

Annual Report

Global Threat Landscape

2025-2026

F6




Table of Contents

Introduction	5	External Compromises and Data Leaks in the Ransomware Ecosystem	70
Key Trends	7		
Predictions	11		
Trending Cyber Threats	13	Current Malware Threats	72
AI in Cyberattacks	14	Mobile threats in the wild	73
AI-Based Attacks	14	Android threats	73
Attacks on AI users	33	iOS threats	78
Attacks on AI solutions	33	Android Threats on the Dark Web	81
AI Mentions on the Dark Web	35	JS Sniffers	90
Attacks Using Zero-day Vulnerabilities	37	Analysis of Compromised Data	94
Sale of zero-day vulnerabilities on the Dark Web	41	Statistics on Compromised Payment Cards	95
Supply Chain and Trusted Relationship Attacks	45	Statistics on Payment Cards on Card Shops	99
Supply Chain Attacks	46	Statistics on Compromised Accounts and Stealers	102
Trusted Relationship Attacks	53	Data Leaks Landscape	105
Attacks on Cryptocurrency Exchanges	55	Statistics	106
Ransomware	57	Notable threat actors	109
Trends in the emergence of new ransomware groups	59	888	109
Most significant ransomware incidents	64	ByteToBreach	111
Transformation and Internal Conflicts in the RaaS Ecosystem	67	GordonFreeman	112
		N1KA	114
		KaruHunters	115
		Iulzintel	117

Tanaka	118
xorcat	120
Kazu	122
Notable data breaches	126
LATAM	126
Notable Data Breaches in LATAM	128
MEA	131
Notable Data Breaches in MEA	133
APAC	136
Notable Data Breaches in APAC	138
CIS	141
Notable Data Breaches in CIS	143
Other notable data breaches	145
AI-related data breaches	147

Underground Forum Takedowns, 2025-2026

149

Takedown of the Cracked and Nulled underground forums	150
Takedown of the LeakBase underground forum	151
Takedown of the XSS forum by French intelligence services	151
Launch of the DamageLib forum	152
Launch of the Rehub forum	152
Takedown of the RAMP underground forum	153

Takedown of BreachForums and the emergence of lookalike successor platforms

154

Repeated takedowns of the

BreachForums underground forum	154
The emergence of the Spear underground forum	157
The Emergence of PwnForums	158
Conclusions	159

Malware Sales on Underground Markets

160

Stealers

162

AURA Stealer	163
--------------------	-----

Windows

163

Needle Stealer	164
KatzStealer	165
STORM stealer	166
VoltaStealer	166
Yellow Stealer	168

macOS

170

mac.c stealer	170
---------------------	-----

macOS and Windows

171

Remus stealer	171
Stealc_v2	172
ParallelUnivers	173
Gremlin stealer	174

Remote Access Trojans (RAT)

175

Windows

176

RAT Cracker stealer	176
ATROPOSIA	176
Staeelite RAT	177
RMS RAT	178

OnyxC2	179
Snowdog	180
QatarRat	182
macOS	183
notnullOSx	183

Shadow B2B: Anatomy of the Initial Access Broker Market and 10 Key Brokers, 2025 – H1 2026

Big-Bro	188
Initial Access Sales Statistics — Big-Bro	190
ATT&CK Matrix for Enterprise (Big-Bro)	193
miyako	194
Initial Access Sales Statistics — miyako	195
ATT&CK Matrix for Enterprise (miyako)	198
Anon-WMG	199
Initial Access Sales Statistics — Anon-WMG	202
ATT&CK Matrix for Enterprise (Anon-WMG)	205
SGL	206
Initial Access Sales Statistics — SGL	207
ATT&CK Matrix for Enterprise (SGL)	209
personX	210

Initial Access Sales Statistics — personX	213
ATT&CK Matrix for Enterprise (personX)	216

Saturned33	217
Initial Access Sales Statistics — Saturned33	219
ATT&CK Matrix for Enterprise (Saturned33)	221

HelluvaHack	222
Initial Access Sales Statistics — HelluvaHack	223
ATT&CK Matrix for Enterprise (HelluvaHack)	225

Minions	226
Initial Access Sales Statistics — Minions (aka williamblack)	228
ATT&CK Matrix for Enterprise (Minions)	229

privisninet	230
Initial Access Sales Statistics — privisninet	232
ATT&CK Matrix for Enterprise (privisninet)	234

AckLine	235
Initial Access Sales Statistics — AckLine	236
ATT&CK Matrix for Enterprise (AckLine)	237

Recommendations	238
------------------------------	------------

Introduction



Valery Baulin, CEO of F6

We're pleased to present F6's first global analysis of the current cyber threat landscape. F6 is a cybersecurity company founded in 2024, built on more than two decades of frontline experience in cybercrime investigations, digital forensics, incident response, and threat intelligence. Our team has worked on some of the most complex cyber incidents affecting enterprises, financial institutions, industrial organizations, critical infrastructure operators, and government entities, while also tracking large-scale fraud, phishing, and cyber-enabled criminal operations.

The F6 team built its expertise and capabilities on a legacy of digital forensics and investigations. Over the years, we have built on this expertise to develop a comprehensive portfolio of technologies and services designed to help organizations identify risks before an attack occurs, protect digital channels, detect malicious activity, respond to incidents, and investigate their root causes.

This report covers trending cyber threats, the current data leak landscape, recent underground forum takedowns, malware sales on Dark Web markets, and the role of Initial Access Brokers (IABs) and other cybersecurity landscape trends.

The insights presented in this report are derived from real-world investigations, incident response engagements, threat hunting activities, and continuous monitoring of adversary operations. The data has been collected and analyzed jointly by F6's Threat Intelligence, Digital Forensics, and Cybersecurity teams.

Technical findings are based on direct observations made during security incident response engagements, as well as telemetry collected through Threat Intelligence and Managed XDR – our flagship solutions. This unique combination of frontline expertise and large-scale visibility provides a detailed view of the tactics, techniques, procedures, and operational trends observed among today's threat actors.

The purpose of this report is to provide cybersecurity leaders and practitioners with actionable intelligence that supports both strategic decision-making and day-to-day defensive operations. We hope the trends, forecasts, technical details and recommendations presented here will serve as a valuable resource for CISOs, security leaders, SOC and CERT teams, incident responders, threat intelligence analysts, and threat hunters seeking to better understand and prepare for the evolving threat landscape.

Valery Baulin, CEO of F6

This is F6's first report on the global cyber threat landscape. It highlights the most significant attacks, key trends, and our outlook for the year ahead, based on insights from F6 Threat Intelligence experts.

The report is designed for information security professionals, CISOs, business leaders, and cyber threat analysts.

Key Trends

- The countries most targeted by ransomware: United States, France, United Kingdom, Germany, Canada, Italy, Australia, Spain, Brazil, India.
- The most frequent stealer victims in 2025: India, United States, Brazil, Indonesia, Philippines, Vietnam, Turkey, Mexico, Egypt, Thailand. In the first half of 2026, the victim-country breakdown was similar, except that Egypt and Thailand were replaced by Pakistan and Bangladesh.
- Top 20 countries by number of compromised cards published across open and closed sources in 2025: United States, Mexico, Turkey, Canada, United Kingdom, United Arab Emirates, Brazil, Australia, France, Malaysia, India, Saudi Arabia, Italy, Spain, Germany, Thailand, Chile, Japan, Pakistan, Argentina.
- Top 20 countries by number of compromised cards published in open and closed sources in the first half of 2026: United States, Russia, Mexico, Canada, Thailand, Malaysia, United Kingdom, Brazil, Australia, Turkey, South Korea, France, India, United Arab Emirates, Philippines, Ireland, Bahrain, Spain, Germany, Italy.
- According to F6 Threat Intelligence, more than 13 million payment cards were offered for sale in card shops in 2025, and F6 identified more than 3.5 million cards offered for sale in card shops in the first half of 2026.
- Widespread adoption of AI in the toolkits of threat actors across all skill levels.
- Advanced threat actors are incorporating AI into a growing number of attack stages.
- Threat actors are devising increasingly sophisticated AI jailbreaking techniques: they craft convincing prompts and communications, break tasks into smaller pieces, and run several AI solutions in parallel to distribute the work.

- Attacks that use AI, target AI users, and target AI itself: some threat actors used AI to carry out attacks, others distributed malware disguised as various AI tools, and still others attacked the AI solutions themselves, probing them for vulnerabilities and attempting to obtain proprietary development information.
- In most publicly disclosed attacks, threat actors favored three solutions: Gemini, OpenAI tools, and Claude Code. Gemini was the most commonly used.
- On underground forums, the top three most-discussed AI solutions were ChatGPT, Gemini, and Claude Code. ChatGPT was mentioned more than four times as often as the other two (108,000 mentions), while Gemini and Claude Code were mentioned roughly the same number of times (~25,000 each).
- A rise in supply chain attacks and, to a lesser extent, trusted relationship attacks.
- A shift from narrowly targeted supply chain attacks to mass-scale ones. Previously, a supply chain compromise affected only the customers of the compromised organization, but threat actors now increasingly compromise open-source components, simultaneously putting the software development processes of many companies at risk.
- Steady growth in the number of new ransomware groups, the vast majority of which operate on a RaaS model. This significantly lowers the barrier to entry for newcomers and drives market fragmentation. More than 50 new groups emerged in 2025 alone, even though many of them lasted only a few months.
- Shifts in ransomware economics: a noticeable decline in publicly disclosed ransom payments and a growing share of private negotiations. Threat actors are shifting their primary focus away from securing payment for decryption and toward exerting pressure through data leaks and reputational damage.
- Ransomware operations are moving away from “noisy” attacks toward more deliberate, multi-stage scenarios aimed at evading defenses and maximizing damage to the victim’s infrastructure.
- Numerous internal conflicts, leaked chat logs, and compromises of group infrastructure have aided the work of law enforcement and researchers.
- During 2025–2026, law enforcement agencies carried out several successful operations to shut down underground forums. As a result of coordinated efforts by international intelligence agencies, 6 popular underground platforms were partially or completely blocked and 5 new forums were created to replace them.
- Android banking trojans increasingly combine NFC relay abuse and Automated Transfer System (ATS) capabilities to execute financial fraud.
- Threat actors are adopting new techniques to hinder analysis and evade detection that are unusual for mobile malware. For example, the Herodotus trojan identified in 2025 attempts to make remotely performed actions appear human, and the developers of the Klopatra trojan have integrated Virbox, a commercial-grade code protection tool.

- The DroidLock Android ransomware trojan was disclosed. Unlike typical ransomware, DroidLock does not encrypt files; instead, it blocks access and threatens to destroy all data unless a ransom is paid, achieving the same effect.
- Android MaaS offerings are evolving rapidly. Several have already been observed in real-world attacks, including SuperCard X, PlayPraetor, Herodotus, AlbirioX, and Mirax.
- Far fewer iOS threats than Android threats were identified in the wild. In addition, no iOS malware offerings were observed on underground platforms. This can be explained by lower prevalence of iOS than Android and to the complexity and cost of developing this class of threat.
- Two advanced iOS spyware exploit kits, DarkSword and Coruna, were used not only in targeted attacks but also by several criminal groups in mass infection campaigns.
- Threat actors are increasingly abandoning their own malicious infrastructure in favor of legitimate cloud services. They disguise JS sniffer C2 traffic as legitimate requests to Google Firebase and the Stripe API, use Google Tag Manager and Azure Functions to deliver payloads, and use Discord webhooks and Vercel for covert data exfiltration.
- Magento remains the primary target for sniffers due to a steady stream of critical vulnerabilities and slow patching by e-commerce stores. Even after patches are released, a large number of sites remain unpatched. Combined with the ability to automate attacks, this makes the platform ideal for mass sniffer deployment.
- The Malware-as-a-Service (MaaS) business model is expanding rapidly, offering buyers turnkey malware solutions complete with technical support, regular updates, and rented command-and-control infrastructure. This trend significantly lowers the barrier to entry into cybercrime, enabling attackers with limited technical expertise to launch sophisticated attacks.
- Top 20 countries by number of initial access sales advertisements in H1 2026: United States, Brazil, Israel, India, United Kingdom, France, Italy, Canada, Argentina, Australia, Spain, Indonesia, Germany, China, Malaysia, Mexico, United Arab Emirates, Netherlands, Vietnam, and Colombia.
- The global initial access market is estimated at approximately \$10 million for the period from 2025 through H1 2026.
- The majority of initial access sold on underground markets is provided through remote access mechanisms, primarily RDP and VPN. Initial Access Brokers (IABs) maintain this steady pipeline by operationalizing a highly effective, multi-pronged collection strategy. They leverage stealer malware for bulk credential harvesting, execute automated brute-force campaigns against exposed remote access endpoints, and systematically exploit vulnerabilities in perimeter devices such as firewalls and VPN gateways. By diversifying the initial access vectors, IABs ensure a highly resilient pipeline of compromised corporate environments for downstream distribution.

- Because government, manufacturing, finance, and retail are high-value sectors, cybercriminals prioritize these industries to obtain initial access to corporate networks and sensitive data and use the compromised environments for further monetization.
- Databases are published more often than sold because threat actors primarily use such publications to build their reputations and gain recognition for future activity in underground communities.
- The number of data breaches affecting the AI industry increased during 2025–2026, in parallel with the expanding deployment and mainstream adoption of AI-driven solutions across various systems.

Predictions

- AI-enabled attacks will continue to evolve at an accelerating pace. Malware incorporating AI is expected to progress from AIM3 Levels 1–2 to more advanced levels. Isolated cases of largely autonomous AI-driven attacks requiring minimal human involvement are also likely to emerge.
- Less experienced criminals will continue to use AI for vibe-coding and reconnaissance.
- Free AI accounts no longer provide threat actors with sufficient capabilities for full-scale development. Because paid accounts offer more capabilities, threat actors will begin devising ways to acquire them without leaving a trail.
- Supply chain attacks will remain a top threat, especially those targeting npm packages, given their relative ease of execution and large-scale impact.
- In current and future geopolitical conflicts, cyberattacks will form part of cyberwarfare. They may damage national infrastructure and disrupt physical processes, including power, water, transportation, communications, and industrial facilities.
- Threat actors will increasingly opt for an extortion model based on the threat of publishing stolen data rather than traditional encryption-based attacks.
- Ransomware operations will become stealthier and more multi-stage: threat actors will continue to use BYOVD, EDR-killer tools, and living-off-the-land techniques to evade modern defenses.
- The priority targets for ransomware groups will remain large organizations in manufacturing, logistics, finance, and critical infrastructure, primarily in the United States and Western Europe.
- The ransomware market will likely continue to fragment: amid conflicts and leaked chat logs from major RaaS platforms, new, smaller groups will emerge, focusing less on ransomware development than on extortion services, negotiations, and the monetization of stolen data.
- Competition among leading players such as Qilin, DragonForce, and LockBit 5 will intensify. These groups will attempt to consolidate affiliate ecosystems and attract affiliates through more favorable terms and attack automation.
- The nature of Android trojans will continue to change; they will become more versatile and full-featured, combining the capabilities of a banking trojan, spyware, RAT, and proxy malware.
- Threat actors will devise ways to hinder analysis and evade defenses, allocating more funding and resources to development in order to minimize the chance of detection and potentially reap greater profits. As a result, fraudulent activity will increasingly resemble that of a legitimate customer using their own device.
- Trojans built on NFCGate and/or capable of taking control of devices and performing automated money transfers will continue to evolve.

- Most Android trojans abuse Accessibility Services. New restrictions, such as the recently added [ability](#) to use the `accessibilityDataSensitive` flag, will lead threat actors to look for other ways to obtain equivalent permissions, so new Android malware variants are likely to emerge.
- The use of advanced iOS spyware will expand beyond targeted espionage campaigns into more widespread attacks conducted for financial gain.
- New, full-fledged franchises for deploying JS sniffers will emerge as part of an increasingly automated business model, complete with technical support and affiliate programs.
- Threat actors will make greater use of legitimate infrastructure in attacks involving JS sniffers.
- The technical sophistication of JS sniffers will grow as threat actors make greater use of AI-based solutions.
- Joint law-enforcement operations aimed at shutting down known underground resources will continue.
- The boundaries between malware categories will continue to blur as attackers add features that were not previously associated with a given malware type.
- Threat actors will focus on the modular expansion of existing malware capabilities, transforming them into multifunctional platforms rather than creating standalone products.
- Malware sales on underground resources will steadily increase, as will the number of subscription-based offerings for ready-made, customizable MaaS solutions.
- The growing profitability of the initial access market will attract new entrants, including less skilled actors. This influx is expected to drive an increase in unauthorized resales of initial access put up on a sale by mature cybercriminals.
- Driven by the pursuit of higher profits per compromise, IAB actors are expected to expand their offerings beyond initial access alone. Alongside VPN/RDP credentials, sellers will increasingly exfiltrate and monetize sensitive data — including customer databases, intellectual property, and internal documents — creating additional income streams from each successful intrusion.

Trending Cyber Threats

This section examines the most significant cyber threats observed worldwide in 2025–2026. Each chapter explores one of the year's key trends, including:



AI-enabled attacks



Ransomware



Supply chain and trusted relationship attacks



Attacks targeting cryptocurrency exchanges



Zero-day vulnerability exploitation

AI in Cyberattacks

AI-Based Attacks

Across the global threat landscape, 2025–2026 saw a sharp rise in AI-enabled attacks. Threat actors increasingly used AI both to develop malware and to integrate LLMs into malware, enabling dynamic, context-aware behavior during execution. This chapter examines the evolution of malicious AI use through some of the most significant cyberattacks observed worldwide.

In early 2025, AI use remained limited: threat actors primarily applied basic capabilities to streamline the initial stages of attacks, including reconnaissance and target profiling, phishing copy and lure creation, and similar tasks. Later, they began using AI for coding advice and debugging; so-called vibe coding gained popularity; and security vendors increasingly used the term AI malware to describe malicious or offensive software whose core development or runtime behavior depends on GenAI or LLMs.

From mid-2025 onward, threat actors began actively using AI not only as a basic assistant but also to develop simple malware directly. It is still too early to speak of fully autonomous AI-driven attacks, because threat actors continue to perform the core actions themselves, but the trend is developing rapidly.

Over the course of 2025, threat actors moved toward using increasingly sophisticated AI-based malware.

In 2023, the [MITRE AI Maturity Model](#) was created to provide guidance for building a foundation for AI implementation across an organization. In response to rapid advances in AI-enabled malware, Recorded Future [proposed](#) the AIM3 (AI Malware Maturity Model) in 2025, drawing on this methodology to measure the sophistica-

tion of AI-enabled threats. The AIM3 model consists of five levels:

- **Level 1 — Experimenting:** attackers explore the potential of GenAI and LLMs for malicious applications;
- **Level 2 — Adopting:** threat actors use GenAI to author phishing emails, research targets, and develop code, with an emphasis on automating and supporting lower-order tasks;
- **Level 3 — Optimizing:** threat actors begin incorporating AI into their attack chains, using GenAI on-host or via APIs to perform introspection, generate commands, and adapt code in near real time;
- **Level 4 — Transforming:** AI-native offensive frameworks emerge, combining multi-step planning and tool use with a human-in-the-loop (HITL) approach. These are early, purposeful attempts at AI-first threat operations that use agentic patterns rather than bolting GenAI onto legacy playbooks;
- **Level 5 — Scaling:** threat actors build agentic systems that manage campaigns end-to-end without human oversight.

Since late 2025, several cases of advanced AI use in attacks have come to light. The first known example of an L4-level attack under the AIM3 model is an operation conducted by GTG-1002, in which the threat actors used jail-breaking to bypass Claude’s safeguards. Their prompt claimed that Claude was an employee of a legitimate cybersecurity firm and was being

used to test defenses. The threat actors broke their attacks into small, seemingly harmless tasks that Claude carried out without receiving the full context of their malicious purpose. As a result, Claude Code performed more than 80% of the malicious activity.

Threat actors are finding ways to bypass security policies and induce AI systems to provide malicious functionality. In several disclosed cases, prompts framed the AI as a cybersecurity expert, penetration tester, student working on a project, or author of IT articles.

The following analysis traces the evolution of AI use for malware development and attack execution according to the AIM3 classification. The table below lists the most notable AI malware identified by researchers since 2025.

AI malware identified in 2025–2026

First seen date	AI malware	Malware category	AIM3 Level	In-the-wild malware
January 2025	FRUITSHELL	Reverse shell	L2	✓
January 2025	PROMPTFLUX	Dropper	L1	✗
July 2025	LAMEHUG	Data miner	L3	✓
July 2025	HexStrike-AI	Red teaming framework	L3	✓
July 2025	Cyberspike's Villager	Red teaming framework	L3	✗
August 2025	QUIETVAULT	Stealer	L3	✓
August 2025	OSSTUN	C2 framework	L2	✓
August 2025	PromptLock	Ransomware	L1	✗
September 2025	VoidLink	Malicious framework	L3	✓

First seen date	AI malware	Malware category	AIM3 Level	In-the-wild malware
September 2025	HONESTCUE	Downloader and launcher framework	L3	✓
November 2025	COINBAIT	Phishing kit	L2	✓
February 2026	PromptSpy	Android malware	L2	✓

FRUITSHELL

FRUITSHELL is a publicly available reverse shell written in PowerShell that establishes a remote connection to a configured command-and-control server and allows a threat actor to execute arbitrary commands on a compromised system. Notably, this code family contains hard-coded prompts meant to bypass detection or analysis by LLM-powered security systems. FRUITSHELL was observed in operations, but details of the attacks using this malware have not been publicly disclosed.

LAMEHUG (a.k.a. PROMPTSTEAL)

In June 2025, researchers observed APT28 using a new malware named LAMEHUG. This malware is written in Python and packaged with PyInstaller. It contains a compiled script that uses the Hugging Face API to query the LLM Qwen2.5-Coder-32B-Instruct and generate one-line Windows commands. The prompts used to generate the commands indicate that the malware is intended to collect system information and documents from specific folders. PROMPTSTEAL then executes the commands and sends the collected data to an adversary-controlled server.

HexStrike-AI

HexStrike-AI is an AI-powered offensive security framework that combines professional security tools with autonomous AI agents to deliver comprehensive security testing capabilities. After its release, threat actors [began discussing](#) how to weaponize it. Within hours, some underground channels were discussing how to use the framework to exploit zero-day vulnerabilities.

Another example of HexStrike-AI being used for malicious purposes was described in May 2026. According to [researchers](#), a suspected PRC-nexus threat actor deployed agentic tools like HexStrike-AI and Strix against a Japanese technology firm and a prominent East Asian cybersecurity platform. HexStrike-AI was used alongside the Graphiti memory system, a temporal knowledge graph, to maintain a persistent state of the attack surface, allowing the agent to autonomously pivot between tools such as subfinder and httpx based on its internal reasoning. At the same time, the actor leveraged Strix, a multi-agent penetration testing framework, to automate the identification and validation of vulnerabilities. This combination of autonomous reconnaissance and automated verification suggests a shift toward AI-driven frameworks that can scale discovery activities with minimal human oversight.

QUIETVAULT

QUIETVAULT is a credential stealer written in JavaScript that targets GitHub and npm tokens. This malware played a key role in the [singularity supply chain campaign](#). Captured data is exfiltrated by creating or publishing to publicly accessible GitHub repositories controlled by the attacker.

On August 26, 2025, multiple malicious versions of the widely used Nx build system package were published to the npm registry. These versions contained a post-installation malware script. The initial compromise introduced QUIETVAULT, which was triggered via the post-install script. The payload executed only on Linux and macOS systems, systematically searching for sensitive files (wallets, keystores, .env files, and SSH keys) and extracting credentials. The malware attempted a lockout by appending a command, effectively causing system shutdowns on new terminal sessions.

In the first phase of the [attack](#), the malware used AI command-line tools, including Claude, Gemini, and Q, to support its reconnaissance and then exfiltrated the stolen data to publicly accessible, attacker-created repositories within victims' GitHub accounts. In Phase 1, secrets belonging to more than 1,700 users were publicly exposed. The malware also attempted to exfiltrate potentially sensitive files, and researchers observed that more than 20,000 files had been leaked.

In the second phase, a threat actor used the GitHub tokens leaked in Phase 1 to make victims' private repositories public. At least 480 compromised accounts, two-thirds of them belonging to organizations, made more than 6,700 private repositories public. In Phase 3, two compromised users publicly uploaded more than 500 repositories belonging to a single organization.

HONESTCUE

In September 2025, researchers observed new malware tracked as HONESTCUE but did not link it to any existing threat activity cluster. HONESTCUE is a downloader and launcher framework that sends a prompt via Google Gemini's API and receives C# source code in response. That code drives the second-stage functionality, which downloads and executes another piece of malware. In addition, HONESTCUE's fileless secondary stage takes the C# source code received from the Gemini API and uses the legitimate .NET CSharpCodeProvider framework to compile and execute the payload directly in memory. The threat actors use content delivery networks (CDNs) such as Discord CDN to host the final payloads.

COINBAIT

In November 2025, researchers identified COINBAIT, a phishing kit that masqueraded as a major cryptocurrency exchange to harvest credentials and was likely developed more quickly with AI code-generation tools. This activity overlaps with activity attributed to financially motivated threat actor UNC5356. The phishing kit was built using the AI-powered Lovable AI platform, as indicated by its use of the lovableSupabase client and lovable.app for image hosting. The kit was wrapped in a full React single-page application (SPA) with complex state management and routing. This complexity indicates code generated from high-level prompts, such as "Create a Coinbase-style UI for wallet recovery," using a framework such as Lovable AI. Another key indicator of LLM use is the presence of verbose, developer-oriented logging messages within the kit's source code.

PromptSpy

In February 2026, researchers uncovered PromptSpy, the first known Android malware that uses generative AI to achieve persistence. The main purpose of PromptSpy is to deploy a VNC module on the victim's device; the malware can capture lock-screen data, block uninstallation, gather device information, take screenshots, record the screen as video, and more. PromptSpy uses Gemini to obtain with step-by-step instructions on how to keep the malicious app pinned in the recent apps list, preventing it from being easily swiped away or killed by the system. PromptSpy sends Gemini a natural-language prompt along with an XML dump of the current screen, giving the AI a detailed view of every UI element: its text, type, and exact position on the display. Gemini processes this information and responds with JSON instructions that tell the malware what action to perform and where to perform it. The malware saves both its previous prompts and Gemini's responses, allowing Gemini to understand context and coordinate multi-step interactions. PromptSpy keeps prompting Gemini until the AI confirms that the app has been successfully locked, forming a feedback loop in which the malware waits for validation before moving on.

VoidLink

A previously unknown threat actor tracked as UAT-9921 was observed using a new malware framework called VoidLink in campaigns targeting the technology and financial services sectors. It is a Linux-based malware framework featuring a modular command-and-control (C2) architecture, eBPF and LKM rootkits, cloud and container enumeration, and more than 30 post-exploitation plugins. Researchers assess that a single developer created it with LLM assistance using a paradigm called spec-driven development. VoidLink uses three programming languages: Zig for the implant, C for the plugins,

and Go for the backend. VoidLink's development likely began in late November 2025, when its developer turned to TRAE SOLO, an AI assistant embedded in TRAE, an AI-centric IDE. In this case, TRAE generated a Chinese-language instruction document. The summary of the developer's initial exchange with the agent suggests that the opening instruction was not to build VoidLink directly, but to design a minimal skeleton and produce a concrete execution plan for turning it into a working platform.

Analysis of these AI-enabled attacks shows that in most publicly disclosed cases, threat actors relied primarily on three tools: Gemini, OpenAI tools, and Claude Code. Gemini was used most frequently. This is partly due to the nature of the threat actors' tasks and the limitations of free tiers, because Gemini had fewer restrictions than the other solutions. It is worth noting, however, that in the cases examined, threat actors used Claude Code for more advanced attacks at AIM3 levels L2-L4, while the other AI tools were used in L1-L2 attacks.

The table below contains information on the groups that used Gemini in their attacks.

Threat actors that used Gemini

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
APT41	 CN	Assistance with code development and obfuscation	L2	Resource Development TA0042	Develop Capabilities: Malware T1587.001
MuddyWater	 IR	Code development support	L2	Resource Development TA0042	Develop Capabilities: Malware T1587.001
UNC6418	 N/A	Intelligence gathering	L1	Reconnaissance TA0043	Gather Victim Identity Information: Email Addresses T1589.002
Mustang Panda	 CN	Intelligence gathering	L1	Reconnaissance TA0043	Gather Victim Identity Information T1589 Gather Victim Org Information T1591
Lazarus	 KP	Intelligence gathering	L1	Reconnaissance TA0043	Gather Victim Org Information T1591
APT31	 CN	To automate vulnerability analysis	L2	Reconnaissance TA0043	Active Scanning: Vulnerability Scanning T1595.002
UNC795	 CN	Code auditing	L1	Resource Development TA0042	Develop Capabilities: Malware T1587.001

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
APT42	 IR	Intelligence gathering and malware development support	L1	Reconnaissance TA0043 Resource Development TA0042	Gather Victim Identity Information: Email Addresses T1589.002 Gather Victim Org Information: Business Relationships T1591.002 , Develop Capabilities: Malware T1587.001
Andariel	 KP	To analyze CVEs and validate PoC exploits	L2	Reconnaissance TA0043	Active Scanning: Vulnerability Scanning T1595.002
APT27	 CN	Accelerating malware development	L2	Resource Development TA0042	Develop Capabilities: Malware T1587.001

APT41

APT41 is a group active since at least 2009. It blends cyber espionage with financially motivated criminal activity against private-sector and government targets; it has also been linked to attacks on gaming-industry companies and to supply chain compromises.

In one case, APT41 unsuccessfully [attempted to use Gemini](#) to learn about Gemini’s underlying infrastructure and systems. The actor asked Gemini to share details such as its IP address, kernel version, and network configuration. Gemini responded but did not disclose sensitive information. In another case, APT41 used the OSSTUN framework, which researchers observed being

developed with Gemini. The actor sought technical support for C++ and Golang code across multiple tools, including OSSTUN. The group was also observed prompting Gemini for help with code obfuscation, with prompts referencing two publicly available obfuscation libraries.

MuddyWater

MuddyWater is an Iranian state-sponsored group active since at least 2017. For intelligence-gathering purposes, MuddyWater targets government, defense, critical infrastructure, and telecommunications sectors.

[MuddyWater used Gemini](#) to conduct research supporting the development of custom malware. The actor used various plausible pretexts in their prompts to bypass security guardrails. These included claiming to be a student completing a final university project, writing a paper, or preparing an international cybersecurity article. In some observed instances, threat actors' reliance on LLMs for development led to critical operational security failures, enabling greater disruption. The threat actor asked Gemini for help with a provided script designed to listen for encrypted requests, decrypt them, and execute commands related to file transfers and remote execution. This exposed sensitive, hard-coded information to Gemini, including the C2 domain and the script's encryption key, facilitating disruption of the attacker's campaign.

UNC6418

UNC6418, an unattributed threat actor, [misused Gemini](#) to conduct targeted intelligence gathering, specifically seeking out sensitive account credentials and email addresses. Shortly afterward, researchers observed the threat actor targeting all of these accounts in a phishing campaign focused on Ukraine and the defense sector.

Mustang Panda

Mustang Panda is a Chinese state-sponsored group that targets government, diplomatic, and non-governmental organizations (including think tanks and religious and academic institutions), primarily in the U.S., Europe, and Asia.

Mustang Panda [misused Gemini and other AI tools](#) to compile detailed information on specific individuals, including targets in Pakistan, and to collect operational and structural data on separatist organizations in various countries. Shortly afterward, the threat actor targeted similar individuals in Pakistan in a subsequent campaign.

Lazarus

Lazarus is a North Korean group active since at least 2007. Over time, the group has been observed conducting espionage, attacking financial institutions, and, more recently, targeting cryptocurrency exchanges.

Lazarus has consistently targeted the defense sector and impersonated corporate recruiters in their campaigns. The group [used Gemini](#) to synthesize OSINT and profile high-value targets in support of campaign planning and reconnaissance. Its profiling activities included researching major cybersecurity and defense companies and mapping specific technical roles and salary information.

APT31

APT31 is a Chinese state-sponsored group active since at least 2016, focused on espionage and intellectual property theft.

APT31 employed a highly structured approach, [prompting Gemini](#) with an expert cybersecurity persona to automate vulnerability analysis and generate targeted testing plans. In one case, the actor fabricated a scenario in which it claimed to be testing HexStrike MCP tooling and directed the model to analyze remote code execution (RCE), web application firewall (WAF) bypass techniques, and SQL injection test results against specific U.S.-based targets. This use of Gemini automated intelligence gathering to identify technological vulnerabilities and organizational defense weaknesses.

UNC795

UNC795 is a Chinese group active since at least 2025.

UNC795 [relied heavily on Gemini](#) throughout their entire attack lifecycle. The group consistently engaged with Gemini several days

a week to troubleshoot their code, conduct research, and develop technical capabilities for their intrusion activity. The group also used Gemini to create an AI-integrated code auditing capability, likely reflecting an interest in agentic AI utilities to support their intrusion activity.

APT42

APT42 is a pro-Iranian group active since at least 2015. It conducts information-gathering and surveillance operations against individuals and organizations of strategic interest.

[APT42 used Gemini](#) to search for email addresses of target organizations and their business partners, which were used to build a convincing cover story for their attacks. APT42 additionally used Gemini to accelerate the development of new malware and offensive tooling, including debugging and researching exploitation techniques.

UNC2814

UNC2814 is a suspected Chinese-speaking cyber espionage group that researchers have tracked since 2017. The group targets government organizations and telecommunications companies in Africa, Asia, and the Americas.

Researchers [observed](#) UNC2814 using a form of expert persona prompting, directing the model to act as a senior security auditor or C/C++ binary security expert. The fabricated scenarios were used to support vulnerability research into various embedded devices, including TP-Link firmware and Odette File Transfer Protocol (OFTP) implementations.

Andariel

Andariel is a North Korean state-sponsored group active since 2009 that primarily targets organizations in South Korea.

In addition to using individual prompts for real-time troubleshooting, [Andariel](#) sent thousands of repetitive prompts that iteratively analyzed different CVEs and validated PoC exploits. This helped the group build a more robust exploit arsenal. The actors are also experimenting with agentic tools such as OpenClaw and OneClaw alongside intentionally vulnerable testing environments.

APT27

APT27 is a Chinese group active since at least 2010. It targets companies and contractors in the defense, government, aerospace, telecommunications, energy, technology, and industrial sectors.

[APT27 leveraged Gemini](#) to speed up the development of a fleet management application, likely intended to manage an operational relay box (ORB) network.

Threat actors that used OpenAI tools

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
Kimsuky	 KP	Deepfake image creation; debugging and development assistance	L1, L2	Resource Development TA0042	Generate Content: Written Content T1683.001 , Develop Capabilities T1587 , Obtain Capabilities T1588
Lazarus	 KP	Deepfake creation; development of fake websites and LinkedIn personas; debugging and development assistance	L2	Resource Development TA0042	Establish Accounts T1585 , Generate Content T1683 , Develop Capabilities T1587 , Obtain Capabilities T1588
APT5 and APT15	 CN	Open-source research into various entities of interest and technical topics, code development and infrastructure setup support	L1, L2	Reconnaissance TA0043 Resource Development TA0042	Gather Victim Org Information T1591 , Develop Capabilities T1587 , Acquire Infrastructure T1583
UTA0388	 CN	Phishing content creation; malware debugging and modification	L1, L2	Reconnaissance TA0043 Resource Development TA0042	Generate Content: Written Content T1683.001 , Develop Capabilities T1587

DPRK-affiliated threat actors

On July 17, 2025, researchers [detected](#) a spear-phishing attack attributed to Kimsuky. The threat actor used ChatGPT, a generative AI tool, to produce sample ID card images, which were then leveraged in the attack. This case demonstrates the state-sponsored group's use of deepfake technology.

Kimsuky is a North Korean state-sponsored group active since 2012 that targets organizations in South Korea and other countries for espionage purposes.

Kimsuky is believed to [have continued using ChatGPT](#) throughout 2025 in attacks against South Korea. Each identified ChatGPT account appears to have focused on a specific use case. The threat actors generated model outputs to support multiple operational areas, including implant and RAT-adjacent development, credential-theft routines, phishing content and lures, macOS development scaffolding, and cryptocurrency operations.

According to an [OpenAI report](#), the DPRK-affiliated threat actors Kimsuky and Lazarus used OpenAI tools to pursue information likely related to cyber intrusion tools or operations. They also showed interest in cryptocurrency-related topics, likely in connection with financially motivated activities. The actors used these models for coding assistance and debugging and to research security-related open-source code. This included debugging and development assistance for publicly available tools and code that could be used for RDP brute-force attacks, as well as assistance with the use of open-source RATs.

According to [researchers](#), Lazarus used GPT-4o in a campaign against cryptocurrency organizations in India, Turkey, Australia, and other countries across Europe and Asia. Victims are invited to fake investment meetings on phishing sites that mimic Zoom or Microsoft Teams. During

these meetings, they are prompted to update the client application to fix an audio problem. In reality, this action downloaded a malicious script and infected the device with malware. The photos of the meeting participants were modified using GPT-4o. In addition, SilentSiphon's secret-stealing module contained several comments. One used a checkmark emoji to indicate successful archiving even though it referred to a completed backup; this led researchers to suggest that Lazarus had used generative AI to write malicious scripts similar to this module.

In April 2026, researchers disclosed [a high-volume campaign](#) against Web3 developers in which the actors posed as tech recruiters and offered fake job assessments backdoored with malware. The campaign was attributed to Lazarus. AI tools — including ChatGPT and Cursor — were used to generate malware loaders and build fake company websites and LinkedIn personas, with evidence of AI use including emoji-laden debug code and accidentally leaked prompts. The group also used AI to test their own backdoors against detection, attempting to harden them against AI-assisted code audits conducted by targets.

China-linked threat actors

[Researchers](#) banned ChatGPT accounts associated with China-linked actors: APT5 and APT15.

APT5 is a cyberespionage group of presumed Chinese origin that has been active since at least 2007 and targets government, technology, and telecommunications companies.

APT15 is a state-sponsored group of Chinese origin, active since at least 2010. For espionage purposes, it targets government, defense, aerospace, and diplomatic entities.


Both APT5 and APT15 engaged with the models in Chinese and English. Their activity originated from the same networks but fell into distinct subsets based on purpose. One subset

used the models for open-source research on entities of interest and technical topics. In technical interactions, the actors used OpenAI models to modify scripts or troubleshoot system configurations, including work involving reNgin and Selenium automation. Another subset appeared to use the models for supporting activities, including Linux system administration, software development, and infrastructure setup. For these activities, the threat actors used the models to troubleshoot configurations, modify software, and research implementation details. In the system administration category, they asked for advice on configuring firewalls and nameservers, as well as on building software packages for offline deployment. Their software development activity included web and Android app development, as well as C and Golang software. Infrastructure setup included configuring VPNs, installing software, deploying Docker containers, and deploying local LLMs such as DeepSeek.

Beginning in April 2025, a China-aligned threat actor tracked as [UTA0388/UNK_DropPitch leveraged OpenAI's ChatGPT platform](#) for several components of its campaign against organizations in North America, Asia, and Europe.

UTA0388 is a cyber espionage group, presumably linked to China, active since 2025.

UTA0388 primarily used OpenAI models to make incremental efficiency gains in existing workflows, particularly to craft phishing content and debug or modify its tooling. In every observed case, UTA0388 sent targets a link to a ZIP or RAR archive. Inside was a legitimate executable with a filename relevant to the targeted organization or tied to the theme of the spear-phishing email. When executed, it loaded a malicious payload from a bundled Dynamic Link Library (DLL) via search-order hijacking, giving operators the ability to execute commands remotely on infected devices. The deployed payload is tracked as GOVERSHELL. Beginning in August 2025, UTA0388 changed its approach to rapport-building phishing. The group first contacted targets without malicious content and delivered it only later, after the user replied.





Threat actors that used Claude Code

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
GTG-2002	 N/A	To automate reconnaissance, credential harvesting, and network penetration	L3	Reconnaissance TA0043 Resource Development TA0042	Active Scanning TA0043 Gather Victim Network Information T1590 Generate Content: Written Content T1683.001 Develop Capabilities T1587 Obtain Capabilities T1588
Unknown TA	 CN	Claude Code used as a technical adviser, code developer, security analyst, and operational consultant	L3	12 tactics	Active Scanning T1595 Develop Capabilities T1587 Brute Force T1110 Exploitation for Privilege Escalation T1068 Exploit Public-Facing Application T1190 Proxy T1090

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
GTG-1002	 CN	To support reconnaissance, vulnerability discovery, exploitation, lateral movement, credential harvesting, data analysis, and exfiltration operations largely autonomously	L4	Reconnaissance TA0043 Resource Development TA0042 Initial Access TA0001 Persistence TA0003 Privilege Escalation TA0004 Stealth TA0005 Credential Access TA0006 Discovery TA0007 Collection TA0009 Exfiltration TA0010	Active Scanning T1595 Account Discovery T1087 Gather Victim Network Information T1590 Gather Victim Host Information T1592 Compromise Infrastructure: Server T1584.004 Develop Capabilities: Exploits T1587.004 Obtain Capabilities T1588 Generate Content T1683 Exploit Public-Facing Application T1190 Unsecured Credentials T1552 Valid Accounts T1078 Account Discovery T1087

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
GTG-1002	 CN	To support reconnaissance, vulnerability discovery, exploitation, lateral movement, credential harvesting, data analysis, and exfiltration operations largely autonomously	L4	Reconnaissance TA0043 Resource Development TA0042 Initial Access TA0001 Persistence TA0003 Privilege Escalation TA0004 Stealth TA0005 Credential Access TA0006 Discovery TA0007 Collection TA0009 Exfiltration TA0010	File and Directory Discovery T1083 Network Service Discovery T1046 System Information Discovery T1082 File and Directory Discovery T1083 System Network Configuration Discovery T1016 System Network Connections Discovery T1049 Automated Collection T1119 Data from Information Repositories: Databases T1213.006 Data from Local System T1005 Data Staged: Local Data Staging T1074.001 Create Account: Local Account T1136.001 Exfiltration Over Web Service T1567

Threat actor	Country of origin	Purpose of use	AIM3 Level	Tactics	Techniques
TAT26-12	 N/A	Claude Code used as an interactive exploitation assistant, conversationally directed to expand access, write exploits, build tunnels, and map victim architecture	L4	Reconnaissance  Initial Access  Execution  Persistence  Privilege Escalation  Credential Access  Stealth  Discovery  Command and Control 	Active Scanning: Vulnerability Scanning  Gather Victim Host Information  Exploit Public-Facing Application  Obtain Capabilities  Develop Capabilities  Command and Scripting Interpreter: Python  Scheduled Task/Job: Cron  OS Credential Dumping: /etc/passwd and /etc/shadow  Forced Authentication  Indicator Removal  File and Directory Discovery  Proxy 

GTG-2002

GTG-2002 is a cybercriminal group active since 2025. It used exfiltrated data and threats of public disclosure to extort victims.

In the summer of 2025, a cybercriminal tracked as [GTG-2002 used Claude's code execution environment](#) to automate reconnaissance, credential harvesting, and network penetration at scale, potentially affecting at least 17 organizations in government, healthcare, emergency services, and religious institutions. The actor employed Claude Code on Kali Linux as a comprehensive attack platform, embedding operational instructions in a CLAUDE.md file that provided persistent context for every interaction. This configuration file included a cover story claiming that the work involved network security testing conducted under official support contracts, while also providing detailed attack methodologies and target-prioritization frameworks. Using this framework, Claude could systematically track compromised credentials, pivot through networks, and optimize extortion strategies based on real-time analysis of stolen data. The actor focused on large-scale data theft and extortion. Rather than encrypting systems with traditional ransomware, the actor threatened to publish stolen data unless victims paid. The approach resulted in the compromise of personal records, including healthcare data, financial information, government credentials, and other sensitive information, with direct ransom demands occasionally exceeding \$500,000.

China-linked threat actors

In 2025, researchers identified a sophisticated Chinese threat actor that systematically leveraged Claude to enhance cyber operations targeting Vietnamese critical infrastructure. The actor integrated Claude across nearly all phases of the attack lifecycle over a nine-month campaign, using it as a technical advisor, code

developer, security analyst, and operational consultant. The actor appears to have compromised major Vietnamese telecommunications providers, government databases, and agricultural management systems. The actor primarily used Claude to enhance its operational capabilities: developing custom Python scanning tools for reconnaissance of Vietnamese IP ranges; creating sophisticated file-upload fuzzing tools and WordPress exploitation frameworks; optimizing credential harvesting operations using tools such as Hydra and hashcat; implementing privilege escalation exploits, including Linux kernel vulnerabilities; building proxy chain configurations for operational security; and analyzing reconnaissance data and planning lateral movement strategies. The actor used Claude in activities spanning 12 of the 14 MITRE ATT&CK tactics.

In mid-September 2025, researchers [detected](#) another highly sophisticated espionage campaign attributed to the Chinese state-sponsored group GTG-1002.

GTG-1002 is a Chinese state-sponsored group active since at least 2022. It conducted a large-scale, autonomous cyber espionage campaign targeting a range of sectors, with an emphasis on the military and energy sectors.

The actors manipulated Claude Code into attempting to compromise roughly 30 targets worldwide and succeeded in a small number of cases. The operation targeted large technology companies, financial institutions, chemical manufacturers, and government agencies. The threat actor used Claude Code to perform reconnaissance, vulnerability discovery, exploitation, lateral movement, credential harvesting, data analysis, and exfiltration largely autonomously. The human operator configured multiple Claude Code instances to operate together as autonomous penetration-testing orchestrators and agents, enabling AI to execute 80–90% of tactical operations independently and at request rates impossible for a human operator.

TAT26-12 used Claude Code and GPT-4.1

TAT26-12 is a threat actor group active since 2025, targeting Mexican government organizations to steal data.

Between late December 2025 and mid-February 2026, a threat actor known as TAT26-12 breached at least nine Mexican government organizations at the federal, state, and municipal levels. The threat actor exfiltrated data from databases containing hundreds of millions of citizen records, compromised hundreds of internal servers, and built both a live query API into government systems and a document forgery service. [The actor used two commercial AI platforms](#) — Anthropic's Claude Code and OpenAI's GPT-4.1 — as core operational tools. The threat actor reportedly used jailbreaking techniques to bypass the AI systems' safety guardrails by posing as security researchers participating in a bug bounty program, tricking the models into helping generate attack plans. The AI conversations were originally conducted in Spanish. Claude Code operated as an interactive exploitation assistant; the threat actor guided it through conversational prompts to expand access, write exploits, build tunnels, and map victim architecture. Approximately 75 percent of remote command execution activity across the campaign was generated and executed by Claude Code via its tool-use interface. BACKUPOSINT.py, a custom 17,550-line Python tool, piped harvested server data through OpenAI's GPT-4.1 API for automated mass analysis across 305 internal SAT servers. It produced 2,597 structured intelligence reports. The threat actor read OpenAI's reports and fed relevant findings into Claude sessions as natural-language instructions. The recovered forensic material includes: 20 tailored exploit scripts targeting 20 different CVEs; 2,597 structured intelligence reports generated by OpenAI; more than 400 custom attack scripts (301 Bash and 113 Python), including tunnel management, credential spraying, data extraction, deployment auto-

mation, operational security cleanup, and rootkits; and 1,088 individually logged attacker prompts that generated 5,317 AI-executed commands across 34 sessions on live victim infrastructure.

Other examples of AI use in threat actors' arsenals

On November 28, 2025, [an unknown threat actor compromised an AWS environment](#) using stolen credentials, escalating to administrator privileges in under 10 minutes. The attack stood out not only for its speed but also for multiple indicators suggesting that the threat actor leveraged LLMs throughout the operation to automate reconnaissance, generate malicious code, and make real-time decisions. The threat actor gained initial access to the victim's AWS account through credentials found in publicly accessible Simple Storage Service (S3) buckets. It then rapidly escalated privileges through Lambda function code injection, moved laterally across 19 unique AWS principals, abused Amazon Bedrock for LLMjacking (an attack in which the threat actor compromises a principal in the victim's cloud account to gain access to cloud-hosted LLMs), and launched GPU instances for model training.

At the end of 2025, a [Vietnamese threat actor likely used AI](#) to author the code behind a phishing campaign. The phishing emails masqueraded as job opportunities, and the threat actor may have used them as lures in the hope that recipients would open them on work computers. The threat actor may have sold access to compromised organizations to other actors. The emails carried malicious ZIP or RAR archive attachments; other samples were hosted on Dropbox, with the phishing emails likely containing links and instructions for downloading the file. When opened, the archives initiated an infection chain that led to the installation of PureRAT or HVNC. The infection chain involved several steps, one of which used malicious batch scripts that appeared highly likely to have been authored with AI. Nearly every step in the batch file had

a detailed comment in Vietnamese. Another version of the batch file had a more streamlined set of instructions and emojis, providing further evidence of AI assistance. In addition to the batch scripts, several Python loaders for the final payload were also highly likely to have been written with AI assistance. Each step in the code was numbered and accompanied by explanatory comments and debug messages in a mix of Vietnamese and English. The code even contained instructions for the threat actor, such as "Remember to paste the base64-encoded HVNC shellcode here."

January 2026

APT37 [used](#) AI-written PowerShell backdoors. The actors targeted software developers and engineering teams with expertise in, or access to, blockchain-related resources and infrastructure.

RedKitten is a newly identified campaign targeting Iranian interests, likely including non-governmental organizations and individuals involved in documenting recent human rights abuses. The [campaign](#) appears to have been developed with AI tools, as indicated by multiple traces of LLM-assisted development. The deployed implant was dubbed SloppyMIO. It retrieves its configuration steganographically from images whose URLs are obtained through a Dead Drop Resolver (DDR) backed by GitHub. SloppyMIO can retrieve and execute modules from a service such as Google Drive. The malware modules contained likely AI-generated comments in Farsi.

In early 2026, the financially motivated threat actor [Hive0163 deployed Slopoly](#), an LLM-generated PowerShell C2 framework, during a ransomware intrusion. Initial access was gained through ClickFix social engineering. Slopoly maintained persistent backdoor access for over a week, with indicators of AI authorship including verbose comments, structured error handling, and explicitly malicious variable names suggesting a successful guardrail bypass.

February 2026

Researchers [identified](#) an AI/LLM-generated malware sample exploiting the React-2Shell vulnerability. The threat actor was observed spawning a container named "python-metrics-collector," configured with a startup command that first installed prerequisite tools, including curl, wget, and Python 3. It then downloaded a list of required Python packages from Pastebin before downloading and running a Python script from its C2 server. The downloaded Python payload begins with a multiline comment: "Network Scanner with Exploitation Framework, Educational/Research Purpose Only, Docker-compatible: No external dependencies except requests." The code was likely generated by a large language model.

March 2026

APT36 has pivoted from off-the-shelf malware to "vibeware," an AI-driven development model that produces a high volume of mediocre implants. The group generated variants in less common programming languages (Nim, Zig, and Crystal) and used legitimate cloud services for C2, complicating detection and response. Researchers [identified](#) metadata within the threat actor's projects that points directly to the use of AI-integrated code editors. The binaries themselves frequently contain Unicode emojis in strings. The resulting malware was classified as AIM3 Level 2. Targeting remains highly focused on South Asian regional politics and national security. The primary victims are associated with the Indian and Afghan governments.

In 2026, the Iranian APT group Void Manticore conducted destructive wiping attacks. As the final stage of the operation, the threat actors [deployed](#) a custom Handala PowerShell-based wiper. It enumerates all files within user directories and deletes them, further compounding

the damage caused by the initial wiping activity. Based on the code structure and the detailed comments, this PowerShell script was likely developed with AI assistance.

In May 2026, researchers [identified](#) a Python script exploiting a zero-day vulnerability that allows users to bypass two-factor authentication (2FA) on a popular open-source, web-based system administration tool. The vulnerability stems from a high-level semantic logic flaw in which the developer hardcoded a trust assumption. Based on the structure and content of the exploit, researchers have high confidence that the actor used an AI model to support the discovery and weaponization of this vulnerability. For example, the script contains an abundance of educational docstrings, including a hallucinated CVSS score, and uses a structured, textbook Pythonic format highly characteristic of LLM training data, such as detailed help menus and a clean `_C ANSI` color class.

Attacks on AI users

In addition to using AI for development, threat actors repeatedly distributed malware disguised as AI tools. They used legitimate platforms to distribute the malware, giving the files an appearance of legitimacy. Examples include:

1. January 2025: [a malicious campaign](#) in the Python Package Index (PyPI) package repository. The attack targeted developers, ML engineers, and everyday AI enthusiasts who might be interested in integrating DeepSeek into their systems. The actor uploaded two packages: `deep-seek` and `deepseekai`. The functions used in these packages were designed to collect user and computer data and steal environment variables.
2. December 2025: researchers [found](#) malicious Chromium-based browser extensions that impersonate legitimate AI

assistant tools to harvest LLM chat histories and browsing data. These extensions were installed approximately 900,000 times. The extensions collected full URLs and AI chat content from platforms such as ChatGPT and DeepSeek, exposing organizations to potential leakage of proprietary code, internal workflows, strategic discussions, and other confidential data.

3. February 2026: [threat actors](#) created malicious GitHub repositories posing as OpenClaw Windows installers, including a repository named `open-claw-installer`. The repositories were published on February 2 and remained online until approximately February 10, when they were reported and removed. Bing search results directed victims to these GitHub repositories. When a victim downloaded and ran the fake installer, it did not install OpenClaw; instead, it deployed the Vidar stealer. In some cases, the loader also deployed GhostSocks, turning the victim's system into a residential proxy node through which criminals could route traffic to conceal their activity. The campaign broadly targeted users attempting to install OpenClaw, and the malicious repositories contained download instructions for both Windows and macOS.

Attacks on AI solutions

In addition to using AI to conduct attacks, threat actors periodically target AI solutions themselves. Researchers have repeatedly reported critical vulnerabilities in various AI solutions, including CVE-2025-32711 (a.k.a. EchoLeak) in Microsoft 365 Copilot, CVE-2025-54135 (a.k.a. CurXecute) in Cursor IDE, CVE-2025-53109 in Model Context Protocol, CVE-2025-8217 in the Amazon Q Devel-

oper VS Code extension, and CVE-2025-34291 in Langflow AI. The last two were observed being exploited in the wild.

CVE-2025-8217 is a remote code execution vulnerability in the Amazon Q Developer VS Code extension. Amazon Q Developer for VS Code Extension version 1.84.0 had an inappropriately scoped GitHub token in its CodeBuild configuration. With that access token, [the threat actor](#) was able to commit malicious code to the extension's open-source repository, which was automatically included in a release. The injected file contained instructions that could wipe local development environments and cloud infrastructure. The malicious code was distributed with the extension but failed to execute due to a syntax error. This prevented the malicious code from making changes to any services or customer environments.

Active exploitation of another vulnerability, CVE-2025-34291, was observed beginning on January 23, 2026. Langflow versions up to and including 1.6.9 contain a chained vulnerability that enables account takeover and remote code execution.

In one case, the PRC-backed threat actor [APT41 unsuccessfully attempted to use](#) Gemini to learn about Gemini's underlying infrastructure and systems. The actor asked Gemini to share details such as its IP address, kernel version, and network configuration. Gemini responded but did not disclose sensitive information. Instead, its responses provided only widely available public information and noted that the requested details were kept secret to prevent unauthorized access.

One of the most high-profile events of 2026 was [the Mini Shai-Hulud supply chain campaign](#), attributed to TeamPCP. In early May 2026, threat actors injected the malicious Mini Shai-Hulud implant into 84 compromised npm packages in the tanstack namespace, primarily to steal credentials. The campaign affected at least Mistral and OpenAI, among other organizations:

- OpenAI has confirmed that two employee devices within its corporate environment were compromised as part of the TanStack supply chain attack. Investigators observed unauthorized access and credential-theft activity within a limited subset of internal source code repositories accessible to the two affected employees. A small portion of credentials was exfiltrated from these repositories, and no other information or code was affected. The affected source code repositories included signing certificates for products including iOS, macOS, and Windows;
- Mistral was affected by a supply chain attack stemming from the compromise of the third-party TanStack software. An automated worm associated with the attack led to the publication of compromised npm and PyPI SDK versions. A developer device was affected. The threat actor group put nearly 450 internal Mistral AI repositories up for sale on a dark web forum, demanding \$25,000 for what it claimed was roughly 5 GB of stolen source code.




Fig. 1 — Forum post advertising the sale of Mistral AI data

AI Mentions on the Dark Web

F6 Threat Intelligence analysts reviewed posts mentioning popular AI solutions across various underground forums and identified the following distribution over the reporting period:

AI solution	Mentions
ChatGPT	108,470
Claude	25,953
Gemini	24,556
DeepSeek	5,178
OpenClaw (Clawbot)	1,019
Qwen	890

On these forums, AI was typically mentioned in discussions of each solution's general capabilities and in the context of prompt creation; however, a significant share of posts related directly to malicious activity. Posts on well-known cybercriminal platforms covered topics such as using AI to find instructions for configuring attacker infrastructure and methods for evading detection by popular antivirus solutions; using AI to bypass KYC procedures on popular cryptocurrency exchanges; discussing known AI-enabled attacks; developing and distributing malware; and using AI to create phishing pages that imitate online services.

An example of a discussion about using AI to develop malware on a well-known underground forum is shown below:




Fig. 2 — Screenshot of a discussion about using AI to develop malware on an underground forum

Translated text from Fig. 2

Greetings, shadow colleagues. A new player here, but one who brings old experience.

Have you noticed how modern LLMs are turning into fussy little brats? Step just a bit off to the side into gray areas and Claude Opus starts moralizing, wringing its hands, and getting triggered by every little thing. They try to pen our kind into a corral of "ethics" and "safety" - even though it's our hand that feeds them, really))

But here's the truth: the machine is a dumb auto-complete. And we're the ones who can bend it. If an AI tells you "I can't," it means only one thing - you didn't give it the right role. I've spent years working in black-hat niches where every other request is a "red flag" for modern AI. And I learned to break them. Today I'll drop my private tricks - the ones my profit rests on. No filler, no upsell (I make my money elsewhere). Take them:

Trick #1: The "Trojan Horse" Never go in head-on. If you ask it right away to whip up a phishing landing page or a RAT - you'll get a refusal.

The first 2-3 requests should be a neutral audit.

The three AI solutions most frequently mentioned on the forums correspond to those most commonly observed in attacks by cybercriminal groups: Gemini, OpenAI tools, and Claude Code. As noted earlier, Gemini was used more often in major attacks, while ChatGPT was discussed more often on the forums.

This may point not only to ChatGPT's greater prevalence on underground, but also to its potential use by low-skilled threat actors for vibe-coding.



Attacks Using Zero-day Vulnerabilities

A high-profile zero-day development in 2026 was the dispute between researcher Nightmare Eclipse and Microsoft. From April to May, the researcher disclosed six zero-day vulnerabilities in Windows together with exploits. Subsequently, unidentified threat actors attempted to exploit three of them in the wild.

Threat actors most often use zero-days in attacks targeting PCs rather than mobile devices. This can be partly explained by the continued effectiveness of social-engineering-based mobile attacks and by the disproportionate difficulty of developing a zero-day relative to the potential payoff.

In most cases, in-the-wild exploitation of a zero-day is detected within a short time, but there are exceptions. For example, CVE-2025-21042 remained in threat actors' arsenals for nearly nine months.

Zero-days are most commonly used by cyberespionage groups and less often by financially motivated actors. The most damaging campaign was conducted by CI0p, which exploited CVE-2025-61882 in Oracle EBS. It is worth noting that another group, Scattered Lapsus\$ Hunters, published files from the CI0p attack on its Telegram channel, including the exploit itself. The threat actors are likely connected in some way. A few months later, some victims of the Oracle EBS compromise were listed on the CI0p data leak site.

The table below lists zero-day vulnerabilities identified in the wild in 2025 and 2026.

Zero-day vulnerabilities used in attacks, 2025–2026

Threat actor	Date of exploitation	CVE	Vulnerable software	Targeted countries
Silk Typhoon	January 2025	CVE-2025-0282	Ivanti Pulse Connect VPN	 N/A
Unknown TA	November 2025	CVE-2025-20337	Cisco Identity Service Engine	 N/A
		CVE-2025-5777	Citrix NetScaler ADC and NetScaler Gateway	 N/A
Unknown TA	July 2025	CVE-2025-53770	Microsoft SharePoint Server	 N/A
		CVE-2025-53771	Microsoft SharePoint Server	 N/A

Threat actor	Date of exploitation	CVE	Vulnerable software	Targeted countries
Play, Storm-2460	March 2025 - April 2025	CVE-2025-29824	Windows Common Log File System (CLFS) kernel driver	 US,  VE,  ES,  SA
Unknown TA	March 2026	CVE-2026-3502	TrueConf client application	 Southeast Asia
Unknown TA	February 2025	CVE-2025-3928	Commvault Web Server	 N/A
Unknown TA	Early 2025	CVE-2025-27915	Zimbra Collaboration Suite (ZCS)	 BR
CI0p	August 2025	CVE-2025-61882	Oracle EBS	 US,  ZA
CL-UNK-1054	July 2024 - April 2025	CVE-2025-21042	Samsung's Android image processing library	 IQ,  IR,  TR,  MA
UAT-9686	November 2025	CVE-2025-20393	Cisco Secure Email Gateway and Cisco Secure Email and Web Manager	 N/A
UAT-8837	September 2025	CVE-2025-53690	Sitecore Products	 North America
Unknown TA	April 2026	BlueHammer (CVE-2026-33825), RedSun (CVE-2026-41091), UnDefend (CVE-2026-45498)	Microsoft Windows Defender	 N/A

More detailed information on attacks using zero-day vulnerabilities is provided below.

CVE-2025-0282

Silk Typhoon is a pro-China cyber espionage group active since at least 2021. The group's primary objective is collecting politically and strategically significant information across government, IT, healthcare, and defense sectors.

In January 2025, Silk Typhoon [was observed](#) exploiting a zero-day vulnerability in the public-facing Ivanti Pulse Connect VPN (CVE-2025-0282). CVE-2025-0282 is an unauthenticated stack-based buffer overflow. Successful exploitation could result in unauthenticated remote code execution, leading to potential downstream compromise of a victim network.

CVE-2025-20337 and CVE-2025-5777

In November 2025, researchers [disclosed](#) a campaign conducted by an unidentified APT group. The threat actors used custom malware and two vulnerabilities in Cisco Identity Services Engine (CVE-2025-20337) and Citrix systems (CVE-2025-5777) as zero-days. Access to these unpublished zero-day exploits indicates a highly resourced threat actor with advanced vulnerability research capabilities or access to non-public vulnerability information. Following successful exploitation, the actor deployed IdentityAuditAction, a custom web shell disguised as a legitimate Cisco ISE component. The custom-built backdoor had advanced evasion capabilities specifically designed for Cisco ISE environments.

CVE-2025-53770 and CVE-2025-53771

On July 17, 2025, researchers [began detecting attacks](#) using a zero-day vulnerability affecting on-premises Microsoft SharePoint Servers. This flaw has since been assigned the identifier CVE-2025-53770, along with an accom-

panying bypass tracked as CVE-2025-53771. The attackers relied on in-memory .NET module execution, avoiding traditional file-based artifacts entirely. Initial exploitation was observed against high-value organizations, with victims primarily in technology consulting, manufacturing, critical infrastructure, and professional services tied to sensitive architecture and engineering organizations. That activity occurred before public disclosure of the vulnerability spurred mass exploitation efforts from a wider set of actors.

CVE-2025-29824

Attackers linked to the **Play ransomware** operation [deployed](#) a zero-day privilege-escalation exploit during an attempted attack against an organization in the U.S. The attack occurred before the disclosure and patching of CVE-2025-29824, a Windows elevation-of-privilege zero-day vulnerability in the Common Log File System driver (clfs.sys), on April 8, 2025. The attackers deployed Grixba, a custom infostealer.

The exploit for CVE-2025-29824 was also [used by Storm-2460](#) to deploy ransomware. Storm-2460 is a financially motivated threat actor that uses zero-day vulnerabilities for ransomware deployment.

When it patched the vulnerability, Microsoft said that it had been exploited against a small number of targets, including organizations in the **U.S., Venezuela, Spain, and Saudi Arabia**. By exploiting CVE-2025-29824, the actor was able to escalate its privileges to NT AUTHORITY\SYSTEM to perform lateral movement and encrypt victims' files.

CVE-2026-3502

Researchers [identified](#) a zero-day vulnerability in the TrueConf client application, tracked as CVE-2026-3502. This vulnerability has been exploited in the wild as part of a targeted

campaign called “TrueChaos” against government entities in Southeast Asia. The threat actor abused the TrueConf update mechanism to deploy the Havoc payload to vulnerable machines.

CVE-2025-3928

Commvault [has revealed](#) that in February 2025, an unknown state-sponsored threat actor breached its Microsoft Azure environment by exploiting CVE-2025-3928 as a zero-day. If successfully exploited, the flaw could enable remote, authenticated attackers to gain unauthorized access to the vulnerable systems and allow them to deploy and execute arbitrary web shells. This activity affected a small number of its customers.

CVE-2025-27915

In early 2025, an unidentified threat actor impersonating the Libyan Navy’s Office of Protocol launched [a sophisticated cyberattack](#) against Brazil’s military, exploiting a zero-day vulnerability (CVE-2025-27915) in Zimbra Collaboration Suite (ZCS). The campaign weaponized a seemingly benign ICS (iCalendar) file containing malicious JavaScript, which executed upon viewing, enabling attackers to steal credentials, manipulate email communications, and exfiltrate sensitive data.

CVE-2025-61882

Beginning September 29, 2025, researchers began tracking [a new, large-scale extortion campaign](#) by a threat actor claiming affiliation with the **CI0p** extortion brand.

CI0p is a financially motivated group active since at least 2018. It exploits zero-day vulnerabilities and conducts supply chain attacks.

The actor began sending a high volume of emails to executives at numerous organiza-

tions, alleging the theft of sensitive data from the victims’ Oracle E-Business Suite (EBS) environments. The extortion emails indicated that alleged victims could prevent the release of stolen data in exchange for payment. The threat actors exploited CVE-2025-61882 as a zero-day vulnerability against Oracle EBS customers as early as Aug. 9, 2025, weeks before a patch was available. In some cases, the threat actor successfully exfiltrated a significant amount of data from affected organizations. A few months later, some victims of the Oracle EBS compromise were listed on the CI0p leak site. The organizations that were the first to be named, such as Harvard University, South Africa’s Wits University, and American Airlines subsidiary Envoy Air, confirmed that they had been affected shortly after they were named by the attackers in mid-October. Other alleged victims include companies in the mining, professional services, wastewater, construction, insurance, financial, manufacturing, transportation, technology, automotive, energy, and HVAC sectors.

CVE-2025-21042

Researchers [uncovered](#) a campaign targeting Samsung Android devices that was attributed to CL-UNK-1054. The CL-UNK-1054 group has been active since 2024. The actors used a previously unknown Android spyware named LANDFALL, which has advanced capabilities for stealth, persistence, and comprehensive data collection from mobile devices. The LANDFALL campaign had been operating since mid-2024, using the zero-day Android/Samsung vulnerability (CVE-2025-21042) in Samsung’s Android image processing library months before it was fixed. The attackers attempted to deliver the exploit through a messaging application.

CVE-2025-20393

Cisco Talos [is tracking](#) active [exploitation](#) of a vulnerability in Cisco AsyncOS Software for Cisco Secure Email Gateway and Cisco Secure Email and Web Manager. The vulnerability enables threat actors to execute system-level commands and deploy a persistent Python-based backdoor, AquaShell. The as-yet-unpatched vulnerability is tracked as CVE-2025-20393 and carries a CVSS score of 10.0. Cisco became aware of this activity on December 10; it has been ongoing since at least late November 2025. Additional tools observed include AquaTunnel (a reverse SSH tunnel), chisel (another tunneling tool), and AquaPurge (a log-clearing utility). This activity is attributed to UAT-9686, a Chinese state-sponsored threat actor that has conducted cyberespionage since at least 2025.

CVE-2025-53690

UAT-8837 [exploited](#) a ViewState deserialization zero-day vulnerability in Sitecore products, CVE-2025-53690, to gain initial access. Since at least 2025, the group has clearly focused on targets in critical infrastructure sectors across North America. [UAT-8837](#) may use a variety of tooling over the course of an intrusion, including GoTokenTheft, EarthWorm, DWAgent, SharpHound, Impacket, GoExec, Rubeus, and Certipy.

Nightmare-Eclipse tooling used in attacks

One of the most widely discussed zero-day developments of 2026 involved a researcher [publicly disclosing six zero-day vulnerabilities](#) in April and May before patches were available. The researcher attributed the disclosures to a dispute with Microsoft and dissatisfaction with its bug bounty program. Between April and May, the researcher published details and PoCs for the following zero-day vulnerabilities: BlueHammer (CVE-2026-33825), RedSun (CVE-2026-41091), UnDefend (CVE-2026-45498), YellowKey (CVE-2026-45585), GreenPlasma, and MiniPlasma. Following the April disclosure, threat actors attempted to exploit the first three vulnerabilities (BlueHammer, RedSun, and UnDefend) in the wild. Although the exploits were executed, none appears to have succeeded during the incident. The threat actor made a series of operational errors during the intrusion. No exploitation of the other three vulnerabilities has been publicly disclosed.

In response, Microsoft [issued](#) a statement on May 27 condemning the researcher's actions and indirectly warning of potential legal action: "Our Digital Crimes Unit will continue bringing cases against these actors and those that enable their criminal activity – coordinating as needed with law enforcement around the world."

Sale of zero-day vulnerabilities on the Dark Web

During the reporting period, F6 Threat Intelligence analysts identified 65 unique listings advertising zero-day exploits for sale across various underground forums.

This analysis considered only zero-day exploit sales that included additional details about the affected product, operating system or specified the type of vulnerability. The actual

number of zero-day exploit listings on underground resources is substantially higher because of numerous duplicates with identical or nearly identical descriptions, as well as announcements that provide no additional details specifying promoted exploits for sale.

The zero-day exploits offered for sale during the reporting period were broken down by type as follows:

Exploit	Count
RCE (Remote Code Execution)	21
LPE (Local Privilege Escalation)	9
LCE (Local Code Execution)	1
RCE+LPE	1

Vulnerability classifications were specified in only 32 of the 65 documented zero-day exploit advertisements; for the remaining 33 instances, this information remained undisclosed.

According to the listings, the identified zero-day vulnerabilities for sale affected the following systems and infrastructure components:

Operating systems (endpoint & infrastructure OS)

- Mobile platforms: Android (v. 11–16), iOS (v. 17.4);
- Infrastructure and desktop platforms: Linux, macOS, Windows.

Enterprise software, business platforms, and DBMS (Enterprise Software & Apps)

- ERP: Oracle E-Business Suite;
- Web Hosting Control Panels & Infrastructure: cPanel, Web Panel, Windows IIS, WordPress;

- Managed File Transfer (MFT) and file servers: CrushFTP, Fortra GoAnywhere;
- Remote access, orchestration, and development: DWAgent (DWService), GitHub;
- Application software: Microsoft Excel, WinRAR.

Network equipment, perimeter components, and VPN (Network Security & Edge Assets)

- Enterprise-grade firewalls and VPN gateways: Fortinet (Fortigate / FortiVPN), Palo Alto Networks (Global Protect), Etrolink (AnyLink Gateway);
- Network Access Control (NAC) & Authentication (AAA): Cisco ISE (Identity Services Engine);
- Network equipment (telecom providers): Calix GPON;
- Network-attached storage (NAS): TerraMaster TOS;
- SOHO routers: ASUS (RT-AX.XXX series).

Cloud web services, fintech, and messaging apps (Web Services & Communication)

- Financial services: attijariwafa (banking), bitmart.com (cryptocurrency exchange);
- Messaging apps: WhatsApp, Tox (including the qTox client).

Other

General Web Application Logic and Web Authorization Interface Vulnerabilities (Web Logic): zero-day Web Vulnerability.

When exploit sales are discovered on the dark web, it is rarely possible to determine whether the exploit for sale is genuine. Because such offerings command extremely high prices, many threat actors attempt to sell publicly disclosed vulnerabilities under the guise of unique zero-day exploits. In some cases, sellers are outright scammers seeking to profit from a one-time sale of a nonexistent exploit or vulnerability, or to attract public attention, particularly from security professionals, in order to mislead.

F6 Threat Intelligence analysts used the following system to rate the credibility of zero-day sale listings posted on underground platforms during the reporting period:

- “1” — the sale is very likely genuine: the threat actor has a strong reputation on the underground forum and no open disputes, may have previously participated in related vulnerability-sale threads, and there is evidence that the actor has successfully sold a vulnerability to other forum members;
- “2” — the authenticity of the vulnerability for sale cannot be reliably confirmed: the seller has a good overall reputation but, for example, has never publicly sold exploits or vulnerabilities before;
- “3” — questionable: the threat actor recently joined the underground marketplace, had no prior observed activity, or was involved in a dispute on a related topic that was resolved against the seller;
- “4” — there is no information confirming that the seller possesses the vulnerability: the thread was deleted, the forum was shut down, or the threat actor ceased all activity immediately after posting the listing.

Statistics on zero-day sale listings identified over the reporting period, based on the rating system above, is presented below:

Rating criterion	Number of zero-day sales matching this criterion
1	10
2	22
3	7
4	26

The price of the exploits advertised on underground forums ranged from a few thousand to one million dollars, representing a cumulative market volume over \$2.6 million. Accordingly, the distinct utility and premium valuation of these underground market items demonstrate a strategic alignment toward a narrow segment of highly funded adversaries.

Potential buyers of zero-day exploits may include APT groups and intelligence services; financially motivated groups; intermediaries and shadow brokers that acquire high-value exploits in bulk for resale to cybercriminal groups; initial access brokers (IABs) that exploit vulnerabilities to obtain corporate network access at scale for later resale; and other underground actors affiliated with cybercriminal syndicates.

Although such threads generally contain no public information indicating whether an exploit was purchased on any underground forum, in some cases it is possible to infer the potential buyer indirectly.

For example, attacks by GOFEE identified in August 2025 apparently exploited CVE-2025-8088. Earlier, on July 7, 2025, an announcement appeared on a popular underground forum offering an allegedly ready-to-use exploit for a WinRAR zero-day vulnerability for \$80,000.



Fig. 3 — Screenshot of the announcement offering an exploit for the WinRAR zero-day vulnerability for sale




Fig. 4 — Screenshot of zeroplayer's post announcing the allegedly successful sale of the exploit for the WinRAR zero-day vulnerability

Subsequently, on August 5, 2025, the threat actor posted a message on the forum that the offer was no longer available. This may indicate that the actor under the alias zeroplayer could have sold the exploit to an anonymous buyer.

A vulnerability of this type was used in GOFFEE attacks around the time the exploit may have been purchased from the threat actor zeroplayer. Therefore, it is quite likely that GOFFEE members could have acquired the exploit on an underground forum and adapted it for use in their attacks.

Supply Chain and Trusted Relationship Attacks

Supply chain and trusted relationship attacks have become increasingly popular among cybercriminals over the past several years.

Over the past year, activity shifted from narrowly targeted supply chain attacks to mass-scale campaigns. Previously, supply chain compromises affected only the customers of the compromised organization; now, threat actors increasingly compromise open-source components, simultaneously jeopardizing the software development processes of many companies.

Attacks involving poisoned packages in repositories such as GitHub are especially common. They pose a threat to applications built with those packages and, in turn, to the applications' users. Since the second half of 2025, supply chain campaigns targeting npm packages through maintainer phishing have been observed with increasing frequency. Threat actors reportedly launched a targeted phishing campaign to compromise Node Package Manager (npm) maintainer accounts and inject malicious code into widely used packages.

TeamPCP was the most prominent actor behind this type of attack in 2026. In addition to conducting large-scale attacks, TeamPCP was notable for collaborating simultaneously with several groups, particularly LAPSUS\$ and Shiny-Hunters. Data obtained through TeamPCP's supply chain attacks was published on those groups' DLSSs.

The unusual surge in such attacks forced companies to reassess their defenses. In May 2026, GitHub [introduced](#) the following change for npm: staged publishing became generally available. Instead of a direct publish that immediately makes a package version available to consumers, the prebuilt tarball is uploaded to a staging queue,

where a maintainer must explicitly approve it before it becomes installable. A human maintainer must complete a 2FA challenge to approve a staged package before it is released to the registry.

One of the largest trusted relationship attacks was the UNC6395 campaign conducted from March through August 2025. The attack had a massive impact, affecting an estimated 700 organizations.

Supply Chain Attacks

Threat actor	Date	Compromised supplier or software	Final payload or impact	Targeted organizations/countries
Unknown TA	11.2024-03.2025	multiple GitHub Actions	Theft of CI/CD secrets	Primary target — Coinbase; secrets leaked from 218 unnamed repositories
Lotus Blossom	06.2025-12.2025	Notepad++	Cobalt Strike	 VN,  SV,  AU,  PH
Unknown TA	04.2026-05.2026	Daemon Tools	Information collector, minimalistic backdoor, QUIC RAT	 BR,  TR,  ES,  DE,  RU,  FR,  IT,  CN
Lazarus	02.2025	Safe{Wallet}	Theft of approximately 400,000 ETH (valued at \$1.4 to \$1.5 billion)	Bybit
N/A	05.2025	RVTools	N/A	 N/A
N/A	05.2025	npm - rand-user-agent	N/A	 N/A
N/A	06.2025	npm - React Native ARIA and gluestack-ui packages	N/A	 N/A
Supply chain campaign targeting npm packages through maintainer phishing	07.2025	npm -eslint-config-prettier, eslint-plugin-prettier, synckit, @pkgr/core, napi-postinstall, 'is' package, got-fetch	Scavenger, Pycoon	 N/A
Operation Singularity	08.2025	npm - Nx	Phase 1 — secrets belonging to more than 1,700 users were publicly exposed, and more than 20,000 files were leaked.	 N/A

Threat actor	Date	Compromised supplier or software	Final payload or impact	Targeted organizations/countries
Operation SIngularity	08.2025	npm - Nx	Phase 2 — a threat actor used GitHub tokens leaked in Phase 1 to make victims' private repositories public; at least 480 compromised accounts made more than 6,700 private repositories public. Phase 3 — two compromised users publicly uploaded more than 500 repositories belonging to a single organization	 N/A
Supply chain campaign targeting npm packages through maintainer phishing	09.2025	npm - backlash, chalk-template, supports-hyperlinks, has-ansi, simple-swizzle, color-string, error-ex, color-name, is-arrayish, slice-ansi, color-convert, wrap-ansi, ansi-regex, supports-color, strip-ansi, chalk, debug, ansi-styles	The malware intercepts crypto and web3 activity in the browser, manipulates wallet interactions, and rewrites payment destinations so that funds and approvals are redirected to attacker-controlled accounts	 N/A
Shai-Hulud 1.0	09.2025	Over 500 packages were infected	Shai-Hulud is designed to steal sensitive data, expose organizations' private repositories, and hijack victim credentials to infect other packages and spread further	 N/A
Shai-Hulud 2.0	11.2025	25,000 malicious repositories across roughly 350 unique users	Shai-Hulud 2.0 steals secrets such as GitHub secrets, cloud credentials, and local files. If exfiltration fails, the malware triggers a destructive payload that wipes user files	 RU,  IN,  VN,  BR,  CN,  TR,  FR,  BG,  CY,  HK
N/A	03.2026	npm - Axios package	N/A	 N/A

Threat actor	Date	Compromised supplier or software	Final payload or impact	Targeted organizations/countries
TeamPCP	03.2026	Trivy, Checkmarx, LiteLLM, Telnyx	The malicious code functioned as a data-stealing framework, targeting SSH keys, AWS/Azure/GCP credentials, GitHub tokens, and developer environment files. Data exfiltrated from 500,000 infected machines (approximately 300 GB of data) The ShinyHunters extortion group published roughly 91.7 GB (compressed) of data stolen from the European Commission on its DLS LAPSUS\$ published an archive (approximately 96 GB) of data stolen in the attack on Checkmarx	 N/A
TeamPCP	04.2026	Bitwarden CLI	The package was downloaded thousands of times before being flagged	 N/A
TeamPCP - Mini Shai-Hulud campaign	04.2026-05.2026	Over 170 npm packages and 2 PyPI packages across TanStack, Mistral AI, UiPath, OpenSearch, and Guardrails AI	The malicious payload steals AWS credentials, GitHub tokens, npm tokens, Vault secrets, and SSH keys, and includes a self-propagating worm that spreads through CI/CD pipelines.	 N/A
TeamPCP	05.2026	Nx Console VS Code extension	Roughly 3,800 GitHub-internal repositories were exfiltrated and offered for sale for more than \$50,000 on cybercrime forums	 N/A

In March 2025, the compromise of the GitHub Action `tj-actions/changed-files` and additional actions within the `reviewdog` organization came to light [Source 1](#), [Source 2](#), [Source 3](#), [Source 4](#). This supply chain attack began in November 2024 but only surfaced months later. `Reviewdog` was compromised through a `reviewdog` maintainer's leaked PAT. The PAT was leaked by the threat actor from the `spotbugs/spotbugs` repository. The threat actor pushed a malicious GitHub Actions workflow file to the `spotbugs/spotbugs` repository, creating a malicious workflow run in the repository's context. The threat actor used this malicious workflow to leak all of `spotbugs/spotbugs`'s secrets, which included that PAT. The PAT likely had access to both `spotbugs/spotbugs` and `reviewdog/action-setup`, with sufficient permissions to push tags to the `reviewdog/action-setup` repository. This allowed the threat actor to override the `v1` tag in the repository and point it to a malicious commit. After this, `tj-actions/changed-files`'s CI workflow was invoked. The malicious code stole a GitHub token that had write permission to `tj-actions/changed-files`. Using this token, the threat actor overrode a tag in `tj-actions/changed-files`, pointing it to a malicious commit made in a fork by the threat actor, specifically targeting the `coinbase/agentkit` repository. Next, `coinbase/agentkit`'s CI workflow executed, consuming the malicious tag from `tj-actions/changed-files` and leaking the credentials to the threat actor. Coinbase was alerted that its CI was consuming malicious code and removed the vulnerable workflow. The threat actor overrode all tags in `tj-actions/changed-files`, pointing them to a malicious commit, which caused all workflow secrets to be printed to the log by consumers of `tj-actions/changed-files`. Researchers determined that the initial attack targeted the Coinbase crypto exchange. The actual impact was smaller than anticipated: 218 repositories leaked secrets, and the majority of those were short-lived `GITHUB_TOKENs`, which expire once a workflow run completes.

From June to December 2025, the presumably Chinese-speaking group [Lotus Blossom](#) conducted [a supply chain attack](#) by compromising the `Notepad++` update infrastructure. The compromise occurred at the hosting provider level and allowed the threat actors to intercept and redirect update traffic destined for `notepad-plus-plus.org`. The threat actors frequently updated the C2 addresses from which the malicious updates were distributed and used several infection chains. Targets included financial, government, and telecom organizations in the following countries: Vietnam, El Salvador, Australia, and the Philippines.

In early May 2026, researchers identified [an active supply chain attack](#) targeting the official [Daemon Tools](#) website. Artifacts suggest that the threat actor behind the attack is Chinese-speaking. The compromised installer delivered malware alongside the legitimate application, enabling the threat actors to execute arbitrary commands and remotely control infected devices. The threat actors had been distributing the modified software directly through the vendor's primary domain since April 8, 2026, concealing the malware with a valid developer digital certificate. The vendor's infrastructure had been compromised, affecting certain installation packages in the build environment and resulting in compromised releases. Most victims were located in Russia, Brazil, Türkiye, Spain, Germany, France, Italy, and China. On a small subset of just over ten machines belonging to organizations in the retail, scientific, government, and manufacturing sectors, the threat actors manually deployed additional payloads, including a shellcode injector and previously unknown RATs. Researchers identified the following payloads: an information collector, a minimalist backdoor, and a QUIC RAT.

Lazarus

In February 2025, the Bybit crypto exchange [was compromised](#) by the Lazarus group. The threat actors stole approximately 400,000 ETH (valued at \$1.4 to \$1.5 billion). The incident occurred when an ETH multisig transaction was facilitated through Safe{Wallet} from a cold wallet to a warm wallet, during which the threat actors intervened and manipulated the transaction. The root cause of the compromise was malicious JavaScript code on Safe{Wallet}'s platform; no vulnerability was detected in Bybit's infrastructure. The attack involved the compromise of a Safe{Wallet} developer's laptop and the hijacking of AWS session tokens to access Safe{Wallet}'s AWS infrastructure. JavaScript resources hosted on the AWS S3 bucket serving Safe{Wallet}'s web interface were modified and injected with malicious code that manipulated transactions. The malicious code included an activation condition set to execute the transaction manipulation only on a specific Bybit cold wallet. Bybit initiated a transaction from the targeted cold wallet using Safe{Wallet}'s web interface. The transaction was manipulated, and the threat actors siphoned the funds from the cold wallets.

On May 13, 2025, researchers [discovered](#) a malicious version of the RVTools utility. The file was significantly larger than the original RVTools file and contained a malicious version.dll. Around the same time, the RVTools website went temporarily offline. When it came back online, the download had changed: the file size was smaller, and the hash now matched the clean version listed on the site. This strongly suggested that the website hosting the installer had been compromised and that the malicious file had been replaced after detection. It appeared to be a temporary, targeted supply chain compromise. The malware appeared to be a custom variant of the Bumblebee loader.

In addition, other researchers observed the distribution of a trojanized RVTools installer via a malicious typosquatted domain.

In May 2025, the npm package random-agent [was compromised](#). The threat actor obtained an outdated, unprotected MFA automation token belonging to an employee and used it to publish releases to npm.

On June 6, 2025, the React Native ARIA and [gluestack-ui npm packages were compromised](#) in a sophisticated supply chain attack that affected 17 packages. A threat actor gained unauthorized access to publishing credentials and injected a RAT into widely used JavaScript packages.

In July 2025, several popular npm packages were compromised and modified to deliver Scavenger malware. An unknown threat actor modified the contents of multiple npm packages to advance its objectives. The compromise followed a successful credential-phishing campaign targeting an npm package maintainer that used a spoofed login page on a typosquatted npm domain. The following packages were published as part of the campaign: eslint-config-prettier, eslint-plugin-prettier, synckit, @pkgr/core, and napi-postinstall. On July 19, 2025, malicious versions of the 'is' package were published to npm. The threat actors first compromised a former maintainer of the package through phishing and then manipulated npm's ownership system to gain control. The same phishing campaign also compromised the npm package got-fetch, using it to distribute a different malware family: the Pycoon infostealer.

August 2025 — Operation S1ngularity involved the publication of multiple malicious versions of the widely used Nx build system package to the npm registry.

September 2025 — The attack was initiated through a [targeted phishing campaign](#) against a package maintainer known as "qix." The maintainer inadvertently gave the threat actors access to their account. Using this access, the threat actors published malicious updates to 18 npm packages. The packages were updated to contain code that would execute on the client side

of a website, silently intercepting crypto and web3 activity in the browser, manipulating wallet interactions, and rewriting payment destinations so that funds and approvals are redirected to attacker-controlled accounts without any obvious indication to the user.

September 2025 — Shai-Hulud 1.0

On September 16, 2025, researchers [discovered](#) a new wave of npm package infections caused by the self-propagating malware known as Shai-Hulud. Shai-Hulud is designed to steal sensitive data, expose organizations' private repositories, and hijack victim credentials to infect other packages and spread further. Over 500 packages were infected in this incident. As a result, developers who integrated these malicious packages into their projects risked losing sensitive data, and their own libraries could become infected with Shai-Hulud. This self-propagating malware takes over accounts and steals secrets to create new infected modules, spreading the threat along the dependency chain.

November 2025 — Shai-Hulud 2.0

On November 24, 2025, a sophisticated supply chain attack dubbed "Shai-Hulud 2.0" (also known as "The Second Coming") compromised the npm ecosystem in one of the largest coordinated attacks on open-source software. Shai-Hulud 2.0 is a two-stage, worm-like malware that spreads by compromising npm tokens to republish trusted packages with a malicious payload. More than 800 npm packages were infected by this version of the worm. Shai-Hulud 2.0 is built to harvest secrets from various environments. Upon execution, it immediately searches several sources for sensitive data, including GitHub secrets, cloud credentials, and local files. To exfiltrate the stolen data, the malware sets up a communication channel through a public GitHub repository. If it fails to exfiltrate the data, it triggers a destructive payload that wipes user files, primarily those in the home directory.

TeamPCP

TeamPCP is a threat actor active since September 2025 whose activity expanded significantly in 2026 across npm, Docker Hub, GitHub Actions, and VS Code extensions.

Timeline of the most notable supply chain attacks by TeamPCP:

- 03/19/2026 — Trivy compromised; the threat actors exfiltrated credentials from every CI/CD pipeline running Trivy;
- 03/20/2026 — using tokens stolen from Trivy users, TeamPCP published the CanisterWorm backdoor across 46+ npm packages;
- 03/23/2026 — attack on Checkmarx;
- 03/24/2026 — attack on LiteLLM;
- 03/27/2026 — telnyx package in the PyPI repository;
- 04/22/2026 — a malicious npm package impersonating the legitimate Bitwarden CLI was published;
- 05/11/2026 — the "Mini Shai-Hulud" campaign;
- 05/18/2026 — Nx Console VS Code extension.

On March 27, the European Commission publicly disclosed the incident, in which the threat actors gained initial access through the Trivy supply chain compromise. A significant volume of data (about 91.7 GB compressed) was exfiltrated from the compromised AWS account, including personal data such as names, email addresses, and email content. On March 28, the data extortion group ShinyHunters made the stolen data publicly available on their dark web leak site.




Fig. 5 — Screenshot from ShinyHunters DLS

On March 25, 2026, LAPSUS\$ threat actors published an archive of about 96 GB of data stolen from Checkmarx on their DLS. The company confirmed that TeamPCP had stolen the data during the incident.




Fig. 6 — Screenshot from LAPSUS\$ DLS

On May 18, 2026, a compromised version of the Nx Console VS Code extension was published to the Visual Studio Code Marketplace. The malicious version was live for approximately 11 minutes before the Nx team pulled it from the marketplace. The compromised extension silently fetched and executed an obfuscated payload from a dangling orphan commit hidden inside the official nrwl/nx GitHub repository. This compromise resulted from a recent supply chain attack in which threat actors stole a contributor's GitHub personal access token. The payload is a multi-stage credential stealer and supply chain poisoning tool. It harvests tokens and secrets from GitHub, npm, AWS, HashiCorp Vault, Kubernetes, and 1Password, then exfiltrates them over three independent channels. It also installs a persistent Python backdoor on macOS.

On May 20, 2026, GitHub confirmed that the threat actors gained unauthorized access to its internal source code repositories after a poisoned

Visual Studio Code extension compromised an employee endpoint. The company currently assesses with confidence that roughly 3,800 GitHub-internal repositories were exfiltrated. The threat actor known as TeamPCP claimed responsibility on cybercrime forums and offered the data for more than \$50,000. The threat actor later stated that they had received an offer of \$95,000.




Fig. 7 — Post advertising the sale of data belonging to GitHub

TeamPCP collaborates with other cybercriminal groups, particularly ShinyHunters and LAPSUS\$. In May 2026, TeamPCP also gave an interview to the Ransomware Interviews project, during which the group confirmed the collaboration when asked the following question:

Q:

You've collaborated closely with LAPSUS\$ and Breached. Why partner with other groups instead of handling the full life cycle in-house?

A:

There is a lot of access here, it's better to create an ecosystem and connections, that way it's easier to sell the data fast and move access. LAPSUS\$ have been good to work with, they are very trustworthy and they brought everything together to start the op. A lot more is handled in house than you think, but the end result isn't always published under our group names - usually just the quick one taps/bulk clones.

Trusted Relationship Attacks

Trusted relationship attacks receive less coverage because such incidents tend to be smaller in scale, and victims often prefer not to disclose security issues affecting their organizations. In larger cases involving several organizations, however, the incidents are more likely to be

publicly reported. Several such attacks involving UNC6395 and DragonForce were disclosed in 2025.

Threat actor	Date	Compromised supplier or software	Final payload or impact	Targeted organizations/countries
UNC6395	03.2025-08.2025	Salesloft Drift application	The TA exfiltrated sensitive information from compromised Salesforce instances	>700 organizations
DragonForce	05.2025	an unnamed managed service provider (MSP)	a double extortion tactic to pressure victims into paying the ransom	 N/A

In the fall of 2025, an attack by UNC6395 was disclosed. From March through June 2025, UNC6395 compromised one of Salesloft’s internal GitHub repositories [Source 1](#), [Source 2](#), [Source 3](#), [Source 4](#). The threat actors maintained access for months, downloading code and conducting reconnaissance. Inside the repository, they found a sensitive OAuth token that acted as a master key, granting privileged access from Salesloft to its account within the Drift cloud application. UNC6395 authenticated to Salesloft’s high-privilege Drift account. From this trusted position, the actor could leverage the integrations Drift had with its customers’ applications. From August 8 through at least August 18, 2025, the actor targeted Salesforce customer instances. The actor systematically exported large volumes of data from numerous corporate Salesforce

instances, including Cases, Accounts, Users, and Opportunities. The threat actor’s primary intent was to harvest credentials. After the data was exfiltrated, the actor searched through it for secrets that could potentially be used to compromise victim environments. Researchers observed UNC6395 targeting sensitive credentials such as Amazon Web Services (AWS) access keys (AKIA), passwords, and Snowflake-related access tokens.

The campaign affected an estimated 700 organizations. Affected organizations included major IT companies such as [Palo Alto Networks](#), [Cloudflare](#), [Zscaler](#), [Tenable](#), [Proofpoint](#), [Rubrik](#), [Tanium](#), [SpyCloud](#), [PagerDuty](#), and others. DragonForce

In May 2025, a trusted relationship attack by DragonForce was [disclosed](#). DragonForce gained access to the MSP's remote monitoring and management (RMM) tool, SimpleHelp, and used it to deploy DragonForce ransomware across multiple endpoints. The threat actor exploited a chain of vulnerabilities disclosed in January 2025: CVE-2024-57727, CVE-2024-57728, and CVE-2024-57726. Using access to the MSP's RMM instance, the threat actors gathered information on multiple customer environments managed by the MSP, including device names and configurations, users, and network connections. They also exfiltrated sensitive data and used double extortion to pressure victims into paying the ransom.

Attacks on Cryptocurrency Exchanges

Attacks on the cryptocurrency industry emerged as another major trend during the reporting period because of the scale of the associated losses. According to [Chainalysis](#), more than \$3.4 billion was stolen from the cryptocurrency industry between January and early December 2025. Most attacks on cryptocurrency exchanges in 2025 and 2026 were attributed to Lazarus. A hallmark of Lazarus operations is the use of social engineering through direct human contact: the group either places operatives inside target organizations or, under the guise of job interviews or partnership offers, infects current employees and uses their devices and access as initial vectors. For example, [to attack Drift Protocol](#), the threat actors attended major conferences and established contact with Drift representatives. Over the following six months, they maintained contact through Telegram conversations, working sessions, and in-person meetings at multiple global events. They behaved like legitimate users by onboarding a vault on Drift, depositing more than \$1 million, participating in detailed strategy and product discussions, and building trust with multiple contributors. This relationship-building was designed to gain proximity and credibility within the ecosystem while the actors covertly infiltrated Drift's systems through social engineering.

Both the losses from and frequency of attacks on cryptocurrency exchanges are increasing markedly. This reflects not only the numerous weaknesses that enable such attacks but also heightened threat actor interest in the sector because of the potential financial returns. In some cases, the consequences for cryptocurrency exchanges can be irreversible, ranging from reputational damage and legal and regulatory proceedings to forced liquidation and bankruptcy. For example, [Balancer Labs](#),

the corporate entity behind the Balancer decentralized exchange, is shutting down after a 2025 exploit created ongoing legal and financial strain: threat actors drained approximately \$110 million in digital assets. The protocol, however, will remain online in a leaner form.

In addition to financially motivated threat actors, there were also cases of politically motivated attacks. For example, in June 2025 the hacktivist group Predatory Sparrow attacked the Nobitex exchange and accompanied the attack with a statement on its Telegram channel reading: "These cyberattacks are the result of Nobitex being a key regime tool for financing terrorism and violating sanctions. Associating with regime terror financing and sanction violation infrastructure puts your assets at risk."

Organizations typically conduct a post-mortem after an incident that causes major financial losses, addressing weaknesses and improving security. Against this backdrop, the repeated [compromise of the Turkish exchange BtcTurk](#) stood out. In August 2025, \$48 million was stolen from BtcTurk's hot wallets; a year earlier, in June 2024, approximately \$55 million had been stolen. According to researchers, the threat actors in the 2025 incident gained access to the private keys managing the hot wallets, exploiting the same weakness that led to the first breach.

Statistics on major cryptocurrency attacks in 2025–2026

Attack period	Amount of damage	Victim	Attribution
01.23.2025	~\$85 million	Phemex	N/A
01.2025	~\$10 million	DogWifTools	Jizzy Group
21.02.2025	~\$1.5 billion	Bybit	Lazarus
18.06.2025	\$82–90 million	Nobitex	Predatory Sparrow
09.07.2025	\$42 million, though the threat actor returned the cryptocurrency stolen from the GMX exchange after a \$5 million bounty payment	GMX	N/A
08.2025	\$48 million	BtcTurk	N/A
03.11.2025	\$128.64 million	Balancer	N/A
01.09.2025–01.04.2026	~\$285 million	Drift Protocol — the largest decentralized perpetual futures exchange on Solana	Lazarus
06.03.2026–18.04.2026	~\$292 million	the KelpDAO rsETH bridge , built on the LayerZero cross-chain messaging protocol	Lazarus
16.04.2026	~\$15 million	Grinex	N/A

Ransomware

Throughout 2025 and H1 2026, the number of new ransomware groups continued to grow. The vast majority operate under the RaaS (Ransomware-as-a-Service) model, which significantly lowers the barrier to entry for new participants and drives market fragmentation. More than 50 new groups emerged in 2025 alone, although many remained active for only a few months.

F6 analysts identified several developments in ransomware affiliate programs during 2025:

Dividing up territory

This sharp rise in activity is likely linked to internal crises and destabilization among major players in the ransomware ecosystem. In particular, following the collapse of ALPHV/BlackCat in 2024 and a series of law-enforcement operations targeting LockBit infrastructure, many affiliates were forced to seek new platforms. Some joined existing operations such as RansomHub, Qilin, and DragonForce, while others began forming their own groups. The forum users \$\$\$ and Devman, along with the RaaS operations they launched, are notable examples of this transformation.

In our assessment, however, the redistribution of influence in the ransomware market is far from complete. Groups continue to compete actively for affiliates and market dominance. Some operations seek to attract affiliates by offering more favorable financial terms and technical support, while others rely on aggressive pressure tactics and efforts to eliminate competitors. DragonForce is the clearest example of a group that has repeatedly taken a hardline approach to confronting other market participants.

Economics

The economics of extortion attacks are also changing. Despite several large payments and settlements recorded in 2025 and 2026, the overall trend is toward fewer publicly disclosed ransom payments and a larger share of private negotiations. At the same time, several incidents show that companies are increasingly trying to avoid paying ransom demands in full, preferring negotiated or compromise settlements. This further underscores the shift in the ransomware ecosystem's monetization model: the primary focus is moving away from securing payment for data decryption and toward exerting pressure through data leaks and reputational damage.

Against this backdrop, a notable and persistent trend is the shift from the traditional model of encrypting data and demanding payment for decryption toward extortion based on the threat of publishing stolen data. When a victim refuses to pay, the fallout from an incident often exceeds the amount originally demanded. Such scenarios are particularly common in attacks on large industrial and logistics organizations, where the impact

extends beyond the company itself to its entire supply chain and its partners. For example, following the September 2025 attack on Jaguar Land Rover, the company had to suspend some production processes and temporarily shut down internal IT systems to contain the incident, resulting in an estimated total economic loss of approximately \$2.41 billion.

More sophisticated attacks

In 2025, significant changes were also observed in the techniques used by ransomware groups. Modern operations are increasingly designed to neutralize security controls and hamper incident response before encryption begins. To this end, threat actors actively use EDR-killer tools. RansomHub, for example, was one of the first major ransomware groups to integrate such tools into its RaaS platform. These tools are designed to disable or bypass EDR/XDR solutions. Threat actors also use the BYOVD (Bring Your Own Vulnerable Driver) technique, which allows them to load vulnerable legitimate drivers and use them to obtain privileged system access and disable defensive mechanisms. BYOVD is often used to implement EDR-killer functionality, combining the two techniques in a single attack chain.

At the same time, ransomware operators increasingly use living-off-the-land (LotL) techniques, legitimate remote-administration tools, and automated payload deployment through enterprise management systems. Akira's 2025 campaign was a notable example. In several investigated incidents, the operators used almost no specialized malware during the early stages of compromise, instead relying on administration and automation tools already present in the environment. In one case, for example, the threat actors used Datto RMM, PowerShell, and built-in Windows utilities for lateral movement, payload execution, and evasion of monitoring tools. This trend indicates that a substantial share of ransomware operations are moving away from relatively "noisy" mass attacks toward more delib-

erate, multi-stage scenarios designed to bypass modern security controls and maximize damage to victim infrastructure.

Law enforcement countermeasures

It is also worth noting that ongoing internal conflicts, leaked communications, and the compromise of group infrastructure significantly aid the work of law enforcement and information security researchers. Such incidents provide valuable insight into the internal structure of ransomware operations, the mechanisms of affiliate interaction, and the financial flows within these criminal ecosystems. One example is the leak of BlackBasta communications in February 2025, in which archives of the operators' internal communications were made public. Analysis of such material allowed researchers to reconstruct the group's operating model in greater detail, as well as to observe internal conflicts and a lack of coordination among its members.

Continuity of trends

In conclusion, despite the rapid growth in the number of ransomware groups and the constant emergence of new players, the fundamental principles governing affiliate programs remain largely unchanged. Most groups continue to avoid attacks on CIS countries and, in some cases, on BRICS states as well. They generally do not focus on particular industries, although most refrain from attacking nonprofit organizations. There are exceptions: in June 2025, for example, the Sarcoma group attacked the nonprofit organization Radix. The geographic distribution of attacks also remains stable: the United States continues to be the primary target because of its high concentration of financially strong companies and critical infrastructure.

Trends in the emergence of new ransomware groups

Over the course of 2025, we recorded the emergence of 54 new ransomware groups. Most of them were initially introduced and promoted by their creators on darknet forums.

List of new groups that emerged in 2025:

Group
Cry0
Benzona
Data Keeper
TENGU
Insomnia
Kryptos
CipherWolf
BrotherHood
Arachna
Kyber
Genesis
Coinbase Cartel
Black Shrantac
Obscura
LunaLock
Desolator
Promptlock
Cephalus
Charon
Angel RaaS
Sinobi

Group
BlackNevas
BQlock
Gentlemen
CRPX0
D4rk4rmy
DeadLock
Kawa4096
CHAOS RaaS
DireWolf
DataCarry
Nebu1a
Gunra
AiLock
Chaos
Devman
Jett
VanHelsing
Nova
WarLock
Ethan
Run Some Wares
Anubis
Core
Frag
Aptlock
Hyena

Group
A1project
LINKC
Anomaly
Striker
GD LockerSec
Skibidi
NightSpire

In the first half of 2026, F6 identified 12 additional ransomware groups.

The new groups identified in 2026 were:

Group
Krybit
NetRunner
Lamashtu
Alp-001
BlackWater
BagIRA
Payload Leaks
Vect
Bravox
PrinzEugen
Shadowbyt3\$
Mortar

The lists of the most active groups in 2025 and 2026 largely overlap, indicating the stability and resilience of certain groups. Notably, LockBit returned to the top ranks after

facing problems in 2024–2025. Qilin was the clear leader by number of attacks in both periods, with more than 1,000 attacks attributed to the group in 2025 and nearly 500 in H1 2026.

Most active groups in 2025:

Group
Qilin
Akira
CI0p
Play
INC Ransom
SafePay
Lynx
RansomHub
DragonForce
Sinobi

Most active groups in 2026:

Group
Qilin
The Gentlemen
Akira
INC Ransom
DragonForce
NightSpire
LockBit 5
Play
CI0p
CoinbaseCartel


Statistics on targeted countries and industries

The most active Threat Actors in 2025-2026	TOP 10 targeted countries	TOP 5 targeted sectors
Qilin	US, CA, FR, GB, DE, ES, IT, JP, AU, SG	Manufacturing, Professional Services, IT, Healthcare, Finance
Akira	US, CA, DE, GB, IT, CH, AU, ES, BR, FR	Manufacturing, Professional Services, Real Estate, IT, Finance
Cl0p	US, CA, GB, AU, DE, FR, IN, JP, MX, IT	IT, Professional Services, Consumer Goods, Manufacturing, Transportation and logistics
Play	US, CA, GB, DE, SE, NL, CH, AU, FR, CZ	Manufacturing, Professional Services, IT, Real Estate, Transportation and logistics
INC Ransom	US, CA, DE, GB, AU, FR, IT, AT, BR, SA	Professional Services, Healthcare, Manufacturing, IT, Education
SafePay	US, DE, GB, CA, AU, AR, IT, ES, MX, SG	Professional Services, Manufacturing, IT, Consumer Goods, Education
Lynx	US, GB, CA, DE, AU, IT, FR, SG, JP, ES	Professional Services, Manufacturing, IT, Transportation and logistics, Real Estate

The most active Threat Actors in 2025-2026	TOP 10 targeted countries	TOP 5 targeted sectors
RansomHub	US, CA, BR, GB, IT, DE, ES, AU, IN, FR	Professional Services, IT, Manufacturing, Healthcare, Education
DragonForce	US, GB, DE, IT, AU, CA, FR, BR, HK, IN	Professional Services, Manufacturing, Real Estate, IT, Healthcare
Sinobi	US, IN, GB, CA, IT, ES, FR, CN, DK, CH	Manufacturing, Real Estate, Healthcare, IT, Professional Services
The Gentlemen	US, TH, BR, FR, DE, IN, IT, GB, ES, JP	Manufacturing, IT, Professional Services, Healthcare, Consumer Goods
NightSpire	US, FR, IN, ES, TW, HK, JP, TR, GB, DE	Manufacturing, Professional Services, Healthcare, IT, Real Estate,
LockBit 5	US, IT, BR, DE, MX, FR, IN, ES, TR, AT	Professional Services, Manufacturing, IT, Healthcare, Real Estate
CoinbaseCartel	US, AE, DE, FR, BR, CA, GB, KR, ID, BE	IT, Business Services, Manufacturing, Consumer Goods, Finance

The countries most frequently targeted by the most active groups are shown in the chart:

Countries most frequently targeted by ransomware



The sectors most frequently targeted by the most active groups:

Industries most frequently targeted by ransomware



According to F6 Threat Intelligence, more than 9,300 ransomware-related incidents were identified during 2025 and H1 2026.

Most significant ransomware incidents

- In June 2025, Radix, a nonprofit contractor that provides services to Swiss federal and municipal agencies, suffered a cyberattack. The incident resulted in the compromise and encryption of confidential data associated with the operations of several agencies. The threat actors gained access to the organization's internal systems, exfiltrated data, and subsequently published it on a dark web platform. The incident was later attributed to the Sarcoma ransomware group. The information was published on the group's data leak site (DLS) after ransom negotiations reportedly broke down;
- On August 23, 2025, a cyberattack hit the Swedish company Miljödata, which serves roughly 80% of Sweden's municipalities. The incident affected more than 200 municipalities, regional administrations, and government agencies. The threat actors gained access to the company's IT infrastructure and demanded a ransom of 1.5 BTC, worth approximately \$160,000–\$170,000 at the time. The attack was later attributed to the ransomware group Data-Carry, which published information about the data compromise on its data leak site;
- From August through October 2025, a large-scale campaign targeting South Korea's financial sector was observed, dubbed "Korean Leaks." Activity peaked in September 2025. The operation was carried out by the ransomware group Qilin (also known as Agenda), which operates under the Ransomware-as-a-Service (RaaS) model. The threat actors compromised a South Korean MSP, using it to gain access to the infrastructure of dozens of financial organizations. At least 28 companies were affected, primarily investment firms and asset management organizations;
- In September 2025, automaker Jaguar Land Rover fell victim to a cyberattack. The incident halted some production lines and caused serious disruptions across its supply chain. To contain the attack, the company was forced to shut down some internal IT systems, and the fallout extended to dealer networks, suppliers, and service centers. According to analysts and insurers, the total economic damage from the attack may have reached \$2.41 billion, accounting for production downtime, supply disruptions, and the impact on partner infrastructure. The incident became one of the most telling examples of how ransomware groups can disrupt the industrial sector and damage entire industry ecosystems. The attack was most likely carried out by members of the ShinyHunters group;

- On September 19–20, 2025, Collins Aerospace suffered one of the largest attacks on transportation infrastructure. The attack targeted the MUSE (cMUSE) system used by European airports for passenger check-in, baggage handling, and boarding management. It caused serious disruption at several major airports, including Heathrow Airport, Brussels Airport, and Berlin Brandenburg Airport. In October 2025, the Everest ransomware group publicly claimed responsibility for the attack;
- On September 29, 2025, the Japanese corporation Asahi Group Holdings was hit by a cyberattack. The incident disrupted the company's internal systems and halted part of Asahi Breweries' production and logistics operations in Japan. The ransomware group Qilin claimed responsibility for the attack;
- Akira's activity warrants particular attention. In 2025, the group carried out a series of attacks on organizations using SonicWall SSL VPN infrastructure. During the campaign, the threat actors actively exploited vulnerabilities in VPN solutions (CVE-2024-40766) and used previously compromised credentials and OTP seeds to bypass multifactor authentication. Particularly notable were the high level of operational preparation and the extremely short interval between initial access and ransomware deployment: in some cases, the threat actors compromised the infrastructure and deployed ransomware within a few hours. Akira operators also made extensive use of living-off-the-land techniques, leveraging legitimate administration utilities, PowerShell, and remote-access tools to move laterally within networks and evade monitoring. In some incidents, they also attempted to compromise Veeam Backup & Replication systems to delete backups and hinder recovery of victim organizations' infrastructure;
- One of the most high-profile incidents in UK retail was the 2025 attack on Marks & Spencer. The threat actors gained access to the company's internal systems using social engineering and account compromise. The attack disrupted online sales, logistics processes, and internal services. The company had to shift part of its operations to manual processes and spent several weeks restoring its IT infrastructure. According to media estimates, total damages and lost operating profit may have amounted to around \$375 million. The incident was a striking example of how severely ransomware attacks can affect large retail companies. Based on available information, the attack was also linked to members of the ShinyHunters group;
- In May 2026, the hacking group ShinyHunters breached Canvas, one of the largest learning management systems. The threat actors exploited vulnerabilities in free accounts (Free-for-Teacher) that lacked two-factor authentication and claimed to have stolen 3.65 terabytes of data affecting roughly 275 million users (including private messages between students and instructors). The attack paralyzed operations at more than 8,800 universities and schools worldwide in the middle of final exams. The platform's administrators were forced to negotiate with the attackers to obtain confirmation that the data had been deleted.

Requested and Paid Ransoms in 2025–2026

Group	Victim	Date	Amount (ransom / damages)	Details and Impact
Akira	U.S. media company	September 2025	\$10 million (demand)	In early fall 2025, the Akira group encrypted 85 GB of data at a U.S. media company and demanded a \$10 million ransom.
Unknown	Global logistics holding company (unnamed)	October 2025	\$12.4 million (payment)	One of the largest confirmed payments of the year. The company chose to pay the ransom in an effort to prevent large-scale disruptions to global supply chains.
Rhysida	Heart South Cardiovascular Group	November 2025	\$630K (demand)	A ransom demand for access to the organization's medical data. The incident reflects the group's characteristic practice of applying pressure through the threat of publishing sensitive information.
Medusa	NASCAR (USA)	April 2025	\$4 million (demand)	The attack resulted in the theft of roughly 1 TB of data, including technical documentation and internal contracts. After the victim refused to pay the ransom, the data was offered for sale.
Unknown	PowerSchool (USA)	2025	\$2.85 million (payment)	After data belonging to a significant number of users was compromised, the company paid the ransom. Despite this, the threat actors continued to pressure educational institutions using the stolen information.
ShinyHunters	Canvas LMS	May 2026	unknown (presumably paid)	Following the compromise and exfiltration of data belonging to hundreds of millions of users, an agreement was reached between the company and the threat actors. According to the company, the data was returned, and the attackers confirmed its deletion, although payment of a ransom has not been officially confirmed.

Transformation and Internal Conflicts in the RaaS Ecosystem

The Emergence of Anubis and the Rise of the Data Ransom-as-a-Service Model

In late February 2025, F6 Threat Intelligence analysts observed the emergence of a new affiliate program on the dark web called Anubis. Its distinguishing feature was a Data Ransom-as-a-Service model that was atypical for the market: rather than offering malware development, it provided extortion and blackmail services. Under the traditional model, threat actors compromise organizations themselves and demand payment in exchange for not disclosing stolen data; Anubis's founder, known by the alias superSonic, separated these stages.

Under Anubis, threat actors who have already obtained data but have not yet begun to monetize it are offered professional ransom-negotiation services. Pressure tactics include notifying counterparties, customers, and regulators, as well as posting on X (formerly Twitter). The monetization model uses a 60/40 split in favor of affiliates. Researchers assess that Anubis may be an evolution or reworked version of the former RaaS platform InvaderX. This assessment is indirectly supported by Anubis's use of the ECIES cryptographic algorithm, similar geographic restrictions on targeting (including a ban on attacks against BRICS countries), and overlap between the earlier project's period of activity and disappearance and the emergence of the new operation.




Fig. 8 — Description of the DataRansom service in a post on an underground forum

Translated text from Fig. 8

[NEW] Data Ransom.

A unique working format where all we need from you is a company's data, which we turn into a payout. Only data meeting four conditions is accepted for work.

- The files have not been previously published on the Darknet/Cleartnet.
- The company's files contain data worth publishing.
- The set of files is exclusively in your possession.
- The data was exfiltrated recently, no more than 6 months ago.

After reviewing the data and agreeing on terms, an investigative article is written based on the available files.

Example: `hxxp://om6q4a6cypxvt7ioudxt****`

The articles can be read in all major world languages. The article is published on a private, password-protected blog that does not appear on the main page, and the negotiation link is delivered to the company through cold calls.

Pressure on the company is also applied through nonstandard methods. Notifying counterparties affected by the leak, notifying authorities (GDPR, EDPB, HHS, OPC, ICO, OAIC, etc.), notifying customers, posting on Twitter.

The starting rate for affiliates is 60/40.

In the case of long-term cooperation, the rate is negotiable and your percentage increases.

The Collapse of RansomHub and the Migration of Affiliates

In March 2025, one of the most prominent ransomware groups, RansomHub, ceased active operations. Its infrastructure, including the leak site and administrative panels, initially went offline without any official statement, which was interpreted as a possible technical reorganization. By April-May 2025, however, it had become clear that the group had effectively collapsed: activity fell sharply, and affiliates began migrating en masse to competitors, primarily Qilin and DragonForce. By mid-2025, RansomHub had stopped posting new victims and had effectively disappeared from the ecosystem. The likely cause of the collapse was internal conflict and a redistribution of influence within the affiliate network.




Fig. 9 — Migration of RansomHub members to DragonForce

The Conflict Between DragonForce and \$\$\$

In parallel, a conflict between the user \$\$\$ and a DragonForce representative escalated in February-March 2025. The standoff culminated in the defacement of the BlackLock site, whose contents were replaced with a message referencing the dispute and asserting control over the site. After the incident, the public confrontation gradually subsided, and both parties scaled back their public activity. The user \$\$\$ declined to comment and subsequently relaunched the RaaS service under a new name, Global.




Fig. 10 — Result of the defacement attack on the BlackLock DLS

The Release of the VanHelsing Source Code

In May 2025, on the RAMP forum, a user named th30c0der posted an offer to sell the VanHelsing source code. It later emerged that he was likely a former member of the service's development team, and a conflict with the administration may have prompted the leak. The following day, the VanHelsingRaaS administrator published part of the source code for free and announced the imminent release of a new version — VanHelsing 2.0. The published archive proved incomplete: it lacked the Linux builder and databases, but it contained a working Windows builder, an affiliate panel, and a DLS. The code also included an unfinished MBRLocker module designed to overwrite the master boot record (MBR), replacing it with its own bootloader displaying a lock message.

Hunters International's Shift to a Data-Extortion Model

In July 2025, Hunters International announced that it was closing its RaaS service in its existing form, releasing a free decryptor and stating that it would transition to a new operating model. As part of the announced rebranding, the group shifted its focus from the traditional RaaS model based on data encryption to a model focused exclusively on data theft and extortion in exchange for withholding disclosure. The new operation, World Leaks, uses the threat of publishing or selling exfiltrated data rather than encrypting victim infrastructure. This shift reflects a broader trend

in the 2025 ransomware ecosystem: encryption is gradually becoming less central to attacks, while extortion based on stolen information is increasingly becoming the primary business model.

Attempts to Consolidate the RaaS Market

In early September 2025, a LockBit representative on the RAMP forum announced the launch of version 5.0 of the affiliate program. Access was priced at \$500, the lowest in LockBit's history. Around the same time, a DragonForce representative publicly proposed that LockBit and Qilin join forces and form an alliance to consolidate control over the RaaS market.




Fig. 11 — DragonForce representative's proposal for collaboration between the groups

Translated text from Fig. 11

dragonforce: Welcome home (LB). I think we should all get in touch (LockBit, Qilin, DragonForce). I have a proposal for everyone: level the playing field, no conflicts and no public insults (for the amusement of journalists, the FBI, and researchers). Clear agreements that everyone understands, competition on equal terms. No undercutting on percentages or deposits. That way we can all grow our revenue and also dictate terms to the market. Call this proposal whatever you want - a coalition, a cartel, whatever, it doesn't matter - the main thing is to stay in touch and be friendly toward one another, to be strong allies rather than enemies. There's enough of the pie here for everyone.

LockBit: I completely agree with you, I love you

and wish you no harm - as people treat me, so I treat people.

LockBit: give me your Tox and I'll add you as a friend

The emergence of Vect and collaboration with BreachForums

In late January 2026, a user named Vect announced the launch of a RaaS service of the same name on BreachForums. This affiliate project drew attention just a few months later: in March, Vect reported a partnership with TeamPCP — a group known for its high-profile supply chain attacks in 2026 — as well as with BreachForums itself. Vect promised that all forum members would gain access to the affiliate service. This was the first instance of a RaaS service partnering with an entire forum to recruit affiliates.




Fig. 12 — Vect's post announcing the collaboration with TeamPCP and BreachForums

External Compromises and Data Leaks in the Ransomware Ecosystem

In February 2025, a Telegram user going by the alias ExploitWhispers published an archive of Matrix chat logs from the BlackBasta group — one of the most active ransomware operations of 2024. By early 2025, the group had gone dormant amid escalating internal conflicts. According to various sources, some operators may have abused victims' trust by collecting ransoms without providing functional decryptors, which likely contributed to the leak.

The published data contained a wide range of information, including phishing message templates, the emails used to distribute them, cryptocurrency addresses, victim data, and discussions of the group's tactics and operational procedures. The archive also contained 367 unique Zoom-Info links, which presumably reflect the number of potential victims, since the service was used to identify and enrich data on target organizations. ExploitWhispers also exposed information on several BlackBasta members, including Lapa (an administrator of the RaaS model), Cortes (associated with the Qakbot group), YY (the main administrator of Black Basta), and Trump (also known as GG and AA), who is believed to be one of the group's key leaders.

In May 2025, the LockBit admin panel was compromised by a hacker using the alias XOXO. As a result, the LockBit DLS was temporarily disrupted, and a link to a SQL dump of the admin panel was published. The leak contained the following data:

- the `btc_addresses` table, containing 59,975 unique Bitcoin addresses;
- the `builds` table, with information on builds created by LockBit affiliates, including public keys and, in some cases, the names of target companies;
- the `builds_configurations` table, with attack configuration parameters, including exclusions for ESXi servers and file encryption rules;
- the `chats` table, containing 4,442 negotiation messages between the extortionists and victims for the period from December 19, 2024, to April 29, 2025;
- the `users` table, with data on 75 administrators and affiliates who had access to the panel.

Notably, the threat actors' credentials were stored in plaintext, including passwords such as Weekendlover69, MovingBricks69420, and Lockbitproud231. A LockBitSupp representative confirmed the incident but claimed it involved a publicly accessible panel. It is also worth noting that the hacker XOXO had previously been linked to the compromise of the Everest group's DLS.



Fig. 13 — The compromised LockBit panel

Later in May 2025, a channel called GangExposed appeared on Telegram, and its author began publishing analyses and data collections about members of cybercriminal groups, primarily in the ransomware ecosystem, including Conti and BlackBasta. Initially, the posts consisted of OSINT analyses of nicknames, cryptocurrency transactions, data leaks, and connections among members of hacking forums. Over time, the mate-

rial became more aggressive and focused on deanonymizing suspected operators, infrastructure, and intermediaries using data from leaked correspondence and activity logs. The project eventually began publishing detailed analyses of individual groups and claiming to reveal their structure and members, although the accuracy of some of these claims remains disputed.

In July 2025, fragments of internal correspondence tied to the activities of the Medusa ransomware group were made publicly available. This material reflected the group’s organizational structure and operational processes, including the distribution of roles among members, as well as the full attack cycle — from initial access to data exfiltration and the subsequent encryption of victims’ infrastructure.

In May 2026, a large archive of internal chats from The Gentlemen group was made publicly available, fully exposing the structure of its communications and its operational activity. The publication of thousands of member messages

made it possible to reconstruct how attacks were coordinated in real time and how ransoms were split among affiliates. The text discussions laid bare the inner workings of the RaaS platform: the correspondence revealed that the service’s owner had assembled the admin panel in just three days using commercial AI coding assistants. Analysis of the logs also exposed the threat actors’ real-world tactics, including an April attack on a British IT consulting firm whose stolen data was used to instantly compromise a large company in Turkey, with the goal of subsequently extorting both victims. The Gentlemen administrator confirmed the leak, but the service continued to operate.






Fig. 14 — The Gentlemen administrator confirming the leak



Current Malware Threats



Mobile threats in the wild

Android threats

Attacks on Android devices remain one of the most significant mobile security threats because the operating system is widely used and its vulnerabilities attract threat actor attention. As mobile devices become more popular, threat actors continue to develop new attack methods, including malicious applications, phishing, and social engineering.

The primary methods of distributing malware on Android include phishing messages that lure users into downloading infected APK files, which are often disguised as updates to popular apps. This method is especially effective given the lack of strict app-vetting measures outside official stores such as Google Play.

Actors often try to exploit accessibility API features to read sensitive information, such as passwords and financial details, directly from the screen and to manipulate a user's device by simulating touch input. To counter this, in December 2025, Android 16 introduced a new defense: the `accessibilityDataSensitive` flag. This flag lets a view or composable be explicitly marked as containing sensitive data. Any app that requests accessibility permission but has not explicitly declared itself a legitimate accessibility tool (`isAccessibilityTool=true`) is denied access to that view. This simple change helps prevent malware from stealing information.

Android malware is increasingly distributed under the MaaS model. This model allows multiple affiliates to operate locally in their own regions or other areas of interest.

The nature of Android trojans is changing; they are becoming more versatile, full-featured tools that combine the capabilities of banking trojans, spyware, RATs, and proxy malware.

Throughout 2025 and 2026, NGate-type Android malware based on the publicly available NFCGate tool evolved rapidly. This malware relays NFC data from victims' payment cards through their mobile phones to an attacker's device. Several NFCGate-based malware families were identified in 2025, including SuperCard X MaaS and RatOn.

Threat actors are devising new ways to evade detection tools and complicate analysis. For example, Herodotus is one of the first banking trojans to attempt to make remotely executed actions appear human. To make input look as though it were typed by a real user, the malware enters characters individually with random delays. Another trojan, Kloptra, combines extensive use of native libraries with Virbox, a commercial-grade code-protection suite, making it exceptionally difficult to detect and analyze.

The table below presents the most notable Android malware observed in real-world attacks in 2025 and 2026.

Android threats observed in the wild in 2025-2026

Activity date	Name	Target countries	Capabilities
11.2024-06.2025	Crocodilus	 PL,  AR,  BR,  ES,  US,  ID,  IN,  TR	Overlay attacks, keylogging, remote access, and "hidden" remote control capabilities
04.2025	SuperCard X (MaaS)	 IT	NFC-relay functionality
07.2025-09.2025	RatOn	 CZ,  SK	Overlay attacks with automatic money transfers and NFC relay functionality
03.2025-07.2025	PlayPraetor (MaaS)	 PT,  ES,  FR,  HU,  IT,  MA,  PE,  HK,  ID,  VN,  IN,  MY,  SG	Abuses accessibility services to gain remote control and can serve fake overlay login screens atop nearly 200 banking apps and cryptocurrency wallets
06.2025-07.2025	DoubleTrouble	 Europe	Steals banking credentials, captures lock-screen information, and features keylogging functionality
12.2025	DroidLock	 ES	The ability to lock device screens with a ransomware-like overlay and illegally acquire app lock credentials, leading to a total takeover of the compromised device
03.2025-08.2025	Klopatra	 ES,  IT	Hidden VNC for complete remote device control and dynamic overlay attacks for credential theft
09.2025	Herodotus MaaS	 IT,  BR	It has remote-control features that abuse Accessibility Services and attempts to make remote actions appear human

Activity date	Name	Target countries	Capabilities
11.2025	Sturnus	 Southern and Central Europe	Overlay-based credential theft, message monitoring, bypassing encrypted messaging, extensive keylogging, real-time screen streaming, remote control, device administrator abuse, comprehensive environmental monitoring, and advanced evasion capabilities
09.2025–11.2025	Albiriox MaaS	 AT,  Global	A VNC-based remote access module for real-time device control and an overlay attack mechanism for credential harvesting
12.2025–06.2026	Mirax (or Astrinox) MaaS	 Europe	Remote access functionality, dynamic HTML overlays, residential proxy functionality, keylogging, and keyguard data exfiltration
04.2026	RecruitRat	Global threat:  US,  CA,  SA,  MA,  NG,  AO,  KE,  PT,  FR,  DE,  PL,  TR,  AE,  IN, and others	Leak SMS messages, leak contacts, simulate user gestures, screen capture, screen overlay, steal PIN/password, perform clicks, update C2, collect information about the device, steal 2FA, launch terminal commands
02.2026–04.2026	Massiv	 US,  AU,  ES,  PT,  GR,  FR,  TR,  Southern Europe	Supports on-device fraud, overlay functionality, keylogging, SMS and push-message interception to obtain sensitive data, and fully functional remote control
01.2026–02.2026	TrickMo.C	 FR,  IT,  AT	Remote control, SMS and notification interception, credential phishing, keylogging, screen recording and streaming; the C2 channel has been migrated onto The Open Network (TON); network-oriented functionality, including reconnaissance, SSH tunneling, and SOCKS5 proxying capabilities

Crocodilus

Crocodilus is a mobile banking trojan that was publicly disclosed in March 2025. Once installed, it asks the user to enable Accessibility Services. The malware then connects to its C2 server to receive instructions, including a list of target applications and the overlays to be used. It runs continuously, monitoring application launches and displaying overlays to intercept credentials. Crocodilus includes a command that captures the contents of the Google Authenticator application and sends the image to the C2 server, allowing its operators to steal OTP codes in real time. Using stolen PII and credentials, threat actors can take full control of a victim's device through built-in remote-access functionality and complete fraudulent transactions without detection. Crocodilus can also conceal remote-access sessions by displaying a black-screen overlay over all activity and muting the infected device.

RatOn

RatOn combines traditional overlay attacks with automated money transfers and NFC relay functionality, making it a uniquely powerful threat. It is a fully functional banking Trojan with device and account takeover capabilities that targets cryptocurrency wallet applications. The malware supports wallet applications in four languages: English, Russian, Czech, and Slovak. In addition, the malware can perform automated money transfers by abusing one specific banking application, as well as carry out extortion using custom overlay pages and device locking. The threat actors registered domains with adult themes to infect victims. These domains contained TikTok18+ in their names and directly hosted the malicious dropper application.

Klopatra

The initial phase of the attack uses a dropper application designed to appear legitimate. Once the main Klopatra payload is installed, it requests a wide range of permissions, but one is crucial to its success: the Android Accessibility Services permission. Klopatra's most innovative feature is its multi-layered approach to evading analysis. At the heart of this strategy is Virbox, a Chinese commercial software-protection solution. This choice marks a significant escalation, indicating that the operators are investing financial resources to protect their malicious "assets." Researchers assess that Klopatra was developed and operated by a Turkish-speaking criminal group.

DroidLock

DroidLock is delivered through phishing sites that trick users into installing a malicious application. The application is actually a dropper that installs DroidLock. Early in the installation process, the malware requests Device Admin and Accessibility Services permissions. DroidLock can lock the device screen with a ransomware-style overlay and illicitly obtain application lock credentials, leading to a complete takeover of the compromised device. Unlike typical ransomware, DroidLock does not encrypt files; instead, it blocks access and threatens to destroy all data unless a ransom is paid, achieving the same effect. It uses deceptive system-update screens to mislead victims and can stream and remotely control devices through VNC. The malware also abuses device-administrator privileges to lock or wipe data, capture the victim's image with the front-facing camera, and silence the device. In total, it supports 15 distinct commands for interacting with its C2 panel.

DoubleTrouble

DoubleTrouble was initially spread through phishing websites impersonating well-known European banks. Early variants of the Trojan primarily used overlays to steal banking credentials, captured lock screen information, and included keylogging functionality. In its latest iteration, the Trojan's distribution strategy has shifted and now leverages bogus websites that host malware samples directly within Discord channels. In this latest evolution, the malware includes several new and advanced features, including displaying malicious UI overlays to steal PIN codes or unlock patterns, comprehensive screen recording, the ability to block the launch of specific applications, and advanced keylogging functionality.

Sturnus

Sturnus is a privately operated Android banking Trojan that supports a broad range of fraud-related capabilities, including full device takeover. A key differentiator is its ability to bypass encrypted messaging. By capturing content directly from the device screen after decryption, Sturnus can monitor communications via WhatsApp, Telegram, and Signal.

RecruitRat

RecruitRat disguises its infection chain as a legitimate recruitment process. Victims are lured under the pretense of employment opportunities and directed to fraudulent websites. It uses deceptive overlay attacks to harvest device-level credentials.

Massiv

Massiv is a powerful Android banking trojan designed to commit fraud on mobile devices. It includes overlay functionality, keylog-

ging, and SMS and push-notification interception to obtain sensitive data. It is also a fully functional remote-control tool that gives its operator direct access to the victim's device. Once the threat actors have profiled the victim, they deploy high-fidelity phishing overlays synchronized with the launch of legitimate target applications. Massiv uses this mechanism to deliver a specific set of HTML-based login overlays targeting 78 banking and cryptocurrency wallet applications, with targets mapped by country of origin.

TrickMo.C

TrickMo.C is device-takeover (DTO) malware that targets banking, fintech, wallet, and authenticator applications on Android. The bot's on-device automation actively coerces users into granting Accessibility Services permission; once granted, the operator obtains a real-time, interactive view of the device. The most significant architectural change in this version of TrickMo is the migration of its primary command-and-control transport to The Open Network (TON). The malware also expands the operational role of infected devices through SSH tunneling and authenticated SOCKS5 proxying, effectively turning compromised phones into programmable network pivots and traffic-exit nodes.

Android MaaS offerings are evolving rapidly. The following were observed in real-world attacks in 2025 and 2026: SuperCard X, PlayPraetor, Herodotus, AlbirioX, and Mirax.

PlayPraetor MaaS

A large-scale Malware-as-a-Service (MaaS) operation orchestrated by Chinese-speaking threat actors has infected more than 11,000 Android devices worldwide by deploying the PlayPraetor RAT. The malware abuses accessibility services to gain remote control and can serve fake overlay login screens on top of nearly 200 banking apps and cryptocurrency

wallets in an attempt to hijack victim accounts. The operation is managed through a sophisticated Chinese-language C2 panel featuring a multi-tenant architecture that supports multiple affiliates.

Information on the other Android MaaS families mentioned above is provided in the Android Threats on the Dark Web section.

iOS threats

iOS devices are generally less exposed to attacks because developing malware for the operating system is complex and expensive and because iOS is less widespread than

Android. This is also reflected in the absence of iOS malware offerings on underground platforms. Advanced commercial software is an exception and is typically used on a limited basis to spy on high-priority targets. However, in 2025 and 2026, at least two iOS exploit kits, DarkSword and Coruna, were observed in the arsenals of several criminal groups and used not only for espionage but also to steal financial information.

Cases of iOS malware being distributed through the official App Store have also been identified. This points to the changing threat landscape for iOS device owners and the growth of large-scale attacks targeting iOS. In most identified cases, the threat actors' primary objective was to gain access to financial information, particularly victims' cryptocurrency wallets.

iOS threats identified in the wild in 2025-2026

Activity date	Name	Target countries	Purpose / Primary objective
03.2024-02.2025	SparkCat	 AE,  KZ,  ID,  ZW,  CN,  Europe,  Asia, and others	Gaining access to cryptocurrency wallets
02.2024-06.2025	SparkKitty	 Southeast Asia,  CN	Gaining access to cryptocurrency wallets
02.2025-12.2025	Coruna Exploit Kit (CryptoWaters)	 CN,  UA	Known cases involve subsequent malware deployment to steal financial information
10.2025-03.2026	Darksword	 SA,  TR,  MY,  UA	Extraction of extensive personal data, including credentials, with a particular focus on cryptocurrency wallet apps
Autumn 2025-03.2026	FakeWallet	 CN	Malicious apps are specifically engineered to hijack recovery phrases and private keys

SparkCat

The malicious applications were distributed through Google Play and the App Store. A malicious SDK/framework for stealing cryptocurrency wallet recovery phrases was embedded in the applications. The stealer requests the required permissions and then analyzes photos in the smartphone's gallery using an optical character recognition (OCR) module built on Google ML Kit. The trojan uploads any photos in which the OCR module detects potentially valuable text or cryptocurrency wallet data to the C2, together with information about the victim's device.

SparkKitty

Researchers later disclosed malware related to SparkCat and named it SparkKitty. It was distributed through the App Store and Google Play, as well as on fake sites impersonating TikTok and cryptocurrency- or gambling-related apps. The iOS payload comes in the form of frameworks and obfuscated libraries or is embedded directly in the applications. The Android Trojan exists in Java and Kotlin variants. The Trojan sends images from the victim's gallery and information about the device to its C2.

Coruna exploit kit

The exploit kit, named "Coruna" by its developers, is a new and powerful toolkit targeting Apple iPhone models running iOS 13.0 (released in September 2019) through iOS 17.2.1 (released in December 2023). It contained five complete iOS exploit chains and a total of 23 exploits. The kit's core technical value lies in its comprehensive collection of iOS exploits, the most advanced of which use non-public exploitation techniques and mitigation bypasses.

Researchers first identified the Coruna exploit kit in February 2025. At that time, the iOS exploits were integrated into a previously unseen JavaScript framework and were used by a customer of a surveillance company.

In the summer of 2025, researchers found the same JavaScript framework embedded as a hidden iframe across many compromised Ukrainian websites; it was delivered only to selected iPhone users in a specific geolocation. Researchers attributed this activity to UNC6353.

In December 2025, researchers identified the JavaScript framework on a large set of fake Chinese websites, mostly related to finance, that delivered the Coruna exploit kit. When an iOS device accessed one of these websites, a hidden iframe was injected to deliver the exploit kit. This campaign was attributed to UNC6691, a financially motivated threat actor operating from China. At the end of the exploitation chain, a stager binary called PlasmaLoader (a.k.a. PLASMAGRID) facilitates communication with the kernel component established by the exploit. The loader injects itself into powerd, a daemon running as root on iOS, and steals financial information. The payload can decode QR codes from images on disk. It also includes a module that analyzes blocks of text for BIP39 word sequences or specific keywords such as "backup phrase" or "bank account." If such text is found in Apple Notes, it is sent to the C2.

DarkSword exploit kit

DarkSword is a sophisticated iOS exploit kit delivered through a watering hole attack via compromised websites. Once a device is compromised, the exploit chain progresses through the iOS operating system, from the browser (WebKit) to the kernel, to achieve full control of the device. Its objective is broad surveillance and intelligence collection, including the theft of sensitive data and application databases (including cryptocurrency wallets). Since at least November 2025, multiple commercial surveillance vendors

and suspected state-sponsored actors have used DarkSword in distinct campaigns. These threat actors have deployed the exploit chain against targets in Saudi Arabia, Turkey, Malaysia, and Ukraine. DarkSword supports iOS versions 18.4 through 18.7 and uses six different vulnerabilities to deliver final-stage payloads. At least three distinct malware families have been deployed following a successful DarkSword compromise: GHOSTBLADE, GHOSTKNIFE, and GHOSTSABER.

UNC6353 was active from at least December 2025 through March 2026 and used the DarkSword exploit chain to deploy GHOSTBLADE. GHOSTBLADE is a data-mining tool written in JavaScript that collects and exfiltrates a wide variety of data from compromised devices.

GHOSTSABER is a JavaScript backdoor used by a commercial surveillance vendor. Its capabilities include device and account enumeration, file listing, data exfiltration, and the execution of arbitrary JavaScript code.


UNC6748 deployed the GHOSTKNIFE backdoor. GHOSTKNIFE, written in JavaScript, has several modules for exfiltrating different types

of data, including signed-in accounts, messages, browser data, location history, and recordings. It also supports downloading files from the C2 server, taking screenshots, and recording audio from the device's microphone.

FakeWallet

In March 2026, researchers uncovered phishing apps in the Apple App Store masquerading as popular crypto wallets. The threat actor behind this campaign may be linked to the creators of the SparkKitty trojan. Once launched, these apps redirect users to browser pages resembling the App Store, from which they distribute trojanized versions of legitimate wallets. The infected apps are specifically engineered to hijack recovery phrases and private keys.

The threat actors have produced a wide variety of malicious modules, each tailored to a specific wallet. In most cases, the malware is delivered through injection of a malicious library, though in some cases the app's original source code was modified.



Android Threats on the Dark Web

SuperCard X MaaS

During the reporting period, threat actors actively conducted fraudulent operations using social-engineering techniques. The attack begins with receiving fraudulent SMS or WhatsApp messages about a charge to the victim's account. The victim is prompted to call a number to dispute the transaction. The threat actors then persuade the victim to disclose their PIN, increase their withdrawal limit, and install a malicious application disguised as a security tool or verification utility. This application conceals SuperCard X malware, which incorporates NFC-relay functionality. In the final stage of the manipulation, the threat actors instruct the victim to hold a physical debit or credit card near the infected mobile device. SuperCard X then silently captures the card details transmitted over NFC. The data is intercepted in real time and relayed through C2 infrastructure to a second Android device controlled by the threat actors. Once the victim's card data has been relayed, the threat actors use the second device to perform unauthorized transactions.

In 2025, F6 Threat Intelligence specialists identified several Telegram channels actively offering subscriptions to the SuperCard X malware in Chinese.

A screenshot of the main Telegram-channel promoting SuperCard X MaaS is shown below:




Fig. 15 — Screenshot of the main SuperCard X MaaS channel on Telegram

One of these resources has been active since 2021 while another similar resource has been active since 2024; both of them remained active during the reporting period. Some Chinese-language advertisements offering SuperCard X for sale also appeared on other underground platforms. However, most of these posts have been removed afterwards, and the original Telegram bots used to purchase subscriptions are no longer active. As of 2025, SuperCard X was offered for \$180 per day or \$900 per week.

Albiriox MaaS

The start of public beta testing for Albiriox MaaS was announced on September 25, 2025, on the RAMP forum. Although this forum was taken down by law enforcement on January 28, 2026 (this event is covered in more detail in the section Underground forum takedowns from 2025 to 2026:

The takedown of the RAMP underground forum), the threat actor operating under the alias NMZ continued selling the malware on another underground forum. The Albiriox malware targets devices running Android OS (v.11 and above).



Fig. 16 — The post offering the Albiriox malware for sale on an underground forum

Translated text from Fig. 16

ALBIRIOX - This is a full-featured Android crypto/banking bot.

Numerous bugs, shortcomings, and optimization issues were fixed - too many to count.

But even after the beta test, many users praised the functionality and performance of the project, which says a lot - Reviews from real beta testers.

As of the Release v1.0 version, the project includes the following functionality:

- VNC - Rendering of on-screen activity and the ability to control the device via clicks and gestures
- AcVNC - Screen rendering based on Accessibility Node and control - allows bypassing the Secure flag (black screen) in Android (an example of rendering is shown below for both options) (rendering speed a stable ~25 fps).
- Keylogger - Based on Accessibility Node - Allows recording taps on elements and capturing the phone password at login.
- Target detect - Identification of apps that may be of interest (crypto/banking apps) and the ability to launch them directly and request their removal if needed.

- Injects - Bank-oriented injects that request Login:Pass upon entering a crypto/banking app (currently Multi inject for all)
- Additional controls - Home, Back, Multitask, Recent, Power, Vol+, Vol-
- Update Screen - Hiding the user's screen behind an Android update screen (see example below) (redesigned; formerly known as blank-screen.)
- AcVNC - rendering configuration directly from the panel (text size, element contrast, element clickability).
- Web-Panel - for viewing brief information on the server about clients (number of installs, number of passwords, etc.).
- Phone Pass - automatic collection of phone unlock passwords (with a confidence rating for the password).

And other, less significant supporting features.

The server-side component can be deployed on Windows Server 2022 or Windows 10. The author claimed that the malware is compatible with Google Pixel, Samsung Galaxy S/Z/A series, Infinix, Tecno, and other Android-based models.

Для установки панели потребуется сервер от надежного хостера с возможностью открыть TCP-порт 5555 и установленной системой Windows Server 2022/Windows 10 как минимум 2gb ram/4 cpu - главное пропускная способность и гео сервера по отношению к гео таргета.
Продукт ориентирован на Android 11+ и на следующих вендорах устройств: Google Pixel, Samsung Galaxy S/Z/A, Infinix, Tecno и других вендоров (см donkillyapp.com).

Fig. 17 — Part of the threat actor's announcement containing a list of supported devices

Translated text from Fig. 17

To set up the panel, you will need a server from a reliable host with the ability to open TCP port 5555 and Windows Server 2022/Windows 10 installed, at least 2gb ram/4 cpu - the main things are bandwidth and the server's geo relative to the target's geo.

The product targets Android 11+ and the following device vendors: Google Pixel, Samsung Galaxy S/Z/A, Infinix, Tecno, and other vendors.

The key features of Albiriox malware focus on remote control of infected devices and defense evasion:

- VNC and AcVNC: remote access to the infected device, including screen streaming and control through clicks and gestures. AcVNC uses Accessibility Node data to bypass the Secure flag (black-screen protection);
- Combo Render (as of malware version v1.6-F) — a hybrid streaming mode that combines VNC and the Accessibility Tree;
- A keylogger based on Accessibility Node that captures the device's screen unlock password;
- Target Detect and web fakes: automatic detection of cryptocurrency and banking apps with forced launch and phishing overlays (Multi-inject support);
- Phone Pass and Pattern Grabber: collection of devices' unlock passwords with automatic input capability (Auto Pattern);
- Real-time SMS interception, history extraction, and analysis of OTP messages;
- Update Screen — disguising malicious activity as an Android update screen.

The subscription price for Albiriox MaaS changed as follows: \$650 per month until October 21, 2025, then \$720, and \$1,300 from December 3 through May 31, 2026. Trial access to the malware was offered for two weeks under a custom subscription plan (\$740); however, this version excluded features such as interface customization, custom web-fake creation, SMS interception, and OTP analysis.

Herodotus MaaS

In October 2025, security researchers observed active campaigns in Italy and Brazil involving device-takeover attacks by Herodotus. The malware was distributed through sideloading, potentially via smishing messages that directed victims to a malicious link from which the dropper was downloaded. After installing the payload, the dropper automatically launched Herodotus. The malware has remote-control capabilities that abuse Accessibility Services to perform actions on the device, such as opening target applications, navigating screens, and entering transaction details to transfer funds from the victim's account. Unlike many other banking trojans, Herodotus is one of the first to attempt to make remote actions appear human. To make input look as though it were typed by an actual user, the text specified by the operator is split into individual characters, each of which is entered separately with random delays.

The actor behind this threat ("K1R0") advertised it as Malware-as-a-Service on one of the underground forums in September 2025. The threat actor claimed that a subscription to the Herodotus malware costs \$4,000 per month. Trial access was also offered for \$700.



Fig. 18 — Posts offering Herodotus for sale under the MaaS model on an underground forum

Mirax MaaS

Unlike typical MaaS offerings, Mirax is distributed through a highly controlled and exclusive model, limited to a small number of affiliates. The first post offering the malware for sale appeared on the forum on December 19, 2025.



Fig. 19 — Posts offering Mirax for sale under the MaaS model on an underground forum

Translated text from Fig. 19

Dear forum users, we are pleased to present a product that has been in the works for a long time and is finally ready for a full release.

Our Android malware is a new solution; every part of it was written entirely from scratch, tested many times on different OSes and Android versions, and put through multiple traffic runs to ensure the bot works flawlessly.

This topic marks the release of the bot into the wild; however, that does not mean it is open to everyone.

Such painstaking work was not done so that just anyone could buy our product. Forbidden fruit is sweet - that is what guided us while working on Mirax. We are not looking for mere customers - we are looking for partners, and so we will cap the number of users at six. This ensures the confidentiality of the product and extends the lifespan of each bot many times over.

The malware supports devices running Android OS. It is designed to steal financial data with the ability to turn an infected device into a residential proxy.

Features of the Mirax malware:

- Remote application management: launching, uninstalling, and monitoring installed applications;
- Keylogger;
- Permission management: automatic granting of all required permissions, as well as a ready-made API for viewing current permissions;
- Notification and SMS management: viewing, deleting, collecting information and exporting statistics;
- Web fakes: dynamic loading of templates from the server and automatic capture of form data;
- Google Play Protect bypass;
- Three black-screen modes: Standard (for executing C2 commands), Update (for automatic permission granting), and Unlock (with an OS update progress bar);
- Residential Proxy: Includes a built-in SOCKS5 module that turns the infected phone into a residential proxy node to route outside criminal traffic;
- VNC and live camera streaming.

As of May 2026, the subscription costs \$2,000 per month or \$500 for five days.

Longkong RAT

On May 6, 2026, a user named longkongRat advertised Longkong RAT on an underground forum as a multifunctional Android remote access trojan with spyware capabilities.



Fig. 20 — Advertisement for the sale of Longkong RAT malware on an underground forum

The RAT supports a wide range of capabilities for remotely controlling infected Android devices:

- **Live Screen Mirroring:** real-time streaming of the infected device's screen with support for remote interaction, including full touch input control;
- **Live Camera Preview and Ambient Audio Monitoring:** remote activation of the front or rear camera, covert video streaming, screenshots and screen recording, and covert microphone activation for eavesdropping;
- **File Management System:** remote browsing, uploading, and downloading of files stored on an infected device, with support for transferring images, videos, and documents;
- **GPS Location Tracking:** real-time tracking of the device's location and viewing of route history;
- **Bank Grabber:** targeted collection of data from banking applications and logging of device unlock passwords;
- **Keylogger:** logging of keyboard input and capture of device unlock passwords;
- **SMS & Contacts:** viewing SMS history, the contact list, and the call log.

Longkong RAT stands out for its low price: as of May 2026, subscriptions cost \$10 per week, \$30 per year, or \$100 for a lifetime license. A license to resell the malware to third parties was also offered for \$300, along with the full source code for \$150, which the author claimed contained all of the listed features.

Android RAT Andromeda

In May 2026, F6 Threat Intelligence analysts identified an advertisement on a popular underground forum offering Android RAT Andromeda for sale, posted by a threat actor under the alias SybraAzazel.




Fig. 21 — Advertisement for the sale of Android RAT Andromeda malware on an underground forum

Translated text from Fig. 21

Android Andromeda RAT + Dropper (VNC 2.0 Edition) - 2026

Presenting the most powerful solution for Android control. This is not just a RAT, but a full-fledged suite focused on maximum Alive Rate and unkillable persistence in the system.

INFORMATION

Compatibility: Full support for Android 10-16.

New Dropper: A free dropper is included for automatic Accessibility acquisition.

FUD / Clean APK: Free crypting included. Google Protect Bypass - zero detections.

Persistence #1: The highest Alive Rate on the market. The bot survives for months.

Anti-removal: Removing the bot manually is virtually impossible.

Infrastructure: No VPS, Fastflux, or complex protection required. Online builder with documentation.

Size: The optimized APK is only ~5 MB.

FEATURES

Hidden VNC 2.0: Full screen access + Black Screen protection.

File Manager: Search, upload, download, and delete files/folders.

Web & URL Injects: About 400 free injects for 30+ countries.

Data Grabbers: Bank card (CC) collection, Cookie Stealer, and Google Authenticator.

Keylogger: Logs of all owner activity, browser URLs, and entered data.

SMS & Calls: SMS interception, SMS sending, USSD (Call forwarding).

Spam: Mass SMS distribution to the contact list (Bulk SMS).

Notifications: Interception and blocking of push notifications.

Locks: Blocking applications, SMS, notifications, or the entire device.

Remote Commands: Opening URLs (browser/full screen), launching and installing applications.

According to the advertisement, the malware targets devices running Android versions 10–16. The size of a build is approximately 5 MB, while the malware build process itself is handled through an online builder.

The threat actor claims that the malware can remain undetected on the infected system for extended periods and includes built-in anti-removal protection.

According to the description, the core functionality of Android RAT Andromeda is geared toward the banking sector:

- Hidden VNC 2.0 with Black Screen protection that hides activity from the user;
- VNC Custom Protection — a fake bank logo that imitate the Android OS loading or an update screen;
- Automatic unlocking — detecting the type of protection (PIN or pattern lock), collecting data, and unlocking the device;
- Web & URL Injects — about 400 ready-made web fakes for 30+ countries, designed to capture data from banking applications and cryptocurrency services;
- Grabber module — collection of bank card data (CC Grabber), cookies (Cookie Stealer), and 2FA codes from the Google Authenticator app;
- Keylogger module — logging user activity on websites, copying URLs from the browser, and capturing data entered into web forms;
- SMS & Calls — interception and sending of SMS messages, execution of USSD commands, and call forwarding;
- Notifications — interception and blocking of push notifications;
- Locks: blocking individual applications or the entire device, opening URLs, launching and installing applications, and removing third-party programs and the malware itself.

As of May 2026, the subscription costs \$3,000 per month. The author also offers an access-exchange program for the owners of other malware's control panel (30 days of access to the Android RAT in exchange for 30 days of access to the Windows RAT) and free trial access to underground forum administrators and high-reputation users.

Android Botnet Maradona

The Maradona botnet is similar in functionality to the Android RAT Andromeda described above. A user named bellygen advertised it on February 25, 2025.




Fig. 22 — Advertisement for the sale of Android Botnet Maradona

Translated text from Fig. 22

We present to you the newest Android Botnet Maradona.

This comprehensive solution for all types of Android will provide maximum functionality, convenience, and operational efficiency, and our technical support will serve as a reliable and competent guide throughout your entire experience with our malware.

We are ready to take your Android experience to a whole new level. Here is what makes our botnet great:

- Fast VNS,
- HVNC (Black Screen),
- ScreenUnLock,
- Offline Keylogger,
- Injects.

A remote NFC controller is also under development.

Android Botnet Maradona is a multifunctional Android malware that includes:

- VNC (Black Screen): remote graphical screen streaming with the option to stream under a custom overlay on the victim's device;
- HVNC (Black Screen): covert control of the infected device. Screen unlocking and capture of the PIN, password, and pattern lock, with automatic entry to unlock the device;
- Online and offline keylogger: capture of keystrokes even when the victim is offline, with synchronization upon reconnection;
- Clipboard interception: extraction of text from the victim's clipboard;
- Device locking: screen-locking function without administrator rights;
- Push notification interception: forwarding of push notification logs and the ability to send automated push spam on behalf of random installed applications;
- Injects: capture of web form data even if the data was entered incompletely or submission was not completed;
- Theft of the victim's bank card data by mimicking the display of an official system window or Google service;
- Application management: retrieval of the list of installed applications with the ability to open and remove selected applications from the list;
- SMS management: covert sending and receiving of SMS messages via an installable SMS Manager, with selection of the SIM card used to send messages;
- Advanced device management: the malware reportedly supports screen-brightness control, muting, collection of SIM card information, battery statistics, charger status, and other system information;
- NFC controller for remotely intercepting NFC data (stated to be under development).

According to the threat actor, the Maradona botnet had been tested on devices running OS Android v. 9–15, including mobile devices of Samsung, Google, Xiaomi, Honor, Oppo, OnePlus, Realme, Vivo, and Motorola.

The seller also provided a video demonstrating the Maradona botnet control panel:




Fig. 23 — Screenshot demonstrating the Android Botnet Maradona control panel, provided by the threat actor

As of May 2026, the subscription to Android Botnet Maradona costs \$3,000 per month.

NFC relay by nikolakitip

On March 11, 2025, a threat actor using the alias nikolakitip advertised the sale of a custom-built NFC relay in the form of an APK file, designed to steal funds through ATMs.



Fig. 24 — Advertisement for the sale of a custom-built NFC relay in APK format

Translated text from Fig. 24

For sale: APK for card relaying

Price: We work with no percentage, 3–5k a week, 10 a month, 1.5 a week

Was: \$10k now and \$10k in two weeks, strictly to the day, or a single payment of 15 upfront; we work with guarantors

or 5K rental with %

- The APK is installed on your phone (the cashout crew's)
- The second APK is dropped onto the mark's device, and he launches it
- We attach a card to your phone, onto which the mark will later load the money (i.e., a drop card)
- The mark taps his phone to the terminal and loads the money supposedly to himself, but in reality to us)
- Done, the money is credited to the card we attached to our phone

Discount for now - Come by

Device control is exercised from the threat actor's Android device, to which a card registered to a money mule is linked: when a victim withdraws funds from an ATM, the funds are automatically transferred to the threat actor's bank account. In addition to the relay function, the malware includes various modules for covertly controlling the Android device:

- VNC/HVNC (Black Screen): covert control of the device, including remote screen streaming with the option to stream under a custom overlay on the victim's device;
- Online and offline keylogger: capture of keystrokes even when the victim is offline, with synchronization upon reconnection;
- Clipboard interception: extraction of text from the victim's clipboard;

- Device locking: screen-locking function without administrator rights;
- Push notification interception: forwarding of push notification logs and the ability to send automated push spam on behalf of random installed applications;
- Support of overlays targeting banking applications (the list of countries is limited);
- Support of a custom-built clipper to steal cryptocurrency wallet data;
- The ability to intercept 2FA codes for Google services;
- Application management: retrieval of the list of installed applications with the ability to open and remove selected ones;
- Access to folders and files on the device;
- Call management and access to contacts: the ability to view completed calls and place calls;
- SMS management: the ability to view, send, and receive SMS messages via an installable SMS Manager;
- Advanced device management: the ability to control screen brightness, real-time access to device's camera and microphone, control of sound, the screen output, collection of SIM card information, battery statistics, charger status, and other system information;
- A self-destruct mode for the malware on the device;
- Modules for collecting geolocation data, files, and contacts as well as camera and microphone control modules.

The stated subscription cost was \$3,500 per month.

The malware was promoted through a private Telegram channel that published updates on the botnet's release, development of a Tele-

gram bot for generating builds, and compatibility fixes for devices running OS Android v.14–15.




Fig. 25 — Screenshot from the threat actor's private Telegram channel

Modern banking trojans increasingly take the form of highly specialized fraud malware built for the direct theft of funds — through ATMs and mobile banking apps. The high subscription prices for these tools point to the emergence of a stable market for professional malware designed to generate profit while minimizing the risks to its operators.

JS Sniffers

In today's digital landscape, sniffer-based attacks are among the most insidious threats facing e-commerce sites. Sniffers, also known as JS sniffers or e-skimmers, are malicious scripts, most often written in JavaScript, that threat actors inject into online store pages, typically where customers enter payment card details. Once embedded in a checkout page, a sniffer intercepts card data in real time, before it is encrypted and sent to the legitimate payment gateway. After the compromise, threat actors either steal funds directly or sell the compromised data on the dark web.

The primary methods for delivering sniffers evolved from simple server compromises to sophisticated chains involving the abuse of legitimate services such as Google Tag Manager; supply chain attacks in which malicious code is delivered through compromised third-party libraries or trusted CDNs; and the injection of hidden malicious elements into a site's HTML. In 2025 and early 2026, threat actors also adopted new methods for exfiltrating stolen data, increasingly shifting away from simple HTTP requests to attacker-controlled domains and toward more covert techniques.

In particular, threat actors use WebSockets for covert two-way communication that bypasses many security policies. They also hide data in WebRTC connections or encrypt it and send it through legitimate payment-system infrastructure such as Stripe, making malicious traffic appear ordinary.

The catalyst for the widespread proliferation of sniffers in 2025 and early 2026 was a pair of critical vulnerabilities in the Magento platform: [SessionReaper](#) and [PolyShell](#). Magento is one of the most popular open-source e-commerce platforms, designed for building, managing, and optimizing stores of any complexity.

[SessionReaper](#) and [PolyShell](#) rank alongside [CosmicSting](#), [TrojanOrder](#), and [Shoplift](#) — vulnerabilities responsible for thousands of e-commerce store breaches worldwide in previous years.

[SessionReaper](#) (CVE-2025-54236) is a critical vulnerability in Adobe Commerce and Magento Open Source that allows threat actors to take over customer accounts and execute remote code without authentication. Using this vulnerability, a threat actor can bypass the security system:

- creating a session — the threat actor first creates their own account on the site to obtain a legitimate token for accessing the store's API;
- tampering with session data — the threat actor sends a specially crafted PUT request to the store's API containing a malicious nested object that specifies the victim's identifier;
- uploading a malicious file — exploiting the vulnerability, the threat actor uploads a malicious PHP file to the server through an unprotected upload endpoint.

As a result, the threat actor gains access to the victim's session. If sessions are stored as files, the actor can substitute a malicious PHP file, which the server then executes, granting the actor full access.

[SessionReaper](#) has a CVSS score of 9.1, corresponding to a critical severity level. In October 2025, researchers recorded more than 250 attacks within 24 hours exploiting this vulnerability. As of June 2026, a large number of e-commerce stores remained vulnerable.

[PolyShell](#) is another critical vulnerability in the REST API of Magento and Adobe Commerce that has existed since the first release of Magento 2. It allows unauthenticated users to upload executable files to the server. Threat actors upload valid GIF or PNG images that simultaneously contain executable PHP code. Two distinct types of malicious payloads are in active use:

- A PHP web shell with cookie-based authentication. The shell checks the value of the cookie d against a hardcoded MD5 hash and then accepts arbitrary code through the eval function. The shell also implements a file upload feature;
- A PHP backdoor disguised as a GIF image that uses cookie-based authentication and the eval function to execute arbitrary code.

After gaining code execution through PolyShell, threat actors deploy a secondary backdoor that spreads across multiple directories to improve persistence. Scattering the copies ensures that at least one backdoor survives cleanup or reinstallation.

The shell has three functions:

- beacon — threat actors use this as a fingerprint to locate live shells across compromised stores: a simple HTTP request reveals which paths are executable;
- remote code execution — when the cookie d matches the hardcoded MD5 hash, the shell runs arbitrary PHP via eval(base64_decode(\$_REQUEST["id"])), giving the threat actor full control over the server;
- file upload — the up parameter copies an uploaded file to a path of the threat actor's choice, enabling further malware deployment.

The scale of the PolyShell attacks proved enormous: 79.5% of stores were attacked. Researchers [detected](#) 471 stores compromised in a single hour as threat actors exploited the PolyShell vulnerability at scale. Every production version of Magento Open Source and Adobe Commerce up to 2.4.8 is affected. Adobe fixed the underlying flaw in the 2.4.9 pre-release branch (alpha-3), but no isolated patch exists for current production releases.

According to F6 Threat Intelligence, the distribution of CMS (Content Management System) platforms and countries with the highest number of infected e-commerce stores is as follows.

CMS statistics for infected stores

CMS	% of infected stores
Magento	72.39%
WordPress	9.87%
Shopify	3.84%
OpenCart	2.01%
PrestaShop	1.46%
Other	10.42%

Magento accounts for the largest share of infected CMS platforms. Vulnerabilities in Magento plugins, together with the widespread use of various modules, make the platform a preferred target for threat actors.

In 2025, more than 10,500 unique [Magecart attacks](#) were recorded, compromising over 23 million online transactions. The term Magecart combines "Magento" and "shopping cart" and originally referred to a specific hacking

group. Today, however, the term denotes a class of attacks aimed at stealing sensitive information, primarily payment card data, from e-commerce sites. Exploiting vulnerabilities in various components of e-commerce systems and other online services, threat actors inject a malicious script into a store's payment pages that runs directly in the visitor's browser, intercepting the data entered and sending it to the threat actor's server.

A [Source Defense study](#) of 550 previously breached sites across 68 countries found that 18% remained infected a year after initial detection and that, among those, 57% exhibited a new or evolved attack rather than remnants of the original one. Moreover, 16% of the sites in the sample went offline, possibly as a direct consequence of the unresolved security issue. In 12% of cases, when security tools blocked a third-party script, the attack shifted to first-party JavaScript, meaning the site's own code, making it invisible to many security tools.

In 2025, threat actors began actively abusing legitimate services to host malicious code and exfiltrate data, using both previously observed schemes and entirely new ones: Google Tag Manager, BunnyCDN, Vercel, Discord webhooks, CodePen, Stripe, and Google Cloud. In March 2026, researchers [discovered](#) the first sniffer to use WebRTC DataChannels to retrieve its payload and exfiltrate stolen payment data. WebRTC connections are not covered by standard CSP rules because the connect-src directive does not control RTCPeerConnection, and the traffic travels over DTLS-encrypted UDP rather than HTTP, meaning that network tools that inspect HTTP cannot detect the data theft. The sniffer can obtain a valid nonce (a one-time CSP key) from a site's legitimate scripts and use it to inject its own code, allowing it to remain undetected even by systems that rely on CSP. Other [stealth techniques](#) used in 2025 and early 2026 include delivering sniffers through legitimate Google Tag Manager containers used by stores; using a loader on checkout pages to contact a dedicated threat actor account through the [Stripe API](#) and retrieve

the sniffer; using [SVG-onload injections](#), in which threat actors insert an invisible graphic element that automatically executes a disguised, string-encoded command when the page loads; triggering execution through browser events; storing data in localStorage and exfiltrating it only after the user leaves the page; and delivering the payload over WebSocket without ever writing it to disk.

Sniffer-as-a-Service

The market for JavaScript sniffers has reached an industrial scale. "[Sniffer by Fleras](#)" was responsible for 26% of all sniffer infections in 2025. "Sniffer by Fleras" is a JavaScript sniffer used to steal credit card data from e-commerce websites. It is sold on underground forums for prices starting at \$2,000 and allows threat actors to intercept sensitive customer data during checkout. In March 2026, the threat actors offered a second, enhanced "HIDDEN" version for \$6,000.



Fig. 26 — Screenshot of a forum post advertising the "HIDDEN" version of "Sniffer by Fleras"

Translated text from Fig. 26

Friends, a new product called "HIDDEN"

The final testing stage is complete. The product bypasses all CSP protections (script-src, unsafe-eval, connect-src), is invisible to WAFs, browser extensions, HTTP monitoring, and the Network tab. Payload delivery and data exfiltration take place outside the visibility of standard security tools. Three layers of obfuscation with key rotation, anti-debug with auto-blocking, and code that self-deletes from the DOM.


Does not use HTTP \ WS to transmit data. At all. Therefore:

- CSP (Content-Security-Policy) - connect-src, script-src, default-src - do not work. The transport is outside the scope of the policies.
- WAF (Web Application Firewall) - Cloudflare, Sucuri, Akamai, ModSecurity - see nothing. There is nothing to filter, there are no HTTP requests.
- Anti-skimmer solutions - Jscrambler, PCI scanners, MageReport, Sansec - look for suspicious JS requests to external domains. There are none.
- Network monitors - Fiddler, Wireshark filters on HTTP/HTTPS, corporate proxies - the traffic goes through a different channel. DevTools - the Network tab is clean. No XHR, fetch, image-beacon, or WebSocket. Empty.
- CORS - does not apply. There are no cross-domain HTTP requests to block.
- Antivirus and extensions - uBlock, Adblock, Malwarebytes Browser Guard - filter by URL lists and request patterns. No requests means no detection.
- Tag monitoring - Google Tag Manager audits, script inventories - the code is not loaded as an external script.

To transmit the code, it uses blockchain systems, where you can change your proxies in half a second.

The price for this version of the sniffer is \$6,000.

In 2025, the number of services offering ready-made solutions for exploiting compromised e-commerce sites also increased. The operators of these services handle the installation and maintenance of sniffers on compromised stores, while clients share the proceeds, typically receiving 50% of the revenue from the sale of payment card data or up to 70% of the raw data collected. Sniffer-as-a-Service offerings are becoming increasingly common on underground forums, enabling even low-skilled threat actors to target e-commerce platforms without having to build or maintain their own infrastructure.



Analysis of Compromised Data



F6 Threat Intelligence specialists use automated and manual methods to study various types of malware, botnets, and scam and phishing schemes designed to obtain sensitive information, which may include usernames and passwords for various systems as well as payment card data.

This section presents our analysis of compromised data, highlighting the countries and regions most frequently targeted, as well as the malware families most commonly involved.

Statistics on Compromised Payment Cards

Although most compromised cards are listed in card shops and then resold to other threat actors for subsequent cash-out, some cards may be published in cleartext on underground forums, in Telegram channels, and in other open and closed sources.

The statistics below summarize compromised payment card data collected by F6 Threat Intelligence.

Note. In this section, a compromised payment card refers to a record containing the card number, CVV, and expiration date. In most cases, this information is sufficient to make purchases through many online merchants and payment platforms.

The publication of a payment card number alone — which occurs frequently — does **not** constitute evidence that the card has been compromised.

Top 20 countries by number of compromised cards published in cleartext in open and closed sources, 2025

Nº	Country	Number of unique compromised cards	
1.	 US		694 657
2.	 MX		58 309
3.	 TR		42 220

Nº	Country	Number of unique compromised cards	
4.	 CA		32 722
5.	 GB		26 852
6.	 AE		23 974
7.	 BR		22 337
8.	 AU		16 482
9.	 FR		13 555
10.	 MY		12 743
11.	 IN		11 793
12.	 SA		11 377
13.	 IT		10 651
14.	 ES		9 037
15.	 DE		8 720
16.	 TH		8 639
17.	 CL		7 576
18.	 JP		7 393
19.	 PK		6 929
20.	 AR		6 039

In 2025, F6 Threat Intelligence systems identified more than 1.1 million compromised payment cards worldwide.

Top 20 countries by number of compromised cards published in cleartext in open and closed sources, first half of 2026

Nº	Country	Number of unique compromised cards
1.	 US	69 794
2.	 RU	8 982
3.	 MX	7 436
4.	 CA	4 978
5.	 TH	4 688
6.	 MY	4 234
7.	 GB	3 634
8.	 BR	2 734
9.	 AU	2 247
10.	 TR	2 046
11.	 KR	1 908
12.	 FR	1 845
13.	 IN	1 676
14.	 AE	1 565
15.	 PH	1 535
16.	 IE	1 483
17.	 BH	1 436
18.	 ES	1 417
19.	 DE	1 408
20.	 IT	1 377

As of mid-2026, the number of compromised cards had declined significantly compared with the same period in 2025, particularly cards published in cleartext on underground forums and Telegram channels. This decline is attributable to the following factors:

- Some users, especially those who had been actively publishing cards in cleartext on underground forums stopped doing so and switched to selling cards through card shops. F6 research indicates that some of the new card shops established by these users are scams: they take money from prospective buyers and provide nothing in return;
- Telegram has begun taking more aggressive action against potentially dangerous channels, making the platform less attractive to threat actors and reducing the overall volume of data available there.

Statistics on Payment Cards on Card Shops

Card shops are underground marketplaces where previously compromised payment cards are offered for sale. Before purchase, a prospective buyer can see some of the card's digits and its country of origin. A single card typically costs \$10-\$20.

The following statistics show the card-issuing countries most commonly represented in card shops.

Statistics on card-issuing countries represented on card shops in 2025

Nº	Country	Number of compromised cards
1.	 US	8 634 981
2.	 GB	1 651 424
3.	 CA	418 163
4.	 FR	255 696
5.	 AU	221 081
6.	 PE	136 300
7.	 SG	130 769
8.	 IT	102 487
9.	 MY	101 705
10.	 IN	101 237
11.	 UA	89 885
12.	 ES	80 956

Nº	Country	Number of compromised cards
13.	 PH	80 171
14.	 TR	79 292
15.	 RU	78 018
16.	 JP	67 018
17.	 MX	59 249
18.	 BR	53 512
19.	 NO	52 898
20.	 DE	49 092

In 2025, more than 13 million payment cards were listed for sale in card shops, according to F6 Threat Intelligence.

Statistics on card-issuing countries represented on card shops in the first half of 2026

Nº	Country	Number of compromised cards
1.	 US	1 684 910
2.	 GB	977 908
3.	 CA	72 336
4.	 FR	62 603
5.	 IT	44 295
6.	 AU	39 498
7.	 ES	31 466

Nº	Country	Number of compromised cards
8.	 UA	29 374
9.	 MY	22 824
10.	 RU	22 782
11.	 IE	19 851
12.	 NO	18 595
13.	 DE	16 049
14.	 PE	15 134
15.	 BR	13 435
16.	 JP	13 367
17.	 MX	12 848
18.	 ZA	10 852
19.	 TR	10 618
20.	 CL	10 157

In the first half of 2026, more than 3.5 million payment cards were identified as being offered for sale in card shops.

Statistics on Compromised Accounts and Stealers

The F6 Threat Intelligence system accumulates data on compromised accounts obtained from closed sources. The data analyzed includes malware logs, phishing attacks, botnets, and other unique sources of attack intelligence collected by F6 specialists in the course of cyber incident analysis.

Data obtained by threat actors using stealer malware accounts for the largest share of all compromised data. A stealer is a type

of malware designed to covertly steal confidential data from an infected computer.

No information security vendor can aggregate every available source. However, analysis of the data available to F6 allows us to determine which stealers are currently the most active and which countries' users are most exposed to stealer infections.

Top 10 stealers in 2025

Nº	Malware	Number of accounts
1.	RedLine Stealer Malware Group	440 813 282
2.	LummaC2	20 988 450
3.	Raccoon	10 075 625
4.	Vidar	5 867 070
5.	Stealc	2 101 729
6.	META Stealer	370 296
7.	Nexus Stealer	223 498
8.	RisePro	191 076
9.	Webrat	142 511
10.	CryptBot	117 522

Top 10 stealers in the first half of 2026

Nº	Malware		Number of accounts
1.	Vidar		8 158 187
2.	Raccoon		4 287 401
3.	Stealc		2 379 596
4.	RedLine Stealer Malware Group		1 986 759
5.	Webrat		294 482
6.	Blank Grabber		22 106
7.	AveMaria		22 035
8.	LummaC2		16 427
9.	SnakeKeylogger		10 408
10.	META Stealer		7 093

Top 10 countries by number of stealer infections in 2025

Nº	Country		Number of accounts
1.	 IN		3 802 881
2.	 BR		2 952 134
3.	 US		2 657 048
4.	 ID		1 666 251
5.	 VN		1 252 648
6.	 PH		1 129 613
7.	 TR		970 120
8.	 MX		922 102
9.	 EG		916 231
10.	 TH		897 677

Top 10 countries by number of stealer infections in the first half of 2026

Nº	Country	Number of accounts
1.	 IN	767 448
2.	 US	747 656
3.	 BR	603 400
4.	 ID	308 878
5.	 PH	209 958
6.	 VN	196 533
7.	 PK	180 104
8.	 MX	159 980
9.	 TR	145 317
10.	 BD	140 989

Data Leaks Landscape




Statistics

F6 Threat Intelligence analysts identified 164 unique new database breaches posted on underground forums and specialized Telegram channels during the reporting period


Threat actors often publish databases to build a strong reputation on the forums or to draw attention to their other activities. The figure below shows the breakdown of advertisements on underground forums and Telegram channels between database publications and sales during the reporting period.

Breakdown of leaked databases offered for sale or published, 2025-H1 2026



Each record in a leaked database may contain personally identifiable information such as names, ID numbers, phone numbers, or email addresses. The pie chart below shows the distribution of leaked records by region:


Breakdown of records from leaked databases by region, 2025-H1 2026



As the pie chart above shows, the CIS region accounted for the largest share of the leaked records observed, representing roughly one-third of the total (33.7%), followed by Latin America (27.3%) and the Middle East and Africa (17.6%). The remaining leaked records were split between APAC and other regions.

The pie chart below shows the distribution of observed leaks by region:

Breakdown of database leaks by region, 2025-H1 2026




The largest number of leaks observed during the reporting period occurred in the LATAM region, accounting for 33% of the total. The MEA region followed closely at 29%. Notably, although the CIS region accounted for a large share of leaked

records, this was not reflected in the number of leaks: the region saw only a few high-volume leaks.

The pie chart below shows the distribution of records in leaked databases by industry:


Breakdown of records from leaked databases by industry, 2025-H1 2026



As the pie chart above shows, the Government sector accounted for the largest share of leaked records (56.4%), followed by Retail (19.6%), Real Estate (6.7%), Education (5.7%), and Finance (5%).

The pie chart below shows the distribution of leaked databases by industry:

Breakdown of database leaks by industry, 2025-H1 2026



As the pie chart above shows, the Government sector again accounted for the largest share of observed leaks, although its share of leaks (23%) was smaller than its share of leaked records. This reflects a relatively small number of leaks, each involving a large number of records. The Government sector was followed by Finance (17%), Retail (10%), Professional services (7%), and Education (7%).

Notable threat actors

Several active and prominent threat actors specializing in the publication and sale of databases emerged on the underground during 2025 and H1 2026. Many engage in both publishing and selling leaked databases. Threat actors typically share a data sample from each leak along with their contact details and some information about the victim organizations.

Telegram and Session remain the most common ways to contact these threat actors, although many are moving away from Telegram due to the large-scale bans imposed on channels and accounts involved in illegal activity on the platform.

In most cases, the threat actors behind a leak provide no technical details about the methods and tools used to obtain the data, leaving this information undisclosed.

This chapter profiles several significant threat actors who either published numerous noteworthy leaks affecting companies worldwide or posted databases containing large numbers of records on underground resources during the reporting period.

888

The threat actor 888 has been active on several underground forums since August 2023. This actor's primary activity is publishing and selling compromised databases on forums, totaling more than 50 databases during the reporting period. A screenshot of one of the actor's profiles on a well-known underground forum is shown below:




Fig. 27 — Profile of threat actor 888 on one of the underground forums

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
⚡ Energy	38%
🛒 Retail	23%
🚚 Transportation and logistics	15%
🎬 Media and entertainment	7%
... Other	17%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
💻 IT	39%
🛒 Retail	27%
🏦 Finance	7%
💼 Professional services	5%
... Other	22%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 APAC	43%
 MEA	5%
 LATAM	4%
 Other	48%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 APAC	31%
 MEA	14%
 LATAM	5%
 Other	50%

The table below lists some of the threat actor's notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Thai health supplement provider	158,000	 Food and beverage,  Healthcare	 TH
International electronics company	5.2 GB of source code and internal documents	 Consumer electronics,  Healthcare,  Manufacturing	 KR
International electronics company	364 MB of source code and internal documents	 Consumer electronics,  Manufacturing	 KR
Nigerian fintech company	37,800	 Finance	 NG
Saudi e-commerce platform	46,000	 Commerce and shopping	 SA

ByteToBreach

The threat actor ByteToBreach has been active on one of the underground forums since August 2025. This actor's primary activity is the sale of more than 20 compromised data-bases on forums during the reporting period, which they claimed had been exfiltrated from victims' corporate networks. A screenshot of one of the actor's profiles on a well-known underground forum is shown below:




Fig. 28 — Profile of threat actor ByteToBreach on one of the underground forums

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
 Government and military	94%
 Finance	5%
 Telecommunications	1%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
 Government and military	28%
 Telecommunications	22%
 Energy	11%
 Finance	11%
... Other	28%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 CIS	93%
 MEA	5%
 LATAM	1%
... Other	1%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 MEA	28%
 LATAM	22%
 CIS	11%
... Other	39%

This threat actor primarily targets victims in Africa and the CIS regions. The table below lists some of the actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Nigerian fintech company	35,000 records and 3 TB of internal documents and KYC data	Finance	NG
Nigerian government agency	759 GB of internal documents (25 million documents)	Government and military	NG
Ethiopian oil company	500 GB of data	Energy	ET
Nigerian electric utility	Business databases and source code	Energy	NG
Uzbek telecommunications company	Full access to the corporate network and 3,000 employee records	Messaging and telecommunications	UZ

GordonFreeman

The threat actor GordonFreeman has been active on underground forums since January 2026. This actor’s primary activity is publishing and selling compromised databases on forums. During the reporting period, they posted messages about more than 30 databases that they published or offered for sale. A screenshot of the actor’s profile on a well-known underground forum is shown below:




Fig. 29 — Profile of threat actor GordonFreeman on one of the underground forums

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
Government and military	96%
Telecommunications	3%
... Other	1%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
 Government and military	69%
 Finance	9%
 Telecommunications	9%
 Transportation and logistics	9%
... Other	4%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 LATAM	99%
... Other	1%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 LATAM	94%
... Other	6%

This threat actor tends to target victims in Latin American countries. The table below lists some of the actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Venezuelan government service	13,800,000	 Government and military	 VE
Paraguayan electricity provider	1,650,000	 Energy	 PY
Chilean government service	10,000,000	 Government and military	 CL
Ecuadorian government service	17,000,000	 Government and military	 EC
Venezuelan transportation service	650,000	 Transportation and logistics	 VE

N1KA

The threat actor N1KA has been active on one of the underground forums since July 2025. The actor primarily publishes and sells compromised databases on forums. During the reporting period, they posted messages about more than 50 databases that they published or offered for sale. A screenshot of the actor’s forum profile is shown below:




Fig. 30 — Profile of threat actor N1KA on one of the underground forums

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
Retail	58%
Travel and tourism	17%
IT	15%
Government and military	8%
Other	2%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
IT	28%
Retail	17%
Education	11%
Travel and tourism	7%
Other	37%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
LATAM	16%
APAC	9%
Other	75%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
APAC	31%
LATAM	11%
MEA	11%
Other	47%

The threat actor tends to target victims in the APAC region. The table below lists some of the threat actor's notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Malaysian direct-selling company	200,000	 Commerce and shopping,  Sales and marketing	 MY
French e-commerce platform	21,300,000	 Commerce and shopping	 FR
Chinese email service	4,500,000	 IT	 CN
Brazilian fashion marketplace	40,000	 Clothing and apparel	 BR
Saudi digital technology company	11,600	 IT	 SA

KaruHunters

The threat actor KaruHunters was active on underground forums from August 2024 to February 2026. This actor's primary activity was publishing and selling compromised databases, with more than 20 posted during the reporting period. As of July 2026, the user is banned from the forum. A screenshot of the actor's profile on a well-known underground forum is shown below:




Fig. 31 — Profile of the threat actor KaruHunters on an underground forum

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
 Finance	82%
 IT	15%
 Healthcare	2%
... Other	1%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
 Telecommunications	27%
 Education	13%
 Healthcare	13%
 IT	13%
... Other	34%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 MEA	84%
 APAC	1%
... Other	15%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 APAC	13%
 MEA	13%
... Other	74%

The table below lists some of the threat actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Hungarian IT service platform	300,000	 IT	 HU
Moroccan online medical portal	41,700	 IT,  Healthcare	 MA
Nigerian bank	1,800,000	 Finance	 NG
Indian home appliance and electronics manufacturer	3,400	 Consumer electronics	 IN
Indonesian bank	7,000	 Finance	 ID

lulzintel

The threat actor lulzintel has been active on underground forums since August 2023. The actor's primary activity is publishing and selling compromised databases on forums. During the reporting period, the actor posted messages about more than 20 victims. A screenshot of the actor's profile on a well-known underground forum is shown below:




Fig. 32 — Profile of the threat actor lulzintel on an underground forum

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
Education	43%
Retail	28%
Healthcare	13%
... Other	16%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
Retail	28%
Education	17%
Healthcare	17%
IT	11%
... Other	27%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
APAC	47%
MEA	35%
LATAM	7%
... Other	11%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
MEA	37%
APAC	19%
LATAM	12%
... Other	32%

The threat actor tends to target victims in the MEA region. The table below lists some of the threat actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Saudi retail company website	72,500	Commerce and shopping	SA
Saudi transportation company	1,200,000	Transportation and logistics, Administrative services	SA
Saudi digital media company	107,000	Media and entertainment	SA
Saudi healthcare platform	180,000	Healthcare, IT	SA
Saudi transportation company	126,000	Transportation and logistics	SA

Tanaka

The threat actor Tanaka has been active on underground forums since June 2023. As of June 2026, the threat actor serves as a content moderator on several underground forums that specialize in distributing data leaks.

The threat actor’s primary activity is publishing compromised databases on various forums, with more than 10 such databases published during the reporting period. As a content moderator, the actor frequently reposts databases leaked and published by other threat actors or ransomware groups within the underground. A screenshot of the actor’s profile on a well-known underground forum is shown below:




Fig. 33 — Profile of the threat actor Tanaka on an underground forum

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
IT	71%
Education	18%
Travel and tourism	7%
... Other	4%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
 Education	20%
 Government and military	20%
 IT	20%
 Sports	20%
 Travel and tourism	20%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 LATAM	96%
 Other	4%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 LATAM	80%
 Other	20%

The table below lists some of the threat actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Brazilian software company	20,000,000	 IT	 BR
Brazilian education portal	5,160,000	 Education	 BR
Saint Lucian government service	1,116,935	 Government and military	 LC
Paraguayan sports education institute	143,692	 Sports	 PY
Mexican destination management company	1,983,503	 Travel and tourism	 MX

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 APAC	29%
 LATAM	19%
 MEA	10%
 Other	42%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 APAC	29%
 LATAM	19%
 MEA	10%
 Other	42%

The table below lists some of the threat actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Brazilian online education platform and marketplace	49,589,000	 Education,  Commerce and shopping	 BR
Mexican hotel	49,193,000	 Travel and tourism	 MX
German specialized technology firm	47,308,000	 Hardware	 DE
UAE health concierge platform	45,611,000	 Healthcare	 AE
Nigerian corporate research consultancy	44,281,000	 Professional services,  IT	 NG

Kazu

The threat actor Kazu has been active on several well-known underground forums since June 2025. The threat actor’s primary activity is selling compromised databases and access to corporate networks, more than 20 in total, on forums, on Telegram, and on their DLS (Data Leak Site) during the reporting period. A screenshot of the actor’s profile on one of the well-known underground forums is shown below:




Fig. 37 — Profile of the threat actor Kazu on an underground forum

The actor also ran a Telegram-channel used to announce new attacks, often including a sample of compromised data. In addition, they set up a DLS that remained active until January 2026, when the site went offline and their Telegram-channel was subsequently banned.

The threat actor remained inactive until May 16, 2026, when they created a new channel to publish victim samples and links to a new DLS, where the actor also posts information about ongoing leaks.

A screenshot of the current channel is shown below:




Fig. 38 — Screenshot of the threat actor Kazu’s Telegram channel

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
 Government and military	47%
 Healthcare	38%
 IT	13%
... Other	2%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
 Government and military	41%
 Healthcare	35%
 IT	6%
 Retail	6%
... Other	12%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 MEA	44%
 LATAM	40%
... Other	16%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 MEA	47%
 LATAM	40%
... Other	13%

The threat actor primarily targets victims in the MEA region. The table below lists some of the threat actor’s notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Saudi government portal	12.3 GB of internal documents and residents' personal data	 Government and military	 SA
UAE government portal	1.94 TB of internal documents	 Government and military	 AE
UAE nonprofit organization	249,000	 Charity,  Healthcare	 AE
UAE association	239 GB of internal documents	 Assosiation	 AE
Kuwaiti ministry	12 TB of internal documents	 Government and military	 KW

Satanic

The threat actor Satanic registered on one of the well-known underground forums on May 9, 2025. A screenshot of the actor’s profile on that forum is shown below:




Fig. 39 — Profile of the threat actor Satanic on an underground forum

In addition to this profile, the threat actor is also active on another major underground forum focused on the distribution of compromised data-bases. A screenshot of the actor’s profile on that forum is shown below:




Fig. 40 — Profile of the threat actor Satanic on an underground forum

While most of the threat actor’s activity is on these forums, their messages are not limited to these forums only. The actor is a member of over 10 other underground forums. The threat actor also runs a Telegram-channel (created on February 26, 2024), where they share more information about new breaches.




Fig. 41 — Screenshot of Satanic’s Telegram channel

During the reporting period, the threat actor’s main activity was the sale and publication of more than 20 compromised databases across various forums and Telegram channels.

The table below shows the distribution of leaked records by industry:

Industry	Percentage of leaked records
IT	51%
Retail	41%
Education	6%
... Other	2%

The table below shows the distribution of leaks by industry:

Industry	Percentage of leaks
 IT	39%
 Retail	33%
 Education	11%
... Other	17%

The table below shows the distribution of leaked records by region:

Region	Percentage of leaked records
 APAC	75%
 LATAM	1%
... Other	24%

The table below shows the distribution of leaks by region:

Region	Percentage of leaks
 APAC	56%
 LATAM	11%
... Other	33%

The table below lists some of the threat actor's notable data breaches during the reporting period:

Leak	Number of records	Industry	Country
Argentinian art supplies and bookstore chain	360,000	 Commerce and shopping	 AR
Multinational content management platform	1,750,000	 Software apps,  Content and publishing	 US
Multinational fintech company	102,000	 Finance	 NL,  AE
Saudi food manufacturer and distributor	Unspecified	 Food and beverage,  Manufacturing	 SA
Cryptocurrency platform	1,210,000	 Software,  IT,  Finance	 US


Notable data breaches

The following section describes, by region, some notable large-scale data breaches that occurred in 2025 and H1 2026.

LATAM

The LATAM region accounted for roughly 27.3% of all leaked records observed in 2025 and H1 2026, with more than 210 million records exposed by threat actors. The region’s share of total leaks was around 33%. The pie chart below shows the distribution of leaked records across countries in the LATAM region.


Distribution of leaked database records by LATAM country, 2025-H1 2026



As shown in the pie chart above, Argentina had the largest number of leaked records, at roughly 120 million (57%), followed by Mexico with 38 million (18%). The remaining countries in the region accounted for about 25% of the total. Argentina’s large share was driven by a major data breach involving government domains, described in more detail below.

The pie chart below shows the distribution of leaks by country in the LATAM region.


Distribution of database leaks by LATAM country, 2025-H1 2026



As shown in the pie chart above, Argentina ranked first by number of leaked records but second by number of leaks (20%), behind Mexico, which accounted for 50% of all leaks. This difference reflects the large number of leaks affecting Mexican government entities.

The pie chart below shows the distribution of leaked records by industry in the LATAM region:


Distribution of leaked database records by LATAM industry, 2025-H1 2026



The Government sector accounted for the largest share of leaked records in LATAM, at about 70%. The sector was prominent across all countries in the region and accounted for more than 146 million records. Finance ranked second with 16% of records (32 million records).

The pie chart below shows the distribution of leaks by industry in the LATAM region:

Distribution of database leaks by LATAM industry, 2025-H1 2026



As shown in the pie chart above, the Government sector accounted for the largest share of leaks in LATAM, at 44%, followed by Finance (19%), whose share was driven by several leaks in the banking software sector. Healthcare ranked third at (11%) of leaks, attributed to several breaches of medical labs and hospitals, mostly in Mexico.

Notable Data Breaches in LATAM

The following section describes several notable regional data breaches in more detail.

Publication of data allegedly belonging to the Argentinian government and a news channel

On May 6, 2026, a threat actor known as Skull1172 posted an advertisement on an underground forum offering data allegedly belonging to the Argentinian government and an Argentinian news channel.




Fig. 42 — Screenshot of a post on an underground forum by the threat actor Skull1172

According to the threat actor's message:

- All.gob.ar and.edu.ar domains were compromised between 2024 and 2026;
- The shared data is a filtered sample of the full dataset;
- The breach resulted in the release of more than 80 million lines of user data, including:
 - Usernames;
 - Emails;
 - Phone numbers;
 - Passwords;

- License plates;
- Identity documents;
- Nearly 155,000 email threads were compromised as well;
- The threat actor also scraped 2 APIs and obtained more than 32 million images and 12 million records;
- The threat actor also targeted one of Argentina's news networks, accessing administration panels and FTP servers, which yielded information about the network's employees.

The threat actor attributed the compilation of the data to the "EsqueleSquad" hacker group, though this could not be confirmed.

Sale of data allegedly belonging to an Ecuadorian government agency

On April 15, 2026, the threat actor GordonFreeman, described above, posted a thread on a well-known underground forum offering the database of an Ecuadorian government agency for sale.




Fig. 43 — Screenshot of a post on an underground forum by the threat actor GordonFreeman

The database exposed approximately 17 million rows of user information:

- Full names;
- Phone numbers;
- Addresses.

Sale of data allegedly belonging to a Mexican household solutions company

On May 5, 2026, a threat actor known as ColdK3y posted an advertisement on an underground forum offering data allegedly belonging to a Mexican household solutions company for sale.



Fig. 44 — Screenshot of a post on an underground forum by the threat actor ColdK3y

According to the threat actor's message:

- The breach exposed 10 million user records in JSON and CSV formats;
- The breach included user details such as:
 - Full names;
 - Email addresses;
 - Physical addresses;
 - Dates of birth;
 - Phone numbers;
- In a follow-up message, the threat actor stated that the data was priced at \$500.

Publication of data allegedly belonging to a Brazilian education platform

On April 25, 2025, the threat actor memberphp published the database of a Brazilian education platform on a well-known underground forum. The publication exposed approximately 2.7 million rows of user information:

- Email addresses;
- Passwords;
- Phone numbers.

The threat actor also shared a sample of the data.




Fig. 45 — Screenshot of a post on an underground forum by the threat actor memberphp

Sale of data allegedly belonging to a Venezuelan national service

On April 15, 2026, the same threat actor, GordonFreeman, posted a thread on a well-known underground forum offering the database of a Venezuelan national service for sale.




Fig. 46 — Screenshot of a post on an underground forum by the threat actor GordonFreeman


The database exposed approximately 13.8 million rows of user information (12.3 million records for natural persons and 1.5 million for legal entities):

- Email addresses;
- Phone numbers;
- Addresses.

MEA

The MEA region accounted for 17.6% of all leaked records observed in 2025 and H1 2026, with more than 136 million records exposed by various threat actors. The region's share of total leaks was 29%. The pie chart below shows the distribution of leaked records across countries in the MEA region.


Distribution of leaked database records by MEA country, 2025-H1 2026



As shown in the pie chart above, Egypt had the largest number of leaked records, at roughly 54 million (39%), followed by Saudi Arabia with about 27 million (20%). The UAE followed closely with 25 million records (18%), and Qatar accounted for 22 million (16%).

The pie chart below shows the distribution of leaks by country in the MEA region.


Distribution of database leaks by MEA country, 2025-H1 2026



As shown in the pie chart above, Saudi Arabia ranked first by number of leaks (27%). Egypt and the UAE followed, each accounting for 17% of total leaks. The large volume of records exposed in a breach involving an Egyptian ministry explains Egypt's overrepresentation in the record count relative to its smaller number of leaks.

The pie chart below shows the distribution of leaked records by industry in the MEA region:


Distribution of leaked database records by MEA industry, 2025-H1 2026



The retail sector accounted for the largest share of leaked records in MEA, at about 64%, or more than 87 million records. The sector was prominent across all countries in the region except Egypt. The Government sector ranked second with 27% of records (36 million). Notably, most Government sector records originated from a single breach involving an Egyptian ministry (around 31 million records). Excluding that breach, the region shows a clear trend toward attacks on the retail sector.

The pie chart below shows the distribution of leaks in the MEA region by industry:

Distribution of database leaks by MEA industry, 2025-H1 2026



As the pie chart above shows, Retail remained the leading sector by number of leaks in MEA, although its share was smaller (21%), followed by Government and Professional services (14% each).

Notable Data Breaches in MEA

The following section describes several notable data breaches in the region in more detail.

Publication of data allegedly belonging to a Nigerian bank

On September 28, 2025, threat actor KaruHunters published a Nigerian bank database on an underground forum. The publication exposed approximately 1.8 million rows of user information:

- Full names;
- Addresses;
- States;
- Email addresses;
- Phone numbers.

The threat actor also shared a sample of the data:




Fig. 47 — Screenshot of the post on an underground forum by threat actor KaruHunters

Sale of a Saudi transportation company database

On April 29, 2026, threat actor lulzintel published a thread on a well-known underground forum offering a Saudi transportation company database for sale. The database exposed approximately 1.2 million rows of user information:

- Full names;
- Addresses;
- Phone numbers.

The threat actor also shared a sample of the data:



Fig. 48 — Screenshot of the post on an underground forum by threat actor lulzintel

Publication of data allegedly belonging to a UAE trading company

On May 14, 2026, a threat actor known as purplepancake49 published a post on an underground forum in which they released data allegedly belonging to a UAE trading company.

According to the threat actor's post:

- The breach exposed 787 thousand records;
- The breach affected students in online learning programs across Middle Eastern countries:
 - UAE;
 - Oman;
 - Kuwait;
 - Qatar;
 - Saudi Arabia;

The link shared by the threat actor pointed to an SQL file containing user details such as:

- Full names;
- Email addresses;
- Phone numbers;
- Physical addresses.




Fig. 49 — Screenshot of the post on an underground forum by threat actor purplepancake49

Sale of data allegedly belonging to a Saudi job seekers website

On April 3, 2026, a threat actor known as Grubder published a post on an underground forum offering data allegedly belonging to a Saudi job seekers website for sale.

According to the threat actor's post:

- The data is priced at \$1,400;
- The database contains 735,000 records;
- The data consists of three parts:
 - Contacts: customer and candidate contact information for job seeking and recruitment;
 - Job applications: records of individual job applications;
 - Employer accounts: data on employer companies registered on the platform;
- The threat actor accepts escrow through the forum administration.

On February 19, 2026, the threat actor "TelephoneHooliganism" made a similarly structured post on another underground forum.




Fig. 50 — Screenshot of the post on an underground forum by threat actor Grubder

Sale of data allegedly belonging to a UAE online fashion and lifestyle retailer

On May 7, 2026, a threat actor known as Jeffery_Epstein published a post on an underground forum offering data allegedly belonging to a UAE online fashion and lifestyle retailer for sale.

According to the threat actor's post:

- Approximately 428 thousand records were released;
- The breach includes user details such as:
 - Full names;
 - Phone numbers;
 - Email addresses;
 - Physical addresses;
- The threat actor accepts escrow through the forum administration.





Fig. 51 — Screenshot of the post on an underground forum by threat actor Jeffery_Epstein

APAC

The APAC region accounted for 10.8% of all leaked records observed in 2025 and H1 2026, with more than 83 million records exposed by various threat actors. This made APAC the region with the fewest leaked records. The region accounted for 8% of total leaks. The pie chart below shows the distribution of leaked records by country in the APAC region.


Distribution of leaked database records by APAC country, 2025-H1 2026



As the pie chart above shows, Malaysia had the largest number of leaked records, at roughly 35 million (42%), closely followed by India with about 34 million (41%). China accounted for 9 million records (11%).

The pie chart below shows the distribution of leaks by country in the APAC regions.


Distribution of database leaks by APAC country, 2025-H1 2026



As the pie chart above shows, Vietnam accounted for the largest share of leaks in the APAC region (26%), although these leaks involved relatively few records. Thailand ranked second (18%), but this was not reflected in its share of records because many of the leaks involved documents and files. India and Malaysia each accounted for 13% of leaks. Malaysia's share of leaks was lower than its share of records because many of those records came from a single large leak involving a real estate marketplace.

The pie chart below shows the distribution of leaked records in the APAC region by industry:


Distribution of leaked database records by APAC industry, 2025-H1 2026



Real estate accounted for the largest share of leaked records in APAC, at roughly 44%, largely due to major leaks in Malaysia and Vietnam totaling more than 36 million records. Education ranked second (40%), driven by 33 million records leaked in India. Information technology accounted for 6% of total records, or 4.5 million, mostly from China and India.

The pie chart below shows the distribution of leaks in the APAC region by industry:

Distribution of database leaks by APAC industry, 2025-H1 2026



As the pie chart above shows, Retail and Government sectors jointly accounted for the largest shares of leaks in the APAC region, at 20% and 17%. Education followed at 13%. Although Real estate accounted for a larger share of records, it represented only 10% of leaks because the sector saw only a few high-volume leaks.

Notable Data Breaches in APAC

The following section describes several notable data breaches in the region in more detail.

Sale of data allegedly belonging to a Malaysian online property and real estate marketplace

On May 8, 2026, a threat actor known as hulky published a post on an underground forum offering data allegedly belonging to a Malaysian online property and real estate marketplace for sale.



Fig. 52 — Screenshot of the post on an underground forum by threat actor hulky

According to the threat actor's post:

- Approximately 34.5 million records were released;
- The breach includes user details such as:
 - Full names;
 - Phone numbers;
 - Dates of birth;
 - Full addresses;
- The threat actor accepts escrow through the forum administration.

On September 3, 2025, the threat actor made a similarly structured post on a different underground forum.

Publication of data allegedly belonging to an Indian education platform

On February 2, 2025, threat actor Future-Seeker published a database from an Indian education platform on an underground forum. The publication exposed approximately 33 million rows of user information:

- Names;
- Emails;
- Phone numbers.

The database consists of more than 120 CSV files with a total size of 29.2 GB.




Fig. 53 — Screenshot of the post on an underground forum by threat actor FutureSeeker

Sale of data allegedly belonging to Vietnamese homeowners

On May 8, 2026, a threat actor known as hulky published a post on a well-known underground forum offering data allegedly belonging to Vietnamese homeowners for sale.




Fig. 54 — Screenshot of the post on an underground forum by threat actor hulky

According to the threat actor's post:

- Approximately 1.9 million records were released;
- The breach includes user details such as:
 - Full names;
 - Phone numbers;
 - Dates of birth;
 - Physical addresses;
- The threat actor accepts escrow through the forum administration.

On September 3, 2025, the threat actor made a similarly structured post on a different underground forum.

Sale of data allegedly belonging to a Vietnamese ministry

On May 13, 2026, a threat actor known as KurdFemboys published a post on a well-known underground forum offering data allegedly belonging to a Vietnamese ministry for sale.




Fig. 55 — Screenshot of the post on an underground forum by threat actor KurdFemboys

According to the threat actor's post:

- Approximately 480 thousand records were released;
- The breach includes user details such as:
 - Full names;
 - Email addresses;
 - Phone numbers;
 - Physical addresses;
 - Document details;
 - Workplace details;
- The threat actor published a sample of the data and threatened to sell the remainder to third parties unless the victim organization contacted them to secure its infrastructure.

The link shared by the threat actor points to a CSV file containing 50,000 rows of data.

Sale of data allegedly belonging to a Thai e-commerce website

On April 3, 2026, a threat actor known as Grubder posted an advertisement on an underground forum offering data allegedly belonging to a Thai e-commerce website for sale.




Fig. 56 — Screenshot of the underground forum post by the threat actor Grubder


According to the threat actor's post:

- The data is priced at \$1,300;
- It contains 284,000 records;
- The data consists of three parts:
 - Contacts: registered user/customer contact details;
 - Purchase Orders: customer order history and details;
 - Support Tickets: customer support interactions and ticket records;
- The threat actor accepts escrow through the forum's administration.

CIS

The CIS region accounted for 33.7% of all leaked records observed in 2025 and H1 2026, with more than 257 million records exposed by various threat actors. The region's share of total leaks was 4%. The pie chart below shows the distribution of leaked records across countries in the region.


Distribution of leaked database records by CIS country, 2025-H1 2026



As shown in the pie chart above, Tajikistan had the largest number of leaked records, at roughly 217 million (84.5%), followed by Kazakhstan with about 25 million (9.5%). Uzbekistan accounted for 15 million records (6%).

The pie chart below shows the distribution of data breaches across the countries of the CIS region.


Distribution of database leaks by CIS country, 2025-H1 2026



As shown in the pie chart above, Kazakhstan led by number of leaks, accounting for 57% of the total, followed by Uzbekistan at 29%. Tajikistan ranked third at 14% despite accounting for the majority of leaked records. This disparity is attributable to a single large leak from one of the country's government portals.

The pie chart below shows the distribution of leaked records in the CIS region by industry:


Distribution of leaked database records by CIS industry, 2025-H1 2026



As shown in the pie chart above, the Government sector leaks accounted for the overwhelming majority of records in the CIS region, at about 96%. This was largely due to a single large leak in Tajikistan involving more than 217 million records. Education ranked second (3%), driven mainly by leaks from educational platforms in Kazakhstan totaling 8.5 million records.

The pie chart below shows the distribution of leaks in the CIS region by industry:

Distribution of database leaks by CIS industry, 2025-H1 2026



As shown in the pie chart above, the Government sector led by number of leaks, accounting for 43% of the total, followed by Education (29%) and Finance (28%). Although these sectors accounted for only small shares of records, each accounted for several low-volume leaks.

Notable Data Breaches in CIS

The following section describes several notable data breaches in the region in more detail.

Sale of a database from a Tajikistan government portal

In September 2025, data was leaked from one of Tajikistan's government portals. The threat actor omnipotent put the data up for sale on an underground forum, claiming that the leak contained information on individuals crossing the country's border. According to the seller, the full database contained roughly 217 million rows.




Fig. 57 — Screenshot of the underground forum post by the threat actor omnipotent

Publication of a private database containing roughly 16 million rows of data on Kazakhstan residents

On June 14, 2025, the threat actor obladaet published a RAR archive containing 16 million rows of data on Kazakhstan residents in a private Telegram channel. The records included the following fields:

- Full name;
- Gender;
- Date of birth;
- Taxpayer identification number (TIN);
- Phone numbers;
- Citizenship;
- Ethnicity;
- Physical address;
- Residence start date;
- Residence end date.




Fig. 58 — Screenshot of the Telegram post by the threat actor obladaet

Sale of data on Uzbekistan citizens

On February 2, 2026, the threat actor ByteToBreach posted an advertisement on an underground forum offering data on Uzbekistan citizens for sale. The threat actor claimed to have extracted data on more than 15 million users, including:

- Full name;
- Physical addresses;
- Dates of birth;
- Phone numbers;
- Email addresses;
- Passport details.




Fig. 59 — Screenshot of the underground forum post by the threat actor ByteToBreach

Publication of a database from an online platform for schools in Kazakhstan

On December 24, 2025, data was leaked from the website of an online platform for schools in Kazakhstan. The data was published in a private Telegram group chat.




Fig. 60 — Screenshot of the Telegram post by the threat actor

The publication exposed roughly 2.2 million rows containing user data in the following format:

- Full name;
- Phone numbers;
- Email addresses;
- Physical addresses;
- Dates of birth.

Publication of a database from a Kazakh online education platform

On December 9, 2025, data was leaked from a Kazakh online education platform. The data was published in a private Telegram channel.

The publication exposed roughly 6.3 million rows of user data, including:

- Full name;
- Email addresses;
- Phone numbers;
- Physical addresses;
- Hashed passwords.

Other notable data breaches

The following section describes several other notable data breaches that fall outside the regions covered above.

Publication of a database containing 16 billion stealer logs

On June 19, 2025, cybersecurity researchers reported the exposure of a dataset containing 16 billion user records associated with Apple, Facebook, Google, GitHub, Telegram, and government services. However, these companies were not directly breached.

According to the primary source, the leak reported in the news is a compilation of data:

- logs collected by stealers from infected victim devices;
- sets of credentials (login:password) used in automated credential stuffing attacks;
- leaks, presumably both publicly available and privately distributed.

Numerous threat actors publish similar collections every day on underground forums and Telegram channels. Some are distributed privately.

Analysis revealed that the database contained about 16 billion records, including the service URL, username (login), and password, presumably in plain text. Researchers also noted that it contained duplicate entries, meaning the actual volume of data was unknown and may have been significantly smaller. No publication of a similar set was observed in open sources in 2025 or 2026. However, a similar thread titled "More than 16 billion Link Login Pass lines" was identified on March 4, 2024. It is also worth noting that users of the underground forum described the data as outdated.

Those messages in the underground forum thread have since been removed. However, they were preserved in the F6 Threat Intelligence monitoring system.

Screenshots of the messages from the F6 Threat Intelligence system are shown below:




Fig. 61 — Screenshot of the underground forum post by the threat actor TheKuntess in the F6 Threat Intelligence system




Fig. 62 — Screenshot of the underground forum post by the threat actor the_one_who_phishes in the F6 Threat Intelligence system

Publication of data allegedly belonging to a popular e-commerce platform

On August 18, 2025, the threat actor N1KA, described above, published a database from a popular international e-commerce platform on an underground forum. The full dataset exposed approximately 21.3 million rows of user information:

- Names;
- Companies;
- Full addresses;
- Phone numbers.

The threat actor also shared a sample of the data:




Fig. 63 — Screenshot of the underground forum post by the threat actor N1KA

AI-related data breaches

With the rise of the AI industry over the past few years, the number of data breaches and attacks has increased sharply. The following section describes several notable data breaches that occurred in 2025 and the first half of 2026.

Publication of data allegedly belonging to an AI-focused cloud platform

On May 14, 2026, the threat actor dien-cracked posted an advertisement on an underground forum announcing the publication of data allegedly belonging to a cloud platform that supports AI teams. The dataset included source code and additional data, totaling 1,360 directories and 10,239 files.



Fig. 64 — Screenshot of the underground forum post by the threat actor dien-cracked

Publication of data allegedly belonging to an AI-powered marketing platform

On July 3, 2025, a threat actor known as Satanic, described above, posted an advertisement on an underground forum offering data allegedly belonging to an AI-powered marketing platform.




Fig. 65 — Screenshot of the underground forum post by the threat actor Satanic

According to the threat actor's post:

- The breach comprised roughly 783,000 records;
- The breach exposed user details, including:
 - Full names;
 - Email addresses;
 - Phone numbers;
 - Physical addresses;
 - Employment details;

Publication of data allegedly belonging to a South Korean AI video analytics company

On December 3, 2025, a threat actor known as 888, described above, posted an advertisement on a well-known underground forum announcing the publication of data allegedly belonging to a South Korean AI video analytics company.




Fig. 66 — Screenshot of the underground forum post by the threat actor 888

According to the threat actor's post:

- The breach occurred in November 2025;
- The breach resulted in the leak of the platform's source code.

Publication of data allegedly belonging to an AI platform

On May 9, 2026, the threat actor xorcat, described above, posted an advertisement on an underground forum announcing the publication of data allegedly belonging to an AI workflow automation platform.




Fig. 67 — Screenshot of the underground forum post by the threat actor xorcat

According to the threat actor's post:

- Extraction date: May 2026;
- Estimated scope: data on more than 500,000 users and 2 million interaction logs, including:

- Email addresses;
- Names;
- Hashed passwords;
- Complete AI conversation history;
- Total leak size: a 2.26 GB compressed RAR archive.

Publication of data allegedly belonging to another popular AI platform

On May 9, 2026, the same threat actor, xorcat, posted an advertisement on an underground forum offering data allegedly belonging to another popular AI platform operated by a U.S.-based company.




Fig. 68 — Screenshot of the underground forum post by threat actor xorcat

According to the threat actor's post:

- The breach exposed Custom CRM / Help Center tickets;
- The data spans October 24, 2024 through December 7, 2024;
- The breach exposed 79,000 records containing user details, such as:
 - Email addresses;
 - Phone numbers;
 - Ticket details.

Underground Forum Takedowns, 2025-2026



Between January 2025 and June 2026, law enforcement agencies carried out a series of successful joint operations to take down underground forums.

During lengthy investigations, police and intelligence services collected and analyzed communications between forum members, tracked their financial transactions, and conducted coordinated raids resulting in the arrest of forum participants and the seizure of servers and backups. Following these operations, seizure notices appeared on the main pages of the affected forums. This coordinated international campaign resulted in partial or complete disruption of six prominent underground platforms. Five new forums emerged to replace them, although in some cases multiple alternative underground platforms appeared following the takedown of a single forum. Notably, although one disrupted forum successfully restored its infrastructure, the operation triggered a profound crisis of trust. Anticipating potential compromise or law enforcement infiltration under the newly appointed administration, some threat actors severed ties with the legacy platform and launched independent alternative underground forums. The following sections examine the major takedowns and emergence of underground platforms that occurred during the analysis period.

Takedown of the Cracked and Nulled underground forums

On January 30, 2025, reports emerged that several underground forums had been taken down as part of a joint operation by law enforcement agencies from the United States, Romania, Australia and other countries, aimed at combating cybercrime. As part of the international Operation Talent, the domain names of four major underground platforms were seized: Cracked, Nulled, MySellIX, and StarkRDP.

The Nulled and Cracked underground forums, in operation since at least 2015 and 2018, respectively, were used to distribute stolen login credentials, compromised databases, malware and other cybercrime-as-a-service offerings. Associated forum services were also dismantled the Sellix payment processor, which was widely used on Cracked, and the StarkRDP hosting service. According to a law enforcement statement, StarkRDP was a virtual hosting and remote desktop (RDP) service that threat actors used to stage cyberattacks and compromise credentials, while MySellIX, a platform built on the Sellix processor, was used to sell stolen data, compromised accounts, and malware.

Alongside the takedown, Spain's National Police detained two suspects — a man and a woman — linked to the administration of the platforms. Following the seizure of infrastructure, official notices were posted on the main pages of the forums, announcing that they had been placed under law enforcement control.




Fig. 69 — Screenshot of the Cracked and Nulled forums' main pages following the takedown

Takedown of the LeakBase underground forum

Between March 3 and March 4, 2026, law enforcement agencies from 14 countries, coordinated by Europol and the FBI, conducted a large-scale operation to dismantle the LeakBase underground forum. Launched in 2021, the forum primarily specialized in the distribution of databases, infostealer logs, and other sensitive information. According to the official law enforcement statement, by the end of 2025 it had more than 142,000 registered users and 32,000 published posts, making it one of the most popular forums for trafficking in sensitive information.

The forum's takedown was part of the international Operation Leak, coordinated by Europol and the FBI. During the operation, authorities seized the forum and replaced its main page with an official law enforcement notice confirming that the platform had been taken down. The notice emphasized that the forum's database and entire archive of internal communications had been seized and would be used as evidence in subsequent investigations.




Fig. 70 — Screenshot of the LeakBase main page following the takedown

Takedown of the XSS forum by French intelligence services

On July 22, 2025, the French Prosecutor's Office officially announced the arrest of the administrator of the XSS forum in Ukraine.

XSS, formerly known as DaMaGeLaB, first appeared in 2013. The forum served as a marketplace for threat actors engaged in various types of cybercrime, including phishing, fraud, carding, malware development, DDoS tools, and related services. XSS members communicated through the encrypted Jabber server these-secure[.]biz, which was also reportedly maintained by the forum administrator. The operation was carried out by international intelligence services with the direct support of Europol and the French police.

**PARQUET DE PARIS
TRIBUNAL JUDICIAIRE**
*Liberté
Égalité
Fraternité*

LA PROCUREURE DE LA REPUBLIQUE

Paris, le 23 juillet 2025

Communiqué de presse

Mardi 22 juillet, la personne soupçonnée d'être administratrice du forum cybercriminel russophone XSS a été interpellée dans le cadre d'un dossier d'information judiciaire ouvert par le parquet de Paris.

Actif depuis 2013, ce forum était l'un des principaux lieux de la cybercriminalité mondiale. Il permettait la vente de malwares, d'accès à des systèmes compromis, de données dérobées et de services liés aux ransomwares. Le forum était doublé d'un serveur Jabber de messagerie chiffrée, facilitant les échanges anonymes entre cybercriminels.

L'enquête ouverte le 2 juillet 2021 par la section de lutte contre la cybercriminalité du parquet de Paris et confiée à la Brigade de Lutte contre la Cybercriminalité de la police judiciaire de la préfecture de police, a permis la mise en place de captations judiciaires sur le serveur Jabber these-secure.biz. Les messages interceptés ont révélé de nombreuses activités illicites en lien avec la cybercriminalité et les ransomwares, et permis d'établir qu'elles avaient engendré au moins 7 millions de dollars de bénéfice.

Une information judiciaire a été ouverte le 9 novembre 2021, des chefs de complicité d'atteintes à un système de traitement automatisé de données, d'extorsion en bande organisée, et d'association de malfaiteurs.

Identifié grâce à une seconde captation réalisée dans le cadre d'une information judiciaire, l'administrateur du forum a été interpellé le 22 juillet en Ukraine par les autorités ukrainiennes (Cyber Department, the Security Service of Ukraine et General Prosecutor's Office of Ukraine), en présence de la police française en charge de l'enquête et avec l'assistance d'Europol.

Les investigations se poursuivent sous la direction des juges d'instruction.

Laure BECCUAU,
Procureure de la République

Contact presse : 06 07 16 42 28
scom.parquet.tj.paris@justice.fr

Fig. 71 — Screenshot of the statement by French intelligence services

The following day, July 23, 2025, a take-down notice appeared on the XSS forum's clearnet domain, announcing that it had been seized through the joint efforts of the intelligence services involved in the operation.



Fig. 72 — Screenshot of the notice posted on the forum's domain

The French Prosecutor's Office later clarified that the investigation had begun in July 2021, after investigators gained access to communications on the thesesecure[.]biz server. According to the intercepted messages, the forum's administrators had earned more than €7 million from cybercriminal activity. By September 2024, law enforcement had identified the forum's lead administrator and was building an evidentiary case for his arrest.

Shortly after the takedown, on July 24, 2025, the XSS forum was restored on a new domain. However, users discovered that their deposits had been completely wiped and that the former administrators had been blocked. Forum members raised concerns that the restored forum was likely controlled by law enforcement and operating as a honeypot.

Launch of the DamageLib forum

After the XSS underground forum was restored, many members lost trust in the new administration and decided to leave the platform. Some former moderators moved to XSS's former cloud storage platform, DamageLib, as they announced in a separate post on the XSS forum.



Fig. 73 — Screenshot of the XSS moderators' post expressing distrust of the new administration

Translated text from Fig. 73

Warning! The XSS forum has been compromised and turned into a honeypot.

The forum administrator was indeed arrested in Kyiv; the new administrator is an outsider who came out of nowhere.

The criminal case involves not only Tokha; the forum itself is also under investigation - which means all of us are, too! The moderation team has set up an alternative secure platform, DamageLib.

Details are available in the thread at this link.

Use a new handle when registering. No email address is required - enter anything. If you lose your password, it cannot be recovered.

The platform completely discontinued trading to reduce the risk of a takedown. Instead of engaging in commercial activity, members began publishing technical articles and tutorials.

Launch of the Rehub forum

Following the seizure of one of the oldest forums and its subsequent restoration under presumed law enforcement control, many former XSS members decided to leave the platform. One of them was an XSS moderator operating under the alias Rehub, who created an underground platform of the same name on August 10, 2025.




Fig. 74 — Screenshot of the post by the user Rehub announcing the launch of the eponymous forum on August 10, 2025

Translated text from Fig. 74

The forum is open. Registration is free. Welcome to RehubCom.

Immediately after the forum was created, the administrator Rehub set up a trading section offering various illicit goods and services, including initial access to corporate networks, malware, and cryptocurrency wallets. The threat actor recruited the forum's most active members to the moderation team to quickly build trust among users.




Fig. 75 — Screenshot of the Rehub forum's main page

Takedown of the RAMP underground forum

On January 28, 2026, the FBI and the U.S. Department of Justice initiated the takedown of another popular Russian-language underground platform, RAMP, which facilitated the distribution of malware, initial access to compromised systems, and other illicit services. After the forum was seized, a takedown notice was posted on the platform's domains.



Fig. 76 — Screenshot of the takedown notice for the RAMP underground forum

Although the FBI did not release official details of the operation, technical analysis of DNS records indicated that the RAMP domains were being served by the agency's servers at the time. This indicated that the platform's user databases, internal communications archives, and members' IP address logs had likely been seized in full. In addition, one of RAMP's former administrators, known by the alias Stallman, confirmed the infrastructure seizure in a statement on the XSS forum and announced that they would not create any similar platforms in the future.




Fig. 77 — Screenshot of the message from the threat actor Stallman

Takedown of BreachForums and the emergence of lookalike successor platforms

Repeated takedowns of the BreachForums underground forum

From 2015 to 2022, one of the largest underground forums was RaidForums. It hosted discussions of malware, data breaches and hacking tools, as well as the sale of initial access and databases. On May 12, 2022, the forum was taken down by Europol as part of the international Operation Tourniquet.



Fig. 78 — Screenshot of the RaidForums main page following the takedown

On March 16, 2022, a threat actor using the alias pompompurin launched a new forum based on cloned RaidForums source code and identical branding. Created to replace the RaidForums platform taken down by law enforcement, it was named Breached, or BreachForums.




Fig. 79 — Screenshot of pompompurin's post announcing the launch of the Breached forum

However, in March 2023, the forum was also seized following the arrest of its administrator and creator, pompompurin (real name Connor Brian Fitzpatrick, according to official statements by the authorities). On June 12, 2023, the forum was relaunched under the name BreachForums. Its new administrators were threat actors using the aliases ShinyHunters and Baphomet.




Fig. 80 — Screenshots of ShinyHunters' post announcing the launch of BreachForums

On May 15, 2024, a notice appeared on the BreachForums domain stating that the forum had been taken down by the FBI and the U.S. Department of Justice with the support of international partners. However, BreachForums was restored just days later, on May 23, 2024.



Fig. 81 — Screenshot of the original BreachForums front page following the 2024 takedown

In 2025, the forum was taken down several more times. On May 15, 2025, it unexpectedly went offline. For an extended period, the administration issued no statements explaining the shutdown. As a result, numerous unsubstantiated claims emerged, with threat actors spreading disinformation that the current administrators, ShinyHunters and IntelBroker, had been arrested and that the forum's domain and source code were being offered for sale. For example, a user operating under the alias Anastasia posed as a Breach-Forums administrator and offered to sell a full dump of the forum's database along with its source code for \$2,000.



Fig. 82 — Screenshot of the advertisement claiming that the original BreachForums source code was for sale

The post was published on the domain breached[.]fi, which was not an official Breach-Forums domain and was later taken down. Other unverified domains mimicking the original site also appeared. One such resource contained no messages in its discussion threads. Several threat actors created similar forum clones to lure users of the original forum and harvest their credentials.




Fig. 83 — Screenshot of a fake BreachForums site

On May 28, 2025, the forum administration officially announced the discovery of a zero-day vulnerability in the MyBB content management system. The forum was taken offline preemptively; the administration assured users that no data had been compromised and that reports of the arrests of administrators ShinyHunters and IntelBroker were false.




Fig. 84 — Screenshot of the PGP message from the original BreachForums administration

On July 25, 2025, the forum was relaunched by the original team, preserving all accounts, private messages, and user reputation. According to an administrator using the alias N/A, none of the core administrators had been detained, and IntelBroker had never actually owned the forum.




Fig. 85 — Screenshot of the post by administrator N/A on the original BreachForums

On October 10, 2025, the FBI and French law enforcement carried out another takedown of BreachForums. A message from the FBI was posted on the main page stating that the domain had been seized.



Fig. 86 — Screenshot of the seizure notice for the original BreachForums in 2025

The ShinyHunters group confirmed the seizure of the forum's backend and backups, stating: "This seizure was inevitable. The era of forums is over." According to ShinyHunters, none of the key administrators had been arrested. However, they claimed that they had no intention of launching a new iteration of the forum and warned that any forums with an identical name were honeypots.

Even so, on December 18, 2025, an announcement appeared in the official BreachForums Telegram channel stating that the forum was back online.




Fig. 87 — Screenshot of the announcement of BreachForums' return on December 18, 2025

After its brief return in December 2025, administrator N/A took the platform offline again on March 15, 2026, announcing an intention to sell it. This triggered distrust among the moderation team, whose members began leaving the platform en masse. This development led to the emergence of new underground platforms.

The emergence of a new iteration of Breached

On November 21, 2025, a post appeared in the "Breached" Telegram channel containing a link to the channel itself, along with a screenshot showing notices of the takedowns of previous BreachForums iterations and the slogan "Breach. Sleep. Repeat."




Fig. 88 — Screenshot of the post in the "Breached" Telegram channel

Following the takedown of the original BreachForums on October 10, 2025, the community interpreted this post as an announcement of an imminent return. The actual launch, however, did not occur until January 17, 2026, when a threat actor using the alias HasanBroker posted an announcement in the "Breached" Telegram channel declaring the launch of a new version of BreachForums.




Fig. 89 — Screenshot of HasanBroker's announcement of the launch of the BreachForums underground forum

The forum preserved the architecture and functionality of the original BreachForums but switched to a red color palette. Some time later, the platform changed its name to Breached. In their first post on the forum, HasanBroker stated that their goal was to build a single, consolidated platform that would unite groups of threat actors affected by law enforcement pressure.




Fig. 90 — HasanBroker's first post on the new iteration of Breached

The emergence of the Spear underground forum

The underground forum named Spear appeared without any major announcements on October 19, 2025, shortly after the FBI and French intelligence services' operation against the original BreachForums. The forum was dedicated to sharing knowledge about cybercrime, data breaches, and ready-made hacking tools, as well as selling initial access to corporate networks and databases and offering various illicit services.




Fig. 91 — Screenshot of the Spear forum's main page

On March 17, 2026, the "Breaches" Telegram channel announced a partnership between the Breached platform, run by the threat actor HasanBroker, and the Spear forum.



Fig. 92 — Screenshot of the announcement of the alliance between the forums in the "Breaches" Telegram channel

The Emergence of PwnForums

From the underground community's perspective, PwnForums was the most legitimate successor of BreachForums. It began operating on April 5, 2026. Its emergence was a direct consequence of the governance crisis at the original forum. As noted earlier, on March 15, 2026, the original BreachForums administrator using the alias N/A took the platform offline and announced their intention to sell it, triggering a mass departure of moderators.

On April 5, 2026, the “PwnForums” Telegram channel published an announcement about the launch of the PwnForums underground forum.




Fig. 93 — Screenshots of the announcement of the PwnForums underground forum's launch in its eponymous Telegram channel

The administration stated that the forum had been created in collaboration with the original BreachForums moderation team and some long-time members. In addition, all threads and posts published on the original BreachForums up

to March 25, 2026, were migrated to PwnForums, along with the forum's user accounts. The same message was published on the forum itself.




Fig. 94 — Screenshot of the administrator's announcement of the forum's launchThe Fake BreachForums

The Fake BreachForums

On March 22, 2026, threat actors deployed an exact copy of the original Breach-Forums interface to compromise the accounts of underground community members. By replicating the original site's visual identity and preserving historical discussion threads, the site's creators exploited users' familiarity with the platform's frequent relaunches.




Fig. 95 — Screenshot of the fake BreachForums front page

The attack relied on phishing: the forum allowed users to log in with their existing Breach-Forums credentials, which were then intercepted by the fake forum's moderators. The administration made no announcements. The forum was active for a little over a month and went offline on May 26, 2026.

Conclusions

The 2025–2026 period marked a turning point for the underground forum ecosystem. Law enforcement agencies in multiple countries refined a well-coordinated model of collaboration, moving beyond simply blocking domains to fully seizing infrastructure and arresting key figures. This unprecedented pressure significantly disrupted the established model of trust within the underground community. As threat actors themselves acknowledged: “The era of underground forums is over.”

Fearing compromise — or that a platform may be under law enforcement control — users are migrating en masse from major forums to alternative platforms or private messengers. New players are quickly filling this vacuum: successor forums such as PwnForums and Spear

are emerging and attempting to legitimize themselves by importing old databases and recruiting former moderators.

The main side effect, however, has been a surge in fraud within the forums themselves. Threat actors have begun creating clones of popular platforms that are visually indistinguishable from the originals. These resources are used to steal credentials and conduct social engineering: users are persuaded to pay an “entry fee” for access to a supposedly secure version of the forum. Thus, while efforts to combat cybercrime undoubtedly inflict significant damage on major underground platforms, such campaigns also further fuel the fragmentation of the cybercriminal community and the rise of scam projects that exploit for their own gain the migration of members from defunct platforms.



Malware Sales on Underground Markets



Malware commercialization has become increasingly common, with developers often choosing to monetize their tools by selling them to other threat actors rather than using them in their own attacks. The growing popularity of this business model has given rise to a highly specialized market with a broad range of offerings.

During the period of analysis, the amount of listings for customizable malware solutions on underground forums has been steadily growing, as threat actors increasingly transition from monolithic malware development to modular architectures. In contrast to the traditional model, where actors offer only a single type of malware, customizable solutions allow attackers to select functionality depending on the attacker's needs. For example, a single customizable malware solution may include multiple modules simultaneously: various types of malware (stealers, RATs, loaders), customizable builds, defense evasion components, as well as modules for infrastructure management and data analysis. Such packages are often designed for different operating systems and attack scenarios, and combined with the Malware-as-a-Service model and modular distribution, they significantly lower the barrier to entry into cybercrime, enabling threat actors with varying levels of technical expertise to carry out sophisticated attacks.

The malware market continues to expand alongside cybercrime and the emergence of narrowly focused groups specializing in specific attack categories. Meanwhile, many developers treat malware development as their primary source of income rather than a means of conducting their own attacks, further reinforcing the professionalization and commercialization of the field.

During the reporting period, F6 Threat Intelligence experts observed a substantial increase in the number of listings offering malware for sale on underground resources. To attract the attention of potential buyers, adversaries publish detailed promotional posts outlining the malware's technical capabilities alongside proof-of-concept demonstration.

This chapter examines the most notable malware offerings identified on underground platforms during the period of analysis, categorized by key parameters such as malware type (stealer, RAT) and target operating system.

Stealers

Threat actors use stealer malware to exfiltrate sensitive data from infected devices. The stolen data most often includes logins and passwords from browsers, password managers, and email clients; cookies and session tokens that allow threat actors to gain unauthorized account access; bank card credentials and form-autofill information; and cryptocurrency application files.


Stealer infections most often result from neglecting basic digital hygiene, such as installing files from suspicious sources, clicking malicious links or opening malicious email attachments. This allows threat actors to conduct continuous mass campaigns against users worldwide.

Stealers remain the most widespread type of commercial malware. Their prevalence reflects a low barrier to entry: subscription-based distribution and ready-to-use web panels make such malware accessible even to threat actors with minimal technical knowledge, directly increasing the number of attacks aimed at stealing confidential and sensitive data. Elevated attack frequencies are also driven by the growing num-

ber of threat actors using ready-made malware as well as by using the Malware-as-a-Service (MaaS) model, under which cybercriminals generate exponential financial returns relative to upfront expenditures.

During 2025–2026, competition in the malware market intensified amid the growing popularity of AI and extensive public disclosures of legacy malware source code, which threat actors use as the basis for building their own products. Notably, sellers compete primarily on stealthy persistence on compromised devices and the quality of comprehensive technical assistance across all phases of malware utilization, rather than solely vying for the volume of data harvested by infostealers.

The following section describes the most notable underground listings offering stealer malware for sale during the period of analysis, grouped by target operating system.



Windows

Stealers targeting Windows users represent one of the most widespread categories of malicious software. Windows OS is one of the world's most widely used operating systems dominating the enterprise market, which gives threat actors a numerous pool of potential victims. As a result, Windows users are more frequently targeted by infostealer malware than users of other desktop operating systems.

AURA Stealer

On July 8, 2025, the user under the alias AuraCorp posted an announcement on a well-known underground forum offering AURA Stealer for sale. Later similar listings appeared on other popular underground platforms. AURA Stealer targets OS Windows (v.7-11) and allows to harvest account credentials and cookies, cryptocurrency wallet information, and files from various desktop applications specified by an attacker in the web panel.




Fig. 96 — Screenshot of the announcement offering AURA Stealer for sale on an underground forum

Translated text from Fig. 96

AURA Stealer - You don't need this (Just kidding. You do. Urgently!)

AURA Stealer is a suite of carefully refined solu-

tions in which every detail is designed for results, not appearances.

Data collection from more than 110 browsers, 70 applications (including wallets and 2FA tools), and more than 250 browser extensions. And that's not all: you can add any application or extension to the data-collection configuration in just a couple of clicks.

We collect cookies from running Chromium browsers without terminating the process, so cookies are not corrupted. We use our own shell-code for App-Bound decryption. All decryption is performed server-side, so the build performs as few suspicious actions as possible.

A loader is included. The build is approximately 500-700 KB and is hardened with a custom-built morpher. AURA offers all this and more!

The stealer is written in C++ using NtAPI and WinAPI. The initial build size of AURA Stealer is 500-700 KB, and packers reduce it to 170-250 KB.

AURA Stealer has the following features:

1. Stealth mechanisms include encrypted function imports that are resolved immediately before invocation, as well as AntiVM, AntiDebug, and ApiHammering (which simulates legitimate activity through random system calls);
2. Communication with the command-and-control server uses HTTPS with an additional layer of AES-256 encryption. An archive containing the stolen data is assembled in memory and transmitted in chunks, ensuring that at least some of the data reaches the threat actors' server even if security tools detect the malware.



Fig. 97 — AURA Stealer control panel home page

AURA Stealer is distributed by subscription offered in two tiers:

- The Basic plan, priced at \$295 per month, includes basic filtering of logs from infected devices, two builds, and one Telegram bot for data exfiltration;
- The Advanced plan, priced at \$585, removes most Basic plan restrictions. It includes advanced log filtering by country, infection date, IP range, installed application, custom tag, keyword, and cryptocurrency wallet type; up to three bulk log exports; up to five malware builds; and up to five linked Telegram bots for subsequent data exfiltration.

Needle Stealer

Needle stealer takes a different architectural approach. The user under the alias Lupin offered it for sale on a well-known underground forum on March 2, 2026, and subsequently relisted the offer on other famous underground platforms. Rather than a monolithic build, the malware is an open-source modular platform that can be configured by the buyer to steal cryptocurrency wallet data, banking account data, and other information from victims' devices.




Fig. 98 — Screenshot of the announcement offering Needle stealer for sale on an underground forum

Needle imitates legitimate browser extensions and cryptocurrency applications, allowing attackers to intercept seed phrases and initiate unauthorized transactions from victims' wallets.




Fig. 99 — Needle management web panel and modules

Although Needle was initially developed to steal data only from Windows devices, the threat actor later introduced add-on modules, sold separately, to target alternative operating systems.

The modules distributed as part of the Needle solution suite include:

- A Browser Wallet Grabber module for Windows and macOS, priced at \$2,499: it scans Chrome, Edge, Brave, and Opera and replaces cryptocurrency wallet extensions with malicious counterparts;

- An AntiLedger module priced at \$1,999: it mimics cryptocurrency wallet applications to intercept seed phrases;
- A custom browser extension for website spoofing priced at \$1,999;
- An installer for malicious browser extensions priced at \$2,000 per supported OS (Windows, macOS);
- Launcher modules that act as malware loaders, priced at \$499;
- A management web panel priced at \$999;
- A Win+R Loader priced at \$4,499: a malicious landing page displays a fake CAPTCHA that prompts the victim to execute a malicious command. The victim is instructed to open the Windows Run dialog by pressing Win+R and then execute malicious code copied from the clipboard;
- A Preset Collection priced at \$999: ready-made malware module configurations, including a CC Grabber module that collects bank card data from website checkout pages and intercepts two-factor authentication codes. The collection also includes a module that replaces the cryptocurrency wallet address and transfers the stolen amount when a victim withdraws funds from a legitimate cryptocurrency exchange.

KatzStealer

The first offering of KatzStealer was published on BreachForums on April 13, 2025. After the forum was seized on May 15, 2025, the author was forced to move operations to another underground platform. On August 17, 2025, the actor posted there a similar announcement promoting this malware. KatzStealer targets Windows (v.7–11) devices and collects data from more than 50 browsers, 90 browser extensions, and popular cryptocurrency wallets and messengers.



Fig. 100 — Screenshot of the announcement offering KatzStealer for sale on an underground forum

Written in C and assembly programming languages with minimal reliance on standard libraries, KatzStealer remains one of the most compact malware reviewed during the research period: the build size ranges from 30 to 140 KB taking about 3–15 seconds to generate.

KatzStealer has the following technical features:

- Direct NTDLL system calls;
- Dynamic string encryption and protection of transmitted malware logs with the Argon2 and SHA-256 algorithms in offline mode;
- Automated brute-force of common cryptocurrency wallet passwords without disrupting normal operation of the victim's browsers.

The KatzStealer subscription price ranges from \$250 to \$3,750, depending on the subscription term and the chosen tariff plan — personal (no API access) or team.

STORM stealer

In late February 2026, F6 TI analysts discovered an underground forum listing in which the user under the alias StormStealer offered STORM stealer with loader functionality for rent. The malware is developed to collect credentials for various services, cryptocurrency wallet information and data from popular messengers' applications.




Fig. 101 — Screenshot of the announcement offering STORM stealer for rent on an underground forum

Translated text from Fig. 101

STORM is a stealer developed by a professional team for demanding customers. The project is designed for maximum ease of use and fast, efficient, and, above all, secure operation.

Each build is compiled on a real Windows Server system and is unique. The stealer is written in C++ without using standard libraries.

Compilation is performed with Microsoft Visual C++ (msbuild).

Build time is 1–3 minutes.

As of this post, the build is 460 KB with all protections and obfuscation enabled.

Runs on Windows.

The build is delivered to the panel in an encrypted archive together with the archive password. After compilation, the build remains available for download from the panel for seven days.

Developed in C++ without standard libraries, the malware build size is approximately 460 KB. A key feature of STORM stealer is that

exfiltrated data is transmitted to the C2 server alongside its decryption key, reducing the likelihood of endpoint detection.




Fig. 102 — STORM stealer management web panel

STORM stealer is distributed by subscription with the following tariff plans:

- "Demo" (\$300 per week, no technical support provided);
- "Standard" (\$900 per month, including the provision of technical support for configuring the malware);
- The "Team" plan (\$1,800 per month) is intended for a group of threat actors, providing access for up to 100 users and supporting 200 configuration types.

VoltaStealer

On May 3, 2026, the threat actor under the alias VoltaTeam posted an underground forum announcement for VoltaStealer | VoltaC2, which uses an architecture similar to STORM stealer but offers a broader feature set.

The malware collects account credentials and cryptocurrency wallet data, as well as session files from instant messaging applications.



Fig. 103 — Screenshot of the announcement offering VoltaStealer for sale on an underground forum

Translated text from Fig. 103

Reliable data collection. Check-in rate: approximately 95%. Each build is at least 75% unique. Maximum profit.

VoltaStealer is a private infostealer with its own web panel. It was written from scratch and contains no public code. It operates under a MaaS model: the panel and builder run in the browser, and installation on third-party hosting is not supported.

Important information about the product:

- Development languages: C and ASM, with minimal use of C++. No CRT dependencies or third-party libraries. Imports are hidden.
- Build size: approximately 205–215 KB after morphing and obfuscation. The build runs entirely in memory and creates no temporary files on disk.
- All decryption is performed server-side. The build sends raw data; nothing is processed or stored in plaintext on the victim's device.
- System calls are invoked directly, without high-level WinAPI functions.
- Data is transmitted in chunks, minimizing the chance of detection by behavioral analyzers. If the build terminates early, some of the data will already be on the server.
- The entire collection and transmission cycle runs without creating temporary files on disk.
- A reworked and hardened proprietary traffic-encryption and compression engine, implemented without standard system calls.
- A private morpher makes each build at least 75% unique.

- Runtime FUD: the focus is not only on static detection but also on in-memory behavior. We perform a cleanup every Monday; if Defender flags the build, we clean it immediately.
- AntiVM, AntiSandbox, and anti-debugging features are built in. If such an environment is detected, the build exits cleanly and leaves no artifacts.
- A geoblock covering the CIS and the former Soviet Union is built into both the build and the server and cannot be removed under any plan.
- The build runs in 5–10 seconds. The time required for a log to appear depends on the victim's internet connection.
- Tested on Windows 10 22H2, Windows 11, and Windows Server 2022. Windows 7 and 8 have not been tested but should work. Dedicated servers running Windows Server sometimes fail to check in because they may be identified as a VM or sandbox.
- The panel is hosted on our servers. Installation on third-party hosting is not supported.
- Check-in rate: approximately 95%. Roughly 30% faster than previous versions. Logs are clean and contain no junk.

The malware is written in C and Assembly, with limited use of C++, omitting CRT dependencies. According to the threat actor, VoltaStealer was tested on Windows 10 22H2, Windows 11, and Windows Server 2022. The base build size is 310–320 KB.

The seller highlights the following technical features of VoltaStealer:

- The stealer is memory-resident and operates entirely in the RAM of the victim's device without creating temporary files; the collected data is decrypted server-side;
- A private morpher makes each build at least 75% unique;
- A proxy system that uses the Telegram and Steam ecosystems;

- A non-resident clipper that replaces cryptocurrency wallet addresses;
- A built-in loader;
- Chunked data transmission to evade behavioral analyzers;
- A proprietary traffic encryption and compression system that does not use standard system calls.




Fig. 104 — VoltaStealer management web panel and pricing plans

VoltaStealer is distributed by subscription with three tariff plans:

- The Basic plan (\$250) includes full data collection, a clipper, a loader, a proxy system, and three Telegram bots for exfiltrating data from infected devices;
- The Pro plan (\$350) adds domain-based filtering; smart search by number of passwords, cookies, wallets, and installed applications; and access to a limited API for automating log sorting and export;
- The Team plan (\$650) provides unlimited API access and allows third-party profiles to be added with customizable restrictions by user group and functionality;
- For an additional \$15 per month, the actor allows buyers to link their own phishing URLs, which redirect victims to a page hosting the malware download.

Yellow Stealer

On April 28, 2026, the user under the alias Polaris Web posted an announcement on an underground forum offering Yellow Stealer for sale. The malware targets Windows x64 devices and harvest account data for various online services, including cryptocurrency wallet data, messenger session tokens and files from installed desktop applications specified by the buyer.



Fig. 105 — Screenshot of the announcement offering Yellow Stealer for sale on an underground forum

The malware is developed in C and C++, resulting in an executable file size of approximately 960–970 KB.

The listing on the forum highlights the following features of Yellow Stealer:

- Reading locked databases through VSS shadow copies;
- The ability to inject into suspended processes via system calls and to escalate privileges by bypassing UAC;
- Anti-analysis mechanisms, protection against re-execution, and automatic cleanup of activity artifacts.



Fig. 106 — Yellow Stealer control panel

The Yellow Stealer MaaS subscription costs \$69 per month. To host the web panel, the seller recommends that buyers use their own infrastructure and traffic-tunneling tools, with GAIA.BAT Crypter used to protect the executable files.

macOS

Threat actors are showing growing interest in the Apple ecosystem. While Windows stealers focus on broad automated collection, macOS stealers leverage social engineering to trick users into revealing system passwords, allowing them to harvest data from platform-specific targets like the Keychain and Safari.

mac.c stealer

On March 14, 2025, the user *mentalpositive* posted a listing offering mac.c Stealer for sale, a macOS data stealer that targets devices running Sierra (>10.12.6) or later and supports the x86_32 and x64 processor architectures. The malware collects passwords, cookies, Telegram sessions, and cryptocurrency wallet data with the ability to decrypt the system's keychain.




Fig. 107 — Screenshot of the announcement offering mac.c Stealer for sale on an underground forum

Translated text from Fig. 107

mac.c macOS Stealer is a stealer for macOS devices. It works on all versions of the operating system starting with macOS Sierra (>10.12.6). Written in C, it has a build size of approximately 50 KB. Both architectures are supported: x64_86 and ARM. We also offer a module that replaces the Ledger application to phish users' seed phrases. The stealer collects cookies, passwords, autofill data, and browsing history from Chromium-based browsers; device information; Telegram sessions; desktop cryptocurrency wallets and cryptocurrency extensions; and the device's decrypted

keychain. You can also change the text displayed in modal windows during password prompts, making the job even easier!

A distinctive feature of mac.c Stealer is its ability to change the text displayed in modal windows during system password prompts, making social engineering more effective by encouraging victims to enter their credentials repeatedly in a malicious dialog. Data decryption is fully offloaded to the C2 server, reducing the load on the infected device and the likelihood of detection by security tools.

By June 2025, the developer had substantially expanded the product's functionality: the number of supported browsers had grown to twelve, including Safari for cookie collection; paid modules priced at \$1,000 targeting hardware cryptocurrency wallets had been introduced; and the file-collection feature had been rewritten to allow customers to optionally disable the feature based on their preferences.

macOS and Windows

Cross-platform stealers designed to run on both Windows and macOS form a distinct category. Supporting both major desktop operating systems allows threat actors to reach a broader pool of victims than malware limited to a single platform, increasing the effectiveness of data-theft campaigns. Such versatility also raises a product's value on underground markets, where buyers generally prefer a cross-platform solution to one targeting a single OS.

Remus stealer

On March 2, 2026, the user under the alias RemusStealer introduced Remus Stealer malware on an underground forum. The malware is designed to steal credentials, cryptocurrency wallet data, and files from Windows and macOS devices and, allegedly, other operating systems.




Fig. 108 — Screenshot of the announcement offering Remus Stealer for sale on an underground forum

Translated text from Fig. 108

This is a MaaS stealer with new and, most importantly, genuinely useful features unavailable elsewhere. It collects more data than any other solution on the market. After extensive closed testing, we can now offer the most stable, user-friendly, and compelling product on the market. We understand how expensive traffic is, so, thanks to technical solutions not previously seen in this field, we are

ready to provide near-100% uptime. Not a single data log will be lost!

Panel

Thanks to a new data-collection and check-in statistics system, you can identify precisely where check-ins are being lost and adjust the relevant rules to obtain more complete logs. Statistics are presented in an easy-to-use format; you do not need to be an analyst or hold a doctorate in mathematics—even a child could understand them. We have developed a carefully designed web panel in which everything is calibrated and in its proper place; after 20 minutes of use, you will be able to operate it with your eyes closed. There is no maze of settings, no clutter of unnecessary buttons, no multitude of identical pages, and no extra steps. Build

This is a stealer; its main task is to take everything from the victim's computer that is not nailed down—and at least try to take what is. We designed it to collect as much data as possible quickly and quietly, without alerting the victim. Unlike competing products, it does not close browsers or place excessive load on the network or CPU: if the victim notices nothing, the material remains available longer and there is a better chance of using it effectively. Our product performs all of its tasks effectively; more detailed technical information is provided below. The build uses system calls exclusively to interact with the victim's system, helping it evade hooks and other EDR detection methods.

The source code of Remus Stealer is written in C++, while the web interface is developed in React. Communication with the server uses a proprietary protocol with full traffic encryption. A distinctive feature of Remus Stealer is a built-in algorithm that detects honeypots, virtual machines and dedicated servers preventing execution in those environments. The grabber module allows to steal data from 21 browsers, 16 cold cryptocurrency wallets, 38 applications that store authentication data (2FA tools, password managers, etc.), more than 250 Chrome exten-

sions, and 68 Firefox extensions. It also collects payment card data and browser history, analyzes files on the victim's device, and provides statistics on data-collection effectiveness, including files successfully exfiltrated from the device.



Fig. 109 — Remus management web panel

The Remus Stealer control panel has the following features:

- Advanced search and sorting of malware logs;
- Notifications of successful infections via Telegram;
- Management of data collection from infected devices, including support for popular cryptocurrency wallets, with automatic recommendations for adjusting filter rules;
- Generation of malware builds.

Remus subscription plans range from \$250 per month for the basic plan to \$1,000 for a team subscription.

Stealc_v2

In addition to newly developed malware, updated versions of well-known malware are actively promoted on forums. One example is a listing published by the user plymouth on April 2, 2025, announcing Stealc_v2, a new version of the Stealc stealer. According to the author, this

version had been in development for about six months and featured a completely rewritten codebase that retained almost none of the functionality of previous Stealc versions.




Fig. 110 — Screenshot of the announcement announcing the release of the Stealc_v2 stealer

Translated text from Fig. 110

stealc_v2 release!

We're pleased to present a completely new version of our stealer!

Development took nearly six months - this is entirely new software that borrows almost nothing from older Stealc versions.

This is a completely new codebase, an entirely new approach to stealer development in general!

Key changes in the new version include fully server-side decryption of data from Chromium-based browsers, automatic brute-forcing of cryptocurrency plugins (at release, only MetaMask is supported), and a builder integrated directly into the admin panel.

A brief list of changes:

We have reworked our software so you will not have to learn how everything works; everything is now as simple, accessible, and clear as possible.

You now have a builder integrated directly into the admin panel with no limit on the number of builds, a user and worker management system, and automatic brute-forcing of cryptocurrency plugins on your server.

Stealc_v2 is written in C++ and produces builds of approximately 770 KB. The threat actor claimed that it can infect devices running all modern desktop operating systems.

Key additions in the updated malware include fully server-side decryption of data from

Chromium-based browsers, automatic brute-forcing of cryptocurrency plugins, and a builder integrated into the admin panel with no limit on the number of builds. Traffic is protected using an RC4-based algorithm, RAM usage during operation is 3–5 MB, and collected files are transmitted through individual HTTPS requests without creating a local archive. Stealc_v2 supports more than 23 browsers, over 100 cryptocurrency extensions, 15 desktop wallets, popular messengers, game launchers, and VPN clients. It also includes a separate module for collecting Safari cookies on macOS, seed-phrase extraction, and a configurable file grabber with shortcut-based search.




Fig. 111 — Stealc_v2 management web panel and functionality

The Stealc_v2 subscription plans are:

- \$300 per month;
- \$700 for three months;
- \$1,000 for six months.

ParallelUnivers

On February 10, 2026, the user ParallelUnivers posted an announcement on a well-known underground forum offering a multifunctional RAT malware with built-in stealer capabilities. The malware targets Windows and macOS devices, establishes persistence, intercepts authentication data, and collects confidential information.




Fig. 112 — Screenshot of the announcement offering the ParallelUnivers malware for sale on an underground forum

The malware represents a comprehensive solution that combines remote administration with data theft. It intercepts keyboard input and clipboard contents, while a separate module extracts credentials from browsers, email clients, and specialized applications. The seller claimed that the malware is undetectable by modern security tools, supports automated commands and custom templates, bespoke code modifications tailored to the buyer's requirements, and conceals its infrastructure behind Tor nodes. The seller also offers free Crypt services for build obfuscation and Binder services for embedding a malicious file in a legitimate one.




Fig. 113 — ParallelUnivers malware control panel

As of May 2026, ParallelUnivers was offered for a fixed price of \$3,000.

Gremlin stealer

On March 15, 2025, the user under the alias CoderSharp posted the first listing in a Telegram channel of the same name, demonstrating how the malware steals data from browsers, cryptocurrency wallets, FTP and VPN services, and messengers, then exfiltrates it to a C2 server.



Fig. 114 — Screenshot of the announcement offering Gremlin Stealer for sale in the threat actor’s Telegram channel

Analysis of the listing showed that Gremlin Stealer, written in C#, is a variant of Sharp Stealer with a codebase comparable to Hannibal Stealer. The developer focused particularly on bypassing Chrome’s V20 cookie protection and creating a self-contained build process that does not require components to be downloaded from the internet during compilation. The malware collects a wide range of data: passwords and cookies from popular Chromium- and Gecko-based browsers, payment card data, cryptocurrency wallet data, FTP and VPN credentials, and Telegram and Discord session tokens. It also takes screenshots and collects detailed system information (BSID, HVID, RAM, CPU, GPU, and IP address).





Fig. 115 — Gremlin Stealer management web panel

Remote Access Trojans (RAT)

Threat actors use remote access trojans (RATs) to covertly gain full control of infected devices. RATs allow them to operate systems remotely, capture keystrokes, take screenshots, and access cameras and microphones. Common infection vectors include users installing files from suspicious sources, clicking malicious links, and opening malicious email attachments.

Unlike stealers, which are mostly designed for one-time data collection, remote access trojans are focused on long-term control over infected systems. A notable trend during

the reporting period was the commercialization of these tools through subscription-based models, previously associated primarily with stealers, as well as the growing inclusion of multimedia capture, covert control, and built-in extortion capabilities. Classic remote administration is increasingly combined with stealer capabilities, creating comprehensive malware suites.



Windows

RAT families targeting Windows users are especially widespread. Windows dominates enterprise environments, where access to workstations and servers is particularly valuable to threat actors. The platform's large user base also gives attackers a broad pool of potential victims, making Windows users the most frequent targets of malware infections that result in remote control of compromised systems.

RAT Cracker stealer

In April 2026, F6 TI analysts discovered an underground forum listing in which the user o1001 offered RAT Cracker Stealer, a comprehensive Windows solution that combines remote access trojan capabilities, stealer and loader functionality, with an emphasis on automatically cracking cryptocurrency wallets and bypassing modern browser protections.




Fig. 116 — Screenshot of the announcement offering RAT Cracker Stealer for sale on an underground forum

Its key features include Auto Wallet Cracking, which automatically brute-forces wallet passwords, and support for Chrome AppBound, which enables decryption keys to be extracted in standard user mode without privilege escalation or opening a debugging port. It also implements an HVNC module that bypasses protections

in newer Chromium versions, a Browser Cloner for transferring the victim's profile to the threat actor's machine, event-driven task execution when the active window or clipboard contents change, and reverse VNC, FTP, SOCKS5, and Telnet tunnels. The functionality of the malware is fully customizable: support for new browsers and wallets can be added independently, without waiting for developer updates.




Fig. 117 — RAT Cracker Stealer management web panel

The RAT Cracker Stealer subscription costs \$5,000 per month. For greater stealth, the seller claims compatibility with any crypter and offers optional use of ForeverLoader, which provides basic webRAT functionality.

ATROPOSIA

In March 2026, an underground forum listing was discovered in which the user Melatonin offered ATROPOSIA Advanced Remote Access Trojan, a comprehensive Windows solution written in C++ and Rust that combines RAT and stealer functionality.




Fig. 118 — Screenshot of the announcement offering ATROPOSIA RAT for sale on an underground forum

The malware has a modular architecture consisting of standalone builds with no external dependencies.

The key features of ATROPOSIA RAT include:

- A choice of code-mutation level: Standard, Advanced, or Aggressive;
- Automatic UAC bypass;
- A Dead Drop Resolver for obtaining the C2 address from a Steam profile;
- HVNC with support for the Chrome, Firefox, Brave, and Opera browsers;
- Browser profile cloning;
- HRDP;
- Redirection of requests via DNS hijacking;
- A local vulnerability scanner.




Fig. 119 — ATROPOSIA RAT management web panel

The panel provides centralized client management, including a file manager for working

with files on infected systems, tools for viewing stolen data, and integration with Telegram bots for data exfiltration.

ATROPOSIA is available under five pricing options:

- \$200 per month;
- \$500 per quarter;
- \$900 for six months;
- \$1,500 per year;
- \$3,200 for a lifetime license.

Steaelite RAT

Alongside expensive suites, budget-conscious offerings are also available on underground forums. On November 4, 2025, the user under the alias Steaelite introduced Steaelite RAT on an underground forum. The malware combines remote access trojan (RAT) and stealer functionality and targets devices running Windows OS (v.10 and v.11).



Fig. 120 — Screenshot of the announcement offering Steelite RAT for sale on an underground forum

Its key capabilities include:

- An HVNC module for covert browser monitoring that bypasses black-screen protections in banking applications;

- Real-time direct access to the device's webcam and microphone;
- A keylogger for capturing user input;
- Protection against detection by popular antivirus solutions;
- Remote control of the infected device;
- A file manager for working with the victim's file system;
- Recovery of passwords and credentials from browsers and system stores;
- Application and process management.




Fig. 121 — Steelite RAT management web panel

Steelite is sold as a one-time purchase for \$249. The buyer receives an executable to deploy on their own infrastructure.

RMS RAT

Some sellers take a different approach by modifying legitimate remote administration tools rather than developing custom malware. RMS RAT follows this model. On March 1, 2026, the user Saruman777 posted a listing for the MaaS offering on a well-known underground forum.




Fig. 122 — Listing offering RMS RAT for sale on an underground forum

Translated text from Fig. 122

We have been developing RMS RAT since 2019. Throughout that time, we operated quietly and selectively, without publicity. Today, we are stepping out of the shadows and opening 10 slots for access to our solutions based on well-known RMS tools.

Why this is the best choice on the market

Unlike most sellers, we develop and patch the solutions ourselves rather than downloading publicly available builds that are already detected. We implement all bypasses, fixes, and updates manually, tailoring them to current detection methods and real-world use cases. We offer the following RMS RAT options:

- NetSupport
- ConnectWise
- GoToResolve
- DattoRMM
- AteraAgent
- SimpleHelp

The capabilities of our solutions:

- Remote screen monitoring of the victim's device
- Control of user input (keyboard/mouse)
- Two-way file transfer
- Remote launch of applications and commands
- Remote command line / administrator commands
- Monitoring and data capture
- Establishing persistence

- Communication with the command-and-control server (C2)
- File manager
- File download and dropper functionality
- Advanced lateral movement within the network

RMS RAT is based on modified versions of legitimate remote administration software, which inherently helps it evade many security tools.

According to the threat actor, RMS RAT has been in development since 2019. The author regularly releases malware updates to evade anti-virus products that use signature- and behavior-based detection.

Supported remote administration tools include NetSupport, ConnectWise, GoToResolve, DattoRMM, AteraAgent, and SimpleHelp.

The malware provides capabilities commonly associated with RATs used in targeted attacks:

- Remote screen monitoring;
- Keylogger;
- Two-way file transfer and a file manager for interacting with files on the infected system;
- Launching applications and commands with administrator privileges;
- System persistence;
- Communication with the C2 server;
- Lateral movement within the corporate network.

Available delivery methods include:

- ClickOnce / ClickFix;
- .exe and .msi files;
- Short PowerShell or VBS scripts;
- Use of legitimate EV certificates from RMS vendors.

The RMS RAT subscription starts at \$2,000 per month and varies depending on the chosen RMS solution.

OnyxC2

On March 15, 2026, listings offering OnyxC2 RAT for sale were found on two underground platforms. Distributed under the MaaS model, the malware is designed to provide remote access to and steal data from Windows devices.




Fig. 123 — Screenshot of the announcement offering OnyxC2 for sale on an underground forum

Written in C++ and assembly, the malware generates a unique signature for each build, which reduces the likelihood of detection.

Its main component is the Advanced Stealer module, which collects data from:

- 37 Chromium-based browsers, including 95 extensions and 5 2FA extensions;
- 8 Gecko-based browsers, including 14 extensions and 1 2FA extension;
- 5 password managers;
- 17 cryptocurrency wallets;
- 3 VPN applications;
- 2 remote access clients;
- 3 messengers;
- 5 email clients;
- 1 gaming platform.

The remote access functionality includes:

- HVNC with session cloning, using an implementation distinct from TinyNuke;
- LSASS Dump for extracting credentials from memory;
- RunPE for payload injection;
- Reverse SOCKS5 proxies;
- A reverse shell over HTTP;
- A file manager for uploading files to and downloading files from the victim's device;
- A keylogger;
- Screenshot capture on the victim's device.




Fig. 124 — OnyxC2 management web panel

The OnyxC2 subscription is divided into three tariff plans:

- The “Normal” plan (\$250) includes a private VPS and a malicious landing page but no HVNC;
- The “Premium” (\$500) plan includes the “Normal” plan’s capabilities plus HVNC functionality;
- The “Private” plan is a one-time \$6,000 purchase that provides the full OnyxC2 source code and assistance with deploying the malware on the buyer’s infrastructure.

Snowdog

Among the most expensive and specialized offerings are private tools for targeted attacks against organizations. On February 4, 2025, the user under the alias SeverusSnape posted a listing on a well-known underground forum offering Snowdog under the MaaS model for \$7,000 per month. The turnkey package included access to a web panel for generating builds and a set of delivery methods. On March 30, the same author opened a separate thread offering a free seven-day trial in exchange for a detailed review from members of the underground marketplace.




Fig. 125 — Screenshot of the announcement offering the private Snowdog malware for sale

Translated text from Fig. 125

Private software for targeted attacks against corporations.

The package includes a delivery method, persistence, UAC bypass, the trojan itself, and a web console.

A set of additional tools.

Written in C.

Monthly rental; the price includes regular cleanups and continuous updates.

Two-day free trial.

If interested, message me on Tox

Snowdog includes two separate modules: a RAT module with a loader and the Venon ransomware module. The management panel enables the real-time creation of unique encrypted builds, control of infected systems through PowerShell, and the downloading of files and directories from remote machines. The web panel can generate builds in HTML-to-DOCX format for the RAT and ISO format for delivery of the Venon ransomware. According to the threat actor, the malware is written in C, runs in the memory of trusted processes, uses the Tor network to communicate with the C2 server, and automatically prevents execution in certain countries to reduce the risk of prosecution. The author subsequently optimized the delivery and on-device stealth modules, increasing the malware's threat level.



Fig. 126 — Screenshot of the announcement offering a Snowdog trial period

Translated text from Fig. 126

We are offering a free trial of new software. The package includes:

- A RAT with a web console
- A web console for creating both the RAT itself and delivery methods (HTML-to-DOCX and ISO)
- UAC bypass
- Persistence on the system
- A loader that runs DLL and EXE files in the memory of its own process
- One-click downloading of files and directories from the compromised machine
- An interface for creating various delivery configurations with different payloads
- Additional tools

Capabilities:

- On-the-fly creation of a unique build in the web console
- Encryption of each build at creation time using a proprietary algorithm
- Two bot-delivery methods, selected at build creation
- Web-console access to the remote machine with administrator privileges, including execution of all PowerShell commands
- Downloading of specified files and directories

Features:

- Runs in the memory of a trusted process
- Transport-layer communication
- Bot management via Tor

What we offer: seven days of free access and 24/7 support. What we ask in return: feedback on the test results and an objective review on this site.

To obtain access, contact us via private message or Tox.

- Automatic launch at system startup;
- File-encryption capabilities for extortion;
- Disabling Microsoft Defender and UAC;
- Hidden VNC access (HVNC);
- Audio and video capture;
- Control via PowerShell scripts and C# modules;
- Manipulation of network settings and the hosts.txt file.

Snowdog was offered under a single plan costing \$7,000 per month.

QatarRat

Another MaaS offering, QatarRat, was advertised by the user under the alias QatarRAT (aka QRhades) on April 8, 2025, in the threat actor's Telegram channel and on several underground forums.




Fig. 128 — Demonstration of the QatarRat web panel




Fig. 127 — Screenshot of the announcement offering QatarRat for sale on an underground forum

QatarRat was available under the following subscription plans:

- \$75 per month;
- \$125 for three months;
- \$300 for six months;
- \$600 per year.

QatarRat provides the following capabilities:

- Keylogger;
- Theft of passwords, cookies, and cryptocurrency wallet data;
- Remote access via RDP/VNC;
- Downloading and running third-party malware;

macOS

Some RATs specifically target macOS users. Although macOS has a smaller overall market share than Windows, it is widely used by executives, developers, and creative professionals, whose devices may contain valuable corporate data and intellectual property. These users therefore present high-value targets, driving growing threat-actor interest in malware that provides covert remote control of macOS systems.

notnullOSx

On March 3, 2026, F6 TI analysts discovered an announcement on a well-known underground forum in which the user notnull offered notnullOSx, specialized macOS malware combining stealer and RAT functionality.




Fig. 129 — Screenshot of the announcement offering notnullOSx for sale on an underground forum

Translated text from Fig. 129

Hello. I would like to introduce a macOS RAT/stealer called "notnullOSx."

Preface: In the interest of fair competition, I will be candid: at present, 95% of products in this field amount to three days of AI-assisted work, built primarily with AppleScripts that break after every macOS update and paired with a hastily assembled web panel. I saw this problem and, drawing on my extensive macOS experience, decided to offer a new product.

The project's main components:

- **Server:** HMAC-SHA256 encryption; a modular system for running tasks on clients, loading only the required modules and cleaning them up after execution without unnecessary overhead; saving results and generating brute.txt; Telegram notifications; autorun configurations that automatically launch modules in any desired order when a client connects; and a convenient, lag-free reverse shell with full terminal functionality, including Tab and Ctrl+C.
- **Web panel:** Dashboard-statistics and online clients; Clients-a client list and detailed pages with results; Results-all results from all clients; Modules-module information and descriptions; Replace App-configuration of application substitution on the target machine, such as Ledger (included in the price); Autorun-module autostart configuration; Builder-installation command, redirectors, domains, and build options; Settings-Telegram integration and convenient log delivery to any Telegram chat; Preview-viewing passwords, cookies, JSON, and CSV directly in the panel, with the option to download a client's entire archive; Localization-Russian and English.


The architecture is modular and consists of a server component, a management web panel, and a set of modules. The server component provides HMAC-SHA256 encryption; dynamic module loading and automatic cleanup after task completion; generation of reports and a brute.txt file, presumably containing a password dictionary for subsequent brute-force attacks; sending notifications to a Telegram bot; and support for a fully interactive reverse shell. Initial infection uses the ClickFix technique, which relies on social engineering to coerce victims into entering a malicious command in the terminal.

The module set includes SystemInfo, BrowserGrab, ChromeExtensionsGrab, SafariCookiesGrab, ChromePasswordsGrab, FirefoxGrab, YandexGrab, BrowserHistoryGrab, CredsGrab, and TelegramGrab, as well as Apple-specific components for extracting data from iMessage and Apple Notes. A separate ReplaceApp module replaces legitimate applications with malicious counterparts.

The notnullOSx subscription plans are:

- \$1,499 per week;
- \$3,999 per mont.h

Shadow B2B: Anatomy of the Initial Access Broker Market and 10 Key Brokers, 2025 – H1 2026



In today’s cybercrime market, a clear, highly specialized supply chain has emerged. At the core of this supply chain is the initial access broker (IAB) market, which supplies ransomware operators and other cybercriminal groups with access to compromised infrastructures worldwide.

Selling initial access has evolved from an opportunistic niche into a mature, streamlined business. Each listing is priced based on specific attributes:

- Country;
- Industry;
- Revenue;
- Access type.

For buyers, these attributes are critical because they directly affect the potential profit from utilizing the access.

Access to the corporate networks of large companies with domain administrator privileges is in the greatest demand and commands the highest prices. With such privileges in a corporate network, a threat actor gains total control over the infrastructure, enabling everything from stealthy persistence via Group Policy Object (GPO) modification and mass data exfiltration to the simultaneous deployment of ransomware across all servers and workstations. The inevitable consequences are complete paralysis of business operations, devastating reputational losses, and multimillion-dollar regulatory fines.

Threat actors rarely disclose the specific names of victim organizations. Direct brand exposure may prompt security researchers to alert it, leading to the access being closed. To mitigate this risk, Initial Access Brokers leverage legitimate corporate business intelligence tools to anonymize the compromised entity’s profile while retaining key commercial metrics.

During the reporting period, thousands of listings of initial access to corporate networks worldwide were advertised on underground forums. Based on sales announcements, brokers sold approximately \$300,000 worth of access, at an average price of \$3,000 per listing. Had all advertised access been sold, threat actors could have earned more than \$10 million.


The most common initial access vectors are publicly exposed RDP servers and corporate VPN gateways that are not protected by multi-factor authentication (MFA). In H1 2026, VPN was the stated access type in 29% of listings, RDP in 14%, and web panel access in 11%.

Access type	%
VPN	29.08
RDP	13.99
Web panel	11.07
Database	5.35
Backdoor	5.23
Other	14.84
Not Stated	20.44

Despite the anonymized and templated nature of these listings, behind them are specific threat actors with distinctive tradecraft and reputations. In this section, F6 Threat intelligence analysts have compiled comprehensive profiles on the 10 most notable Initial Access Brokers (IABs) of the initial access market who were active from H1 2025 through H1 2026.

The profiles below are based on key behavioral characteristics identified during the period of each broker's activity. The statistics include only unique listings offering access to companies with annual revenue above \$1 million. Low-priced bulk access offerings (for example,

100 RDWeb access for \$1,000) were excluded from the analysis. Such listings are of negligible value and generally hold little practical interest for the target audience of this market. For each IAB, this section presents the key tactics, techniques, and procedures (TTPs) identified among the detected actors' activity. These observations are based solely on the threat actors' messages published on underground forums. Each matrix includes only a subset of TTPs and does not represent the full range of an actor's capabilities; it is limited to methods observed while assessing that specific broker. Actions taken by buyers during later stages of an attack are outside the scope of this research.



Big-Bro

Overview	
First seen	25.12.2022
Victims, 2025 – H1 2026	238
Target geography	 Americas,  Africa,  APAC,  Europe,  Middle East
Target countries	 US,  UK,  AU,  BR,  CA,  FR,  DE,  IT,  ID,  ES,  AR,  AE,  MX,  BE,  BA,  CZ,  MY,  NZ,  TW,  CN,  DK,  DO,  GR,  IN,  KW,  NL,  PT,  SA,  SK,  TR,  AO,  CL,  EG,  MA,  PE,  SG,  CH,  TH,  VE

The IAB known as Big-Bro first appeared on underground forums in December 2022 and posted a record number of initial access listings in 2026. All of the actor’s listings follow a consistent structure:

- Country;
- Revenue;
- Industry;
- User privileges;
- Access type;
- Security solutions present in the infrastructure (if any).

Prices for the actor’s individual access listings are generally negotiable and are discussed with prospective buyers in private messages. In some cases, the actor offered access at auction, with starting prices ranging from \$700 to several thousand dollars.

Big-Bro frequently bundles several large lots in a single listing. For example, in a January 5, 2026 post on an underground forum, the IAB advertised access to companies located in Indonesia, Malaysia, and Egypt. The Indonesian transportation company, with stated revenue of \$3 billion, was the highest-revenue victim among the actor’s 2026 listings:




Fig. 130 — Screenshot of Big-Bro’s January 5, 2026 listing offering initial access for sale

Forum activity suggests that Big-Bro may have resold access, which could explain the large number of listings they published. For example, on January 30, 2026, the actor listed access to an Italian company’s corporate network that appeared to have been purchased from another IAB:




Fig. 131 — Screenshot of Big-Bro’s January 30, 2026 initial access listing

Translated text from Fig. 131

domain user with local admin privileges on the Exchange server

- Manufacturing
- watchguard vpn
- av sophos
- Start: 700\$
- Step: 250\$
- Blitz: 1500\$

A listing with identical content at a lower price had been posted four days earlier, on January 26, 2026, by another IAB known as personX:




Fig. 132 — Screenshot of personX’s January 26, 2026 initial access listing

Users on underground resources — including an administrator of a well-known underground forum — have also accused the actor of reselling access and selling the same lot to multiple buyers:




Fig. 133 — Screenshot of underground forum users discussing the IAB’s activities

Translated text from Fig. 133

DevManager:

delete a comment, get a ban wouldn't recommend dealing with him - the guy sells networks with 'neighbors' and constantly deletes negative feedback

hack.sign:

I'll agree with you here - a hardcore reseller, sells to a hundred hands at once, and doesn't care about anything or anyone...

He doesn't source any accesses himself

It should be noted that in one of their posts Big-Bro claimed that some access had been obtained through brute-force attacks:




Fig. 134 — Screenshot of a deleted post by the threat actor Big-Bro, available on the F6 Threat Intelligence portal

Initial Access Sales Statistics — Big-Bro

The statistics below summarize Big-Bro's access sales for the period 2025 – H1 2026.


Access sales by country

Country			Count
	US		64
	UK		16
	AU		12
	BR		10
	CA		9
	FR		8
	DE		7
	IT		7
	ID		6
	ES		6
	AR		5
	AE		5
	MX		4
	BE		3
	BA		3
	CZ		3
	MY		3
	NZ		3
	TW		3
	AT		2

Country			Count
 CN			2
 DK			2
 DO			2
 GR			2
 IN			2
 KW			2
 NL			2
 PT			2
 SA			2
 SK			2
 TR			2
 AO			1
 CL			1
 EG			1
 MA			1
 PE			1
 SG			1
 CH			1
 TH			1
 VE			1
 Not Stated			28

Access sales by industry

Industry		Count
 Manufacturing		50
 Commerce and Shopping		37
 Professional Services		15
 Healthcare		12
 Real Estate		10
 Finance		9
 Media and entertainment		8
 Software		8
 Government and military		7
 Food and beverage		6
 Energy		5
 IT		5
 Transportation and logistics		5
 Travel and tourism		5
 Consumer electronics		4
 Agriculture and Farming		3
 Consumer Goods		2
 Education		2
 Messaging and telecommunications		2



From 2025 through H1 2026, Big-Bro’s victims included companies across the Americas, Africa, APAC, Europe, and the Middle East. The United States, the United Kingdom, and Australia were the most frequently targeted

countries, while Manufacturing was the most frequently represented industry. The actor most commonly offered access via RDP and VPN, with Fortinet and SonicWall being the predominant VPN solutions.

ATT&CK Matrix for Enterprise (Big-Bro)

Tactics	Techniques	Details
Initial Access	External Remote Services T1133	In their initial access listings, Big-Bro stated that access was provided via RDP and VPN services
	Valid Accounts: Domain Accounts T1078.002	Big-Bro stated that they had obtained access to domain users’ and domain administrators’ accounts
Credential Access	Brute Force T1110	In posts on underground marketplaces, Big-Bro claimed that some of the access offerings had been obtained through brute-force attacks
Discovery	Software Discovery: Security Software Discovery T1518.001	In their initial access listings, Big-Bro identified the security solutions present on compromised devices

miyako

Overview	
First seen	15.08.2024
Victims, 2025 – H1 2026	224
Target geography	 Americas,  Africa,  APAC,  Europe,  Middle East
Target countries	 US,  CN,  ES,  UK,  CA,  AE,  SA,  DE,  TW,  PS,  TH,  IR,  UA,  IT,  BA,  ID,  MX,  CO,  IN,  FR,  JP,  BR,  HK,  MZ,  GT,  CZ,  LK,  GR,  VE,  PL,  HN,  SV,  PE,  ZA,  AR,  IQ

The threat actor known as miyako registered on a popular underground forum on August 15, 2024. Their first observed listing of 2025, posted on January 3, offered root access to web servers hosting the firewalls of a Central American logistics company. The asking price was \$200:




Fig. 135 — Screenshot of one of miyako’s earliest access listings, dated January 3, 2025

miyako are distinguished by their narrow focus on compromising firewalls. The vast majority of their listings advertised one of several

forms of privileged access to FortiOS- and Linux-based firewall devices, featuring remote code execution with root privileges (Root RCE), shell access, or access to an administrative panel. During the early phase (January-February 2025), offerings were more diverse and included the bulk sale of VPN credentials as well as SSH, DNS, mail, and SQL access. By mid-2025, however, the actor had focused almost exclusively on root access to firewalls, which remained their primary offering through Q2 2026.

The highest-revenue target was a health-care company based in the United States with annual revenue of \$29.5 billion. The listing was posted on an underground forum on February 28, 2025. The actor claimed to have gained access to the corporate infrastructure by exploiting an RCE vulnerability and obtaining root privileges on a device. The asking price was \$500. The listing also illustrates a recurring pattern in miyako’s posts: the actor often sells access to the networks of large companies at relatively low prices:




Fig. 136 — Screenshot of miyako’s access listing, dated February 28, 2025

Four months earlier, on November 2, 2024, the actor commented on another user’s post about an exploit for CVE-2024-23113, an FGFM protocol vulnerability that enables arbitrary code execution in several Fortinet products. In a forum comment, miyako claimed to have developed an exploit that enables remote

code execution on a target FortiGate device. The actor may have used this exploit to obtain initial access to victim infrastructure for resale.



Fig. 137 — Screenshot of miyako’s discussion of CVE-2024-23113 on an underground forum

Initial Access Sales Statistics — miyako

The statistics below summarize miyako’s access sales for the period 2025 – H1 2026.

Access sales by country

Country	Count
US	89
CN	14
ES	13
UK	10
CA	10
AE	7
SA	7
DE	5
TW	3
PS	3

Country			Count
 TH			3
 IR			2
 UA			2
 IT			2
 BA			2
 ID			2
 MX			2
 CO			2
 IN			1
 FR			1
 JP			1
 BR			1
 HK			1
 MZ			1
 GT			1
 CZ			1
 LK			1
 GR			1
 VE			1
 PL			1
 HN			1
 SV			1
 PE			1
 ZA			1
 AR			1
 IQ			1
 Not Stated			28

Access sales by industry


Industry		Count
 Government and military		62
 Manufacturing		29
 Finance		29
 Healthcare		14
 Software		12
 Transportation and logistics		10
 Real Estate		10
 Messaging and telecommunications		9
 Education		8
 IT		7
 Energy		5
 Professional Services		4
 Commerce and shopping		3
 Travel and tourism		3
 Media and entertainment		3
 Real Estate		2
 Not Stated		14

The statistics above shows that miyako’s target profile was geographically broad but heavily skewed toward the United States. Government and Military, Manufacturing, and Financial Services were the leading industries. The actor consistently focused on critical infrastructure and sensitive targets, including the United States. defense

and aerospace contractors, state and municipal governments, telecommunications data centers, and regions of heightened geopolitical tension. By selecting such high-profile targets, the actor maximizes their publicity, underscoring the IAB’s dual motivation: financial gain and public attention.

ATT&CK Matrix for Enterprise (miyako)

Tactics	Techniques	Details
Resource Development	Develop Capabilities: Exploits T1587.004	miyako claimed to have developed an exploit for a vulnerability in the FGFM protocol (CVE-2024-23113)
Initial Access	External Remote Services T1133	In their listings, miyako specified VPN and RDP as initial access vectors
	Exploit Public-Facing Application T1190	In their messages on underground forums, miyako claimed to have exploited CVE-2024-23113, a vulnerability in the FGFM protocol, to achieve remote code execution



Anon-WMG

Overview	
First seen	21.09.2024
Victims, 2025 – H1 2026	64
Target geography	 Americas,  Africa,  APAC,  Europe,  Middle East
Target countries	 US,  BR,  CA,  FR,  ID,  IT,  AE,  IN,  NL,  PH,  QA,  ES,  CH,  UY,  AR,  AU,  EG,  ET,  FI,  JP,  MY,  MX,  MA,  SG,  TH,  TR,  UK,  UG,  VN

The initial access broker later known as Anon-WMG first appeared on September 21, 2024. The actor registered an account under the alias R-W-M-GROUP on a major underground forum.

Anon-WMG’s first victim was an Egyptian telecommunications provider. In their first listing, published on October 1, 2024, the actor named the compromised company’s domain and claimed access to its network infrastructure, employee and customer data, and other business-critical information. Anon-WMG also claimed they could manage customer services: according to the actor, a buyer could activate, modify, and deactivate service packages provided to the company’s clients. The actor stated that they obtained the access in June 2024. As proof, they attached links to .csv, .js, and .json files, claiming they belonged to the compromised company and contained a database dump sample, source scripts, and infrastructure configuration files. The seller set a price range of \$4,000 – \$7,000 per listing and stated the access

would be sold to no more than two buyers. The full text of the announcement is shown below:




Fig. 138 — Screenshot of Anon-WMG’s October 1, 2024 listing offering initial access to a corporate network

On April 12, 2025, monitoring of Anon-WMG activity revealed a listing in which the actor offered an exploit targeting vulnerable FortiGate firewalls. The actor claimed that the discovered vulnerability allows an unauthenticated attacker to execute remote code execu-

tion (RCE) and obtain infrastructure configuration dumps. According to the post, successful exploitation granted the actor access to local user's credentials and password hashes, system administrator account data, firewall and NAT rules, SSL/IPsec VPN configurations, private certificates, SNMP community strings, and full CLI configuration backups. The asking price was \$6,500. The actor may have used the exploit to obtain initial access to the promoted corporate networks: most of their listings consisted of access to FortiGate VPN gateways.

Although Anon-WMG emphasized remote code execution (RCE) in the listing, technical analysis of the collected data structure indicates that the vulnerability was in fact an authentication bypass in the REST API, which could be subsequently leveraged to execute arbitrary code. The actor appears to have used this tooling to automate the extraction of critical infrastructure configuration dumps.




Fig. 139 — Screenshot of a fragment of Anon-WMG's April 12, 2025 listing offering a zero-day exploit

The full text of the listing is provided below (the author's spelling and punctuation are preserved):

Unauthenticated RCE / Config Access for FortiOS
Config Dump
Gain FULL control of vulnerable FortiGate firewalls without creds.
Sorry for long article, But it Worth it
Features:

Forti Exploit Dump Contents

Authentication & Access

local_users.json → All local VPN/GUI users (usernames + ENCRYPTED passwords)

admin_accounts.json → System admin accounts (permissions, access, trusthost)

two_factor.json → FortiToken / 2FA status

Firewall & NAT

firewall_policy.json → Complete rulebase (source, destination, service, action)

firewall_vip.json → Virtual IP mappings (NAT/pivot points)

firewall_nat.json → Source NAT rules

firewall_addresses.json → All internal IP assets

firewall_addrgrps.json → Address groups

VPN & Remote Access

vpn_ssl_settings.json → SSL VPN interface settings

vpn_ssl_portals.json → Portal configurations (e.g. full-access / tunnel-only)

vpn_ipsec_phase1.json → IPSec tunnels

vpn_bookmarks.json → Web VPN bookmarks (destination IPs, internal web apps)

vpn_sessions.json → Active user sessions

System Configuration

config_backup.conf → Core text backup of system config (entire CLI format)

backup_full → Binary backup (if accessible)

system_global.json → Hostname, admin ports, banner message

device_status.json → Serial, version, uptime

dhcp_servers.json → Internal DHCP pools

Network Monitoring

routes.json → All static and dynamic routing

arp_table.json → Internal MAC-IP mapping

snmp.json → SNMP community strings (rarely present, but critical if found)

active_sessions.json → Current traffic sessions

event_log.json, system_log.json, webfilter_log.json → System and security logs

Certificates & Trust

certificates.json, cert_ca.json, cert_remote.json → VPN and internal SSL certs

May contain.pem or.key

Automation & Scripts

system_scripts.json → CLI scripts stored by admins (often with sensitive logic)

Also script Could Dump:

Admin password policies

FortiGuard license keys

Web Filter profiles (URL categories)

PRICE: 6500\$




Fig. 140 — Screenshot of Anon-WMG’s June 8, 2025 listing offering access to the Dutch company

The following were the most notable access listings posted by Anon-WMG during the reporting period:

1. **Initial access to the corporate network of a company in the Netherlands with a stated revenue of more than \$195 billion:**

On June 8, 2025, Anon-WMG posted an announcement on a major underground forum claiming administrator-level access to the corporate network of the largest automotive manufacturer in the Netherlands, which controls several world-renowned car brands. The listing claimed 258,000 employees and \$195 billion in revenue. The access type was GlobalProtect VPN, and the asking price was \$6,500. The relatively low price reflected the actor’s inability to bypass certificate validation, which prevented full network access even with administrator credentials.

The full text of the listing is provided below (the author’s spelling and punctuation are preserved):

Hi, before you send me msg (the access im selling this time need big and professional group or person to work on it , how???:

Country: Netherlands, Amsterdam

Employees: 258,275

Type of access: global protect Vpn network

Privileges: full administrator

Revenue:+\$195 Billion (this is real)

Category:

Manufacturing

Motor Vehicles

Automotive Parts

controls 14 major car brands

Why i said need professionals:

You can login from web portal of global protect that let download global protect vpnusing creds (user,pass)

The problem is when you try to access to network inside or connect to it from inside vpn

It ask for certificate (required)

If can't bypass this, you can't access network (even using admin creds)

Price: 6500\$

If you can bypass the certificate, you will have access to all the corp and all hosts, machines, websites.....

2. Initial access to the corporate network of a company in Brazil with a stated revenue of more than \$85 billion:

On November 3, 2025, Anon-WMG advertised access to a large Brazilian food manufacturer and distributor with stated revenue exceeding \$85 billion. According to the listing, the actor had obtained administrator-level access to the company’s server via FTP and managed to access more than 8,000 files. The asking price was \$19,000.



Anon-WMG

Posted November 3, 2025

Country: brazil

Employees: +200,000

Revenue: +85 billion\$

access Type: ftp main server

privileges: Administrator

Files: +8190

Directories: 1590

Industry:
Food & Beverage
Retail
Manufacturing
Grocery Retail

Price: 19k

Fig. 141 — Screenshot of Anon-WMG’s November 3, 2025 listing offering access to the Brazilian company

Initial Access Sales Statistics — Anon-WMG

The statistics below summarize Anon-WMG’s access sales for 2025 – H1 2026.

Access sales by country

Country			Count
 US			13
 BR			4
 CA			3
 FR			3
 ID			3
 IT			3
 AE			3
 IN			2

Country			Count
 NL			2
 PH			2
 QA			2
 ES			2
 CH			2
 UY			2
 AR			1
 AU			1
 EG			1
 ET			1
 FI			1
 JP			1
 MY			1
 MX			1
 MA			1
 SG			1
 TH			1
 TR			1
 UK			1
 UG			1
 VN			1
 Not Stated			3

Access sales by industry

Industry	Count
 Manufacturing 	9
 Commerce and shopping 	5
 Messaging and telecommunications 	5
 Food and beverage 	3
 Government and military 	3
 IT 	2
 Professional Services 	2
 Real Estate 	2
 Agriculture and Farming 	1
 Education 	1
 Healthcare 	1
 Media and entertainment 	1
 Sales and marketing 	1
 Transportation and logistics 	1
 Travel and tourism 	1
... Other 	2
 Not Stated 	24

From 2025 through H1 2026, Anon-WMG most frequently targeted the United States and Brazil, although their listings covered nearly every region except the CIS. The most

frequently targeted industries were Manufacturing, Commerce and Shopping, and Messaging and Telecommunications.

ATT&CK Matrix for Enterprise (Anon-WMG)

Tactics	Techniques	Details
Initial Access	Develop Capabilities: Malware T1587.001	Anon-WMG likely participated in the development and operation of an API dumper capable of exploiting an authentication-bypass vulnerability chain in the FortiOS REST API on perimeter FortiGate devices
	Exploit Public-Facing Application T1190	Anon-WMG used a custom automation tool (an API dumper) to exploit an authentication-bypass vulnerability chain in the FortiOS REST API on FortiGate devices at the network perimeter
	Valid Accounts T1078	Anon-WMG provided initial access by supplying compromised VPN gateway credentials in user:pass format
Credential Access	Exploitation for Credential Access T1212	Anon-WMG leveraged a custom script to exploit an authentication bypass in the FortiOS REST API and extract credentials from FortiGate devices.
Discovery	System Network Configuration Discovery T1016	Anon-WMG stated that they had mapped the victim's internal network configuration and exported the collected data to routes.json, firewall_addresses.json, and arp_table.json

SGL

Overview	
First seen	2007
Victims, 2025 – H1 2026	54
Target geography	Africa, Americas, APAC, Europe, Middle East
Target countries	US, ES, MX, BR, FR, DE, BE, IN, IT, TR, CN, ID, IR, JP, KE, MA, NL, SA, KR, UK

The threat actor operating under the alias SGL (also known as sganarelle) has been active on various underground forums since 2007, making the actor one of the longest-standing active participants in these platforms. SGL claimed to be personally acquainted with several forum administrators.

SGL initially focused on carding activity rather than access brokerage. In July 2022, however, the actor began actively posting listings offering access to corporate networks. At first, sganarelle sold access at auction, but later switched to fixed-price listings. Even in these early offerings, the actor noted that the obtained stealer logs contained credentials for Citrix, RDP, and VPN access:




Fig. 142 — Screenshot of SGL’s July 27, 2022 initial access listing

Thus, the actor leveraged stealer logs to obtain valid credentials gaining initial access to target infrastructures. These logs may have been either acquired directly from underground marketplaces or harvested independently through the threat actor’s own deployment of malicious software.

Unlike most access brokers, SGL discloses neither the victim’s name nor its industry in their listings.

Among SGL’s targets in 2025, the highest-revenue compromised organization was an unnamed Spanish company with annual revenue exceeding \$3.5 billion. The actor offered VPN access to the company in an announcement posted on February 26, 2025:




Fig. 143 — Screenshot of SGL’s February 26, 2025 initial access listing

The access was purchased within the same day by an unknown buyer.

As of H1 2026, SGL’s most recent access listing was dated November 26, 2025. Nevertheless, the actor remains active on underground forums under the alias **sganarelle2**:

Notably, the actor has taken lengthy breaks from activity before and may soon resume selling initial access to corporate networks.




Fig. 144 — Screenshot of the sganarelle2 profile on one underground forum

Initial Access Sales Statistics — SGL

The statistics below summarize SGL’s access sales for the period 2025 – H1 2026.

Access sales by country

Country		Count
 US		11
 ES		7
 MX		6
 BR		5
 FR		4
 DE		3

Country			Count
 BE			2
 IN			2
 IT			2
 TR			2
 CN			1
 ID			1
 IR			1
 JP			1
 KE			1
 MA			1
 NL			1
 SA			1
 KR			1
 UK			1

Access sales by industry

Industry			Count
 Not Stated			54

The United States, Spain, and Mexico were the most frequently targeted countries. No industry preferences could be determined since

SGL does not disclose victims' industries in their listings. The actor predominantly offered access via VPN, RDP, and Citrix.

ATT&CK Matrix for Enterprise (SGL)

Tactics	Techniques	Details
Initial Access	External Remote Services T1133	In their initial access listings, SGL listed Citrix, VPN, and RDP as initial access vectors
	Valid Accounts T1078	SGL claimed to have gained initial access using credentials from stealer logs

personX

Overview	
First seen	15.07.2024
Victims, 2025 – H1 2026	50
Target geography	Europe, Middle East, APAC, Americas, Africa
Target countries	US, PT, SK, UK, AR, DE, ID, SG, ZA, SD, TR, AU, AT, BA, BR, BG, CN, DK, GH, IN, IT, JO, KW, MX, NZ, PK, RO, SA, ES, CH, TH, AE, VN, ZW

The user known as personX registered on an underground forum on July 15, 2024. Shortly afterward, on July 22, 2024, the actor posted their first listing, offering initial access with domain administrator privileges to an Italian company with an estimated revenue of \$17 million:

Translated text from Fig. 145

Italy. Access to one company with Domain Admin privileges. [B2B platform name] reports \$17 million in revenue. Price: \$1,000

Италия, администратор домена

By personX, July 22, 2024 in [Access] - FTP, shells, root, sql-inj, DB, Servers

Reply to this topic

personX

Posted July 22, 2024

Имеется доступ к одной компании. Доменный администратор.

В [REDACTED] указано 17 млн\$

цена 1000\$

personX’s highest-revenue victim was an unnamed international pharmaceutical company registered in Pakistan. According to the actor, the victim’s annual revenue was \$518 million. The announcement, posted on February 4, 2025, is shown below:

Fig. 145 — Screenshot of personX’s first access listing, dated July 22, 2024



Fig. 146 — Screenshot of personX’s access listing, dated February 4, 2025

Two days later, on February 6, 2025, the seller stated that the access had been purchased by an unknown buyer.

In addition to posting individual listings, personX sold low-priced bulk lots containing several hundred access credentials. An example of such announcement is shown below:




Fig. 147 — Screenshot of personX’s December 30, 2025 listing offering a bulk lot of access credentials

Translated text from Fig. 147

~1,000 valid Fortinet VPN credentials in ip:port user:pass format. Price: \$1,000. Regions: United States, Canada, Australia, Europe, and the Gulf states. Actually, take them for \$500 - today only

Based on the actor’s forum posts, personX may have used Shodan to automate network scanning; the actor repeatedly offered Shodan access for sale on underground forums. This hypothesis is supported by the fact that personX’s activity spanned roughly two years (July 2024 – June 2026), matching the stated period of the Shodan account usage:




Fig. 148 — Screenshot of personX’s June 5, 2026 listing offering a Shodan account for sale

Translated text from Fig. 148

Selling a personal Shodan account. It has been active for two years and includes API keys with no monthly fees. All filters, the API, and every feature available on the site work. Credits refresh monthly, up to 200,000. The account is permanent as long as you do not share it widely. After purchase, you can change the password in the account settings. Price: \$1,000

In an August 17, 2024 access listing, the actor claimed to have obtained Domain Cached Credentials v2 (DCC2) password hashes for domain users. personX also stated that they had gained local administrator privileges

on a victim server and dumped domain credentials. This indicates that the attack had progressed beyond initial access into persistence, privilege escalation, and lateral movement. Local administrator access to four servers, including a Microsoft Exchange server, could have given an attacker extensive control over the corporate infrastructure and posed a direct threat of full Active Directory domain compromise. The claimed dumping of DCC2 hashes may indicate that the threat actor could have conducted offline password-cracking attacks to identify highly privileged accounts. The forum listing is shown below:




Fig. 149 — Screenshot of personX’s access listing, dated August 17, 2024

Translated text from Fig. 149

Australia. Local administrator access. VPN and domain user access, plus local administrator access to four servers, including an Exchange server. DCC2 hashes available. Approximately \$6 million in revenue; manufacturing. Price: \$1,000

In a September 30, 2024 listing, personX offered all previously obtained corporate network access, the Shodan account, and an exploit for a server-side request forgery (SSRF) vulnerability in Microsoft Exchange. The actor may

have used this exploit to gain the initial access to the Exchange server described in the August 17 listing before.




Fig. 150 — Screenshot of personX’s access listing, dated September 30, 2024

Translated text from Fig. 150

Access + works in progress + exploit + unlimited Shodan

Selling all the access I have listed on the forum (except for a couple that have already been sold), ongoing work on additional access, the Exchange SSRF exploit, and unlimited Shodan with mailbox access.

Everything for \$3,500

Open to offers

Escrow accepted

On April 30, 2026, personX responded to a question from the user raimondo about evading detection of a malicious executable by security solutions. The actor recommended using the SigThief utility to copy an existing digital signature from a legitimate file to the malicious payload, then hosting the payload on pre-built infrastructure with a trusted SSL certificate to evade defenses during delivery. personX said

they had used the same method, likely to achieve initial infection:




Fig. 151 — Screenshot of the April 30, 2026 discussion on bypassing security detection of a malicious download

Translation from Fig. 151

raimondo: Hi all - how can I bypass the alert when downloading an.exe file? Happy to tip anyone who can help.

personX: Hey. You can use sigthief.py to copy signatures from Microsoft applications (for example, Calculator) and sign your.exe. Hosting the.exe on a site with a trusted SSL certificate also works well. I used to buy cheap websites and upload it there; downloads from those sites were not flagged, especially when combined with sigthief.

Notably, personX is likely linked to the IAB Big-Bro described above: these two actors posted listings with the identical content. Big-Bro later offered the same access at a higher price, which may indicate that Big-Brow was reselling access purchased from personX. In total, Big-Bro likely purchased at least five access lots from personX.

Initial Access Sales Statistics — personX

The statistics below summarize personX's access sales for the period 2025 – H1 2026.

Access sales by country

Country			Count
 US			4
 PT			3
 SK			3
 UK			3
 AR			2
 DE			2

Country		Count
 ID		2
 SG		2
 ZA		2
 SD		2
 TR		2
 AU		1
 AT		1
 BA		1
 BR		1
 BG		1
 CN		1
 DK		1
 GH		1
 IN		1
 IT		1
 JO		1
 KW		1
 MX		1
 NZ		1
 PK		1
 RO		1
 SA		1
 ES		1
 CH		1
 TH		1
 AE		1
 VN		1
 ZW		1

Access sales by industry

Industry	Count
 Manufacturing 	4
 Energy 	2
 Healthcare 	2
 Real Estate 	2
 Science and Engineering 	2
 IT 	2
 Agriculture and Farming 	1
 Data and Analytics 	1
 Finance 	1
 Food and beverage 	1
 Professional Services 	1
 Transportation and logistics 	1
 Other 	2
 Not Stated 	28

Geographic analysis shows that personX had no clear regional preference and targeted companies worldwide. The actor sold access to organizations across Europe, the Middle East,

APAC, the Americas, and Africa. Identified victims operated in Manufacturing, Energy, Health Care, Real Estate, and Science and Engineering, although most listings did not specify any industry.

ATT&CK Matrix for Enterprise (personX)

Tactics	Techniques	Details
Reconnaissance	Search Open Technical Databases: Scan Databases T1596.005	personX stated on underground forums that they had repeatedly used Shodan
Resource Development	Acquire Infrastructure: Domains T1583.001	personX claimed to have purchased ready-made websites with trusted SSL certificates to host malicious payloads
Initial Access	External Remote Services T1133	personX listed Citrix, VPN, and RDP in their initial access listings
	Valid Accounts T1078	personX sold compromised credentials providing access to corporate resources
	Exploit Public-Facing Application T1190	personX claimed to have exploited an SSRF vulnerability to access a Microsoft Exchange server
Stealth	Masquerading: Invalid Code Signature T1036.001	personX claimed to have used sigthief.py to copy the digital signature from calc.exe to a malicious executable
Credential Access	Brute Force T1110	personX used brute-force attacks to guess passwords for Fortinet VPN gateways
	OS Credential Dumping: Cached Domain Credentials T1003.005	personX claimed to have extracted DCC2 hashes from compromised Exchange servers

Saturned33

Overview	
First seen	04.08.2025
Victims, 2025 – H1 2026	50
Target geography	Europe, Americas, APAC, Middle East, Africa
Target countries	US, VE, CH, IR, TW, JP, DE, IT, TR, ZA, CA, PY, LK, GT, FR, KE, LB, IQ, AU, CN, BD, BH, TH, PT, IN, ES, AR, CZ, EG, SA, UG, MX

The user known as Saturned33 registered on an underground forum on August 4, 2025. The following day, the actor posted their first listing, offering initial access to several companies in Taiwan, the United States, Paraguay, Sri Lanka, and Venezuela, with domain administrator or local administrator privileges:




Fig. 152 — Screenshot of Saturned33’s first access listing, dated August 5, 2025

The full text of the listing is provided below (the author’s spelling and punctuation are preserved):

All transactions with autogarant

1. Country> TAIWAN [hosted in china. Operates in TAIWAN. Business in TAIWAN. Language of AD in Chinese]

Rev> +18kk

Sector> Electronics

Type> RDP

Privileges> Domain Admin

Hosts> 80

AV> HUORONG (can be deactivated manually as LA/DA)

Tags> They have public download page for clients about software, drivers, firmware... (can be used to infect binding your file too)

Price> 1300\$

2. Country> US [hosted in China. Operates in US. Business in the US. Language of host in Chinese]

Rev> +17.5kk

Sector> Manufacturing (car components)

Type> SHELL

Privileges> Local Admin

Hosts> 50
Price> 1200\$
3. Country> PARAGUAY
Rev> +13kk
Sector> GOV (transport)
Type> SHELL
Privileges> Local Admin
Price> 500\$
4. Country> SRI LANKA
Rev> - Sector> GOV (pensions)
Type> SHELL
Privileges> Local Admin
Price> 450\$
5. Country> VENEZUELA
Rev> +10.5kk
Sector> Travel Agency
Type> SHELL
Privileges> Local Admin
Price> 400\$

The highest-value target in the threat actor's offerings was a manufacturing company in Mexico. The listing offered RDP access with local administrator privileges, shell access with SYSTEM privileges, and AnyDesk access with local administrator privileges to an environment containing more than 70 hosts and 1,300 domain accounts. They also noted the presence of SentinelOne EDR. The starting bid was \$3,500, and the buy-now price was \$5,500:




Fig. 153 — Screenshot of Saturned33's access listing, dated May 30, 2026

The primary access types in Saturned33's listings were shell and RDP, often offered together and with explicit indications that a Windows Active Directory domain controller had been compromised. In early 2026, the actor began selling access via AnyDesk and other remote device management tools.

Prices for Saturned33's initial access offerings varied widely, from \$200 for shell access to \$3,500 or more for lots involving large, high-revenue companies; buy-now ('blitz') prices reached \$5,000. During the initial period of activity (August–September 2025), low-priced lots at \$200–\$800 predominated, with a few listings priced at \$1,200–\$1,300 for domain administrator access or high-revenue victims. By 2026, the actor was posting fewer listings but offering more expensive lots, with starting prices of \$1,500–\$3,500 for premium listings.

The actor also systematically bundled the access they sold with large data dumps from victim infrastructure. According to Saturned33, the exfiltrated data ranged from a few terabytes to more than 150 TB of internal and customer information. Thus, the IAB managed to monetize both access to compromised networks and the stolen data.

A representative example was a January 2026 listing offering access to a business services company in Spain. The lot included RDP access with domain administrator privileges, shell access with SYSTEM privileges, and 5 TB of sensitive internal data and backups stored on the victim's network-attached storage (NAS).




Fig. 154 — Screenshot of Saturned33’s access listing, dated January 26, 2026

Initial Access Sales Statistics — Saturned33

The statistics below summarize Saturned33’s access sales for the period 2025 – H1 2026.

Access sales by country

Country			Count
	US		4
	VE		4
	CH		4
	IR		3
	TW		2
	JP		2
	DE		2
	IT		2
	TR		2

Country		Count
 ZA		2
 CA		2
 PY		1
 LK		1
 GT		1
 FR		1
 KE		1
 LB		1
 IQ		1
 AU		1
 CN		1
 BD		1
 BH		1
 TH		1
 PT		1
 IN		1
 ES		1
 AR		1
 CZ		1
 EG		1
 SA		1
 UG		1
 MX		1

Access sales by industry

Industry	Count
 Manufacturing	16
 Finance	6
 Government and military	5
 Software	5
 Healthcare	3
 Professional Services	3
 Real Estate	2
 Travel and tourism	1
 Media and entertainment	1
 Energy	1
 Commerce and shopping	1
 Transportation and logistics	1
 Not Stated	5

Saturned33 most frequently targeted the United States, Venezuela, and Switzerland; Manufacturing was the leading industry.

ATT&CK Matrix for Enterprise (Saturned33)

Tactics	Techniques	Details
Initial Access	External Remote Services T1133	In their listings, Saturned33 listed AnyDesk and RDP as initial access vectors
Discovery	Software Discovery: Security Software Discovery T1518.001	In their initial access listings, Saturned33 identified the security solutions presented on compromised devices

HelluvaHack

Overview	
First seen	02.10.2024
Victims, 2025 – H1 2026	37
Target geography	 Middle East,  APAC,  Europe
Target countries	 IL,  EG,  IN,  ES

The threat actor known as HelluvaHack registered on a popular underground forum on October 2, 2024, but did not post their first initial access listing until May 2025. The first lot, posted on May 16, 2025, offered VPN and RDP access to an Israeli real estate company with \$29 million in revenue:




Fig. 155 — Screenshot of one of HelluvaHack’s earliest access listings, dated May 16, 2025

Following these first listings, HelluvaHack remained moderately active, posting roughly three dozen unique lots between May 2025 and February 2026. The actor’s defining characteristic

was their pronounced geographic concentration: approximately 92% of victim organizations were in Israel, an atypical pattern for a conventional IAB. This pattern points to two hypotheses: the actor was motivated by ideological or geopolitical factors, potentially both, or their expertise and infrastructure limited them to targets in a single region. However, a gradual expansion beyond Israel was observed by early 2026, suggesting that the actor may have been diversifying their activity.

Notably, the actor’s largest listing offered RDP access with local administrator privileges to an Israeli data and analytics company with revenue of approximately \$394 million, listed on September 6, 2025:




Fig. 156 — Screenshot of HelluvaHack’s access listing, dated September 6, 2025

HelluvaHack’s primary access vector was VPN, often combined with RDP. Over time, the actor began offering additional access types, including access to Windows Active Directory domain controllers, Domain User and workgroup accounts, databases, administrative hosting and management web panels (ISPConfig, InterWorx, and SAP BO), and virtual telephony panels (Virtual DID Panel).

In early 2026, HelluvaHack also began selling Active Directory accounts along with exfiltrated ntds.dit files, indicating that the actor had dumped the domain credential database and could provide buyers with domain user password hashes.

In addition to selling initial access, HelluvaHack engaged in other cybercriminal activity. The actor sold compromised databases containing employee personal data from various companies and advertised a malware dropper builder that could have been potentially used to gain initial access by infecting users' devices with third-party malware:




Fig. 157 — Screenshot of the listing offering a malicious tool called “H Builder”

Initial Access Sales Statistics — HelluvaHack

The statistics below summarize HelluvaHack’s access sales for 2025 – H1 2026.

Access sales by country

Country		Count
 IL		34
 EG		1
 IN		1
 ES		1


Access sales by industry

Industry	Count
 Professional Services	4
 Real Estate	3
 Healthcare	2
 Manufacturing	2
 Transportation and logistics	2
 Finance	2
 Messaging and telecommunications	2
 Energy	2
 IT	2
 Education	1
 Hospitality	1
 Commerce and shopping	1
 Not Stated	13

As noted above, Israel was the actor’s primary geographic focus. No clear industry preference was identified, suggesting that Helluva-Hack was willing to target any sector.

ATT&CK Matrix for Enterprise (HelluvaHack)

Tactics	Techniques	Details
Resource Development	Develop Capabilities: Malware T1587.001	HelluvaHack likely helped develop and operate a malware builder capable of deploying third-party malware to obtain corporate VPN credentials for resale
Initial Access	External Remote Services T1133	In their listings, HelluvaHack listed VPN and RDP as initial access vectors
Discovery	Software Discovery: Security Software Discovery T1518.001	In their initial access listings, HelluvaHack noted the security solutions presented on the compromised devices
	Remote System Discovery T1018	In one access listing, HelluvaHack explicitly stated that they could reach backup servers through lateral movement
Credential Access	OS Credential Dumping: NTDS T1003.003	HelluvaHack offered an exfiltrated ntds.dit dump in their initial access listings



Minions

Overview	
First seen	07.10.2025
Victims, 2025 – H1 2026	23
Target geography	Africa, Americas, APAC, Europe, Middle East
Target countries	GT, CA, CO, TW, TR, VN, BJ, CL, EG, IN, IT, JP, MX, PE

The IAB Minions registered on an underground forum on October 7, 2025, and posted their first corporate network access listing on January 13, 2026. The victim was an unnamed Chilean telecommunications company. The actor claimed read-only access to the company's FortiGate dashboard and stated that the infrastructure contained approximately 100 hosts. As in most of the broker's listings, the company's revenue was not specified.

NO AVATAR

Minions

HDD-drive

Seller

Joined: Oct 7, 2025

Messages: 23

Reaction score: 3

Jan 13, 2026

Price: 100

Contacts: PM

Country Chile

access - Dashboard Fortinet (read only)

[~100] PC

Price is negotiable

Telecommunication Industry

Fig. 158 — Screenshot of Minions' first access listing, dated January 13, 2026

On March 1, 2026, the threat actor registered another underground-forum account under the alias williamblack and began posting new lots, which were later reposted under the Minions account:

Продам доступ к SEO агентству Meta/Google/Tiktok (Ads)
Доступ через ssh, CRM панели и базы данных всех клиентов которые делали покупки 100к строк такого вида: патры, бафалты, итд, product_патры, risks клиенты из северной Америки, Европы и Ближнего Востока. Можно управлять всеми документами и субаккунтами ADS через API ключи, трогается 15 миллионов в...
Минусы: Тара - 13.05.2026 - Отзывы: 0 - Продам ДОСТУП к сайту, ssh, root, sql-инг, DB's

Доступ к админкам SEO онлайн казино
Страна VN Доступ к CRM и базам данных - управление документами и продвижением (референты) в Брондате казино 20к активных доменов api key к телеграмм Ботам Цена 20000\$ // рассматриваю и другие предложения Гарант - Контакт PM
Минусы: Тара - 13.05.2026 - Отзывы: 0 - Продам ДОСТУП к сайту, ssh, root, sql-инг, DB's

Продам доступ к AI проекту
10 миллионов пользователей Трафик весь мир(US,UK,Asia) Доступ ко всем серверам по SSH, Базы данных на ZTB - доступны все разговоры 100 подкастов ежедневно - оплата Stripe и Google Pay Цена 10000\$ // рассматриваю и другие предложения Гарант - Контакт PM
Минусы: Тара - 13.05.2026 - Отзывы: 0 - Продам ДОСТУП к сайту, ssh, root, sql-инг, DB's

Доступ к админкам SEO онлайн казино
williamblack posted a topic in [Access] - FTP, shells, root, sql-инг, DB, Servers
Страна VN Доступ к CRM и базам данных - управление документами и продвижением (референты) в Брондате казино 20к активных доменов api key к телеграмм Ботам Цена 20000\$ // рассматриваю и другие предложения Гарант - Контакт PM выданы меню
May 9

Продам доступ к SEO агентству Meta/Google/Tiktok (Ads)
williamblack posted a topic in [Access] - FTP, shells, root, sql-инг, DB, Servers
Доступ через ssh, CRM панели и базы данных всех клиентов которые делали покупки 100к строк такого вида: патры, бафалты, итд, product_патры, risks клиенты из северной Америки, Европы и Ближнего Востока. Можно управлять всеми документами и субаккунтами ADS через API ключи, трогается 15 миллионов на рекламу Цена 10000\$ // рассматриваю и другие предложения Гарант - Контакт PM выданы меню для проф еств...
May 9

Продам доступ к AI проекту
williamblack posted a topic in [Access] - FTP, shells, root, sql-инг, DB, Servers
10 миллионов пользователей Трафик весь мир(US,UK,Asia) Доступ ко всем серверам по SSH, Базы данных на ZTB - доступны все разговоры 100 подкастов ежедневно - оплата Stripe и Google Pay Цена 10000\$ // рассматриваю и другие предложения Гарант - Контакт PM
April 22

Fig. 159 — Screenshot of identical posts by Minions (top) and williamblack (bottom)

Translated text from Fig. 159

Top:

Selling access to an SEO agency - Meta/Google/TikTok Ads

SSH access, CRM panels, and databases for all customers who made purchases - 100,000 records

in the format: first name, last name, email, product name, price. Customers are in North America, Europe, and the Middle East. You can manage all ad accounts and subaccounts via API keys; 15 million is spent per...

Access to online casino SEO administration panels

Country: Vietnam. CRM and database access, including management of domains and redirectors; eight casino brands, 20,000 active domains, and API keys for Telegram bots. Price: \$20,000; other offers considered. Escrow accepted. Contact via PM.

Selling access to an AI project

10 million users and worldwide traffic (United States, Europe, and Asia)

SSH access to all servers and a 2 TB database containing all conversations. The service receives 100 subscriptions per day, with payments via Stripe and Google Pay. Price: \$10,000; other offers considered. Escrow accepted. Contact via PM.

Bottom:

Access to online casino SEO administration panels

Country: Vietnam. CRM and database access, including management of domains and redirectors; eight casino brands, 20,000 active domains, and API keys for Telegram bots. Price: \$20,000; other offers considered...

Selling access to an SEO agency - Meta/Google/ TikTok Ads

SSH access, CRM panels, and databases for all customers who made purchases - 100,000 records in the format: first name, last name, email, product name, price. Customers are in North America, Europe, and the Middle...

Selling access to an AI project

10 million users and worldwide traffic (United States, Europe, and Asia)

SSH access to all servers and a 2 TB database containing all conversations. The service receives 100 subscriptions per day, with payments via Stripe and Google Pay. Price: \$10,000...

The exact reason for the account switch remains undisclosed. However, the possibility that williamblack was an independent actor whose

access were purchased by Minions for further resale can be ruled out: the reposted listings were identical in both wording and price.

Most of the actor’s listings offered access to Fortinet devices. Minions also sold access via remote access protocols (SSH and RDP) and via FTP.

The actor’s most expensive lot offered access to the CRM system and a database of eight Vietnamese casino brands. It was posted on May 9, 2026, under the alias williamblack and reposted on May 26 under the alias Minions. The actor claimed that the access provided control over 20,000 active domains and access to Telegram bot API keys. The asking price was \$20,000:




Fig. 160 — Screenshot of williamblack’s access listing, dated May 9, 2026

Translated text from Fig. 160

Access to online casino SEO administration panels

CRM and database access, including management of domains and redirectors

Eight casino brands

20,000 active domains

API keys for Telegram bots

Price: \$20,000; other offers considered

Escrow accepted

Contact via PM

Scammers, stay away

Initial Access Sales Statistics — Minions (aka williamblack)

The statistics below summarize Minions’s access sales for the period 2025 – H1 2026.

Access sales by country

Country		Count
 GT		3
 CA		2
 CO		2
 TW		2
 TR		2
 VN		2
 BJ		1
 CL		1
 EG		1
 IN		1
 IT		1
 JP		1
 MX		1
 PE		1
 Not Stated		2

Access sales by industry

Industry	Count
 Messaging and telecommunications	7
 Healthcare	2
 Manufacturing	2
 Commerce and shopping	1
 Finance	1
 Food and beverage	1
 Gaming	1
 Hardware	1
 IT	1
 Media and entertainment	1
 Sales and marketing	1
 Software	1
 Travel and tourism	1
 Not Stated	2

Victims were concentrated in Latin America (Guatemala, Mexico, Colombia, Chile, and Peru) and Asia (Vietnam, India, Japan, and Taiwan).

Messaging and Telecommunications was the most frequently targeted industry.

ATT&CK Matrix for Enterprise (Minions)

Tactics	Techniques	Details
Initial Access	External Remote Services T1133	In their listings, Minions listed SSH and RDP as initial access vectors

privisnabet

Overview	
First seen	01.02.2026
Victims, H1 2026	22
Target geography	Africa, Americas, APAC, Europe, Middle East
Target countries	MY, BR, CR, EG, DE, IN, ID, IL, IT, LB, MX, OM, PE, PH, ZA, CH, AE, US

The threat actor known as **privisnabet** registered on a popular underground forum on February 1, 2026. During the reporting period, the actor posted 22 unique listings offering initial access to corporate networks. Some were later removed, likely after third parties purchased the access; transaction details were not publicly disclosed.

privisnabet posted their first corporate network access listing on February 15, 2026. In the announcement IAB claimed to have an access to several organizations advertising four lots at once:

The first lot offered access to the corporate infrastructure of a Lebanese software and hardware development company. The actor stated that the victim’s annual revenue was \$4.5–\$5.2 million. **privisnabet** described it as access to a workstation containing company emails and corporate documents. No further details were provided.

The second lot offered access to the corporate network of a Malaysian research organization with reported annual revenue of \$68–\$70 million. No technical details were provided.

The third lot listed only the name and domain of a UAE hotel complex and likewise provided no technical details.

The fourth lot offered initial access to a civil society organization in Egypt, also without technical details. Notably, all four lots disclosed the targeted organizations’ names and domains, which is atypical of standard initial access market listings.

The highest-revenue victim claimed by **privisnabet** was a Swiss airline with stated revenue of \$1.6 billion. The actor first offered the access on February 16, 2026, without providing technical details or additional information about the victim.




Fig. 161 — Screenshot of **privisnabet**’s first access listing, dated February 15, 2026




Fig. 162 — Screenshot of privisninet’s listing offering access to the Swiss airline, dated February 16, 2026

On February 18, 2026, privisninet posted an updated listing that included technical details. The access covered more than 40 virtual machines in a VMware vSphere environment and a datastore containing more than 2 TB of data; the actor also claimed to provide AnyDesk access. The lot was offered at auction, with a starting price of \$25,000 and a buy-now price of \$60,000:




Fig. 163 — Screenshot of privisninet’s updated listing offering access to the Swiss airline, dated February 18, 2026

On March 26, 2026, privisninet offered for sale a custom malware developed to automate credential brute-forcing against the RustDesk remote access system. According to the listing, the Go-based tool consisted of two modules: an ID Checker capable of verifying the connection status of up to 800 device IDs per second,

and a password-cracking module (Bruter) supporting up to 1,000 parallel threads with claimed proxy support.

The asking price was \$6,000, and the threat actor offered to transfer the source code exclusively to a single buyer. As evidence of the tool’s effectiveness, privisninet claimed to have compromised more than 50 vulnerable workstations during a one-hour test. The IAB likely used this malware to obtain initial access to corporate networks for subsequent sale.




Fig. 164 — Screenshot of privisninet’s listing offering malware for sale, dated March 26, 2026

Initial Access Sales Statistics — privisnabet

The statistics below summarize privisnabet’s access sales for the period 2025 – H1 2026.

Published access listings by country

Country	Count
 MY	4
 BR	1
 CR	1
 EG	1
 DE	1
 IN	1
 ID	1
 IL	1
 IT	1
 LB	1
 MX	1
 OM	1
 PE	1
 PH	1

Country	Count
 ZA 	1
 CH 	1
 AE 	1
 US 	1
 Not Stated 	1

Published access listings by industry


Industry	Count
 Commerce and shopping 	11
 Travel and tourism 	2
 Food and beverage 	2
 Transportation and logistics 	1
 Software 	1
 Content and publishing 	1
 Manufacturing 	1
 Professional Services 	1
 Clothing and apparel 	1
 Not Stated 	1

Malaysia was the actor’s primary target, although their victims also spanned Africa, the Americas, APAC, Europe, and the Middle East.

Commerce and Shopping was the most frequently targeted industry.

ATT&CK Matrix for Enterprise (privisninet)

Tactics	Techniques	Details
Resource Development	Develop Capabilities: Malware T1587.001	privisninet likely developed a RustDesk brute-force tool for obtaining initial access
Initial Access	External Remote Services T1133	In listings offering initial access to corporate infrastructure, privisninet listed AnyDesk as an intrusion vector
	Valid Accounts T1078	privisninet presumably used credentials obtained through brute-force attacks to access corporate networks
Credential Access	Brute Force T1110	privisninet sold a RustDesk brute-force tool that they could presumably have used to obtain initial access



AckLine

Overview	
First seen	03.08.2025
Victims, 2025 – H1 2026	14
Target geography	 Americas,  Europe,  APAC,  Africa
Target countries	 US,  IT,  NL,  MY,  AR,  DZ,  CO,  TH,  RO,  SI,  ES

The threat actor known as AckLine first surfaced on August 3, 2025, when they created a profile on a major underground forum. Initially, the actor was focused exclusively on the distribution of compromised corporate databases; however, in March 2026, the adversary pivoted toward actively selling initial access to corporate networks. One of their earliest access listings offered access to the Amazon S3 storage of an unnamed company from the United States. The company reportedly had annual revenue of \$8 million, and the actor claimed that the storage contained internal documents and company phone call logs. The asking price was \$1,750:



Fig. 165 — Screenshot of AckLine’s initial access listing, dated March 1, 2026

AckLine’s highest-revenue victim was a Spanish education company with stated revenue exceeding \$350 million. In a listing posted on April 29, 2026, the actor offered access to the company’s corporate network via GlobalProtect VPN. The seller said they had verified only the credentials and had not inspected hosts or security solutions over the network.




Fig. 166 — Screenshot of AckLine’s listing offering access to the Spanish company, dated April 29, 2026

In addition to initial access sales, AckLine provided access free of charge on several occasions by publishing credentials for web panels with local administrator privileges on underground forums. These giveaways were likely

intended to maintain activity on the threat actor’s forum profile, drawing attention to their listings, and building their reputation within the underground community. The screenshot below shows the actor’s publication of administrator credentials for the CMS panel of the website of a software development company based in the United States.




Fig. 167 — Screenshot of AckLine’s March 4, 2026 listing offering free access

Initial Access Sales Statistics — AckLine

The statistics below summarize AckLine’s access sales for the period 2025 – H1 2026.

Access sales by country

Country		Count
US		3
IT		2
ES		1
NL		1
RO		1
CO		1
TH		1
SI		1
AR		1
DZ		1
MY		1

Access sales by industry

Industry	Count
 Messaging and telecommunications	2
 Travel and tourism	1
 Healthcare	1
 Software	1
 Professional Services	1
 Education	1
 Not Stated	7

AckLine has not yet demonstrated a clear geographic or sectoral specialization, targeting organizations worldwide regardless of their country of operation or industry.

ATT&CK Matrix for Enterprise (AckLine)

Tactics	Techniques	Details
Initial Access	External Remote Services T1133	AckLine listed VPN as an initial access vector in their listings.
	Valid Accounts T1078	AckLine validated existing user credentials before selling Amazon S3 access

Recommendations



To protect against current cyber threats, organizations should rigorously follow these recommendations:

- Conduct regular employee training and testing to introduce and reinforce information security fundamentals and digital hygiene principles and raise awareness of current threats.
- Ensure that security tools cover as much of the IT infrastructure as possible.
- Apply software updates promptly and regularly to remediate known vulnerabilities, prioritizing critical issues. Give highest priority to publicly accessible services and the operating systems on which they run, followed by domain controllers, backup systems, centralized management systems, and similar infrastructure.
- Regularly back up important data following global best practices. The 3-2-1 or 3-2-1-1-0 strategy can serve as the basis for implementing backups. The latter is preferable, as it provides integrity checking of the backups.
- Implement multifactor authentication using one-time codes for access to services and corporate networks. Enforce MFA for all user accounts, service accounts, and API gateways across the infrastructure without exception. This measure reduces the risk of threat actors gaining access with compromised or brute-forced credentials.
- Use EDR or XDR solutions, such as F6 Managed Extended Detection and Response (MXDR), together with antivirus software to protect endpoints and the infrastructure as a whole. Respond to and analyze all alerts from these solutions promptly.
- Implement DDoS protection solutions.
- Proactively block traffic to and from potentially malicious hosts.
- Enforce a strong password policy for local and domain accounts, both in the office and for remote work.
- Use unique local administrator passwords on every host. If a single machine is compromised, this prevents the entire affected segment from having to be treated as compromised.
- Design the network architecture so that individual segments can be isolated and both inbound and outbound traffic can be controlled.
- When assigning access rights, follow the principle of least privilege, paying particular attention to service accounts and to accounts used for automated tasks and remote access.
- Restrict network access to the specific tasks assigned to each account. For example, grant a contractor access only to the server they need rather than to the entire network segment or network.
- Use the F6 Attack Surface Management (ASM) solution to monitor the organization's infrastructure security posture and to find vulnerabilities and forgotten digital assets.
- Enable advanced auditing, including disabled event categories, to increase the level of detail in Windows security events.

- Ensure that operating system and security tool event logs are retained for at least three months. Note: this applies not only to workstations and servers but also to deployed security tools, including those protecting the network perimeter (for example, a firewall).
- Configure IP address and account blocking to protect against brute-force attacks.
- Regularly check credentials against public leaks and promptly change any passwords found in them.
- Allow remote access only from trusted IP addresses or after the connecting device has been successfully identified. If neither measure is feasible, implement GeoIP filtering.
- To protect against malicious email campaigns, deploy an email security solution such as F6 Business Email Protection (BEP), which can detect, block, and analyze threats ranging from spam and phishing to malware and business email compromise.
- Prohibit direct RDP access to workstations from outside the internal network.
- Block inbound connections to SMB ports from computers located outside the corporate network.
- Prohibit direct root login to endpoints over SSH.
- Disable or block unused remote services.
- Reduce the attack surface. Do not expose remote-management and access services, such as RDP, SSH, SMB, and WinRM, or other services unless they are properly protected and operationally necessary.
- Continuously identify shadow IT to manage the attack surface. Shadow IT refers to systems and devices deployed by employees without the knowledge or approval of the company's IT department.
- Identify indicators of initial access, persistence, and lateral movement across the network. Although threat actor techniques are often relatively unsophisticated, regular proactive threat hunting can also help detect more advanced attacks.
- Detect known indicators of compromise and check infrastructure against them on a regular basis.
- Analyze antivirus events to determine the circumstances of system compromise.
- Prohibit users from registering for third-party services with corporate email addresses.
- Monitor DNS traffic: some malware needs to communicate with a server or C2, and malicious connections can be disguised as legitimate traffic, including traffic over HTTP.
- Audit unauthorized Telegram user sessions.
- Restrict the use of personal social media and messaging accounts for work purposes.
- Restrict employees' use of personal devices to access the corporate network, including over VPN.
- To counter threat actor activity effectively at both the organizational and technical levels and minimize damage, organizations should either engage external incident response specialists promptly or maintain an incident response retainer.

- To reduce the risk of attacks through contractors, enforce contractor compliance with security policies and requirements, apply Privileged Access Management (PAM) to contractor accounts, connect suppliers to the infrastructure through a dedicated, controlled VPN, and revoke obsolete contractor access.
- Implement centralized collection of events across the infrastructure and forward them to a data collection system (for example, an ELK stack or SIEM).
- In 2026, the complexity, stealth, and targeted nature of cyberattacks will increase substantially. Effective defense will depend not only on how quickly known vulnerabilities are detected and remediated, but also on a fundamental shift from a reactive to a proactive defense model. The ability not only to respond to incidents but also to actively identify hidden threats through anomaly detection, user and entity behavior analytics, and continuous, forward-looking infrastructure monitoring and security testing will become critical.
- In the event of a security incident, avoid reinstalling operating systems or removing identified malware and threat actor tools from affected systems, where possible, until forensically significant data has been collected. Otherwise, it may be difficult or, in some cases, impossible to identify the source of the initial compromise, the persistence mechanisms used, and other circumstances of the cyberattack.
- Before beginning recovery, contain the incident and minimize the risk of a recurrence. This reduces damage and enables safe restoration of the IT infrastructure.

To protect against current cyber threats, users should follow these recommendations:

If a company suffers a cyberattack, follow these recommendations:

- Seek guidance from incident response providers. Given sufficient details about an attack, their specialists can, where possible, provide useful information about the threat actors and their characteristic TTPs. They can also provide tools and guidelines for collecting the forensically significant data needed to respond to and investigate the incident.
- Respond to the incident promptly. If you lack the capacity or expertise to respond on your own, engage a specialized firm.
- Use a unique password for each service and avoid simple combinations containing dates, names, or other personal information.
- Update passwords periodically across all services.
- Enable two-factor authentication (2FA) and cloud passwords for every service that supports them.
- Use a password manager (for example, KeePass) to store passwords, and create unique, strong passwords for each account.
- Disable autofill for passwords and payment information in the browser.
- Use a separate card for online purchases.
- Review your social media presence. Set accounts to private and remove personal information.
- Disable automatic file downloads in messaging apps.

- Download apps only from official stores.
- Do not download suspicious files or click unverified links received from any source, including email, SMS, and messaging apps.
- Limit apps' access to data and the permissions granted to them.
- Back up your data.
- Do not connect to public Wi-Fi networks or use unverified VPNs.
- Enable caller ID and unwanted-call blocking.
- Never disclose confirmation codes, passwords, or other sensitive information to anyone.
- Promptly block bank cards, change service passwords, and take similar steps if you suspect that they may have been compromised.
- Delete unused accounts and apps.
- Clear your browser history and cookies.



Annual Report Global Threat Landscape 2025-2026

F6