ADVANCED MOBILE
THREAT REPORT 2026

iVerify.

From Nation-States to Credential Thieves

How iOS exploit kits and Android malware went from targeted to widespread

Contents

Executive summary 3

ANALYSIS

01 Overview: state-grade tools, criminal scale 6

02 Prevalence: nearly 1 in 900 iOS consumer scans shows signs of spyware 10

03 Who is being impacted: criminal adversaries cast a wide net 12

04 Delivery: attacks arrive through ads and forgotten web code 14

05 The App Store: an exploitation toolkit inside an approved app 16

06 Blast radius: each compromise now exposes more 17

07 Android: the year of the RAT 19

08 Detection: attackers leave less evidence, so detection must happen on the device 22

09 Recommendations: what security teams should do now 25

10 Outlook: the assumptions that no longer hold 26

About the data 27

From Nation-States to Credential Thieves

This year, adversaries put leaked iOS exploit toolkits to work at scale, delivering attacks through online ads, reaching ordinary employees, and erasing their traces before defenders could respond. The leaks of DarkSword and Coruna placed that capability in the hands of criminal organizations worldwide. Attackers no longer need a government's budget to break into iPhones — AI-based vulnerability hunting levels the playing field. In iVerify's detections, the people they reached were not executives but rank-and-file employees swept up in broad, untargeted campaigns delivered through websites, ads and, in at least one case, an App Store app. Android has long faced the same reality, with a steady supply of cheap, ready-made kits that let criminals take over phones remotely.

1.07
infections per 1,000 iOS consumer scans (0.7 with enterprise included)

75%
of detections were criminal, not state-sponsored

100%
of enterprise detections from criminal actors were rank-and-file employees

~65
estimated infected phones at an average-sized Fortune 500 company, at the consumer rate

Below are the top-line takeaways from our 2026 detections and research. Read on to understand implications for the mobile security landscape and our actionable insights for defenders as they set budgets and priorities for 2027.

- **1.07 infections per 1,000 devices.** This is the prevalence of iOS spyware across iVerify's consumer user base, or at least 1 in 900. We report the conservative lower end of our estimate. Our early 2025 study found about 1.5 per 1,000, and our consumer user base has more than tripled since, which strengthens the finding. The rate has barely moved and we're seeing confirmed attacks weekly, many targeting on-device crypto wallets and delivered via watering hole attacks.

- **75% of detections were criminal.** Three in four of this year's detections were criminal adversaries stealing credentials and cryptocurrency, not state-sponsored actors. Criminal attacks have overtaken the abuse of mercenary spyware (commercial spyware sold to governments) as a primary mobile threat.
- **100% of enterprise detections from criminal actors were rank-and-file.** Every enterprise detection from campaigns launched by criminal actors involved an employee at or below managing-director (MD) level. Broad campaigns, some leveraging adtech, sweep up ordinary staff rather than singling out leaders or "high-value targets". Mercenary spyware is also moving beyond national-security targets.
- **Each compromised phone exposes more.** Between raiding iPhone keychain stores, messages, notes files and crypto wallets, DarkSword and Coruna go way beyond stealing 2FA codes and session tokens. This year marked the first demonstration of wormable malware that can expand the blast radius even further by automatically spreading to other devices in a victim's contact list.
- **Over 160 vulnerabilities in one Apple update.** iOS 27.2 contained 163 security fixes, the most ever. Mobile software is not exempt from the Bugpocalypse and the marked increase in the vulnerabilities fixed in each iOS release demonstrates AI's impact on accelerating how quickly flaws are found.
- **AI is lowering the bar for mobile exploitation.** Frontier LLMs favor the attacker and vulnerabilities are exploited faster than patches can be deployed. Not every AI-built attack works; there's limited training data for mobile exploits and models aren't quite good enough to make up for a lack of mobile-specific exploit development skills but that isn't stopping criminal actors from trying and in some cases, succeeding.
- **0 Apple Threat Notifications.** To our knowledge, not one device we detected under DarkSword or Coruna attack received an Apple Threat Notification (ATN), the warning Apple sends to users it believes are being successfully or unsuccessfully targeted. ATNs appear overly focused on nation-state threats like Pegasus and Paragon, well-known commercial spyware tools.
- **Forensic evidence is disappearing.** Adversaries now clean up more aggressively, and the traces that remain age off the device faster, so the evidence effectively goes dark. By the time an ATN is received, high-confidence data has disappeared.
- **Android-focused threat actors lean on social engineering.** Where DarkSword and Coruna exploit flaws in iOS, every Android RAT we observed gets in by persuasion and navigating victims through the "allow restricted settings" step Android implemented as an extra security layer since version 13. That said, other advanced Android exploits like LandFall did represent a step change in sophistication targeting that platform.

KEY INSIGHT

“In short, mobile exploitation has gone mass-market. Criminals now reach ordinary employees at scale, silently, through ads, the App Store, and cheap kits.”

Why it matters

The old threat model no longer holds. Victims are overwhelmingly ordinary employees at their desks and in their homes, not just executives or travelers. Executives are still targets, but the risk now reaches well beyond the VIP list. Confirmed attacks on enterprise employees’ phones are consistently routine in iVerify’s detections.

Infections are silent. The attacks typically run in the background with no visible sign to the user. Without detection on the phone itself, an organization has no reliable way to know a device was compromised, and after the fact forensics are unlikely to see anything due to more aggressive cleanup behavior.

Corporate access is exposed. Exploits can pull work passwords and sign-in tokens from the iPhone keychain, and personal phones that skip updates leave the door open longer. On Android, 22% of enterprise devices iVerify monitors allow sideloading, 60% run a security patch more than 90 days old, and 68% are not on the latest version. Furthermore, containerization won’t save an enterprise. Exploits like DarkSword can puncture right through common data governance tools on mobile.

After-the-fact signals arrive too late. Attackers clean up more aggressively and traces age off devices faster, so investigating after the fact turns up little. Even when iVerify investigated a dozen Apple Threat Notifications this year, the evidence that remained was limited. To our knowledge, no device we confirmed as infected with DarkSword or Coruna ever received an Apple Threat Notification.

What this means for defenders

Mobile devices now warrant the same endpoint detection and response coverage as laptops and servers. The findings point to three priorities: keep iOS and Android devices updated, restrict sideloading, and detect attacks while they are running via true mobile EDR, before the evidence ages off.

CHAPTER 01 · OVERVIEW

State-grade tools, criminal scale

For more than a decade, the security of the iPhone rested on an economic argument as much as a technical one. A working iOS exploit chain took rare skills, months of work and millions of dollars to build, so it was used sparingly against a short list of high value targets, where the risk of exposure was worth the cost of exploit development. Everyone else could reasonably assume that installing updates was enough, and most enterprises left phones out of their endpoint detection and response programs. As a result, few organizations were looking for attacks on phones at all.

This year, attackers put leaked iOS exploit toolkits to work at scale. They delivered attacks through online ads, reached ordinary employees, and erased their traces before defenders could respond. The leaks of DarkSword and Coruna, combined with AI vulnerability hunting, put that capability in the hands of criminal organizations worldwide. Attackers no longer need a government's budget to break into iPhones, and the people they reach are not executives but rank-and-file employees swept up in broad, untargeted campaigns. Exploitation toolkits turned up inside at least one App Store app, and Android faces the same reality, with cheap, ready-made kits that let criminals take over phones remotely.

AI is accelerating the trend. Frontier models that favor the attacker are getting better by the day at both finding vulnerabilities and writing code that turn them into exploits. As a result, software vendors including mobile OEMs are increasing both the frequency of updates and the number of vulnerabilities patched in each. For now, LLMs are not good enough to allow anyone with the resources and desire to retool these kits or develop new exploits, but some of the more knowledgeable criminals have been successful thanks to the headstart they got from two leaked nation-state grade iOS exploit kits.

The exploit ecosystem: DarkSword and Coruna opened Pandora's box

iOS exploits are proliferating fast. Since the leaks, criminal groups have copied and modified both toolkits and turned them into N-day campaigns against devices that haven't been updated.

Desktop computers went through the same transition years ago; once powerful attack tools spread, they proliferate and cannot be pulled back.

From state tools to criminal tools

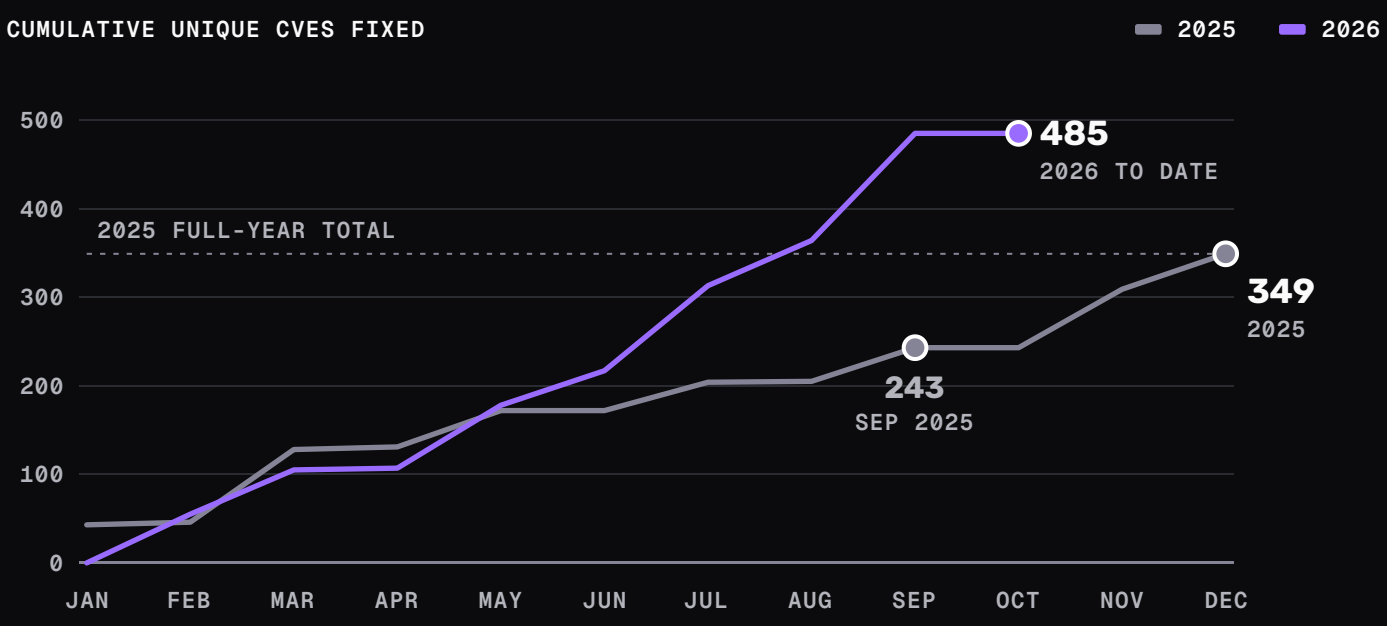
Google Threat Intelligence Group (GTIG) and iVerify disclosed Coruna on March 3, 2026: "five full iOS exploit chains and a total of 23 exploits" for iOS 13.0 through 17.2.1. GTIG traced it from "a customer of a surveillance vendor" to a suspected Russian espionage group and then to a financially motivated actor in China serving it from fake finance and crypto sites. Fifteen days later, GTIG, Lookout and iVerify disclosed DarkSword, covering iOS 18.4 through 18.7, and within a week a newer version appeared on GitHub making them extremely easy for anyone with no expertise to repurpose.

The leak mattered for more than the code. What had kept iOS exploitation out of reach was the set of private techniques needed to turn bugs into a working chain. With Coruna and DarkSword circulating, those techniques are now, as iVerify's researchers put it, "in the training data." By September of 2026, Coruna attacks against devices on vulnerable iOS versions reached 5% of devices, up from 1.5% in August, and iVerify had seen multiple clusters of variants of DarkSword. The most capable is P7, a fork that strips out the original's debug logging, adds two-way command and control and adds commands to steal additional wallet data, Apple Notes and photos. Attackers no longer need to find new bugs (though with AI even that is much easier). They need to find phones that haven't been patched, and there are plenty.

Windows went through this after the Shadow Brokers leaked EternalBlue in April 2017, a month after Microsoft had patched it. WannaCry used it to cause "tens of thousands of infections in over 150 countries," and NotPetya followed weeks later. The evolution of nation-state tooling into ransomware helped push enterprises from antivirus to endpoint detection and response on desktops and servers. On the phones in the same employees' pockets, comparable coverage is still uncommon.

AI acceleration

Apple's patch volume has doubled. By September 28, 2026, its iOS and iPadOS updates had fixed 485 unique CVEs, against 243 by the same point in 2025. iOS 27 alone fixed 145 vulnerabilities.



2026 runs to Oct 16. CVEs deduplicated within each year across all iOS and iPadOS branches.

Figure 1. Apple fixed more iOS vulnerabilities by August 2026 than in all of 2025. Cumulative unique CVEs in iOS and iPadOS security releases, by month.

AI-assisted research appears to be a major driver of that increase. Apple's security notes now credit AI tools for findings, and in June 2026, Apple capped how many open reports a single researcher could have, citing a surge in AI-generated submissions that included both real flaws and low-quality reports. Every fix also tells an attacker where to look. iVerify has seen AI applied to the leaked code directly, including multiple unsuccessful, likely LLM-assisted attempts to update DarkSword for iOS 26.x.

An AI-assisted port that failed

After iVerify published the P7 research, claims spread online of a DarkSword variant that could infect iOS 26.5. Our researchers tested the published samples against devices running iOS 26.0 and later, and none worked. The key files, named for the versions they claimed to target, were byte-identical copies of the iOS 18.7 files, which could never run on newer iOS. The pattern suggests someone used AI to try to port DarkSword without the expertise to guide it. P7 shows the other outcome; its changes reflect skilled operator work. AI can make skilled operators faster, but in the samples we've examined, it hasn't yet produced a working chain for current iOS on its own.

Attackers have the upper hand

Mobile is one of the most impacted vectors from AI because security teams have traditionally had the least visibility there. AI has collapsed traditional patching cycles and despite vendors' best effort, vulnerabilities are exploited faster than patches can be deployed. For phones, this turns one of the most ubiquitous enterprise endpoints into its most vulnerable.

Because LLMs excel at tasks with deterministic success criteria, such as whether a bug is triggered or an exploit works, they are well suited to offensive work. An exploit's utility matters more than its quality, making the offensive use case much more effective than defensive work such as patching bugs. There are no perfect oracles for code or patch quality, and we cannot verify state correctness across every code path. Defenders will use AI for patching, but quality controls and the need for a human to take responsibility for each patch and ship it will remain bottlenecks.

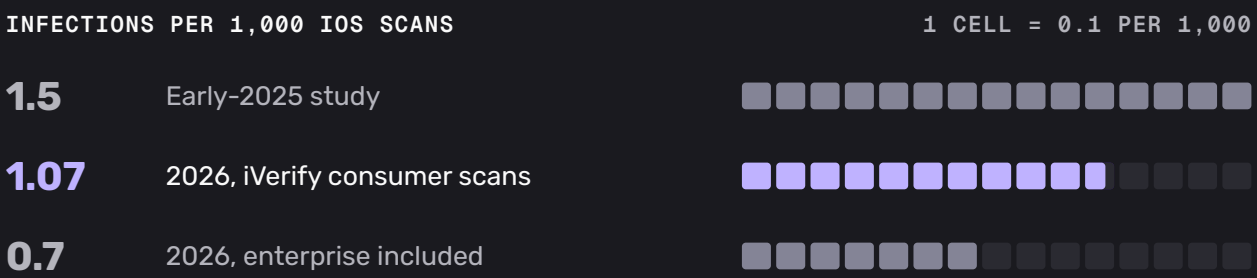
There are millions of existing bugs in critical attack surfaces including mobile software, and humans and LLMs continually introduce new ones. Foundational software projects on which modern society depends were written in memory-unsafe programming languages, with low unit-test coverage and rapid development-to-delivery cycles. This is exactly what led to the PHP meltdown in October 2026 that impacted 70% of websites around the world and also a factor in the development of the first-known, zero-click wormable mobile malware that leveraged such a vulnerability to spread through WeChat.

Continuing this cat-and-mouse game will not work. Technical debt must be addressed, software vendors need to reduce attack surfaces, implement testing, establish KPIs for test and security-test coverage, shorten patch cycles, and support anomaly monitoring with EDR software that has long been standard across desktops. On-device visibility provides the only meaningful assurance against compromise in the gap between vulnerability discovery, patch release and deployment, which can be weeks for large enterprises.

CHAPTER 02 · PREVALENCE

Nearly 1 in 900 iOS consumer scans shows signs of spyware

1.07 infections per 1,000 iOS consumer scans in 2026



Source: iVerify telemetry. The 1.5 and 1.07 rates describe iOS consumer scans; the 0.7 rate also includes devices monitored for enterprise customers. All rates are at the conservative lower end of our estimate.

iVerify's 2026 data puts the prevalence of iOS spyware at 1.07 per 1,000 devices across iVerify's consumer user base, nearly 1 in 900, at the conservative lower end of our estimate. The 2026 figure includes newer threats such as Coruna and DarkSword. Our earlier studies found higher rates in smaller samples that skewed toward people who already suspected a problem. As our consumer user base has more than tripled since early 2025, the rate has settled near 1 per 1,000 rather than continuing to fall.

With enterprise customers included, the rate drops to 0.7 per 1,000 devices monitored.

That offers some reassurance that enterprise security policies, such as enforced patching, on-device detection and network security controls, can meaningfully reduce the risk of iOS spyware, though not eliminate it.

Targeted then, indiscriminate now

iVerify's detections prior to 2026 were mainly mercenary spyware, aimed at specific chosen targets. This year's Pegasus detections are overwhelmingly older indicators, mostly from 2021 to 2023, that our updated detection logic surfaced during 2026. Every detection we could date to a 2026 infection month was Coruna or DarkSword, kits deployed against whoever happened to visit a compromised page or see a malicious ad.

These numbers are likely an undercount. Some DarkSword indicators exist on a device for only a very short time, and a scan that runs after they disappear may miss them.

CHAPTER 03 · WHO IS BEING IMPACTED

Criminal adversaries cast a wide net

Attacks became mass-scale and opportunistic this year, so these dragnet campaigns now routinely sweep up employees and the devices they use for work. We observed attacks reaching employees in a variety of industries including finance, technology and defense sectors.

Criminals lead

Three in four (75%) of this year's detections were criminals stealing logins and cryptocurrency, making them a primary mobile threat, ahead of mercenary spyware. When an exploit chain cost millions, it was rationed for targets someone had carefully chosen. Once it is cheap or free, there is no reason to ration it. Coruna was served from compromised Ukrainian websites and then "a very large set of fake Chinese websites mostly related to finance", and DarkSword was served to every visitor of compromised news and government sites whose iPhone ran a vulnerable version of iOS.

Ordinary employees are the ones getting hit

Every enterprise detection of Coruna and DarkSword iVerify recorded in 2026 involved an employee at or below managing-director level. In none of our confirmed cases have we found evidence that an executive or VIP was targeted. That does not mean executives are safe—they are still in the crosshairs for mercenary spyware as proved by the increasing frequency of Apple Threat Notifications—but they are more likely to already have managed devices, enforced updates and closer attention from security teams, which leaves the rest of the workforce as the easier and less-watched target. Only focusing on "high-value targets" or scrutiny during sensitive travel will cause security teams to miss the more statistically significant threat posed by ordinary employees.

The organizations where these attacks landed lean heavily on bring-your-own-device (BYOD) programs.

Traditional mobile security tools tie protection to device management, which means enrolling personal phones, controlling settings and, in many cases, seeing more of the device than employees are comfortable with. Faced with that, many organizations choose a lighter touch. Their mobile device management (MDM), if it's even used, does not force iOS updates on personal phones, and employees can stay out of compliance for months without consequence.

The result is a weaker security posture on a large share of the devices that touch corporate data, and those are the same phones these campaigns keep landing on. The trade-off comes from the tools, not from privacy itself.

Valid credentials are a prime target for criminal groups including Scattered Spider and ShinyHunters that most often pull them from mobile phones via social engineering as mobile devices sit at the nexus of authentication. Coruna and DarkSword give cyber criminals another credential-pulling tool in their arsenal, with both kits having the ability to pull passwords, notes and messages.

Ordinary people, not obvious targets

Conventional guidance on spyware risk focused on people with a public profile, a sensitive job or a role that put them on a government's radar. Some Apple Threat Notification cases iVerify saw this year don't fit that pattern.

In one case, Apple Threat Notifications reached the children of a city bus driver and a schoolteacher. In Australia, one went to a limousine driver. None of them had a public profile, a sensitive job or an obvious connection to anyone worth spying on.

These cases matter because they show the threat reaching people who had no reason to expect it.

> **For an enterprise, the employees most likely to be hit look much more like the limousine driver than like the CEO.**

CHAPTER 04 · DELIVERY

Attacks arrive through ads and forgotten web code

Attackers increasingly deliver iOS exploits through online advertising and the web's tracking and analytics plumbing. A targeted display ad can carry a half-click attack in which the exploit starts as soon as the ad loads, so the victim never has to tap anything. iVerify observed one served on Instagram while the user was simply scrolling.

Why attackers are turning to ads

A watering hole only reaches people who visit that site, and many of them will be on iOS versions the exploit can't touch. Ad platforms let a buyer choose who sees an ad by device, OS version, location and interests, so an attacker with an exploit for a specific iOS version can show it mainly to iPhones running that version, and to people interested in crypto, as draining crypto wallets was a primary objective of many DarkSword and Coruna campaigns we observed this year. This technique combines the reach of a mass campaign with the precision of a targeted one.

Citizen Lab's October 2026 AdFection report showed the technique in mercenary use. Ads from the Israeli ad tech firm Discover-Tech returned spyware-delivering code only to chosen devices, targeting at least 45 journalists at Al Jazeera, in attacks Citizen Lab links "with high confidence" to Candiru. iVerify's findings confirm that the technique has moved well beyond journalists and mercenary vendors, into criminal campaigns at mass scale.

CASE STUDY

2

Separate attacks on one ordinary person, each with its own DarkSword variant

One person, two attacks

iVerify telemetry recorded two separate attacks on the same individual, an ordinary person browsing the web with nothing that would make them a high-value target. The first came during an Instagram browsing session, through an ad served from ad-track[.]live. The infection began without the person clicking anything, and the payload was a DarkSword variant. The second came later, after the person clicked a link, and involved a different threat actor using a second, distinct DarkSword variant. In both cases, we assess the goal was cryptocurrency theft. The domains behind the ads were themed around this year's World Cup, bait chosen for reach rather than for any one person.

We assess that an ordinary person scrolling a social feed can be compromised by criminal actors, and that other ad-supported platforms may carry similar risk.

Forgotten web plumbing is being hijacked too

Ads are not the only shared infrastructure being abused. In October 2026, Scott Helme of Report URI documented how a group dubbed DeadTag purchased an expired e-commerce analytics domain, ecomtrack[.]io, on September 15, 2026 and turned into a vast campaign co-opting thousands of legitimate sites that combined receive 135 million visits a month to serve DarkSword and Coruna. Its tracking script was "still sitting in the theme of stores that had long since stopped thinking about it," so those online stores began loading code controlled by the new owner. A small redirector filtered visitors and passed real shoppers to traffic brokers, who sold them on to an investment scam, an online casino and a fake crypto trading site. That last site checked for Safari on iOS 18.4.0 to 18.7.2 and served DarkSword, in a build Report URI later matched to iVerify's P7 variant. The shoppers did nothing riskier than browse a store whose owners had forgotten a script in their theme.

CHAPTER 05 · THE APP STORE

An exploitation toolkit inside an approved app

App Store review has long been one of the strongest pillars of Apple's security model. Most iPhone users, and most enterprise policies, treat "only install from the App Store" as a sufficient rule. This year, an app on the store carried a working exploitation toolkit.

SEP 9
v1.1 on the App Store

SEP 12
v1.2 on the App Store

SEP 17
v1.3 removes the toolkit

FomoPeek Case Study

DISTRIBUTED
Apple App Store

BILLED AS
A "read-only monitor" for large crypto transfers

TARGETS
iOS 12.0–18.7.2 and 26.0–26.1

ATTACK METHODS
8

TRACED TO ATTACKER
About $580,000 in USDT

ANALYSIS
SlowMist and OKX security team

FomoPeek billed itself as a "read-only monitor" for large crypto transfers. Versions 1.1 and 1.2, distributed through the App Store on September 9 and 12, carried a hidden toolkit that, according to SlowMist and OKX's security team, used eight attack methods targeting iOS 12.0 through 18.7.2 and 26.0 through 26.1. Once it succeeded, it could escape the app sandbox and per SlowMist, "read and decrypt Keychain items, and pull files belonging to other apps," including crypto wallets. Investigators traced about $580,000 in USDT to the attacker. Version 1.3 removed the toolkit on September 17.

It is one case, not yet a trend, but it breaks a core assumption of mobile security policy. Victims did nothing more than install an app from the official store, and nothing in its listing signaled an exploit. Allow-listing App Store apps does not replace visibility into what actually runs on the device.

CHAPTER 06 · BLAST RADIUS

Each compromise now exposes more

A single infected phone now gives an attacker more to work with, and attackers' capabilities keep expanding. The criminal groups running these kits want what can be resold or spent.

Corporate logins

DarkSword extracts "device keychains" and "signed in accounts," and P7 packages keychain contents on the device before sending them out. A personal phone with corporate email, single sign-on, chat and VPN apps keeps credentials for all of them in that keychain, so an attacker who reads it can sign in to corporate services from anywhere, looking exactly like the employee. iVerify's guidance: "if a device used for work is compromised, treat every credential stored on it as compromised too."

Messaging apps

End-to-end encryption protects a message in transit, but once the operating system is compromised, an attacker can read conversations where they are stored on the phone. DarkSword collected Telegram and WhatsApp data from the start, and in September 2026, Coruna gained a Telegram implant module, with its decrypted configuration also listing WhatsApp and WeChat.

Crypto

Cryptocurrency is the fastest payout. A wallet can be emptied in minutes, and the transfer can't be reversed. Coruna's criminal payload scans for recovery phrases in text and images and targets 18 wallet apps, so a phrase saved as a screenshot or note on the same phone gives the attacker the wallet itself.

P7 added dedicated wallet commands, and the newer build Report URI recovered targets more than 25 crypto apps.

Self-spreading attacks

In September, iVerify analyzed WeWorm, “the first-known, mobile-based, zero-click worm,” developed by Calif and exploiting a flaw in WeChat. An attacker’s phone calls a WeChat contact and takes over their account while the call is still ringing, and the compromised phone then calls the next contact, across both iOS and Android. It was a research demonstration and has been patched, but Calif’s team, working with AI, went from bug to working exploit in about two days. Inside an organization, a worm that moves through contact lists moves through the team.

New protections, new bypasses

In May 2026, Calif detailed a working exploit chain against macOS with Memory Integrity Enforcement (MIE) enabled in five days. Just a few months later in October, press reports described its untethered jailbreak of an iPhone 18 Pro on iOS 27.0.1, the newest hardware and software. Techniques from the leaked kits have also revived public jailbreaks on older iOS versions.

CHAPTER 07 · ANDROID

The year of the RAT

On iOS, the story of 2026 is leaked exploits. On Android, it is the remote access trojan (RAT), malware that lets an attacker control a phone from afar. Criminals can now rent one off the shelf.

Cheap, turnkey kits are widely available

Most of these kits are sold as malware-as-a-service. A subscriber gets a builder that produces the malicious app and a web panel to run attacks, so buyers need little technical skill, and the prices are set for criminals, not governments. iVerify's threat research team has documented a steady stream of them.

Kit	How it's sold	What it does
HyperRat	Russian-language cybercrime forums	Custom malicious app plus a panel offering live screen control (VNC), call and SMS log theft, and mass phishing messages sent from the victim's own phone
Cellik	Dark web, about $150 a month	Wraps its payload inside real Google Play apps so a trojanized app looks legitimate; streams the screen in real time and intercepts notifications, including one-time passcodes
ZeroDayRAT	Openly on Telegram	Advertised to work on Android 5 through 16 and iOS up to 26; live camera and microphone, screen recording, keylogging with app context, and bank and crypto stealer modules
Oblivion RAT	Subscription, from about $300 a month	Generates fake Google Play update pages to trick users into enabling Accessibility Services, then adds remote control and keylogging; a "Wealth Assessment" feature sorts the victim's apps into banks, crypto and government so operators can prioritize
Octagon	$1,400 a month from a Russian-speaking developer	Intercepts SMS one-time codes, captures PINs, passwords and unlock patterns on supported devices, and overlays fake login screens on banking and crypto apps; it went undetected by Google Play Protect during testing.

iVerify threat research. Prices and capabilities as advertised by sellers.

The common mechanism: Accessibility Service abuse

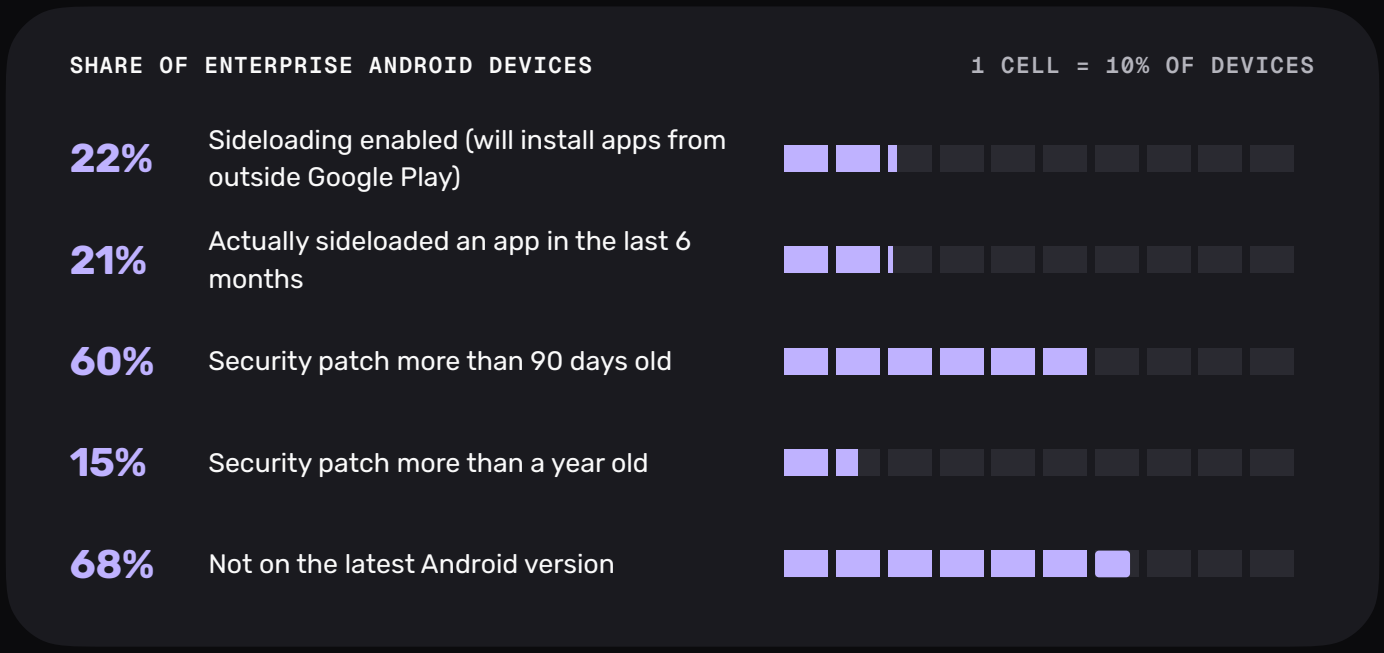
None of these RATs break Android. Where DarkSword and Coruna exploit flaws in iOS, every Android RAT above gets in by persuasion: a smishing link or Telegram post delivers an app from outside Google Play, a fake update page talks the victim through enabling sideloading, and a fake system screen talks them through granting Accessibility Services. Android has required an extra “Allow restricted settings” step for sideloaded apps since version 13, so the kits simply add that step to the script.

Once Accessibility is granted, the app can read everything on screen, tap on the victim’s behalf, draw fake login screens over real banking and crypto apps, and read the one-time passcodes that arrive by text. The builders ship lists of which financial apps to overlay, and Oblivion goes further by ranking victim value based on what they have installed.

Because nothing is exploited, there is nothing to patch, and Google Play Protect, which looks for known malicious code rather than for apps abusing permissions the user granted, reported no harmful apps while Octagon ran. In our 2026 telemetry, Play Protect flagged fewer than half of the Android devices on which iVerify’s signals found a malicious app.

Enterprise Android fleets leave the door open

iVerify's own fleet telemetry shows why these kits work.



Sideloading is a main delivery path for these kits, and more than 1 in 5 enterprise Android devices we monitor will accept an app from outside Google Play. Security patches are Android's fixes for known flaws, and most of the fleet runs a patch level that is months behind. These kits must be caught on the device at install time. Every app needs to be checked against known malicious packages and external malware intelligence, permissions assessed against surveillance and financial-fraud profiles, and have provisions to flag anything that arrives from outside Google Play, so a sideloaded "update" asking for Accessibility gets flagged before the operator ever connects.

Some phones ship compromised

Some Android phones are compromised before anyone opens the box. In October 2026, Bitdefender documented Midnight Mimosa, malware preinstalled in the firmware of budget MediaTek-based phones on "thousands of unique devices across more than 150 countries." Running as a trusted system app from first power-on, it fetched plugins from a server "disguised as a weather API," silently installed apps and briefly disabled the Play Store app during each install, likely to avoid detection by Google Play Protect, all to run ad fraud and turn the phone into a residential proxy. Bitdefender couldn't say where in the supply chain it was added, which makes device sourcing a security decision.

CHAPTER 08 · DETECTION

Attackers leave less evidence, so detection must happen on the device

Forensics, the investigation of a phone after an attack, has lost much of its value. Most of what the public knows about iPhone spyware came from that kind of work. Someone received a warning or noticed something odd, researchers collected a backup, a full file system extraction or a sysdiagnose, and then went looking for traces. It exposed Pegasus and much of the mercenary spyware industry. iVerify's researchers have described it as "slow, unscalable forensic investigation," and it worked largely because early implants were careless, leaving persistent artifacts with little obfuscation or cleanup. That is no longer true.

The going dark problem

Apple sends threat notifications to users it believes "may have been individually targeted" by mercenary spyware either successfully or unsuccessfully. Its October 2026 wave included users in 99 countries and came just 11 weeks after it alerted users in 110 countries in August, the first time we've seen notifications batches come in such quick succession. We investigated a dozen of these notifications this year and found limited evidence of exploitation, with no forensically sound, high-confidence evidence in any of them. Several of the people involved fit the profile of past mercenary spyware targets, some had received more than one notification, and some correlate strongly with earlier Pegasus activity. We have good reason to believe some were targeted, but the evidence that would have given us confidence was gone by the time we looked. The problem was not the warning. It was how little the attack left behind.

The reverse gap matters too.

To our knowledge, none of the devices we confirmed as infected with Coruna or DarkSword received a notification, indicating there are certain constructions of attacks that Apple cannot or is not flagging and reinforcing the need for continuous on-device monitoring to provide meaningful assurance against spyware compromise.

Mercenary spyware has not gone away. In October 2026, Citizen Lab tied attacks on Al Jazeera devices to Candiru. What has changed is how much these tools leave behind.

Many of the Pegasus indicators in our data date from 2021 to 2023 and survive only in fragments, because operators learned to clean up after themselves. iVerify has recorded no novel Pegasus detections since 2023. We read that as a visibility problem, not as evidence that Pegasus, Predator, Candiru or their peers have stopped operating. The tools are going dark, still in use, but much harder to see after the fact.

The 2026 kits push this further. Our threat operations team sees the same patterns again and again:

- DarkSword, as iVerify documented in March, has no persistence. It runs, takes what it wants and leaves, so a reboot or an update removes the running implant along with many of the traces an investigator would look for.
- The P7 variant removes the debug logging that made the original easy to study, cuts down its process injections and deletes some of its own working files once they have been uploaded. File-based indicators were never a strong foundation, since an attacker can rename or move a file, and now the files are being removed outright.
- The delivery infrastructure hides from investigators. Citizen Lab found that Discover-Tech's ad server returned an empty response to anyone who was not a chosen target. In the chain that served P7 through hijacked store scripts, Report URI found that datacenter and VPN addresses consistently received empty responses, and some residential visitors got a payload once and nothing after. By the time an investigator tries to reproduce an infection, the page serves nothing.
- What residue remains is usually there by accident. DarkSword was discovered partly because of "significant operational security (op-sec) failures" by its operators. Defenders cannot plan around attackers' mistakes.

A race against the clock

Some DarkSword indicators exist on a device for only a very short time. That turns detection into a race. Continuous on-device monitoring has to observe the indicator during the brief window it is present, or it may miss it entirely, and an investigation that begins days later has almost no chance of finding it.

In our assessment, our DarkSword numbers are a floor, and the real number is likely higher.

CASE STUDY

P7

DarkSword variant traced from a single alert, August 2026

From one alert to P7

In August 2026, iVerify received a DarkSword alert from the iPhone of a financial institution employee who had clicked a malicious ad. The detection differed slightly from the usual DarkSword detections, enough for our threat hunters to look closer. With the customer's consent, we collected artifacts from the phone, including suspicious files written to /private/var/tmp around the time of infection, and used them to hunt for related infrastructure in Validin. That led to a domain hosting another build of the same spyware family, a recovered sample and the P7 variant.

We published indicators of compromise in the full write-up, and a day later Report URI matched an implant served through hijacked store scripts to the same family, with two modules "byte-for-byte identical to iVerify's samples." That chain, from a small anomaly on one phone to a public analysis other defenders could use within days, depended on catching the attack while the evidence still existed. The customer's SOC was alerted before any lateral movement into its network.

Why OS-level visibility matters

Infections typically run in the background with no visible sign to the user, so without detection on the phone itself, an organization has no reliable way to know a device was compromised. Apple's Endpoint Security framework gives desktop EDR that visibility on macOS, but iOS does not offer a comparable framework to third-party tools. iVerify collects and correlates various on-device signals to achieve on-device detections to catch advanced attacks as they happen but the rest of the security community would benefit greatly from a plug-in that allows telemetry collection in real time. Speed decides the outcome for modern spyware investigations.

CHAPTER 09 · RECOMMENDATIONS

What security teams should do now

Defenders should treat employee phones as a front-line target, on par with laptops and servers.

01

Plan for mass campaigns, not only targeted ones.

Prepare for criminal campaigns against ordinary employees, not just state-sponsored attacks on executives or other high-risk staff.

02

Extend mobile security beyond a short list of leaders.

Ordinary staff are getting caught on iPhone and Android alike, so coverage limited to travelers or executives leaves most exposure unaddressed.

03

Close the patching and sideloading gaps.

Require iOS and Android updates and security patches, and restrict sideloading on Android. Tolerating out-of-date personal devices raises risk.

04

Detect on the device, not after the fact.

Post-incident forensics and after-the-fact signals miss attacks that leave little behind. Detection needs to run continuously on the phone.

05

Collect telemetry in real time.

If you cannot, collect it within hours, before the evidence ages off.

06

Treat a compromised phone as compromised credentials.

If a device used for work is compromised, rotate every credential stored on it, including single sign-on, email, VPN and messaging accounts.

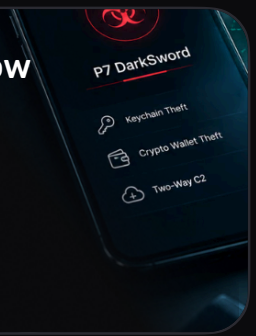
07

Make device sourcing a security decision.

Buy phones from trusted channels. Some budget Android devices arrive with malware preinstalled in their firmware.

From one alert to P7: how on-device detection caught a DarkSword variant

Read the research



CHAPTER 10 · OUTLOOK

The assumptions that no longer hold

For more than a decade, enterprise mobile security rested on three assumptions. Advanced attacks were rare and reserved for a short list of high-value people. Platform protections and app review were enough for everyone else. And if something did go wrong, an investigation would find it. This year's data challenges all three.

Attackers reached rank-and-file employees through ads, compromised websites and, in one case, an App Store app. Leaked iOS kits turned already-patched flaws into live risk for any device behind on updates, and on Android, rented trojans compromised phones without needing an exploit at all. And as attackers cleaned up after themselves and indicators grew shorter-lived, the evidence investigators once relied on was often gone before anyone looked.

Security updates continue to close gaps, but what has changed is the economics and the visibility. Capabilities once limited to governments are now cheap enough to use widely, and the traces defenders depended on are disappearing faster.

Several questions will shape the year ahead, and we will be tracking each of them in our research:

- Whether AI-assisted attempts to adapt leaked exploit code begin producing working chains for current iOS, rather than the failed ports we have seen so far.
- How far ad-based and web supply-chain delivery spreads, and whether ad platforms and site owners find ways to close those paths.
- Whether the Android trojan market keeps lowering its prices and skill requirements, and how much of it continues to depend on Accessibility Service abuse.
- How often compromised devices reach organizations through the supply chain, before they are ever switched on.

KEY INSIGHT

"For security teams, the practical question has shifted. It is no longer whether employee phones will be targeted, but whether anyone will be in a position to see it when they are."

About the data

Detection figures come from iVerify telemetry on consumer and enterprise devices, reviewed by our threat hunting team. Prevalence rates describe the population of devices scanned by iVerify users or monitored for enterprise customers and are reported at the conservative lower end of our estimate. Test devices, known false positives and consensual use of commercial data extraction tools are excluded. Apple vulnerability counts are unique CVEs listed on Apple's iOS and iPadOS security pages, deduplicated within each year. Android figures come from iVerify fleet telemetry on enterprise devices.

ADVANCED MOBILE THREAT REPORT 2026

iVerify.

iverify.com