystem. In  
 1366 *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM  
 1367 '15)*. Association for Computing Machinery, New York NY USA, 87–98. <https://doi.org/10.1145/2808117.2808118>
- 1368 [117] Caroline Tice, Tom Roeder, Peter Collingbourne, Stephen Checkoway, Úlfar Erlingsson, Luis Lozano, and Geoff Pike.  
 1369 2014. Enforcing Forward-Edge Control-Flow Integrity in GCC & LLVM. In *23rd USENIX Security Symposium (USENIX  
 1370 Security 14)*. USENIX Association, San Diego, CA, 941–955. [https://www.usenix.org/conference/usenixsecurity14/  
 1371 technical-sessions/presentation/tice](https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/tice)
- 1372 [118] Sami Tolvanen. 2017. Hardening the Kernel in Android Oreo. [https://android-developers.googleblog.com/2017/08/  
 1373 hardening-kernel-in-android-oreo.html](https://android-developers.googleblog.com/2017/08/hardening-kernel-in-android-oreo.html)
- [119] Sami Tolvanen. 2018. Control Flow Integrity in the Android kernel. [https://security.googleblog.com/2018/10/  
 1374 posted-by-sami-tolvanen-staff-software.html](https://security.googleblog.com/2018/10/posted-by-sami-tolvanen-staff-software.html)
- [120] Sami Tolvanen. 2019. Protecting against code reuse in the Linux kernel with Shadow Call Stack. [https://security.  
 1375 googleblog.com/2019/10/protecting-against-code-reuse-in-linux\\_30.html](https://security.googleblog.com/2019/10/protecting-against-code-reuse-in-linux_30.html)

- 1373 [121] Victor van der Veen, Yanick Fratantonio, Martina Lindorfer, Daniel Gruss, Clementine Maurice, Giovanni Vigna,  
1374 Herbert Bos, Kaveh Razavi, and Cristiano Giuffrida. 2016. Drammer: Deterministic Rowhammer Attacks on Mobile  
1375 Platforms. ACM Press, 1675–1689. <https://doi.org/10.1145/2976749.2978406>
- 1376 [122] Jeff Vander Stoep. 2015. Iocctl Command Whitelisting in SELinux. <http://kernsec.org/files/lss2015/vanderstoep.pdf>  
1377 Linux Security Summit.
- 1378 [123] Jeff Vander Stoep. 2016. Android: Protecting the Kernel. [https://events.static.linuxfound.org/sites/events/files/slides/  
1379 Android-%20protecting%20the%20kernel.pdf](https://events.static.linuxfound.org/sites/events/files/slides/Android-%20protecting%20the%20kernel.pdf) Linux Security Summit.
- 1380 [124] Jeff Vander Stoep. 2017. Shut the HAL up. <https://android-developers.googleblog.com/2017/07/shut-hal-up.html>
- 1381 [125] Jeff Vander Stoep and Sami Tolvanen. 2018. Year in Review: Android Kernel Security. [https://events.linuxfoundation.  
1382 org/wp-content/uploads/2017/11/LSS2018.pdf](https://events.linuxfoundation.org/wp-content/uploads/2017/11/LSS2018.pdf) Linux Security Summit.
- 1383 [126] W3C. [n. d.]. Web Authentication: An API for accessing Public Key Credentials. <https://webauthn.io/>
- 1384 [127] R. Watson. 2012. *New approaches to operating system security extensibility*. Technical Report UCAM-CL-TR-818.  
1385 Cambridge University. <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-818.pdf>
- 1386 [128] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015.  
1387 Android Permissions Remystified: A Field Study on Contextual Integrity. In *24th USENIX Security Symposium (USENIX  
1388 Security 15)*. USENIX Association, Washington, D.C., 499–514. [https://www.usenix.org/conference/usenixsecurity15/  
1389 technical-sessions/presentation/wijesekera](https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/wijesekera)
- 1390 [129] Shawn Willden. 2018. Insider Attack Resistance. [https://android-developers.googleblog.com/2018/05/  
1391 insider-attack-resistance.html](https://android-developers.googleblog.com/2018/05/insider-attack-resistance.html)
- 1392 [130] Yuan Zhang, Min Yang, Bingquan Xu, Zhemin Yang, Guofei Gu, Peng Ning, X. Sean Wang, and Binyu Zang. 2013.  
1393 Vetting Undesirable Behaviors in Android Apps with Permission Use Analysis. In *Proceedings of the 2013 ACM  
1394 SIGSAC Conference on Computer & Communications Security (CCS '13)*. ACM, New York, NY, USA, 611–622. <https://doi.org/10.1145/2508859.2516689>
- 1395
- 1396
- 1397
- 1398
- 1399
- 1400
- 1401
- 1402
- 1403
- 1404
- 1405
- 1406
- 1407
- 1408
- 1409
- 1410
- 1411
- 1412
- 1413
- 1414
- 1415
- 1416
- 1417
- 1418
- 1419
- 1420
- 1421

Table 1. Application sandboxing improvements in Android releases

Release	Improvement	Threats mitigated
≤ 4.3	Isolated process: Apps may optionally run services in a process with no Android permissions and access to only two binder services. For example, the Chrome browser runs its renderer in an isolated process for rendering untrusted web content.	[T10] access to [T5][T8][T9][T12][T13]
5.x	SELinux: SELinux was enabled for all userspace, significantly improving the separation between apps and system processes. Separation between apps is still primarily enforced via the UID sandbox. A major benefit of SELinux is the auditability/testability of policy. The ability to test security requirements during compatibility testing increased dramatically with the introduction of SELinux.	[T8][T14]
5.x	Webview moved to an updatable APK, independent of a full system OTA.	[T10]
6.x	Run time permissions were introduced, which moved the request for dangerous permission from install to first use (cf. above description of permission classes).	[T7]
6.x	Multi-user support: SELinux categories were introduced for a per-physical-user app sandbox.	[T3]
6.x	Safer defaults on private app data: App home directory moved from 0751 UNIX permissions to 0700 (based on targetSdkVersion).	[T9]
6.x	SELinux restrictions on ioctl system call: 59% of all app reachable kernel vulnerabilities were through the ioctl() syscall, and these restrictions limit reachability of potential kernel vulnerabilities from user space code [122, 123].	[T8][T14]
6.x	Removal of app access to debugfs (9% of all app-reachable kernel vulnerabilities).	[T8][T14]
7.x	hidepid=2: Remove /proc/<pid> side channel used to infer when apps were started.	[T11]
7.x	perf-event-hardening (11% of app reachable kernel vulnerabilities were reached via perf_event_open()).	[T8]
7.x	Safer defaults on /proc filesystem access.	[T7][T11]
7.x	MITM CA certificates are not trusted by default.	[T6]
8.x	Safer defaults on /sys filesystem access.	[T7][T11]
8.x	All apps run with a seccomp filter intended to reduce kernel attack surface.	[T8][T14]
8.x	Webviews for all apps move into the isolated process.	[T10]
8.x	Apps must opt-in to use cleartext network traffic.	[T5]
9.0	Per-app SELinux sandbox (for apps with targetSdkVersion=P or greater).	[T9][T11]
10	Apps can only start a new activity with a visible window, in the foreground activity 'back stack', or if more specific exceptions apply [25].	[T8][T9][T10][T11]
10	File access on external storage is scoped to app-owned files.	[T7][T9]
10	Reading clipboard data is only possible for the app that currently has input focus or is the default IME app.	[T12]
10	/proc/net limitations and other side channel mitigations.	[T7]
11	Legacy access of non-scoped external storage is no longer available.	[T7][T9]

Table 2. System sandboxing improvements in Android releases

Release	Improvement	Threats mitigated
4.4	SELinux in enforcing mode: MAC for 4 root processes <code>installd</code> , <code>netd</code> , <code>vold</code> , <code>zygote</code> .	[T7][T8] [T14]
5.x	SELinux: MAC for all userspace processes.	[T7][T8]
6.x	SELinux: MAC for all processes.	
7.x	Architectural decomposition of <code>mediaserver</code> .	[T7][T8] [T14]
7.x	<code>ioctl</code> system call restrictions for system components [122].	[T7][T8] [T14]
8.x	<i>Treble</i> Architectural decomposition: Move HALs (Hardware Abstraction Layer components) into separate processes, reduce permissions, restrict access to hardware drivers [49, 124].	[T7][T8] [T14]
10	Software codecs, the source of approximately 80% of the critical/high severity vulnerabilities in media components were moved into a constrained sandbox	[T8][T14]
10	Bounds Sanitizer (BoundSan): Missing or incorrect bounds checks on arrays accounted for 34% of Android's userspace security vulnerabilities. Clang's BoundSan adds bounds checking on arrays when the size can be determined at compile time. BoundSan was enabled across the bluetooth stack, and in 11 software Codecs.	[T8][T14]
10	Integer Overflow Sanitizer (IOSAN): The process of applying IOSAN to the media frameworks began in Android 7.0 and was completed in Android 10.	[T8][T14]
10	Scudo is a dynamic heap allocator designed to be resilient against heap related vulnerabilities.	[T8][T14]

Table 3. Kernel sandboxing improvements in Android releases

Release	Improvement	Threats mitigated
5.x	Privileged eXecute Never (PXN) [3]: Disallow the kernel from executing userspace. Prevents 'ret2user' style attacks.	[T8][T14]
6.x	Kernel threads moved into SELinux enforcing mode, limiting kernel access to userspace files.	[T8][T14]
8.x	Privileged Access Never (PAN) and PAN emulation: Prevent the kernel from accessing any userspace memory without going through hardened <code>copy-*-user()</code> functions [118].	[T8][T14]
9.0	Control Flow Integrity (CFI): Ensures that front-edge control flow stays within a precomputed graph of allowed function calls [119].	[T8][T14]
10	Shadow Call Stack (SCS): Protects the backwards edge of the call graph by protecting return addresses [120].	[T8][T14]

Table 4. Network sandboxing improvements in Android releases

Release	Improvement	Threats mitigated
6.x	usesCleartextTraffic in manifest to prevent unintentional cleartext connections [45].	[T5][T6]
7.x	Network security config [30] to declaratively specify TLS and cleartext settings on a per-domain or app-wide basis to customize TLS connections.	[T5][T6]
9.0	DNS-over-TLS [81] to reduce sensitive data sent over cleartext and made apps opt-in to using cleartext traffic in their network security config.	[T5][T6]
9.0	TLS is the default for all connections [46]	[T5][T6]
10	MAC randomization is enabled by default for client mode, SoftAp, and Wi-Fi Direct <sup>14</sup>	[T5]
10	TLS 1.3 support	[T5][T6]