GLOBAL MOBILE THREAT REPORT

2026

ZIMPERIUM®

Table of Contents

Executive Summary

AI: Supercharging the Exploitation of Enterprise Mobile Environments and Rewriting the Playbook of Trust and Fraud

We have entered the era of Mythos, a world where AI can find a vulnerability, chain it with others, and build a working exploit before a human analyst finishes their coffee. While we are led to believe that CISOs are waking up in the middle of the night worried about the Mythos 'boogeyman' that's coming for them, the true impact of weaponized AI is already here and is silently wreaking havoc on what is the enterprise's most vulnerable attack surface: **mobile**.

While the industry hardened every other layer, mobile became the blind spot. Mobile isn't one attack surface, it's four — the operating system (OS), the apps, the network it connects to, the human using it. And none of them are static. New apps are installed daily. New WiFi networks joined without a thought. OS updates are pushed silently overnight and often ignored by users for days or weeks. Humans are clicking on links that appear to be genuine. The mobile environment is a playground for AI, whether weaponized or simply as an unchecked pipeline for confidential company data loss via 'Shadow AI'.

Employees use an average of nine work apps daily[1], **and 46% of U.S. frontline workers can't complete their core job functions without a mobile device**[2]**.** Every mobile device is a unique, mutating attack surface — no two alike, and almost all of them living outside the enterprise perimeter and its control, despite holding the keys to user credentials and authentication.

AI thrives in this environment and attackers have taken notice. Combining this mostly unprotected and largely unmonitored medium with what is traditionally the most effective attack technique, social engineering, we are seeing what is arguably historic scale and efficacy of data and credential theft, banking fraud and enterprise compromise - all enabled by AI that is easily accessible by even the most under-resourced cybercriminals.

The problem of AI-driven risk to the organization isn't limited to motivated and clever attackers. Unintentional data compromise due to employee AI usage on mobile is a flourishing compliance gap that most security teams are not equipped to regulate or even monitor.

Mobile was already the most vulnerable and under assessed attack surface in the enterprise, as few organizations have invested in real visibility and protection. The scale and sophistication of the threats we see today were always going to happen. The consensus was that security teams and vendors would not expect them to arrive for another five or more years from now. The availability and capability of AI has accelerated that timeline. The threat is here and many organizations have no plan for defense.

Zimperium's zLabs research has identified four primary AI-powered threat vectors that have the broadest impact across enterprise mobile environments:

MOBILE PHISHING (MISHING)

AI-assisted phishing campaigns have reportedly **grown at a rate of more than 1,200%** according to some sources [3]. AI is used to generate content estimated to be 4.5x more convincing than anything human-written[4] while almost 86% of phishing attacks now contain AI-generated elements[5].

Phishing events detected on employee mobile devices grew 380% since January 2025, and the number of devices where employees clicked a malicious link grew 110% in 2025 over the previous year. Mobile-targeted phishing, which is inclusive of various messaging tools (SMS, Telegram, WhatsApp, etc), QR codes (quishing), and device-aware email, generally sits outside the organization's security perimeter and has a success rate that is 40% higher than traditional PC-based phishing attacks. AI didn't invent mobile phishing. It just made every text, every email, every PDF convincing enough to fool even the most vigilant employees. AI has provided the scale and efficacy that was once only available to the most sophisticated criminal organizations.

MOBILE MALWARE TARGETING EMPLOYEES & CUSTOMERS

The average employee device carries around 80 apps, only nine are for work. The rest are personal, unmanaged, and invisible to standard corporate security tools; which is exactly where malware lives. AI-driven social engineering is tricking users into downloading visually perfect imitations of legitimate mobile apps which carry a silent, sinister malware payload. Spyware is now on nearly one in 10 devices, up over 4x, sold commercially by more than 500 vendors worldwide. AI is accelerating both the scale of distribution, as well as previously unseen sophistication in both capability and obfuscation.

Mobile malware driven fraud is a unique and growing economic threat to the banking and payments sector. **Zimperium's zLabs team tracked 34 active malware families targeting more than 1,200 financial brands across 90 countries**. Banking malware has moved past stealing credentials. It now takes over the transaction itself. Three families alone: TsarBot, CopyBara, and Hook, targeted more than 60% of banking and fintech apps in the Zimperium Mobile Banking Heist report. AI now runs the entire attack chain. It reverse-engineers the legitimate app, regenerates the malware's code on each run to dodge signature detection, and scripts the fraudulent transaction, all with minimal manual intervention by the attacker. Fraud now starts on the mobile device by manipulating legitimate apps and their sessions.

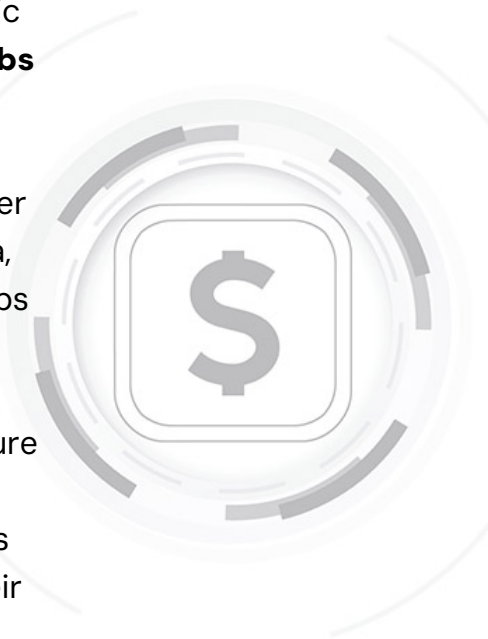

SHADOW AI

Whether intentionally or unknowingly, employees are inviting AI to feast on confidential corporate data in two different ways. Firstly, employee devices often already have multiple GenAI apps installed, most of them never approved, vetted, or even seen by enterprise security teams (unlike their PC-based counterparts). While not intentionally malicious, employees are feeding these apps sensitive enterprise data in the form of code, data and confidential information in order to take advantage of productivity gains available. But neither the employee or the business have visibility or control over where that data is transmitted or stored.

Nearly 1 in 12 employees have used a China-based AI tool in the past 30 days[6], and due to cheaper token costs, more apps will increasingly switch to Chinese open-source models, which already make up 17% of Hugging Face downloads over the past year.

Secondly, many mobile apps, including corporate approved productivity apps, contain embedded AI capabilities, which may not be apparent to the user or enterprise. In the rush to compete on advanced features, mobile app developers are looking for any advantage, including AI sourced from countries and organizations with little concern for data residency regulations. Without the ability to vet mobile apps on the device, security and compliance teams are left with no governance. Enterprise data leakage used to mean a file landing on a U.S. server, under U.S. law. Now it can mean that file leaving U.S. jurisdiction entirely, with nobody signing off and nobody finding out until long after it's gone. Shadow AI on mobile is a compliance nightmare that security teams and regulators had not seen coming, but they both will soon be compelled to act.

AI-ASSISTED MOBILE APP DEVELOPMENT RISKS

The development of mobile apps has up until now been only the domain of highly trained and governed development teams. That time is no more. AI has made the development and launch of new mobile apps both fast and relatively simple for teams including professional developers and clever, motivated amateurs. However, the more AI writes an enterprise's code, the more exploitable its apps become. Veracode[7] found 45% of AI-generated code samples fail basic OWASP Top-10 benchmarks. Apiiro found AI-assisted developers at Fortune 50 scale ship code 3x to 4x faster, but introduced security findings at 10x the rate of their peers[8]. Gartner projects 40% of new enterprise production software will be built via vibe coding by 2028, with 25% of production defects by 2027 tracing to AI-generated code with insufficient review, up from under 1% in 2023[9]. The most exploitable generation of enterprise and “vibe coded” mobile apps in history is already shipping.

Conclusion

It is clear that cybersecurity, and mobile in particular, is going through a seismic shift in regard to the speed, scale and sophistication of attacks that are being enabled by this current generation of AI.

What the data are showing is not a theoretical threat. Mobile has been left largely unprotected by the enterprise, while leveraging employee devices (both company owned and employee owned) as the primary hub of identity and authentication. That is a contradiction that provides opportunity for attackers. And with the use of AI, the bar has never been lower to fool users into taking action that compromises both themselves and the business.

The mobile threat that organizations should be facing five years from now is here today thanks to the unprecedented acceleration enabled by AI. The gap in mobile visibility, much less protection, is being exposed by even the most unsophisticated attackers and fraudsters, creating what could be described as the greatest risk to cybersecurity and compliance that organizations are facing today.

How AI Is Weaponizing Mobile

The sections that follow take each of these threats further. Where the executive summary establishes the scale of the problem, the sections ahead unpack the mechanics behind it: how the attacks work, what the data shows across device types and regions, and what each threat means for the leaders responsible for stopping it.

AI-Assisted Mishing Skyrockets on Employee Devices

Zimperium Labs (zLabs) detected over **2.5 million phishing attacks** on employee devices in the last 12 months. Phishing events detected on employee mobile devices grew 380% since January 2025, and the number of devices where employees clicked a malicious link grew 110% in 2025 over the previous year.

Mobile-centric phishing already succeeds at a rate **40% higher than email** [10] and AI is now widening that gap further. **86% of phishing campaigns are now AI-generated** [11], producing content estimated to be 4.5x more convincing than anything human-written [12].

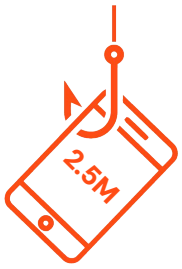

Zimperium Labs (zLabs) detected over

2.5 MILLION PHISHING ATTACKS on employee devices

Why this matters: The old defense, training employees to spot bad grammar, generic greetings, and obvious red flags, assumes the lure is distinguishable from the real thing. AI-generated lures are built specifically to erase that distinction. Awareness training calibrated to yesterday's phishing doesn't transfer to what's arriving today.

The channel breakdown shows where to focus: smishing (SMS phishing) drives 39% of all mobile threats and 70% of mobile-based phishing attacks, fueled by large-scale campaigns like fake toll-authority texts, growing 320% and costing U.S. consumers $470 million in losses [13]. Quishing (QR phishing) grew 146% quarter-over-quarter in Q1 2026 [14]. And malicious PDFs are climbing fast as a payload vector — from 19% of payloads in January to 29% in March [15]. We are seeing PDFs used in SMS or mobile email phishing at a rate four times higher than those phishing messages that only utilize a URL. This PDF tactic is used to bypass existing filtering mechanisms in place on iOS and Android.

Brand Impersonation Tactics Increasingly Used in Mishing

Successful mishing occurs when recipients engage with the mishing message. And, use of well known brands is seeing increased usage because recipients find it compelling. Our data supports this trend:

- Zimperium's zLabs found that brand-identified targeting precision grew 5.4x year-over-year signaling a deliberate concentration of resources around high-fidelity, high-value campaigns.
- Messaging app impersonation surged 258% (WhatsApp, Telegram), reflecting a strategic pivot away from email-delivered payloads toward channels that bypass enterprise email gateways.
- Telecom provider impersonation increased 77% (AT&T, Verizon), consistent with SMS-based delivery targeting employees on personal mobile devices with corporate access.

Malware

Stealing Enterprise Credentials & Data

Malware on employee devices isn't a fringe risk anymore — it's mainstream and accelerating. At the start of 2026, roughly **29% of the iOS fleet** was exposed to specific known threats like Coruna/Darksword, on a platform many security teams still treat as "low risk by default."

The growth curve is just as alarming as the exposure. Google's Threat Intelligence Group reports that adversaries "…are deploying novel AI-enabled malware in active operations."[16] Verizon reports that "…the median threat actor researched or used AI assistance across 15 different documented attack techniques, with some using 40-50."[17]

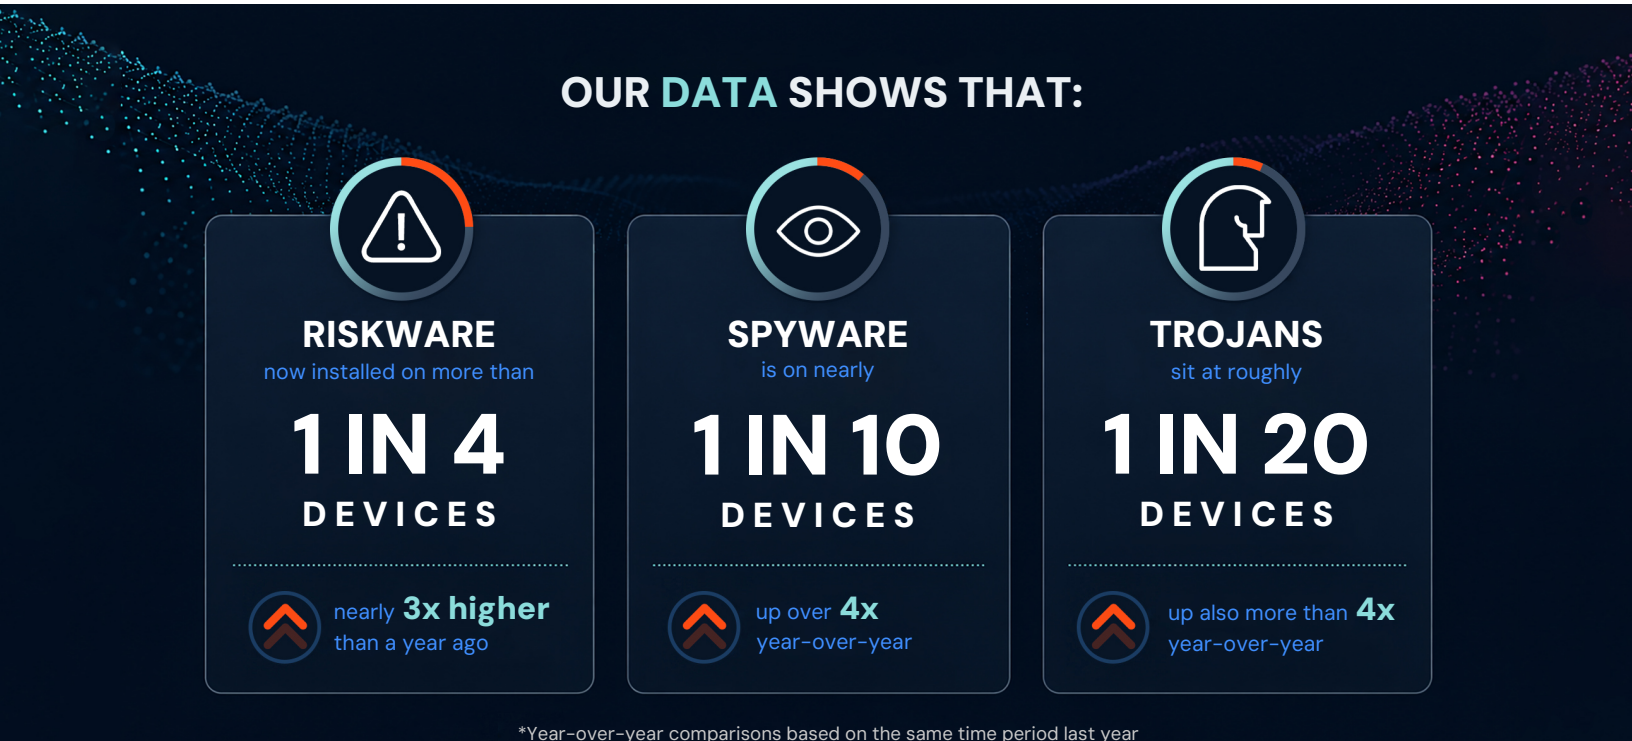

Riskware

Riskware can be a work app or a personal app — the label doesn't change the risk. What makes it "riskware" isn't who installed it. It's the app's access to data, broad permissions, weak data handling, and data sharing with third parties and foreign-hosted backends that go beyond what the user actually understood or agreed to.

Spyware

Spyware doesn't need to be the app. It just needs to be inside it. Spyware's growth isn't a spike — it's sustained. Android spyware has grown at a 22.3% CAGR over four years. The drivers here are structural. Commercial spyware has matured into a real market — more than **500 vendor entities across 40+ countries** now supply surveillance capability[18]. In 2025, for the first time, **more zero-day exploits were attributed to commercial spyware vendors than traditional nation-state groups**[19]**.** Spyware isn't a niche capability anymore, it's a supplied, commercialized product, and zero-click, mobile-messaging exploits command a premium price in that market. It is sold to a widening customer base and reaching more devices, including the ones your employees carry.

Why this matters: This isn't concentrated in one industry. Sideloaded apps and risky malware rank in the top 10 threats we observed across nearly every vertical we track, Financial Services, Retail, Manufacturing, Automotive, and Transportation among them. Zimperium's Global Threat Data shows why: **22% of Android and 14% of iOS devices have apps sideloaded from outside official app stores, bypassing Apple and Google's vetting entirely**. Whatever industry you're in, this is already on your employees' devices.

Today's sophisticated malware now requires nothing from the employee at all. **Zero-click attacks** exploiting flaws like the ImageIO memory corruption issue have made messaging apps like **WhatsApp and WeChat** primary targets.

No link, no download, no click. Just a message arriving on the device.

Mobile Malware-Driven Fraud

Sophisticated Mobile App Fraud at Scale

The Zimperium Mobile Banking Heist [20] reported 34 active malware families targeting 1,243 financial brands across 90 countries in 2025. These families are scalable fraud infrastructure, distributed through Malware-as-a-Service platforms, updated continuously, and sold to fraudsters who never have to write a line of code themselves.

Every control an enterprise relies on to secure a banking session already has a malware capability built specifically to defeat it.

APP SECURITY CONTROLS	MALWARE CAPABILITY THAT BYPASSES CONTROLS
Multi-Factor Authentication (SMS/OTP)	OTP Hijacking - Push Notification & SMS Interception
Device Binding	VNC - Remote Control of Device
Data Encryption	Screen Overlays & Keylogging
Biometric Logins	Cookie Stealing & Deep Fakes
Traditional API Security	ATS on Compromised Mobile Devices
Traditional Fraud Management Systems	On-Device Bot Automation

Why this matters: Every layer of traditional app defense an institution has invested in has a corresponding malware bypass already deployed in the wild. Adding another server side security control to this stack doesn't close the gap. These are client-side runtime attacks happening on the device enabling credential theft, account takeovers (ATO) and fraudulent transactions.

Malware Leverages AI to Rewrite Itself

Automation isn't new to fraud. What's new is what is being automated. Verizon's 2026 DBIR[21], working with Anthropic, analyzed 793 unique threat actors under enforcement action for policy violations, whose activity spanned malware development, capability building, and tasking. The median actor used AI assistance across 15 different documented attack techniques, with some reaching 40 to 50.

The next generation is already visible. **LameHug, an APT28 campaign, used an LLM to generate polymorphic malware code on demand, rewriting itself on each run so no two infections share the same signature**. PromptLock, the first AI-powered ransomware identified, generates its encryption scripts dynamically rather than shipping a fixed payload. VoidLink, a malware framework built entirely by an AI agent in six days with no human developer, is what DBIR calls a point of no return for automated threat development.

The trajectory is clear, malware that reverse-engineers the target app, rewrites its own code to evade whatever detection it encounters, and scripts the fraudulent transaction itself, with a human directing less and less of the process at each step.

Why this matters: Signature-based detection won't work against malware that rewrites itself every run. You need layered runtime detection that catches each part of the attack chain, reverse engineering, code injection, session manipulation, as it evolves, not a single control looking for a signature that no longer exists by the time you see it.

Shadow AI adoption in mobile apps isn't creeping in — it's exploding. **Android apps embedding AI grew 14x; while iOS app embedded AI use grew 7x**[22] (Figure 1). Our research data shows that total adoption in May 2026 had exceeded eight percent.

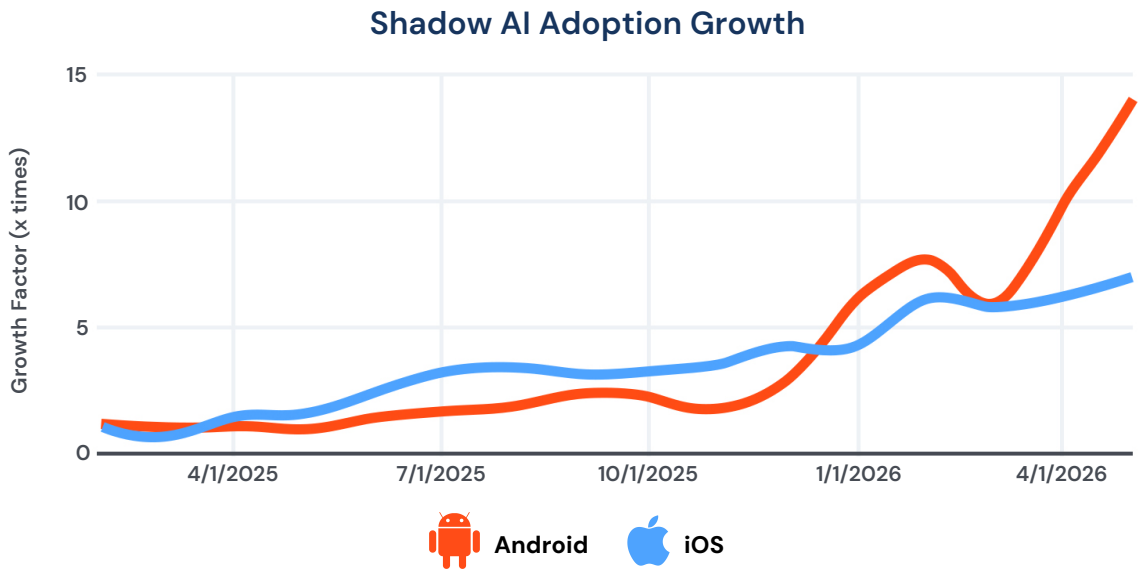

Figure 1: AI adoption growth compared to 2025 baseline

You Approved the Work App, Not the AI Inside

Gartner found employees use an average of nine apps daily at work, while those who use 16-25 apps reporting the highest satisfaction (92%) and productivity[23]. IDC reports that 46% of frontline workers in the U.S. are now considered mobile-dependent[24]; meaning they cannot complete their core job functions without the use of a mobile device.

Work applications on employee mobile devices have become a critical part of the enterprise attack surface, and AI is changing what "work app" means. zLabs data shows adoption of AI-embedded apps is up 14x on Android and 7x on iOS overall, and business apps specifically are growing AI adoption at 136%. A growing share of the business apps that employees rely on daily now ship with a GenAI feature built in. These apps were likely approved once, but never re-reviewed when the AI capability was added. You don't control how those apps were built, and increasingly, AI wrote part of them. You're still accountable for the devices they run on.

On managed devices, work apps often don't run in a separate compartment at all, personal and corporate apps share the same OS, storage, clipboard, and network stack. MAM (mobile application management) or Android Enterprise containers improve on this by isolating corporate data in an encrypted workspace. But that isolation is only as strong as the OS underneath it. If the device is rooted, jailbroken, or running spyware with system-level access, the container's walls no longer hold. Neither approach is safe on a compromised device.

With access to sensitive data, device resources, and corporate systems, weaknesses in these applications create significant security, privacy, and compliance risk.

We assessed the top business apps across the App Store and Google Play. Here's what we found:

1. Your Biggest Exposure: Insecure Data Transmission

63% of Android and **52%** of iOS business apps make network connections without proper encryption or certificate validation. This means they are sending data over plain HTTP, accepting expired or invalid SSL certificates, or skipping certificate pinning entirely.

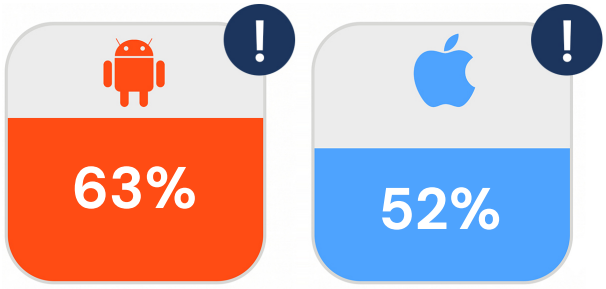

Why this matters: These apps fail to properly encrypt their own data or verify who they're actually talking to. That's the exact protection HTTPS is supposed to guarantee. Without it, any attacker in a position to intercept traffic can read or tamper with enterprise data flowing through the app, regardless of your MDM policy or device compliance score.

2. Sanctioned-Country API Exposure

17% of Android and **27%** of iOS business apps call APIs hosted within sanctioned or trade-restricted countries.

Why this matters: You likely approved the app by evaluating the functions you need and the vendor's reputation. At no point in that process did you make note of every destination the app is capable of sending data to. Most mobile apps are built on a stack of third-party SDKs: analytics, ads, crash reporting, authentication, and those SDKs are closed source. Increasingly, they're AI-generated too, assembled faster than any human reviewer can verify. You can't review what they do, only observe what they call out to. That means the destination exposure isn't just hidden in code you didn't write, it's hidden in code nobody meaningfully reviewed.

3. Sensitive Data Leakage

17% of Android and **38%** of iOS business apps leak PII: sending data like contacts, location, device identifiers, or account details off the device, often over unencrypted connections or to third parties never disclosed to the user.

Why this matters: On a corporate-connected device, "contacts" isn't a personal address book. It's colleagues, clients, partners, and prospects synced from corporate directories and CRM tools. When an approved app leaks that data, and increasingly that same app has an AI feature processing it, it's not just the employee's privacy at risk. It's enterprise relationship data leaving your control through a channel your approval process never tested for, and an AI model you never vetted.

Personal Apps on BYO Devices: Invisible and Increasingly AI

Personal apps are an enterprise problem, not just an employee one. You didn't approve them, you can't patch them, and you have no visibility into whether they pose a threat to enterprise data, but they sit on the same device. Vetting must extend beyond the apps that IT distributes. The device is the actual perimeter, and everything installed on it, approved or not, is inside that perimeter.

These aren't static risks, either. The personal app categories least likely to get any security scrutiny are the ones adding AI the fastest.

APP CATEGORY	GROWTH IN AI-EMBEDDED APPS
Entertainment	179%
Finance	224%
Food & Drink	433%
Lifestyle	278%
Travel	416%

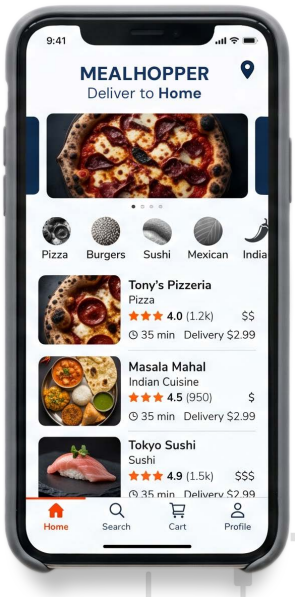

The apps that most enterprises treat as low-stakes, such as those designed for food delivery or travel, are accelerating past the ones assumed to carry the most risk. That shared surface is exploited differently depending on the platform. Android's openness and iOS's API/privacy model each create their own version of the same problem.

We assessed personal apps across both platforms and categories. Here's what we discovered and how that shared surface puts enterprise data and access at risk.

PLATFORM	FINDING	WHY IT MATTERS
	**65%** of apps can write to external storage; **63%** can read from it	Personal apps can access the same storage space corporate apps use for cached files and attachments, no vulnerability in the work app required
	**56%** of apps request dangerous and unexpected permissions (accessibility, storage, camera, contacts)	Can be used to steal credentials, read what's on screen, and turn the device into a surveillance tool
	**71%** of apps fail nutrition label accuracy	Privacy labels reflect developer disclosure, not independent verification. They can't be relied on as a safety signal.
	**66%** of apps request dangerous and unexpected permissions (**49%** microphone, **48%** tracking)	Once granted, the app can listen in and track employee movements in the background
	**45%** of apps are vulnerable to leaking PII data	Leaked PII gives attackers what they need to build a targeted Mishing campaign for employee's corporate access.

Why this matters: None of this requires a flaw in the work app you approved and pushed to the employee devices. The exposure comes from a personal app on the same device that the enterprise never reviewed, never approved, and can't see into. The MDM might tell you which apps are installed and what version they're running. It doesn't tell you what those apps actually do, what they access, or what they leak. Visibility into an app's name isn't the same as vetting its behavior, and on this device, that gap is where the risk actually lives.

Unapproved GenAI Apps on Work Devices

Verizon's 2026 DBIR found 45% of employees are now regular users of AI on their corporate devices, authorized or not, up from 15% the year before. Of those users, 67% are doing it through non-corporate accounts on corporate devices, meaning the enterprise has no visibility into the session, the account, or where the data goes.

Verizon calls this Shadow AI, and it's no longer a minor policy gap. It's now the third most common non-malicious insider action detected in DLP datasets, a fourfold increase in a single year. The most common data type submitted to these unauthorized tools is source code, followed by images and other structured data. In 3.2% of violations, employees uploaded research and technical documentation, a direct intellectual property exposure.

Most enterprises have no policy governing AI use on personal devices at all, and no visibility into it even where a policy exists on paper. MDM covers the device. It doesn't cover what an app on that device sends to a model hosted somewhere else. Shadow AI isn't slipping past governance. For most enterprises, there is no governance to slip past.

The mobile device makes this worse, not better. A GenAI app installed on a phone doesn't need a browser extension or a corporate login to access the same clipboard, storage, and camera that a work app on the same device uses. Every one of Verizon's DBIR findings applies just as easily to a mobile device as a laptop, with one difference: there's less tooling built to catch it on mobile at all.

Why this matters: Enterprises can push MAM or MDM policies with some Data Loss Prevention (DLP) controls to block copy-paste between managed work apps and personal apps, and that control genuinely works within its scope. But it only covers the managed side. An employee can still photograph a screen, upload a research file from the device's storage, or copy code into a GenAI app. And on unmanaged BYOD devices with no MAM enrollment, none of this applies at all. The data doesn't need to be copied to leave. It just needs to be seen.

AI Governance Must Include Mobile

AI-embedded work apps, personal AI apps, unapproved GenAI: none of it is covered by a policy most enterprises actually have. It's not a tooling gap you patch with better DLP. It's a visibility gap: most enterprises can't see which apps have added a model, what data reaches it, or when a feature they already approved quietly started using one.

Every app carrying an AI model, embedded, personal, or unapproved, has to be vetted the same way. Waiting to build that governance is itself a decision, and it's the wrong one.

YOU CAN'T GOVERN WHAT YOU CAN'T SEE.

Vibe Coding

Source code, APIs, network traffic, you've got controls there. But apps keep evolving after release: they run on devices you don't fully trust, pull in third-party SDKs, and can be reverse-engineered or tampered with in the wild. The risk you're not securing isn't in the repo, it's on the device.

And it's compounding. AI now writes much of the code that ships. Gartner projects **25% of production software defects by 2027** will trace back to insufficient human review of AI-generated code — up from less than one percent in 2023[25].

APPS ARE SHIPPING FASTER. NOT SAFER.

To evaluate how today's mobile applications withstand real-world threats, we assessed the top 1,000 mobile apps in official app stores across Finance, Lifestyle, Travel, Food & Drink, and Entertainment — examining protections against reverse engineering, tampering, credential exposure, insecure communications, and data leakage once apps are running on end-user devices.

Financial Services Apps: The most permission hungry category.

Financial apps demand deep access to the device — camera, biometrics, contacts, location, SMS — to do their job. But the data shows the apps themselves aren't well protected.

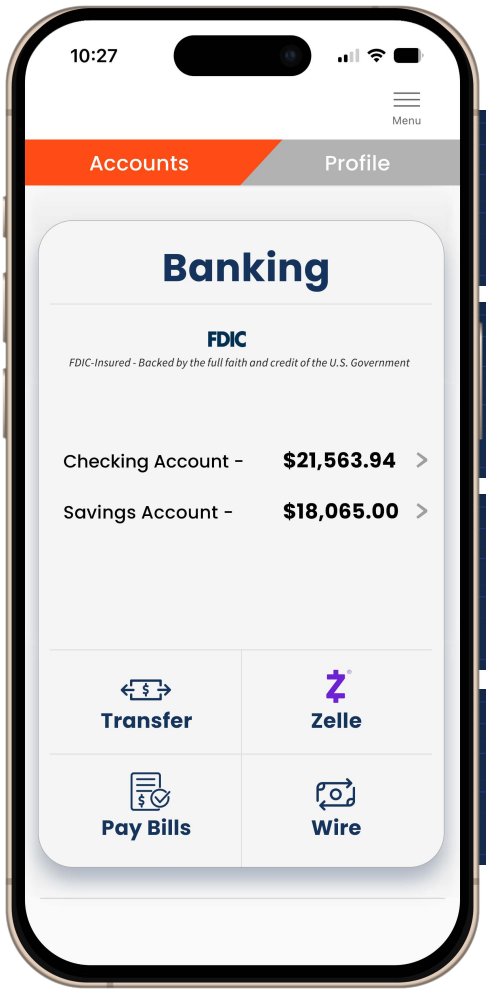

50% have no device trust or fraud-detection checks at all, meaning they can't tell if they're running on a compromised device or if they are being manipulated.

30% of Android financial apps lacked commercial code protection required to withstand modern mobile threats.

28% failed to properly validate SSL/TLS connections, the check meant to confirm they're talking to a real server, not an attacker impersonating one.

21% had secrets or credentials, like cryptographic keys, API keys or passwords, hardcoded directly into the app, readable by anyone who inspects it.

Why this matters: These are build-time decisions, not runtime accidents. No code protection means the app is reverse-engineerable on day one. Weak TLS validation lets it talk to attacker-controlled endpoints without complaint. Exposed secrets sit in the binary for anyone who decompiles it. No device trust signal means the app can't tell a clean device from a compromised one and grants both the same access. App store review and static analysis don't catch most of this; it only shows up once the app is running on a real device, outside your control.

Travel & Hospitality Apps: The most location-aware app category.

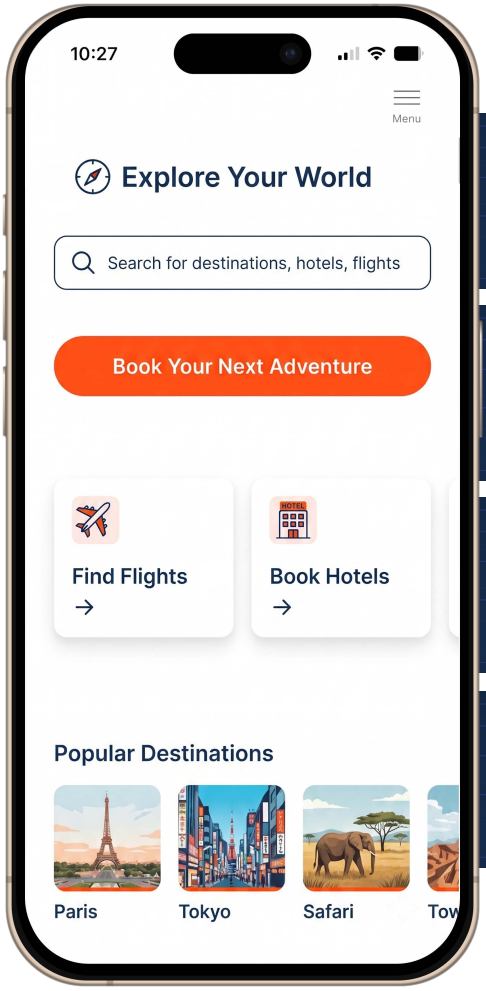

84% of travel applications can access precise location data—the highest rate among all categories analyzed.

48% lack commercial code protection, making them more susceptible to reverse engineering and tampering.

62% exhibit weaknesses in how they protect data during network communications.

16% contain API endpoints hosted in countries subject to U.S. sanctions or trade restrictions.

Why this matters: Travel applications possess a unique combination of location intelligence, itinerary information, travel patterns, and user identity data. Yet nearly half lack meaningful protection against reverse engineering, while many continue to exhibit communication security weaknesses. If compromised, these applications can expose sensitive information about employee movements, customer visits, executive travel, and business operations.

Entertainment Apps: Collect the most; protect the least.

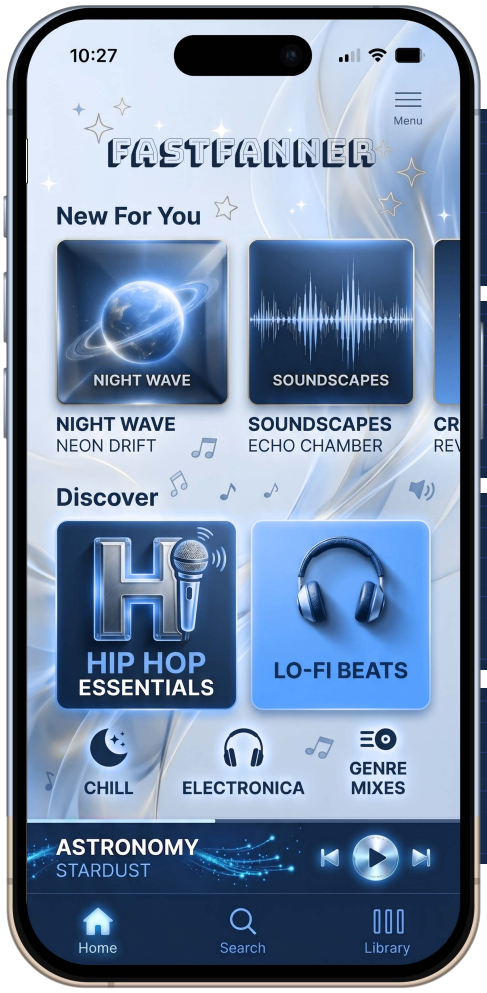

68% of entertainment applications request dangerous permissions, providing access to device resources and user data

47% lack commercial code protection, making them more susceptible to reverse engineering and tampering

71% exhibit weaknesses in how they protect data during network communications

23% contain API endpoints hosted in countries subject to U.S. sanctions or trade restrictions, the highest rate among all categories analyzed.

Why this matters: Entertainment apps collect significant user and device data while relying on extensive third-party ecosystems for advertising, analytics, and content delivery. Yet nearly half lack meaningful protection against reverse engineering, making it easy for attackers to analyze app logic, extract embedded assets, and identify backend integrations. These apps also carry a responsibility most others don't: protecting the copyrighted content they stream from piracy.

Food & Drink Apps: Low-risk perception, high-value data.

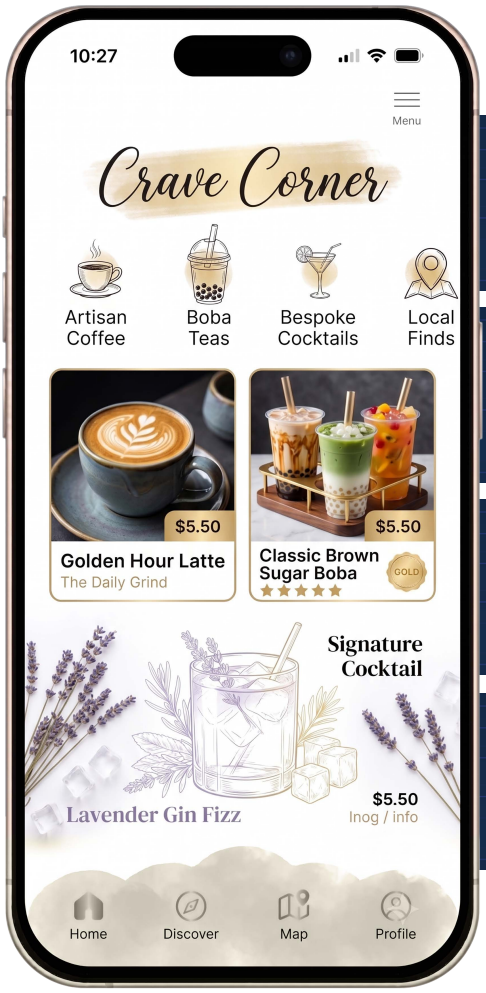

69% of food and drink applications exhibit weaknesses in how they protect data during network communications

98% contain database-related findings, the highest rate among all categories analyzed

14% contain API endpoints hosted in countries subject to U.S. sanctions or trade restrictions

20% exhibit behaviors associated with device intelligence collection or potentially risky activity.

Why this matters: Food and drink apps are often treated as low-risk consumer conveniences, but they routinely handle payment details, delivery addresses, loyalty accounts, order history, location data, and personal preferences. Because these apps are not typically viewed as enterprise risk priorities, they may receive less scrutiny from security teams. The data shows that even everyday applications can introduce exposure through weak communications, backend data handling issues, third-party dependencies, and access to personal information that can be useful for fraud, phishing, or customer profiling.

Lifestyle Apps: The business of personal data.

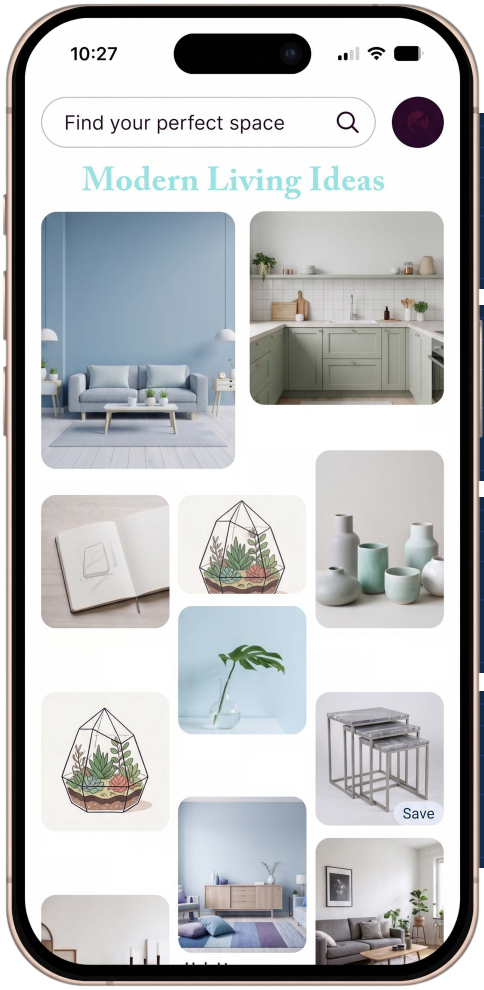

46% lack commercial code protection

66% of lifestyle applications request dangerous permissions

64% exhibit weaknesses in how they protect data during network communications

19% exhibit cryptographic weaknesses—the highest rate among all categories analyzed.

Why this matters: Lifestyle apps derive their value from user behavior, location, preferences, and routines, the more they collect, the more precisely they personalize recommendations and advertising. Yet this category shows some of the highest rates of cryptographic weakness and among the lowest levels of application protection. The apps collecting the most personal data are among the least resilient to reverse engineering, extraction, and tampering.

Customer-facing apps built by enterprises run largely on code the enterprise didn't write. 60% of third-party SDKs inside these apps are closed source, and most ship with no SBOM, no inventory of what's actually in them. This isn't a source code scanning or SCA problem anymore. Those tools need visibility to work; yet that visibility does not exist.

Enterprises are rapidly shipping their most exploitable apps ever, knowing less about their contents than ever before.

Regulatory Changes

In the past year there have been a number of publications from government agencies worldwide that range from strong suggestions to mandates for securing mobile devices. In November 2025 the US government's DISA issued an updated STIG mandating use of MTD on iOS and Android devices. Days later the French government issued guidance on securing mobile devices, which we wrote about here.

In June 2025 the UK government issued a comprehensive report on anti-fraud strategy for the remainder of this decade. In that guidance, they cite the need to secure mobile devices and the apps they run as priorities. They noted that the threat is accelerated by attacker use of AI and strongly suggested three actions that security leaders must implement:

1. **Detect fraud in mobile finance apps in real time**
2. **Know if the mobile app has been repackaged or cloned**
3. **Defend against AI-generated social engineering on mobile devices**

Finally, the Five Eyes[26] governmental security agencies recently published a call to action on Frontier AI because of how its use is accelerating cyber attacks.

Conclusion

Mobile is now the largest, most vulnerable attack surface in the enterprise. Global threat trends show cybercriminals have a mobile-first strategy, moving to devices and apps because that's where enterprise defenses are weakest and least monitored. AI didn't create this shift. It industrialized it, making mobile attacks faster, easier to scale, and harder to detect.

What You're Defending Against

Every app on an employee's device is a potential path to credential theft and data loss. Work and personal apps have the same exposures: insecure data transmission, regulatory compliance violations, AI-embedded apps leaking data. Adoption of AI within these apps is up 14x on Android and 7x on iOS. There is no safe category left.

The same exposures are seen within the apps that enterprises build in-house. AI ships code faster than security reviews can keep up, so apps launch well before runtime protections can be added. Each build brings new risks, with little to no security review taking place.

Attacker tools are evolving just as fast. Malware analyzed this year rewrites its own code on every run, generates its payloads dynamically, and in at least one documented case, was built entirely by an AI agent. The malware isn't waiting for you to catch up; it's already rewritten itself twice since you started reading this.

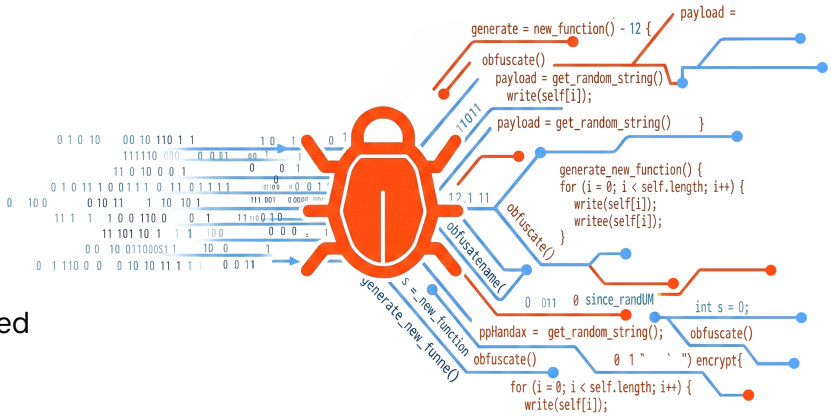

Where Governance and Compliance Have to Catch Up

None of this is covered by the policies most enterprises actually maintain. Visibility stops at the device name and OS version, not at what an app does, what it leaks, or whether it quietly added an AI model nobody reviewed. You can't govern what you can't see, and right now, most enterprises can't see any of it.

Regulators are moving regardless. DISA, the UK, and France have already shifted from suggesting mobile security to mandating it, covering both device protection and mobile app fraud.

Mishing, malware, shadow AI, and vibe coding all point to the same conclusion. AI has made mobile attacks faster, more convincing, and easier to scale, while the governance enterprises rely on hasn't kept pace. AI-driven attacks require a new paradigm of mobile visibility and protection.

Contact us today to schedule your mobile security assessment. For more information on how Zimperium can help you protect against these AI-generated threats, please visit zimperium.com for the latest mobile threat intelligence.

Authors and Acknowledgements

Gianluca Braga
Nico Chiaraviglio
Rajat Goyal
Ignacio Montamat
Pablo Morales
Asaf Peleg
Vishnu Pratapagiri
Tim Roddy
Santiago Rodriguez
Fernando Sanchez
Esteban Tissot
Krishna Vishnubhotla
Aazim Bill SE Yaswant

All data cited in this report is from zLabs unless otherwise noted.

About Zimperium

Zimperium is the world leader in AI-empowered mobile security. Purpose-built for mobile environments, Zimperium provides unparalleled protection for mobile applications and devices, leveraging the power of AI to deliver autonomous security that counters evolving threats including mishing, malware, and zero-day attacks.

ZIMPERIUM®

Disclaimer

Sources

1 "The 4 Work Modes Framework to Enhance the Digital Employee Experience", Gartner 13 May 2025 - ID G00829032
2 IDC MarketScape: Worldwide Rugged Mobile Devices 2025–2026 Vendor Assessment
3 https://www.strongestlayer.com/blog/ai-generated-phishing-enterprise-threat
4 Microsoft Digital Defense Report 2025
5 https://www.knowbe4.com/hubfs/Phishing_Threat_Trends_Report_Vol7_en-US.pdf
6 Harmonic Systems Research, January 19, 2026
7 "GenAI Code Security Report," Veracode, October 2025.
8 https://www.theregister.com/software/2025/09/05/ai-code-assistants-improve-production-of-security-problems/329883?ref=dailydev
9 "Innovation Insight for AI-Native Software Engineering," 13 January 2025, - ID G00822558, Gartner.
10 Verizon DBIR 2026
11 KnowBe4 Phishing Threat Trends Report, Vol 7
12 Microsoft Digital Defense Report 2025
13 https://alluresecurity.com/blog/smishing-statistics/
14 https://www.microsoft.com/en-us/security/blog/2026/04/30/email-threat-landscape-q1-2026-trends-and-insights/
15 https://www.microsoft.com/en-us/security/blog/2026/04/30/email-threat-landscape-q1-2026-trends-and-insights/#qr-code-phishing-attacks
16 https://cloud.google.com/blog/topics/threat-intelligence/threat-actor-usage-of-ai-tools
17 Verizon 2026 DBIR
18 https://www.atlanticcouncil.org/wp-content/uploads/2026/03/Mythical-Beasts-Investigating-the-Role-of-Intermediaries.pdf
19 https://cloud.google.com/blog/topics/threat-intelligence/2025-zero-day-review
20 https://lp.zimperium.com/2026-mobile-banking-heist-report
21 https://www.verizon.com/business/resources/reports/2026-dbir-data-breach-investigations-report.pdf
22 Zimperium Global Threat Research
23 The 4 Work Modes Framework to Enhance the Digital Employee Experience", Gartner 13 May 2025 - ID G00829032
24 IDC MarketScape: Worldwide Rugged Mobile Devices 2025–2026 Vendor Assessment
25 "Innovation Insight for AI-Native Software Engineering," 13 January 2025, - ID G00822558, Gartner.
26 Australia, Canada, New Zealand, the United Kingdom, and the United States