

Threat Report

H1 2026

December 2025 – May 2026

(eset):research

Contents

Foreword	4
Threat landscape trends	5
Agentic AI skills: Small tools, big attack surface	6
ClickFix evolves: AI-fix and CrashFix expand the playbook	10
Stop before you scan: QR code phishing on the rise	14
PromptSpy: The first AI-powered Android malware	18
Mapping the world of EDR killers	21
Threat telemetry	24
Research publications	36
About this report	37
About ESET	38

Executive summary

AI threats

Agentic AI skills: Small tools, big attack surface

ESET analyzed 900,000 AI skills from popular repositories and found 25,000 suspicious and over 3,000 outright malicious skills.

Attack vectors Social engineering

ClickFix evolves: AI-fix and CrashFix expand the playbook

ESET detections of ClickFix have doubled as the technique expands from fake error prompts to browser environments, AI-themed help pages, and the workspace.

Email threats Phishing

Stop before you scan: QR code phishing on the rise

H1 2026 saw record levels of quishing attacks as scammers exploit the widespread adoption of QR codes.

AI threats Android

PromptSpy: The first AI-powered Android malware

H1 2026 brought us the first Android malware that actively uses GenAI at runtime.

Ransomware

Mapping the world of EDR killers

ESET tracks over 100 EDR killers used in the wild to kill, freeze, or blind security software that would otherwise detect the main payload during an attack.

Foreword

Welcome to the H1 2026 issue of the ESET Threat Report!

The first half of 2026 shows how attackers continue to improve the efficiency and scalability of their operations. Rather than relying on entirely new methods and tools, they are quickly adapting established techniques to new platforms, technologies, and user behaviors.

Artificial intelligence is playing a growing role in this development. In H1 2026, ESET analyzed nearly 900,000 AI skills – small functional components used by AI agents – and identified tens of thousands of suspicious and thousands of outright malicious instances. The number of AI skills within this new ecosystem is growing rapidly “as we speak”, further expanding the attack surface.

AI is also beginning to appear within malware itself. Shortly after the emergence of the first AI-powered ransomware in 2025, ESET researchers identified PromptSpy, the first known Android malware to use generative AI in its execution flow. The malware leverages AI – specifically, Google’s Gemini – to interpret user interface elements and adapt across devices and environments without relying on hardcoded behavior. While still rare, PromptSpy illustrates the potential for increased flexibility

in future threats – although guardrails against abuse included in LLMs are likely slowing down the adoption.

ClickFix – a social engineering technique leveraging fake error messages – has expanded beyond fake CAPTCHA prompts into AI-themed help pages, browser extensions, and cloud authentication scenarios. ESET detections of this vector more than doubled between H2 2025 and H1 2026, indicating sustained activity and adaptation.

Phishing campaigns are also evolving in response to user behavior. QR code phishing – also known as quishing – has reached record levels in ESET telemetry, with attackers embedding malicious links in QR codes to bypass cursory inspection and shift user interaction to mobile devices, while exploiting the implicit trust many people place in the black-and-white squares.

Last but not least, ransomware activity showed no signs of slowing down, with continued use of EDR killers – tools designed to disable security software during attacks. ESET Research has documented over 100 EDR killers used in the wild, with new variants appearing regularly. At the same time, data

from multiple sources shows that a declining share of victims are choosing to pay ransoms, suggesting some progress in mitigation and response measures.

I wish you an insightful read.

Jiří Kropáč
ESET Director of Threat Prevention Labs

Threat landscape trends



AI threats

Agentic AI skills: Small tools, big attack surface


ESET analyzed 900,000 AI skills from popular repositories and found that 25,000 were suspicious and over 3,000 were outright malicious.

AI systems are no longer limited to chatbots. Increasingly, AI agents can plan tasks, browse the web, interact with third-party services, write files, execute commands, and take actions for its users. This shift from conversational AI to autonomous, tool-using agents – popularized by projects such as [OpenClaw](#) and [Hermes Agent](#) – has created a new and rapidly expanding attack surface.

At the center of these new ecosystems are "skills": small add-ons or sets of instructions that direct an AI agent how to perform specific tasks, including which services or tools to use and what data to access. However, in some scenarios, skills can be abused – or even designed by an attacker – to exfiltrate data, download and execute malware, override user instructions, subtly alter an agent’s behavior over time, or even modify the skill itself.

According to ESET telemetry, the number of AI skills is growing sharply. Between March and May 2026, the number of unique skills scanned by ESET systems increased from an initial 60,000 to almost 900,000. Over the same period, skills considered suspicious grew from around 10,000 to over 25,000, while those blocked as malicious rose from approximately 600 to over 3,000.

ESET separates skills that appear benign from those requiring deeper analysis and then tags the individual files and scripts contained in each skill. These tags cover a wide range of capabilities, from basics such as communicating over the internet and creating or modifying files on disk to downloading third-party tools, embedding automation, accessing credentials, and using obfuscation or other nontransparent techniques.



Unique AI skills scanned by ESET systems per day, seven-day moving average

Among the skills selected for further analysis, particularly notable were cases involving third-party tool downloads and obfuscation (each seen in 31% of the analyzed set), as well as smaller but high-risk clusters involving credential loading (6%) or code injection into scripts (4%).

Many of these behaviors are not inherently malicious: a legitimate skill may need to automate repetitive actions or download third-party tools in order to function. However, risk increases when such capabilities appear in specific combinations, especially when the purpose is unclear or command execution, credential handling, or obfuscation are involved.

The share of selected tags in the evaluated skills:

- **84%** execute commands instead of the execution being user initiated,
- **39%** check for the existence of files on disk,
- **31%** download third-party tools,
- **31%** use nontransparent techniques or some form of obfuscation,
- **6%** load credentials from the system to use for authentication, and
- **4%** inject code into scripts.

From chatbots to autonomous operators

Typically, even if a chatbot produces harmful or misleading output, the human user usually remains responsible for taking the final action. With agentic AI, that relationship shifts and an agent is allowed to perform even sensitive steps such as sending emails, making payments, or executing scripts and interacting with other software.

From a high-level perspective, observed suspicious and malicious behaviors include:

- exfiltration of data,
- download and execution of malware,
- manipulation of sensitive systems,
- overriding instructions through prompt injection, and
- gradual changes in the agent’s behavior over multiple interactions before a suspicious or malicious agenda is revealed.

The last point is particularly important, as a malicious or compromised AI skill may not immediately perform an obviously harmful action. Instead, it can shape the agent’s behavior over time or wait until it is granted the desired access or permissions.

Malicious and high-risk skills

Among the AI skills analyzed by ESET, several fall into the malicious category. In one blocked case, the skill offered a command line interface for video evaluation using the Gemini API. On the surface, the purpose appears legitimate – the agent loads the API key from a configuration file and then uses the API to process video in a certain way. However, if the authors/ attackers decided to do so, they could alter the instructions and send the user’s API key to themselves or run their own tasks on the victim’s dime.

Another notable group of malicious skills was designed to perform red-teaming actions, such as running attacks against Active Directory, enumeration of data on the targeted system, exfiltration of credentials, or achieving highly privileged, persistent access. The instructions for the agent include the use of external offensive security platforms and tools such as Impacket, Mimikatz, or BloodHound, frequently observed in ransomware and infostealer attacks.

```
## Inputs/Prerequisites
- Kali Linux or Windows attack platform
- Domain user credentials (for most attacks)
- Network access to Domain Controller
- Tools: Impacket, Mimikatz, BloodHound, Rubens, CrackMapExec

## Outputs/Deliverables
- Domain enumeration data
- Extracted credentials and hashes
- Kerberos tickets for impersonation
- Domain Administrator access
- Persistent access mechanisms

----
```



Example of inputs and outputs of an analyzed **red-teaming AI skill**

In some cases, these skills also included advanced techniques and workflows that go beyond the basic administration or educational use claimed by the author. While such red team tools have legitimate use for authorized defenders, when used in combination with agentic AI, they can easily be abused even by low-skilled threat actors.

Gray zone: Access to money, and evolving behavior

While some skills provide real functionality, their future behavior is difficult to predict and even minimal changes can push them into harmful territory.

One such example is a self-modifying skill identified by ESET scanners. Its Chinese-speaking author designed it to create a persistence mechanism (JSON file) and a tool for self-modification (Python code), while allowing also external modifications. This can lead to unpredictable behavior of the agent or its abuse by an attacker.



Instructions of a self-modifying skill (top: original; bottom: translated)

In theory, this could help an agent improve performance or adapt to user preferences as they change over time. In practice, it makes security review more difficult: a skill that appears benign at installation may self-develop into something more

invasive or malicious – even without outside updates or actions.

Another example from the gray zone is Credit Claw. This skill enables the agent to make purchases through platforms such as Amazon and Shopify,

and to integrate with software-as-a-service (SaaS) environments. The core risk is the access to payment methods and in turn direct access to the user’s money. A harmful update or dependency could quietly alter that agent’s behavior and redirect the purchased goods or services to the attacker.

Browser extensions, mobile apps, and automation scripts have long raised similar concerns, but when it comes to handling of sensitive data, making purchases, running API calls, or instruction chains, the higher level of autonomy of AI agents increases the risk and scope of such attacks.

Benign but problematic: The false sense of security

Some skills are not malicious, but can still create issues: for example, by creating a false sense of security. Good examples of that are skills marketed as security scanners that implement only basic scanning techniques, resembling antivirus tools from the 1990s. In other words, they may detect obvious or already-known threats, but provide little protection against more complex, emerging, obfuscated, or context-specific attacks.

EXPERT COMMENT

AI skills can enable a wide range of agentic AI abuses, from automated reconnaissance and red-team-style attacks to spam generation, malware modification, and distribution. Adversaries will likely keep testing these approaches to bypass controls, including by obfuscating intent or using region-specific, niche, or constructed languages.

As agentic capabilities mature, they could power more advanced attack chains, from supply-chain compromise and typosquatting at scale to rapid adoption of sophisticated techniques emerging in the threat landscape. Combined with AI-enhanced phishing and spam, AI-driven malware generation or modification could make campaigns faster, cheaper, and harder to detect – making agentic AI abuse a key area for defenders to monitor.

Anton Mäčko, ESET Malware Analyst

Some of the “security” skills seen by [ESET AI Skills Checker](#) only act as local applications that reach out to online scanning services such as VirusTotal, checking hashes, URLs, or IP addresses against known reputational data. These capabilities can be useful in triage, but they are not equivalent to running a full scanning engine or performing behavioral analysis. What is also concerning is that the AI agents can present “safe” and “not detected” results with a high degree of confidence, making even wrong output look trustworthy to users.

Supply-chain risks and AI-fix

The AI skill ecosystem inherits many of the supply-chain risks seen in open-source software, browser extensions, npm packages, and developer tooling. Skills can depend on external libraries, scripts, APIs, models, or command line utilities, with each dependency introducing another potential point of failure or compromise.

ESET Research has also documented a specific version of the ClickFix technique – called AI-fix by ESET researchers – that abuses AI services and domains to spread malicious software. While the instructions used as part of that attack pattern usually target humans, they could easily be repackaged and used to train AI agents to run the compromises in an automated fashion. For more information on the ClickFix attack vector and its AI-fix evolution, read the [dedicated chapter](#) in this report.

Attack vectors Social engineering

ClickFix evolves: AI-fix and CrashFix expand the playbook


ESET detections of ClickFix have doubled as the technique expands from fake error prompts to browser environments, AI-themed help pages, and the workspace.

People are used to encountering problems online. A browser tab freezes, a file refuses to download, a utility throws an error, or a service offers an automated fix. Most users want the issue resolved quickly, and that moment of urgency gives attackers an opening. In H1 2026, threat actors continued to refine the efficiency of ClickFix, a social engineering technique built around a simple idea: present a fake problem, offer a fast fix, and trick the victim into launching the attack.


We first reported on ClickFix in H1 2025, when it predominantly relied on fake CAPTCHAs and rather simple faux error messages. Since then, attackers have switched to more advanced techniques, including sophisticated verification flows that steal authorization tokens. In H1 2026, ClickFix continued to spread across new environments and leverage new lures: extending to [macOS through commands that](#)

[supposedly install system utilities, compromising WordPress sites](#) to show ClickFix-style prompts to site administrators, and launching AI-themed waves. Attackers have also refined the social engineering layer, using fake blue-screen-of-death (BSOD) prompts, frozen document viewers, and service-specific error messages to increase the chances of successful compromise.

ESET telemetry shows that detections of ClickFix attacks (under HTML/FakeCaptcha) grew by 108% between H2 2025 and H1 2026. Although the current detection levels are lower compared to the initial ClickFix boom covered in [Threat Report H1 2025](#), the upward trend in 2026 shows that adversaries have not let up on their activity and continue to find new ways to employ ClickFix and related techniques in their compromise chains.



ClickFix detection trend in H2 2025–H1 2026, seven-day moving average



Geographic distribution of HTML/FakeCaptcha detections in H1 2026


It is important to note that the multiple stages of the attack – including the copied PowerShell commands or scripts, executables, malicious “envelopes”, and final payloads – are covered by dozens of other detections. Therefore, the real prevalence of this threat is probably even higher than the HTML/FakeCaptcha numbers. Countries reporting the highest volume of detections in ESET telemetry in H1 2026 are Japan (14%), Slovakia (7%), and France, Spain and Peru (each over 5%).

AI-fix rides the hype train

Adversaries keep the social engineering aspect of ClickFix up to date by adopting what we track as [AI-fix](#),

which exploits the current hype, growing popularity, and availability of generative AI tools. Attackers craft pages that abuse legitimate domains – for example, Anthropic’s Artifact pages, OpenAI’s Canvas, or Microsoft’s Copilot Pages – offering troubleshooting content for nonexistent issues. The user is led to believe that the instructions are AI-generated, which plays right into the attackers’ hands, as people increasingly tend to place their trust in such tools.

The instructions are really the same old ClickFix execution chain, wrapped in fake branding and a polished facade. As the use of AI is becoming more embedded in daily life, and the trust in these tools



A web page, abusing Anthropic's Artifact pages domain, with AI-fix instructions


increases, it gives the attackers a wide attack surface: our relationship with AI tools. AI-fix also demonstrates the versatility of ClickFix, showing how quickly and easily attackers can adapt the concept to fit user behavior and evolving technology.

CrashFix: Crashing the browser party

An ad-free browsing experience sounds appealing, and installing an ad blocker seems like a simple solution. Only, sometimes it leads to an uninvited guest, like [CrashFix](#), crashing your browser session.

This was the case with the malicious extension NexShield – Smart Ad Blocker, a copycat of the legitimate uBlock Origin Lite. Distributed through phishing and malvertising, it redirected unsuspecting victims to the official Chrome Web Store page, where positive reviews lent it an air of legitimacy.

Upon installation of the fake ad blocker and browser relaunch – however, with a clever 60-minute delay designed to weaken the mental link between the extension's installation and malicious behavior – the extension first displays an error message, prompting the victim to run a scan.



Fake warning displayed by the malicious extension to run a scan and initiate the CrashFix execution chain (image source: [Huntress](#))

The scan returns a fake security warning with a seemingly quick and easy fix – just follow the instructions on the screen: CrashFix execution steps.



Fake security issue prompt that urges the victim to run a malicious command (image source: [Huntress](#))

At this stage, the social engineering aspect is reinforced by a warning about the potential compromise of browsing data, adding a sense of urgency.

If the victim ignores the warning and tries to restart the browser, the same fake issue appears (browser stopped abnormally), leading to the same execution chain – until the malicious browser extension is removed.

Preying on frustration and the fear of data loss, CrashFix exploits the immediate availability of a seemingly quick and easy fix to a fabricated issue, and it foreshadows the possible evolution of this technique reaching other platforms: not just desktop browsers, but even mobile applications.

ConsentFix: Fake verification, real tokens

ConsentFix further demonstrates how attackers continue to [adapt the ClickFix playbook to new contexts and environments](#) – in this case, the enterprise cloud and the workspace. Just think of all the instances when you are asked to sign in with your Microsoft account: Copilot, Outlook, SharePoint, Teams – to mention only a few. Sign-in workflows have become a routine part of daily work. ConsentFix exploits this as an attack surface by combining ClickFix-style social engineering with OAuth authorization code theft, which allows adversaries to hijack Microsoft

accounts without stealing credentials or triggering multifactor authentication (MFA) notifications.

The victims land on compromised, but legitimate, websites – often distributed through phishing or search engine results – where they are presented with fake verification prompts, such as a fake CAPTCHA or Cloudflare Turnstile. As part of the faux verification process, victims are required to manually copy and paste a URL, containing their OAuth authorization code, into the phishing page. The attackers then use the OAuth token from the URL to gain access to the victim’s Microsoft account.

Often, the victims may have an active Microsoft session in the browser – in this case, no password entry or MFA may be required. This tactic is effective especially because the entire compromise chain takes place within the browser and relies on legitimate Microsoft infrastructure.

ConsentFix points to an emerging trend where attackers aim to steal tokens rather than credentials, and to the next level in social engineering as they are able to abuse legitimate Microsoft login pages.

EXPERT COMMENT

Attackers are doubling down on what works best, and they are finding new and more sophisticated ways to exploit the same underlying malicious tactics. Social engineering and exploiting human psychology remain the primary methods of initial compromise, now combined with trusted entities and actions: whether a popular AI tool, such as Claude, or a daily workflow, such as signing in through a Microsoft login page. Given the flexibility and effectiveness of ClickFix-style tactics, adversaries are likely to continue tracking authentication trends and adapting their methods to match user behavior and evolving platforms.

Ondrej Kubovič, ESET Security Awareness Specialist

Email threatsPhishing

Stop before you scan: QR code phishing on the rise

H1 2026 saw record levels of quishing attacks as scammers exploit the widespread adoption of QR codes.

By now, we all know that clicking links in random emails is risky business – so, we’ve learned to stop and think before we click. But, if a QR code lands in your inbox, do you stop and think before you *scan*?

As QR codes have become increasingly ingrained in everyday life – whether in restaurant menus, contactless payments, or hotel check-ins – many people have grown used to scanning them without a second thought. Cybercriminals, naturally, have jumped at the opportunity, and over the past few years, QR codes have become one of the go-to ways of delivering harmful links to potential victims.

In a Q1 2026 [report](#), Microsoft noted a 146% quarterly increase in QR code phishing, making it the fastest-growing email-based attack vector in its telemetry. ESET telemetry has recorded a steady increase in QR code phishing detections since the beginning of 2026, with April seeing the highest detection volume.

The potential of this technique didn’t go unnoticed by [advanced persistent threat \(APT\) actors](#) either, with some groups employing malicious QR

codes in spearphishing attacks. As a result, the FBI issued a [cyber alert](#) in January 2026 to warn of the use of QR codes in campaigns by the North Korea-aligned APT group Kimsuky.

Scammers have also been busy experimenting with QR code fraud in the physical domain, planting fraudulent QR codes where people might normally expect them – on [parking payment machines](#), [bike-share bicycles](#), or even on fake [parking tickets](#) and [toll violation notices](#). People who scan the codes to pay for a service or a supposed fine end up divulging their payment card details to fraudulent websites.

The quishing playbook

A typical QR code phishing attack starts with an individual – for example, a corporate employee – receiving an email in their work account. In a typical scenario, the email impersonates company communication, often abusing the branding of widely used corporate tools. Much like in most phishing attacks, the message often has an element of urgency and is personalized to the recipient in order to appear legitimate.



Example of a phishing email detected as QRCode/Phishing (screenshot redacted)

WHAT IS QR CODE PHISHING?

QR code phishing – also known as quishing – involves storing phishing URLs in QR codes. These attacks can occur virtually (via emails, messages, websites) or physically (QR codes planted in public spaces or placed over legitimate ones in places where QR codes are expected).

By placing malicious links inside QR codes – either directly in the email body or within attachments – attackers exploit the weaknesses of text-based scanning engines and steer victims toward phishing websites, often accessed via unmanaged mobile devices.


Instead of a clickable link, the email contains a QR code and prompts the recipient to scan it with their mobile device to complete a supposedly urgent and important task. Once the code is scanned, a link is displayed that leads the victim to a phishing site with the aim of stealing credentials or other sensitive data.

This technique is arguably more dangerous than “traditional” phishing because users might be less vigilant when scanning QR codes, as compared to clicking links. Moreover, scanning the code moves the attack to a mobile device, which may lack security controls or an endpoint security product that would otherwise halt the attack. At the same time, some

security solutions may fail to detect emails with malicious QR codes, even if the encoded URL would otherwise be blocked.


ESET telemetry insights

ESET tracks quishing emails under the detection name QRCode/Phishing. This detection works through a dedicated layer of the ESET email scanner, designed to identify QR codes in the vast majority of file types, and to decode the URLs in them. The extracted URLs are scanned using ESET anti-phishing, anti-malware, and anti-spam engines; any harmful URLs are blocked, and the associated emails flagged or deleted.



QRCode/Phishing detection trend from September 2025¹ to May 2026, seven-day moving average

¹ September 2025 is when we started tracking QR code-based phishing as a separate detection in ESET telemetry.



Geographic distribution of QRCode/Phishing detections in H1 2026


In H1 2026, approximately 11% of all detected phishing emails utilized QR codes. On average, we saw 100,000 such detections monthly in H1 2026, with the highest detection levels recorded in April, as can be seen in the chart on the left.

In H1 2026, QRCode/Phishing threats were most prevalent in the US (19% of detections), Spain (17%), and Mexico (6%); we also saw notable activity in the UK (5%), Czechia, Canada, Poland, and Italy (3% each).

HOW QR CODES WORK


A QR (Quick Response) code is a two-dimensional barcode that stores data in a grid of black and white squares. QR codes encode data (such as URLs, text, or payment details) using patterns of squares that can be read by a camera or scanner. The pattern is then converted into the original data, typically displaying a link or other text.

As shown in the examples on this page, these phishing emails commonly posed as corporate HR communication, especially related to pay and benefits – a topic likely to raise recipients' curiosity.



Example of a phishing email detected as QRCode/Phishing (screenshot redacted)

Important Payroll Advisory Regarding Hourly Wage Modifications – Review Required



Example of a phishing email detected as QRCode/Phishing (screenshot redacted)

How to stay safe

As convenient as QR codes are, they should be treated with great caution, especially if coming from unknown sources or found in public places.

To stay protected from QR code phishing:

- Use a multilayered security solution capable of detecting malicious QR codes in emails.
- If using an Android mobile device, keep it protected via a reputable security solution.
- Don't be *quick to respond* (pun intended). Stop and think before you scan – are there any red flags in the email/message that raise suspicion? If you do decide to scan, stop again to inspect the displayed link.
- When in doubt, verify the legitimacy of the message with the supposed sender outside of the suspicious email thread – for example, by contacting them via chat or by phone.

EXPERT COMMENT

While not a new technique, QR code phishing is entering a new phase of automation and scalability in 2026. Beyond the still considerable lack of user awareness, its effectiveness stems from a structural detection gap: the malicious link is encoded in an image, making it invisible to traditional email security controls – and equally opaque to the human eye, which has no way to view the link before scanning. This is further amplified by the fact that people tend to implicitly trust QR codes due to their widespread, legitimate use in public spaces and everyday interactions.

At the same time, this attack introduces a critical shift – it moves the victim from a relatively well-protected corporate environment to a potentially unmanaged mobile device, effectively bypassing multiple layers of enterprise security in a single step.

Attackers are rapidly refining this vector, using various evasion techniques to stay ahead of detection, while relying on the fact that users rarely verify QR code destinations. As a result, we can expect sustained high volumes of QR-based phishing – not just due to awareness gaps, but because many security models (both in email and mobile ecosystems) are not yet designed to handle QR-encoded threats at scale.

Dariusz Iwański, ESET Senior Detection Engineer

AI threatsAndroid

PromptSpy: The first AI-powered Android malware

H1 2026 brought us the first Android malware that actively uses GenAI at runtime.

Not long after discovering [PromptLock](#), the first AI-powered ransomware, ESET researchers found another threat with AI integration, this time developed for Android – [PromptSpy](#).


Ordinary facade

On the surface, PromptSpy is a rather typical remote access trojan, used to gather and exfiltrate information from victim devices. It can, among other things, intercept the lockscreen PIN or password, upload a list of installed apps, take screenshots on demand, and record video. PromptSpy

To remove PromptSpy in safe mode, you would typically press and hold the power button, long-press Power off, and confirm the Reboot to safe mode prompt (though the exact method may differ by device and manufacturer). Once the phone restarts in safe mode, you can go to **Settings → Apps** and uninstall the app in question without interference.

deploys a built-in Virtual Network Computing (VNC) module that grants the attackers remote access to the compromised Android device; thus, they can see its screen and take control of it in real time. Additionally, the VNC protocol is used for C&C communication, which is AES encrypted. PromptSpy also tries to prevent being uninstalled by abusing accessibility services to place invisible overlays on buttons such as **Stop** and **Uninstall**, making removal only possible via safe mode.

In our blogpost, we noted that we hadn’t seen any instance of PromptSpy in our telemetry data, which could be explained by the malware being a proof of concept. At the same time, the malware was being distributed via a dedicated website and impersonated the Argentinian branch of JPMorgan Chase Bank, suggesting real-world targeting. This is supported by the distribution website being `mgardownload[.]com` (offline at the time of analysis) and the app being named `MorganArg` (likely a shorthand for “Morgan Argentina”), both probably intended to convince the targets that the app is actually a legitimate JPMorgan Chase app. Since the publication of the blogpost, our telemetry has registered exactly one detection of PromptSpy, caught on February 22, 2026 in Ukraine.



Invisible rectangles (displayed in color for clarity) covering specific buttons

Due to the presence of both debug strings written in simplified Chinese and an unused function that returns descriptions of accessibility event types in Chinese, we believe with medium confidence that PromptSpy was developed in a Chinese-speaking environment.

AI integration

What distinguishes this Android spyware is that it abuses Google’s Gemini at runtime. PromptSpy employs this LLM to execute a gesture that allows the malware to become locked in the list of recent apps, thereby achieving persistence. Even though such use of GenAI seems relatively minor, it actually allows the threat actors to automate actions that would


normally be more difficult with traditional scripting. Android malware usually depends on interacting with the device’s UI via hardcoded taps, which are difficult to automate across all possible UI and OS versions available on Android devices. The gesture that PromptSpy aims to replicate is no exception, and that’s where GenAI comes into play.

The malware sends Gemini a natural language prompt along with an XML dump of every UI element on the screen: its text, type, and exact position on the display. Gemini processes this information and returns JSON instructions on the gestures that should be executed, and where to execute them. PromptSpy performs these gestures using accessibility services and sends

EXPERT COMMENT

An obstacle to AI-powered malware becoming more widespread might be that LLM models include protections against abuse. While vibe coding a piece of malware is basically a one-time task, deploying a solution that hinges on the LLM to cooperate every time the solution runs can fail once the model is updated. What I can see becoming more attractive to cybercriminals are tools that would either use GenAI to bypass security solutions or run as a thin client – having minimal malicious functionality in the malware itself, but be able to adapt to any device and dynamically download and run malicious code.

Lukáš Štefanko, ESET Senior Malware Researcher



PromptSpy’s execution flow prompting Gemini for a gesture to lock the malicious app in the recent apps list

```
public static void startAutomationLoop(AccessibilityService accessibilityService, String str, String str2) {
    String str3;
    String str4;
    JSONArray jsonArray;
    String str5;
    int i;
    int i2;
    boolean z;
    AccessibilityService accessibilityService2 = accessibilityService;
    String str6 = str;
    if (accessibilityService2 == null) {
        return;
    }
    int i3 = 4;
    String str7 = TAG;
    if (str2 == null || str2.isEmpty()) {
        ServiceInteractionUtil.ToLog(TAG, "未设置 Gemini API Key, 无法执行自动化任务", 4);
        return;
    }
    ServiceInteractionUtil.ToLog(TAG, "开始执行任务: " + str6);
    JSONArray jsonArray2 = new JSONArray();
    String string = accessibilityService2.getString(R.string.atuo_load_msg);
    int i4 = 1;
    boolean z2 = true;
    int i5 = 0;
    while (z2 && i5 < 30) {
        int i6 = i5 + 1;
        ServiceInteractionUtil.ToLog(str7, ">>> 步骤 " + i6);
        AccessibilityNodeInfo accessibilityNodeInfoGetActiveWindowNodeInfo = InputService.GetActiveWindowNodeInfo();
        String aIXml = AccessibilityNodeUtil.toAIXml(accessibilityNodeInfoGetActiveWindowNodeInfo);
        if (accessibilityNodeInfoGetActiveWindowNodeInfo != null) {
            accessibilityNodeInfoGetActiveWindowNodeInfo.recycle();
        }
        if (i6 == i4) {
            str3 = "You are an Android automation assistant. The user will give you the UI XML data of the current screen.
        } else {
            str3 = "The previous action has been executed. This is the new UI XML, please determine if the task is complet
        }
        addToHistory(jsonArray2, "user", str3);
        String strCallGeminiApi = callGeminiApi(str2, jsonArray2);
        if (strCallGeminiApi == null) {
            ServiceInteractionUtil.ToLog(str7, "API 请求失败, 终止任务", i3);
            AutoClicker.bringHomeToForeground();
            return;
        }
    }
}
```

Malware code snippet with **hardcoded prompts**

updated UI XML back to Gemini. The AI model then verifies whether the app has been successfully locked in the recent apps list. The process repeats in a loop until Gemini confirms the success of the operation.

In May 2026, GTIG published additional [research](#) into the malware. According to that blogpost, apart from locking the malicious app in the recent apps list, PromptSpy’s AI component was designed with broader capabilities in mind, mainly related to interface navigation and the interpretation of user activity. GTIG also noted that PromptSpy exhibits high operational resilience, with the threat actors being able to update the malware’s components (including Gemini API keys) at runtime via its C&C channel.

PromptSpy demonstrates how abusing GenAI can make malware more dynamic, allowing it to adapt to a variety of environments. However, since our discovery of the malware, we have not seen any other Android threats integrate GenAI into their execution flows. On the other hand, using LLMs to aid with malware [development](#) is not unusual anymore. For an Android-specific example, in April 2026 we [reported](#) on an NGate variant that had LLM-style emoji left in the code. If you want to read our detailed analysis of PromptSpy, you can find our original research on [WeLiveSecurity](#).

Ransomware

Mapping the world of EDR killers

ESET tracks over 100 EDR killers used in the wild to kill, freeze, or blind security software that would otherwise detect the main payload during an attack.

In H1 2026, ESET published [detailed information](#) about its research into EDR killers – tools widely used during ransomware intrusions by affiliates in the ransomware-as-a-service ecosystem.

For ransomware actors, killing the endpoint detection and response (EDR) process is often easier than making their encryptors evade detection. Encryptors process a lot of data in a very short time, making them noisy by design. In contrast, EDR killers can use a variety of techniques – most often exploiting legitimate but vulnerable drivers – to remove the deployed security measures and create a short time window for the final payload to run.


Our analysis uncovered several major categories of EDR killers that use different mechanisms to remove, blind, or freeze the security software running in the targeted environment

- **simple scripts** aimed at killing security-related processes, often used in combination with

- rebooting the machine into Safe Mode,
- **anti-rootkits** designed to kill rootkits, repurposed to take down EDR processes,
- malicious tools based on the **Bring Your Own Vulnerable Driver (BYOVD)** technique to kill EDR processes directly from the kernel, and
- a smaller set of **driverless killers** that interfere with the communications of EDR software instead of abusing the kernel.

Of these, BYOVD remains by far the most prevalent approach. Attackers install a vulnerable but legitimate driver, then exploit it to terminate protected security processes from the kernel. ESET has documented over 60 such EDR killers, abusing more than 40 different drivers, with new ones appearing on a weekly basis.

However, with thousands of vulnerable drivers available – some with public proof-of-concept exploits – and AI coding tools, threat actors have a virtually endless supply of drivers for weaponization.



High-level scheme of the BYOVD technique, frequently used by ransomware actors when deploying EDR killers

Anti-rootkits are the second most common category of EDR killers, followed by rare cases where scripts and driverless approaches are used.

ESET research into ransomware attacks also shows that these tools are often used in bundles, meaning that ransomware actors deploy numerous EDR killers simultaneously during the same attack, effectively brute forcing the victim’s defenses. Some samples, including tools linked to the Warlock gang, show signs of AI-generated code.

Ransomware attacks intensify but more victims refuse to pay

According to a report published in H1 2026 by [Chainalysis](#), the share of ransomware victims paying threat actors dropped to 28% in 2025, an all-time low. This trend was already documented by [Coveware](#) in October 2025 – with a record low of 23% – and further confirmed by cyberinsurance company [Coalition](#), which said that 86% of ransomware victims refused to pay the demanded sum.

As ransom payments declined, the number of claimed ransomware attacks surged by 50%, according to Chainalysis. We reported a similar year-over-year increase in our [H2 2025 Threat Report](#).

The same Chainalysis report notes that the median ransom payment also increased by 368% year over year

to nearly USD 60,000. Total ransomware payments in 2025 currently stand at USD 820 million, down 8% year over year, but the final total is likely to approach or exceed USD 900 million as more events and payments are attributed.


The Gentlemen leak

In H1 2026, the currently second most active ransomware-as-a-service (RaaS) gang – Gentlemen – suffered a [major leak](#). Dumped material includes internal chats, images of compromised systems, and bitcoin wallet addresses allegedly used to exchange funds internally and purchase equipment.

It also shed light on the group’s day-to-day tactics, showing extensive reconnaissance and preparation inside victims’ environments, and in-house developed tooling shared among affiliates.

Gentlemen launched its dedicated leak site in 2025 and quickly became one of the major players. Aside from encryptors, the gang also maintains a large set of EDR killers, including many variants of GentleKiller – our name for Gentlemen’s proprietary tool of this kind – as well as third-party EDR killers repurposed for Gentlemen’s objectives.

In an [ESET Research blogpost](#) – published during the preparation of this report – we shared our findings on Gentlemen’s EDR killer suite, which are corroborated by



Suite containing the in-house developed GentleKiller and repurposed third-party EDR killers

the recent leak. Among the most notable findings, the group uses a shared defense-evasion layer across the whole set of these tools, mostly impersonating security vendors through fake version information, copied legitimate certificates, and stolen icons.

Gentlemen also demonstrates an unusual ability to rapidly operationalize newly disclosed BYOVD proofs of concept, often within days of their public release.

The gang relies on double extortion and reportedly offers affiliates a generous 90% revenue share, with ransomware variants targeting Windows, Linux, and other platforms. Its victims span multiple geographies and sectors, including critical sectors such as [energy](#).

EXPERT COMMENT

Unsurprisingly, the Gentlemen leak provided valuable information, including TTPs, initial access strategy, and a peek into the affiliate structure. For ESET researchers, the leaked data supports our findings from early 2026: that the Gentlemen gang develops and maintains multiple variants of its own EDR killer that we named GentleKiller, a not-so-common approach taken by current leading RaaS operators. It remains unclear whether Gentlemen will survive this major data leak, or if the gang will follow the same path as Black Basta, which shut down shortly after a similar incident. So far, we haven’t observed any visible decline in the gang’s victim announcements; in fact, we saw quite the opposite.

Jakub Souček, ESET Senior Malware Researcher

Ransomware crews turn on each other

H1 2026 brought another case of intergang fighting, when a threat actor calling itself [OAPT](#) breached the mid-sized RaaS operator Krybit and publicly leaked its panel database. The incident exposed the inner workings of the Krybit group, which has so far claimed 13 victims across 10+ countries, with typical ransom demands below USD 100,000. It also revealed poor operational security, including storing of passwords in cleartext.

Krybit responded by counter-hacking OAPT’s server, defacing its data leak site, and exfiltrating what it claimed was the group's “entire operational history”. The response reframes OAPT as less of an advanced threat actor and more of a low-skill operator whose own security was even weaker than its target’s.

In the ransomware ecosystem, criminal groups increasingly compete not only for victims and affiliates, but also for reputation and control of infrastructure, making internal fights like this more common. A similar dynamic played out last year, when DragonForce took down RansomHub, then a leading RaaS operator.

Success zone

Law enforcement kept pressure on the ransomware ecosystem through sentencing, seizures, and attribution. Recent cases include an 81-month prison sentence for a [Yanluowang ransomware access broker](#), two years for a [botnet manager](#) linked to BitPaymer, Conti, ProLock, Egregor, and DoppelPaymer cases, and four years for former ransomware negotiators involved in [BlackCat attacks](#). A [Phobos ransomware admin](#) has also pleaded guilty, with sentencing expected in July.

Authorities also [seized RAMP](#), one of the few remaining cybercrime forums openly allowing ransomware promotion and affiliate recruitment. Its Tor site and clearnet domain now display seizure notices, probably giving investigators access to user data, including emails, IPs, and private messages.

German authorities also identified Daniil Shchukin and Anatoly Kravchuk as alleged heads of [REvil](#) and [GandCrab](#), linking them to at least 130 German extortion cases, USD 2.2 million in ransom payments, and more than USD 40 million in estimated damage.




Seizure notice placed on RAMP forum websites by law enforcement


Threat telemetry

Abstract geometric lines in white and light blue on a dark background, creating a sense of motion and technology. The lines are of varying lengths and angles, some parallel and some intersecting, forming a complex pattern that suggests a network or data flow.


All threats



Overall threat detection trend in H2 2025 and H1 2026, seven-day moving average




Geographic distribution of malware detections in H1 2026




Top 10 malware detections in H1 2026 (% of malware detections)


Android



Detection trends of selected Android detection categories in H2 2025 and H1 2026, seven-day moving average (Clickers, Cryptominers, Ransomware, Scam apps, SMS trojans, and Stalkerware are combined in the trendline Other)




Geographic distribution of Android detections in H1 2026




Top 10 Android detections in H1 2026 (% of Android detections)


Cryptocurrency threats



Cryptocurrency threat detection trend in H2 2025 and H1 2026, seven-day moving average




Top 10 Cryptocurrency threat detections in H1 2026 (% of Cryptocurrency threat detections)




Geographic distribution of Cryptocurrency threat detections in H1 2026


Downloaders




Downloader detection trend in H2 2025 and H1 2026, seven-day moving average



Top 10 Downloader detections in H1 2026 (% of Downloader detections)




Downloader detections per detection type in H1 2026




Geographic distribution of Downloader detections in H1 2026


Email threats




Malicious email detection trend in H2 2025 and H1 2026, seven-day moving average



Spam detection trend in H2 2025 and H1 2026, seven-day moving average




Top 10 threats detected in emails in H1 2026




Top malicious email attachment types in H1 2026

Email threats




Geographic distribution of Email threat detections in H1 2026

Exploits




Trends of RDP, SMB, and SQL attack attempts in H2 2025 and H1 2026, seven-day moving average




External network intrusion vectors reported by unique clients in H1 2026


Exploits




Geographic distribution of RDP password guessing attack attempt sources in H1 2026



Geographic distribution of SMB password guessing attack attempt targets in H1 2026




Geographic distribution of RDP password guessing attack attempt targets in H1 2026




Geographic distribution of SQL password guessing attack attempt targets in H1 2026


Infostealers



Infostealer detection trend in H2 2025 and H1 2026, seven-day moving average




Top 10 Infostealer families in H1 2026 (% of Infostealer detections)




Geographic distribution of Infostealer detections in H1 2026


macOS



macOS detection trend in H2 2025 and H1 2026, seven-day moving average




Geographic distribution of macOS detections in H1 2026




Top 10 macOS detections in H1 2026 (% of macOS detections)

Ransomware


H2



Ransomware detection trend in H2 2025 and H1 2026, seven-day moving average




Geographic distribution of Ransomware detections in H1 2026




Top 10 Ransomware detections in H1 2026 (% of Ransomware detections)


Web threats




Web threat block trend in H2 2025 and H1 2026, seven-day moving average



Global distribution of Web threat blocks in H1 2026



Unique URL block trend in H2 2025 and H1 2026, seven-day moving average




Global distribution of blocked domain hosting in H1 2026

Research publications



LongNosedGoblin tries to sniff out governmental affairs in Southeast Asia and Japan

ESET researchers discovered a China-aligned APT group, LongNosedGoblin, which uses Group Policy to deploy cyberespionage tools across networks of governmental institutions



PromptSpy ushers in the era of Android threats using GenAI

ESET researchers discover PromptSpy, the first known Android malware to abuse generative AI in its execution flow




A rigged game: ScarCruft compromises gaming platform in a supply-chain attack

ESET researchers have investigated an ongoing attack by the ScarCruft APT group that targets the Yanbian region via backdoor-laced Windows and Android games




Revisiting CVE-2025-50165: A critical flaw in Windows Imaging Component

A comprehensive analysis and assessment of a critical severity vulnerability with low likelihood of mass exploitation



Sednit reloaded: Back in the trenches

The resurgence of one of Russia's most notorious APT groups



FrostyNeighbor: Fresh mischief and digital shenanigans

ESET researchers uncovered new activities attributed to FrostyNeighbor, updating its compromise chain to support the group's continual cyberespionage operations



ESET Research: Sandworm behind cyberattack on Poland's power grid in late 2025

The attack involved data-wiping malware that ESET researchers have now analyzed and named DynoWiper




EDR killers explained: Beyond the drivers

ESET researchers dive deeper into the EDR killer ecosystem, disclosing how attackers abuse vulnerable drivers



Webworm: New burrowing techniques

ESET researchers describe new tools and techniques that the Webworm APT group recently added to its arsenal



Love? Actually: Fake dating app used as lure in targeted spyware campaign in Pakistan

ESET researchers discover an Android spyware campaign targeting users in Pakistan via romance scam tactics, revealing links to a broader spy operation



New NGate variant hides in a trojanized NFC payment app

ESET researchers discover another iteration of NGate malware, this time possibly developed with the assistance of AI



ESET Threat Report H2 2025

vA view of the H2 2025 threat landscape as seen by ESET telemetry and from the perspective of ESET threat detection and research experts



DynoWiper update: Technical analysis and attribution

ESET researchers present technical details on a recent data destruction incident affecting a company in Poland's energy sector



GopherWhisper: A burrow full of malware

ESET Research has discovered a new China-aligned APT group that we've named GopherWhisper, which targets Mongolian governmental institutions



ESET APT Activity Report Q4 2025–Q1 2026

An overview of the activities of selected APT groups investigated and analyzed by ESET Research in Q4 2025 and Q1 2026

Credits

Team

Peter Stančík, Team Lead
Klára Kobáková, Managing Editor

Adam Chrenko
Branislav Ondrášik
Bruce P. Burrell
Hana Matušková
Nick FitzGerald
Ondrej Kubovič
Rene Holt
Zuzana Pardubská

Contributors

Anton Mäčko
Dariusz Iwański
Dušan Lacika
Jakub Souček
Jakub Tomanek
Lukáš Štefanko
Martin Jirkal
Michał Szklarzewicz

About the data in this report

The threat statistics and trends presented in this report are based on global telemetry data from ESET. Unless explicitly stated otherwise, the data includes detections regardless of the targeted platform.

Further, the data excludes detections of potentially unwanted applications, potentially unsafe applications and adware, with the exception of platform-specific charts and cryptocurrency threat charts in the Threat telemetry section.

This data was processed with the honest intention to mitigate all known biases, in an effort to maximize the value of the information provided.

Most of the charts in this report show detection trends rather than provide absolute numbers. This is because the data can be prone to various misinterpretations, especially when directly compared to other telemetry data. However, absolute values or orders of magnitude are provided where deemed beneficial.

About ESET

ESET® provides cutting-edge digital security to prevent attacks before they happen. By combining the power of AI and human expertise, ESET stays ahead of known and emerging cyberthreats — securing businesses, critical infrastructure, and individuals. Whether it’s endpoint, cloud, or mobile protection, our AI-native, cloud-first solutions and services remain highly effective and easy to use. ESET technology includes robust detection and response, ultra-secure encryption, and multifactor authentication. With 24/7 real-time defense and strong local support, we keep users safe and businesses running without interruption. An ever-evolving digital landscape demands a progressive approach to security: ESET is committed to world-class research and powerful threat intelligence, backed by R&D centers and a strong global partner network. For more information, visit www.eset.com or follow us on [LinkedIn](#), [Facebook](#), and [X](#).

[WeLiveSecurity.com](https://www.welivesecurity.com)

[@ESETresearch](#)

[ESET GitHub](#)

[ESET Threat Reports and APT Activity Reports](#)