CSIS
REST ASSURED

Threat Matrix
Spring 2026

csis.com

Executive summary

Stealth exfiltration extortion campaign

A ransomware group conducted extortion without encryption, leveraging long-term VPN-based access, domain administrator compromise, and RDP lateral movement. Data was staged using kopia.exe and exfiltrated to S3-compatible storage. Rapid provider takedown disrupted attacker access, likely preventing data publication and neutralising extortion leverage despite a months-long undetected intrusion.

Akira exploits VPN patch gap

Akira ransomware operators exploited CVE-2024-40766 in an SSL-VPN appliance during a patch delay window, gaining access via non-AD-authenticated sessions. They established persistence with Chrome Remote Desktop, exfiltrated data using rclone, and deployed encryption. Despite takedown efforts, data replication likely occurred, highlighting risks associated with delayed response and edge-device exposure.

ClickFix lures drive intrusions

H2 2025 saw widespread adoption of "fix-it" social engineering techniques such as FileFix and ClickFix, tricking users into executing malicious commands via trusted interfaces. These methods bypassed MOTW protections and enabled delivery of RATs such as NetSupport and Interlock, reflecting a shift towards user-assisted execution and scalable initial access.

Self-spreading supply chain worm

The Shai-Hulud campaign highlighted escalating software supply-chain risk, using phishing to hijack npm maintainers and inject malicious code into trusted packages. Leveraging tools such as TruffleHog, the worm harvested credentials and propagated across dependencies and CI/CD pipelines, amplifying compromise through transitive trust relationships and developer ecosystem exposure.

Identity-driven ransomware intrusions

Ransomware operations in H2 2025 shifted towards identity-centric, hands-on intrusions, with actors such as Scattered Spider (UNC3944) exploiting IT support workflows and social engineering to gain access. Use of legitimate administrative tools enabled stealthy lateral movement, while attacks increasingly targeted VMware and backup systems to maximise operational disruption and extortion leverage.

Ransomware ecosystem fragmentation

The ransomware ecosystem experienced sustained churn, with groups rebranding or dissolving while new families such as Devman, Bert, and Ailock emerged using recycled codebases. Multi-platform encryptors targeting Windows, Linux, and ESXi became standard, while MaaS offerings such as Matanbuchus 3.0 and MSP-focused campaigns (e.g. SafePay) expanded operational scale and accessibility.

100%

increase in malicious npm packages in 2025

Source: ReversingLabs

China expands cyber espionage

China increased cyber espionage against European and Asia-Pacific government and advanced technology sectors, driven by technology sovereignty disputes and supply chain competition. Targeting focused on semiconductor ecosystems, AI, and STEM industries, combining cyber and HUMINT methods to secure intellectual property and strategic advantage amid tightening global technology controls.

Lumma resurgence and identity risk

Despite a May 2025 law enforcement takedown, Lumma Stealer rapidly re-emerged, adopting stealthier delivery and infrastructure. Its core capability—large-scale credential and session cookie theft—enables account takeover, MFA bypass, and downstream threats such as BEC and ransomware. Persistent session tokens extend exposure windows, requiring organisations to prioritise identity containment beyond password resets.

Industrialised malware operations

Android malware in H2 2025 shifted towards industrialised delivery via Google Play droppers masquerading as utilities. Attackers leveraged staged updates and dynamic payloads to evade detection, then abused Accessibility Services for full device control. Banking trojans such as Anatsa and RedHook enabled automated fraud, OTP interception, and session hijacking across multiple global regions.

Russian hybrid pressure intensifies

Russia sustained a multi-domain campaign across Europe, blending cyber espionage, disinformation, hacktivism, and deniable drone incursions to test NATO thresholds. Activity targeted governments, critical infrastructure, and financial systems, with continued initial access operations and influence efforts likely as Moscow exploits geopolitical tensions and EU support to Ukraine.

Election-timed DDoS campaigns

Pro-Russian group NoName057(16) leveraged its DDoSIA platform to conduct persistent DDoS campaigns against Denmark, clustering activity around the November 2025 municipal elections and March 2026 snap general election. Targets expanded from political parties to ministries, transport, and media, indicating deliberate timing to maximise visibility and disruption during sensitive democratic periods.

Phishing market industrialisation

Chinese-speaking phishing actors have evolved into a structured service economy, mirroring legitimate SaaS models with subscriptions, bundled offerings, and customer support. Seasonal Chinese New Year promotions highlight competitive dynamics, customer acquisition strategies, and market expansion, enabling rapid scaling by lowering technical barriers and continuously onboarding less-skilled operators into organised fraud ecosystems.

81%

SMBs that suffered a breach within the last year

Source: ITRC - 2025 Business Impact Report

GLOBAL STATISTICS

82% malware-free detections in 2025
Source: CrowdStrike

97% increase in risky AI prompts in 2025
Source: Check Point

64% third-party applications accessing sensitive data without justification in 2025
Source: reflectiz

GLOBAL COMPANIES

77% CISOs seeing third-party risk as a major threat in 2026
Source: Panorays

15% increase in Business Email Compromise (BEC) attacks in 2025
Source: LevelBlue

90% identity weaknesses playing a material role for attacker success in 2025
Source: Unit 42

PHISHING

83% email threats related to phishing in H2 2025
Source: Acronis

36% increase in email threats in H2 2025 over H1
Source: ESET

CSIS Threat Matrix Report Spring 2026

This latest Threat Matrix report provides an overview of the cyber threat environment as it continues to evolve in both scale and complexity, and one theme stands out clearly: attackers are becoming more effective because they are getting better at exploiting trust.

March 2026

Legitimate services, common workflows, and everyday user behaviour are increasingly used as intrusion paths. As a result, cybersecurity risk now extends well beyond technical vulnerabilities into identity, human behaviour, and the broader geopolitical environment.

A major shift observed in recent months is the growing reliance on user-assisted compromise. Social engineering techniques such as ClickFix and FileFix illustrate how attackers can bypass automated defences by persuading users to execute malicious commands themselves. Instead of exploiting software flaws, attackers often rely on convincing pretexts and trusted interfaces. This approach allows campaigns to scale quickly while evading traditional security controls.

Identity compromise has, therefore, become a central driver of modern cyber operations. Information-stealing malware, phishing campaigns, and session-token theft are now common entry points for fraud, espionage, and ransomware attacks. In many cases, attackers prioritise persistent access rather than immediate disruption, allowing stolen credentials or authenticated sessions to be reused later for deeper compromise or financial gain.

Supply-chain exposure has also expanded significantly. Threat actors are increasingly targeting developer ecosystems, open-source repositories, and software automation pipelines to reach organisations indirectly. A single compromised package, maintainer account, or exposed credential can propagate across multiple environments through dependency chains and automated build systems. This reflects a broader shift in attacker strategy: compromise the platforms that organisations rely on rather than the organisations themselves.

Artificial intelligence (AI) is beginning to play a larger role in this ecosystem as well. AI-assisted tooling is accelerating malware development, enabling more convincing phishing lures, and increasing the speed with which attackers can test and refine campaigns. At the same time, defensive teams are beginning to adopt AI to improve detection, analysis, and response. The result is a rapidly evolving technological race that will likely shape both offensive and defensive capabilities in the coming years.

The cyber threat landscape is continuing to be influenced by geopolitical developments. State-aligned activity such as cyber-espionage, influence campaigns, and politically motivated hacktivism continue to intersect with criminal operations. Over the past months, pro-Russian hacktivist campaigns have targeted Danish institutions and public services with waves of DDoS attacks timed around elections. These attacks aim to create visibility for the Russian threat, disrupt public trust, and amplify political narratives.

Notably, Greenland has also emerged as a target in these campaigns. Its growing geopolitical significance, combined with Denmark's role in European and Arctic policy, has made it a symbolic pressure point within the broader hybrid threat landscape.

Taken together, these developments illustrate a threat environment that is adaptive, persistent, and increasingly interconnected. Criminal groups, state actors, and hacktivist networks operate in overlapping ecosystems where disruption rarely eliminates capability for long. This report is intended to provide clarity in that environment. By combining threat intelligence, incident response insights, and strategic analysis, this Threat Matrix report aims to help organisations translate awareness into resilience.

At CSIS, our goal remains to ensure that organisations not only understand the risks ahead but are prepared to respond with confidence.

Rest assured.

DANIEL SHEPHERD
CEO, CSIS Security Group

Extortion without encryption

Containing a leak-site threat after stealth data exfiltration

Problem:

In late autumn 2025, an organisation (the "Victim") learned through law enforcement that its logo had appeared on a ransomware group's public leak site. This was unexpected: there were no reports of encrypted systems, no ransom note, and no visible operational disruption. The situation nonetheless presented a serious business risk because leak-site postings are typically used to pressure victims into paying by threatening public disclosure of stolen data.

CSIS was engaged to lead the incident. Early triage quickly shifted the incident from a "possible extortion attempt" to a confirmed network compromise with confirmed data exfiltration that had occurred earlier in the month. The most immediate concern became twofold: (1) determine how deep the attacker had penetrated and whether they still had access, and (2) contain and neutralise the data-leak threat before any publication or further attacker action could occur.

Despite the absence of encryption, the indicators pointed to a classic extortion playbook: gain privileged access, identify and stage sensitive files, exfiltrate them to attacker-controlled infrastructure, and then use the leak site to force payment. The organisation needed rapid containment to prevent follow-on actions (such as ransomware deployment, backup destruction, or additional exfiltration) and an evidence-based assessment of what happened, when it happened, and what data was at risk.

Investigation:

CSIS initiated evidence collection immediately upon engagement. The investigation focused on endpoint artefacts and Windows event telemetry from critical servers, supported by targeted triage across multiple hosts using CSIS tooling. The objective was to reconstruct attacker activity across the environment, validate whether privileged compromise occurred, and identify any infrastructure used for exfiltration.

A consistent early theme in the evidence was suspicious activity originating from the SSL VPN segment, suggesting the attacker leveraged remote access as a staging point into the internal network. The earliest confirmed unauthorised events dated back to the summer of 2025, where failed authentication attempts and SMB session failures were observed—behaviour consistent with probing, password guessing, or exploratory lateral movement from a VPN-connected device. Shortly thereafter, the attacker achieved interactive access to an internal terminal server using a compromised legacy domain user account. From this foothold, they executed common administrative and reconnaissance tooling associated with hands-on-keyboard intrusions, including utilities used to enumerate persistence points and query Active Directory.

Over time, the attacker reappeared intermittently, establishing RDP sessions to increasingly sensitive systems. By mid autumn, evidence showed successful access to the domain controller and, critically, later use of a domain administrator account—marking a major escalation and confirming the attacker obtained full administrative control of the Windows domain. Around this time, artefacts consistent with credential access were observed, including the creation of an archive named in a way that strongly suggested an LSASS memory dump, a common method to extract credentials for broader compromise and lateral movement.

As the timeline progressed, attacker behaviour shifted from reconnaissance and privilege escalation to operational execution. Logs showed rapid RDP-based lateral movement across multiple systems with administrative privileges. On the file server, the attacker executed, a legitimate backup tool, in a manner consistent with bulk collection and staging for export. This activity aligned with the confirmed data exfiltration window.

CSIS then pivoted to identifying the destination of the stolen data. Execution parameters revealed the attacker was uploading to an S3-compatible object storage endpoint hosted by a third-party provider in Europe. This was a crucial turning point because it provided an external control point to disrupt the extortion chain. CSIS immediately initiated a takedown/suspension request with the hosting provider, supported by collected evidence linking the storage endpoint to the intrusion.

Two investigation constraints impacted certainty about initial access. Log retention limitations on the VPN appliance and domain controller meant key historical records were unavailable for the earliest stage of the compromise. As a result, CSIS could not conclusively determine whether initial access was achieved via a compromised VPN account, stolen user credentials used for remote access, exploitation of an exposed service, or another pathway. However, the evidence did conclusively establish: long-term unauthorised presence, domain administrator compromise, broad lateral movement, and data exfiltration to attacker-controlled storage.

Solution:

Containment and disruption were executed in two parallel tracks: stopping attacker access to the environment and neutralising the exfiltration/leak leverage.

First, the Victim, guided by CSIS, made the decision to take decisive containment action by shutting down the entire network immediately. This step aimed to immediately sever any active attacker sessions, halt ongoing command execution, and prevent additional lateral movement or destructive activity. Given the confirmed domain administrator compromise, the organisation treated the environment as fully untrusted until proven otherwise.

Second, CSIS focused on cutting off the attacker's ability to access the exfiltrated data. After identifying the attacker's object storage destination, CSIS submitted an evidence-backed request to the hosting provider to suspend the storage server/bucket environment. The provider confirmed suspension quickly, effectively preventing further uploads and, critically, preventing the attacker from retrieving or managing the data stored there.

Follow-on communications from the hosting provider indicated that the attacker later attempted to regain access to the suspended storage environment and was denied. In parallel, the negotiation track included a request to the attacker for proof of possession of the stolen data; the attacker did not provide proof in a manner consistent with having ready access to the dataset. Taken together, these signals strongly suggested that the attacker may not have maintained an alternate copy of the exfiltrated files and that the primary dataset remained contained within the now-suspended hosting environment. As a result, the most significant business risk—the credible threat of public data release—was materially reduced before any publication occurred.

With containment achieved, CSIS recommended remediation actions aligned with a domain-level compromise: rebuild affected hosts from known-good baselines, rotate credentials broadly (including service accounts), and harden authentication controls. Special attention was placed on improving centralised logging and retention so that future investigations would not be limited by short log windows. Although the initial access vector could not be proven due to telemetry gaps, the remediation approach assumed worst-case compromise of credentials and privileged systems to ensure the attacker could not re-enter using previously captured access.

Lessons learned:

- A leak-site posting can indicate a serious compromise even when no ransomware encryption or ransom note is present; extortion-only operations can still involve significant data theft.
- Long dwell time (months) is possible when visibility is limited; inconsistent attacker "check-ins" can mask a sustained campaign until the extortion phase begins.
- Domain administrator compromise changes the incident posture entirely—response should assume full environment compromise and prioritise rapid isolation and credential reset.
- Limited log retention on critical infrastructure can prevent determining initial access and delays confidence-driven remediation decisions.

Recommendations:

- Implement centralised logging with sufficient retention for VPN, domain controllers, and critical servers to support investigations across months, not days/weeks.
- Expand MFA and conditional access policies for remote access paths and privileged accounts; reduce reliance on legacy accounts and enforce least privilege.
- Establish a rebuild-and-recover playbook for domain compromise (gold images, rapid credential rotation procedures, and validated restore processes).
- Adopt 24/7 detection and response coverage (internal SOC or MDR) with alerting tuned for RDP lateral movement, credential dumping indicators, and anomalous bulk file access/exfiltration tooling.

The weekend patch gap

Akira ransomware via SSL-VPN compromise

Problem:

An organisation (the "Victim") experienced a ransomware incident affecting critical Windows file shares. The event began when staff discovered that at least one server had been encrypted and a ransomware note was present. In response, the Victim worked with its hosting provider to rapidly contain the situation by shutting down systems and restricting inbound/outbound traffic to prevent further spread. While these containment actions limited additional encryption, the organisation faced two urgent uncertainties: how the attackers gained access and whether data had already been stolen.

The incident was assessed as an intrusion by the Akira ransomware group, a ransomware-as-a-service (RaaS) actor known for "double extortion"—stealing data before encrypting systems to increase pressure on victims. The risk profile therefore included both operational disruption (encrypted shares) and a potential data breach (exfiltration and possible later publication).

A major complication was timing. The Victim needed insurer approval to start a full investigation, and the insurer/partner response introduced delays. Those delays mattered because for Akira-style operations, the best chance to reduce impact is to quickly identify and disrupt the actor's exfiltration infrastructure before stolen data is replicated elsewhere. Additionally, limited log retention on the edge VPN device reduced visibility into the earliest phase of the attack, making it harder to conclusively prove the initial compromise path.

Investigation:

CSIS took the lead on incident response and performed an in-depth investigation focused on reconstructing the attack chain, determining impact, and identifying evidence of exfiltration. The investigation leveraged endpoint and server artefacts collected from the Windows environment, along with available VPN telemetry and backups/snapshots to recover deleted attacker tooling configuration.

Because edge device logging was limited, CSIS could not definitively replay the initial access sequence end-to-end. However, the weight of evidence supported a high-confidence assessment that the attackers entered through exploitation of a public-facing SSL-VPN/edge appliance vulnerability (CVE-2024-40766). The Victim had auto-updating enabled, but updates were scheduled for weekends, creating a short exposure window between patch release and installation. Akira is known to exploit such windows quickly, particularly against remote access appliances. A key investigative clue was that unauthorised SSL-VPN-connected activity was observed from devices that did not correspond to legitimate Active Directory (AD) VPN authentications. This mismatch aligned with a known post-exploitation behaviour: attackers can create or enable local SSL-VPN accounts on the appliance that authenticate locally (not via AD) and may not appear in AD authentication logs.

Once inside the network via SSL-VPN, the threat actor attempted to access the backup server using network logons. Those attempts failed, which CSIS assessed was likely due to the backup server not being AD-joined—an architectural choice that reduced the attacker's ability to compromise backups. The attacker then pivoted to more consequential targets and achieved interactive access to the Domain Controller via RDP from an internal IP associated with the VPN session.

On the Domain Controller, the attacker established persistence by installing Chrome Remote Desktop as a Windows service, enabling continued access even if initial credentials were changed. The actor performed discovery using a network scanning tool (Advanced IP Scanner) and proceeded to data theft. Evidence showed use of the legitimate tool rclone to stage and exfiltrate data from file shares. Although the attacker deleted rclone configuration files to hinder response efforts, CSIS recovered the configuration from snapshot backups, which contained the exfiltration method (FTP) and the endpoint details. This enabled CSIS to identify the exfiltration server and initiate a takedown request with the relevant hosting provider, and to share details with Danish law enforcement. Despite those actions, the timing indicated that the attackers likely copied the stolen data to an additional location before the exfiltration server could be disrupted, consistent with Akira's operational patterns.

Finally, CSIS confirmed the impact phase: Akira ransomware ("win_locker.exe") was executed, encrypting available file shares and dropping ransomware notes across directories. In parallel, evidence suggested compromise of Active Directory credential material, including exposure of the AD database (NTDS.dit). This elevated the severity because it implied broad credential compromise and the potential for rapid re-compromise if the environment were simply "cleaned" rather than rebuilt.

Solution:

The response strategy focused on stopping attacker activity, reducing the likelihood of data publication, and restoring business operations safely. Early containment by the Victim and hosting provider—shutting down systems and restricting traffic—helped prevent further spread of encryption and limited continued interactive attacker access. CSIS then prioritised actions that materially reduce risk in Akira-style incidents: identify exfiltration mechanisms, disrupt attacker infrastructure where possible, and determine whether key security boundaries (backups and identity systems) were compromised.

A critical operational success was that backups were not compromised. The backup server's separation from Active Directory prevented straightforward credential-based takeover, preserving a viable recovery path. This allowed restoration planning without the additional burden of rebuilding from incomplete or attacker-corrupted backups.

To address ongoing risk, CSIS recommended rebuilding the compromised Domain Controller from a known-good state predating the incident, rather than attempting piecemeal remediation. Because evidence indicated compromise of AD credential data, CSIS advised treating all directory accounts—including service accounts—as compromised. This included rotating all AD passwords, and either resetting the Kerberos service account password twice (to invalidate forged ticket scenarios) or rebuilding AD entirely depending on the organisation's tolerance for risk and the extent of observed compromise. On the data theft front, CSIS used recovered rclone configuration to locate the exfiltration server and initiated a takedown request with the hosting provider, while also sharing indicators with law enforcement. This step was intended to prevent or delay publication and reduce the attacker's leverage. However, due to unavoidable timing (including delays prior to investigation start), the actor likely replicated the stolen data elsewhere. In practical terms, this shifted the solution from "prevent exfiltration" to "prepare for downstream impacts," including validating what data was at risk through crisis management support and ensuring the organisation could notify affected parties if needed.

Restoration and hardening were positioned as parallel tracks: restore critical services from clean sources while implementing controls to prevent reinfection—especially on the VPN edge, identity layer, and logging/monitoring coverage that would be crucial for detecting any attempted return by the threat actor.

Lessons learned:

- In double-extortion ransomware cases, speed is a decisive factor; delays in authorising investigation can increase impact, especially around exfiltration disruption.
- Limited edge-device log retention can prevent definitive root-cause proof, even when the broader attack chain is well supported by endpoint evidence.
- Separating backup infrastructure from Active Directory meaningfully increases resilience by reducing credential-based compromise paths.
- Once a Domain Controller and AD credential material are compromised, rebuild/credential reset strategies must be treated as identity-recovery operations, not just malware cleanup.

Recommendations:

- Reduce patch exposure windows for internet-facing edge devices (accelerated patching for critical CVEs; consider out-of-band updates rather than weekend-only schedules).
- Implement centralised logging and longer retention for edge devices, identity systems, and critical servers to support fast, high-confidence investigations.
- Expand MFA/conditional access coverage and explicitly enforce controls for local VPN accounts and administrative access paths.
- Establish pre-approved incident response playbooks with insurers/partners (including authority to investigate lateral movement) to avoid time lost during the most critical early hours.

General cyber threats

From exploits to execution – social engineering and evasion redefined initial access

In H2 2025, the cyber threat landscape was characterised by a marked increase in the sophistication of evasion techniques, a continued surge in multi-platform information stealers, and persistent activity from state-aligned espionage actors.

Threat actors increasingly leveraged paste-and-run social engineering techniques such as FileFix and ClickFix to bypass traditional security warnings and trust indicators, enabling delivery of a wide range of malware, including remote access trojans (RATs) and loaders. The malware-as-a-service (MaaS) model continued to lower barriers to entry for cybercriminals, with widely available commodity tools such as Katz Stealer and AsyncRAT enabling rapid operational scaling.

Trust-exploitation as default

H2 2025 saw adversaries consistently weaponise trusted interfaces, services, and ecosystems—turning routine user actions and "allowed" enterprise traffic into reliable intrusion paths. The defining pattern was not a single malware family, but repeated exploitation of trust anchors: developer registries, CI/CD workflows, SaaS integrations, and mainstream cloud services used for staging, command-and-control, and exfiltration.

Supply-chain compromise remained a significant concern, with multiple high-profile incidents involving open-source ecosystems. Shai-Hulud was a self-replicating npm supply-chain incident in which attackers compromised maintainers via phishing, injected malicious code into trusted packages, and used those publication rights to propagate laterally across the ecosystem. The worm's payload combined credential harvesting with propagation logic: it leveraged tooling such as TruffleHog to scan developer environments and repositories for secrets (including npm tokens and cloud keys), exfiltrated those credentials, and then used them to compromise additional packages—creating a compounding blast radius through transitive dependencies and CI/CD runners.

Malicious packages were identified in the npm registry and the Arch Linux User Repository (AUR), introducing malware directly into developer environments and downstream build pipelines. In parallel, advanced actors exploited both zero-day and rapidly weaponised vulnerabilities, such as the SharePoint "ToolShell" flaw and CitrixBleed 2.

On the defensive side, coordinated international law-enforcement operations resulted in notable disruptions to cybercrime infrastructure, including action against the pro-Russian DDoS group NoName057(16) and the Badbox 2.0 IoT botnet.

A defining trend during the period was the maturation of "fix-it" style lures that manipulated users into executing attacker-supplied commands. FileFix campaigns persuaded victims to paste malicious commands into the Windows File Explorer address bar, exploiting a trusted navigation interface to achieve execution while bypassing Mark-of-the-Web (MOTW) protections.

ClickFix used a similar psychological pretext, presenting fake browser or system errors and instructing users to paste and execute commands—typically PowerShell—via Run, Terminal, or PowerShell interfaces. These techniques were observed delivering a range of payloads, including NetSupport RAT, Latrodectus, Interlock RAT, and other commodity loaders.

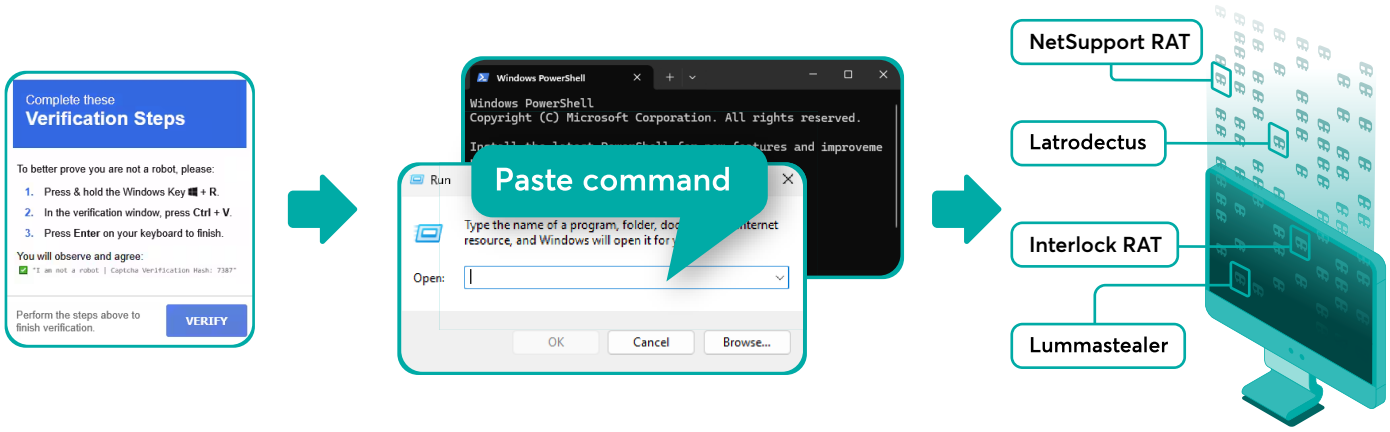

User-assisted initial access

"Fix-it" paste-and-run lures matured into a high-yield access mechanism, with ClickFix/FileFix-style techniques repeatedly bypassing security prompts and delivering commodity loaders and RATs. The psychological pretext (fake errors/CAPTCHA) coupled with direct user execution shifted the intrusion surface from attachment handling to endpoint scripting interpreters and interactive shells, increasing cross-platform reach and reuse across both cybercrime and state-linked targeting.

Several other remote access trojans also made a mark. A PHP-based variant of Interlock RAT was delivered by the KONGTUKE loader via FileFix social engineering. EDSKManager RAT was observed as a multi-stage implant incorporating Hidden VNC (HVNC) functionality and multiple evasion layers. PureRAT, a .NET-based RAT, was delivered using the Ghost Crypt crypter, which employed process re-imaging to evade endpoint detection. XWorm RAT version 6.0 was distributed using steganography, hiding its payload within image files to bypass security scanners.

The use of sophisticated evasion frameworks was widespread. The SHELLTER commercial loader was repeatedly abused to package and deploy commodity infostealers such as Lumma and Rhadamanthys while evading static and behavioural detection. The SLOW#TEMPEST campaign employed advanced obfuscation techniques, including control-flow graph (CFG) manipulation, to hinder reverse engineering.

In the financial malware space, Coyote targeted Brazilian banking users and represented the first documented in-the-wild abuse of Windows UI Automation (UIA), using the legitimate Electra framework to automate fraudulent transactions.

Information-stealer activity remained intense. Lumma Stealer continued to evolve following disruption efforts, with newer versions incorporating enhanced anti-analysis features and diversified distribution channels. Batavia emerged as a multi-stage spyware targeting Russian industrial organisations from mid-2024 onward, exfiltrating documents and detailed system information.

Octalyn Stealer, written in C++ and Delphi and distributed via GitHub-hosted lures, targeted credentials, browser data, and cryptocurrency wallets. Shuyal Stealer was identified as a new advanced stealer with layered evasion techniques. macOS threats also expanded: Atomic macOS Stealer (AMOS) introduced backdoor functionality for persistent access, while the macOS.ZuRu backdoor resurfaced embedded in a trojanised version of the Termius application.

Equally, state-aligned and advanced threat activity persisted. The NightEagle APT was assessed by multiple researchers to be conducting a long-running campaign targeting high-tech, military, and quantum-technology organisations in China, leveraging custom malware and access gained via zero-day exploitation. China-aligned actors were also linked to phishing campaigns targeting Taiwan's semiconductor sector using industry-event lures to deliver bespoke backdoors, as well as enhancements to the DeedRAT backdoor to improve espionage capabilities.

Supply chain targeted

Open-source compromise broadened from typo-squatting into maintainer compromise, malicious updates, and propagation through dependency graphs and automation pipelines. The key takeaway is the expansion of the attack surface from package managers into IDE ecosystems and CI/CD (including workflow injection and secrets harvesting), making developer tooling and build infrastructure a recurring "first breach" vector with downstream enterprise-wide consequences.

A DPRK-linked group targeted Web3 and cryptocurrency organisations using NimDoor, a custom backdoor written in the Nim programming language. The GoldMelody group exploited leaked ASP.NET Machine Keys to forge authentication tickets and gain access to Microsoft IIS servers, subsequently monetising that access via resale to other threat actors. Meanwhile, the pro-Russian DDoS group NoName057(16) was disrupted in a coordinated international action known as "Operation Eastwood," resulting in significant impact on its infrastructure.

Threat actors also demonstrated increased focus on the software development lifecycle. Besides the previously mentioned Shai-Hulud incident, one notable campaign published 67 malicious packages to the npm registry as part of a "contagious interview" social-engineering scheme targeting developers. Separately, malicious AUR packages were identified that installed Chaos RAT on Linux systems. Another novel technique involved embedding a persistent Linux backdoor within image files featuring panda imagery, leveraging steganography and AI-assisted code generation to evade detection; this activity was associated with multi-stage malware capable of long-term persistence.

Botnet activity remained relevant in H2 2025. Rondobox was assessed as a botnet primarily used for DDoS operations and proxy-style services, employing a multi-stage infection chain and WebSocket-based command-and-control. Badbox 2.0 continued to operate at significant scale, compromising millions of uncertified Android devices and monetising access through advertising fraud and traffic proxying; Google initiated legal action aimed at dismantling its infrastructure.

Exploitation activity throughout the period was shaped by rapid attacker adoption of high-impact vulnerabilities. CVE-2025-53770 ("ToolShell") in Microsoft SharePoint's ToolPane component enabled unauthenticated remote code execution and was actively exploited prior to patching by Microsoft. CitrixBleed 2 (CVE-2025-5777) affected NetScaler ADC and Gateway appliances and was exploited to retrieve session tokens and hijack authenticated sessions. SonicWall SMA vulnerabilities facilitated deployment of the OVERSTEP backdoor using living-off-the-land techniques.

Leaked ASP.NET Machine Keys were leveraged by initial access brokers to forge authentication tickets and gain remote code execution on IIS servers. A CrushFTP zero-day (CVE-2025-54309) was also identified as a critical unauthenticated vulnerability, enabling attackers to access sensitive files and ultimately achieve remote code execution.

Stealth-by-design

Evasion shifted from "obfuscate the payload" to "abuse the platform," including steganography/metadata smuggling, policy and driver trust boundary manipulation, and multi-layered anti-analysis patterns. Across campaigns, there was a clear emphasis on persistence and survivability—ranging from multi-stage implants and HVNC-capable RATs to stealth frameworks and packaging techniques designed to survive takedowns and degrade endpoint visibility.

CVE-2025-55182 ("React2Shell") was disclosed on 3 Dec 2025 as a critical unauthenticated RCE in React Server Components and was exploited almost immediately at scale across multiple threat clusters, ranging from opportunistic criminal activity to suspected espionage; activity.

Financially motivated actors rapidly monetised exploitation via XMRig cryptomining and ecosystem chatter accelerated tool and proof-of-concept diffusion, while early exploit "noise" (including AI-generated or non-functional code) increased validation risk and underscored the need to rely on trusted technical sources.

Exploits as a main intrusions-driver

Rapid adoption of critical internet-facing vulnerabilities shaped exploitation tempo throughout the period, with repeated waves targeting widely deployed enterprise platforms and frameworks. The React2Shell cycle in December captured the broader dynamic: fast weaponisation, multi-cluster exploitation spanning espionage and monetisation (cryptomining), diverse post-compromise tooling, and significant operational risk from exploit "noise" and unreliable PoCs that complicated validation and response.

MongoBleed (CVE-2025-14847) was characterised during December 2025 as a high-impact data exposure issue affecting internet-facing MongoDB deployments, driven by unauthenticated reads of uninitialised memory associated with wire-protocol handling (including the OP_COMPRESSED path). The practical risk was leakage of sensitive in-memory artefacts—credentials, tokens, and PII—under certain conditions, with public scanning narratives indicating a large, exposed population on the public internet (commonly described as on the order of tens of thousands).

The n8n automation platform RCE (CVE-2025-68613) was highlighted as a critical remote code execution condition arising from expression/code injection that can lead to RCE in specific authenticated contexts, rather than a generic sandbox escape. The incident was notable in the same period because it aligned with a broader pattern of rapid adoption of high-impact vulnerabilities affecting widely deployed internet-facing automation and integration tooling, with adjacent disclosures in the same timeframe also covering sandbox-bypass mechanics (e.g., CVE-2025-68668) as distinct issues.

Main vulnerabilities exploited

CVE-2025-53770 – "ToolShell" (MS SharePoint ToolPane)

Unauthenticated remote code execution (RCE) in SharePoint's ToolPane component.

CVE-2025-5777 – "CitrixBleed 2" (NetScaler ADC / Gateway)

Lets attackers steal valid session tokens from NetScaler, enabling session hijacking.

SonicWall SMA vulnerabilities

Used to deploy the OVERSTEP backdoor by "living off the land" techniques.

Leaked ASP.NET Machine Keys

Stolen machine keys let attackers forge tickets and run code on IIS servers.

CVE-2025-54309 – CrushFTP zero-day

Enables access to sensitive files and escalation to full remote code execution.

CVE-2025-55182 – "React2Shell"

Widely exploited by criminal and state-linked groups to deploy tunnels and backdoors.

CVE-2025-14847 – "MongoBleed" (MongoDB)

Can leak in-memory credentials, tokens, and PII from exposed MongoDB instances.

CVE-2025-68613 – n8n automation platform RCE

Can cause remote code execution in certain authenticated scenarios, not a full escape.

CVE-2025-68668 – Sandbox-bypass mechanics

Bypasses sandbox protections, adding risk beyond expression injection RCE.

Ransomware

Identity abuse, multi-platform encryption, and the rise of hands-on intrusions

The ransomware threat landscape in H2 2025 was shaped less by the dominance of a single group and more by a convergence of aggressive intrusion tradecraft, ecosystem volatility, and continued professionalisation of access and extortion operations.

While Scattered Spider (UNC3944) remained one of the most operationally impactful actors during the period, its activity should be understood as illustrative of broader shifts rather than as the sole defining force. Across the ecosystem, ransomware operations demonstrated increased willingness to conduct hands-on-keyboard intrusions, leverage social engineering against identity and support workflows, and exploit trusted third-party relationships to bypass perimeter controls.

The ransomware ecosystem itself remained highly fluid. Multiple established brands either ceased operations or rebranded, often in response to law-enforcement pressure, internal disputes, or reputational degradation. Hunters International announced the cessation of ransomware activity, with its operators claiming a transition toward a data-extortion-only model under the "World Leaks" brand. SatanLock similarly announced a shutdown. These exits were offset by the emergence of new or rebranded families—often leveraging recycled tooling or leaked codebases—demonstrating the continued churn typical of the ransomware economy rather than a net reduction in capability.

Multi-platform ransomware continued to mature as a standard design goal. Threat actors increasingly deployed encryptors capable of targeting Windows, Linux, and VMware ESXi environments, enabling full-environment disruption rather than workstation-centric impact. This trend aligned with attackers' focus on virtualised infrastructure and backup systems as primary leverage points in enterprise environments.

Law-enforcement actions produced tangible, if localised, disruption. Infrastructure associated with the BlackSuit ransomware operation was seized, and a separate action resulted in the seizure of approximately USD 24 million in cryptocurrency linked to the Chaos ransomware operation. While these actions temporarily degraded specific operations, they did not materially reduce overall ransomware activity, as operators rapidly migrated infrastructure and affiliates redistributed across competing platforms.

Scattered Spider (also tracked as UNC3944) significantly escalated its operational scope during the period. Initially known for data theft and extortion without encryption, the group increasingly incorporated destructive ransomware deployment into its campaigns. Its operations relied heavily on social engineering, including impersonation of IT support and abuse of identity recovery workflows, rather than vulnerability exploitation. The group made extensive use of legitimate administrative, remote access, and security tools to blend malicious activity into normal enterprise telemetry, enabling prolonged hands-on-keyboard intrusions. In several cases, Scattered Spider activity culminated in targeted attacks against VMware vSphere and backup infrastructure, maximising operational disruption.

Ransomware targeted identity

Leading ransomware actors—most notably Scattered Spider (UNC3944)—demonstrated that high-impact intrusions can be achieved without reliance on zero-day exploits, instead abusing identity recovery processes, IT support impersonation, and trusted relationships to obtain persistent access and conduct hands-on-keyboard attacks and degrade endpoint visibility.

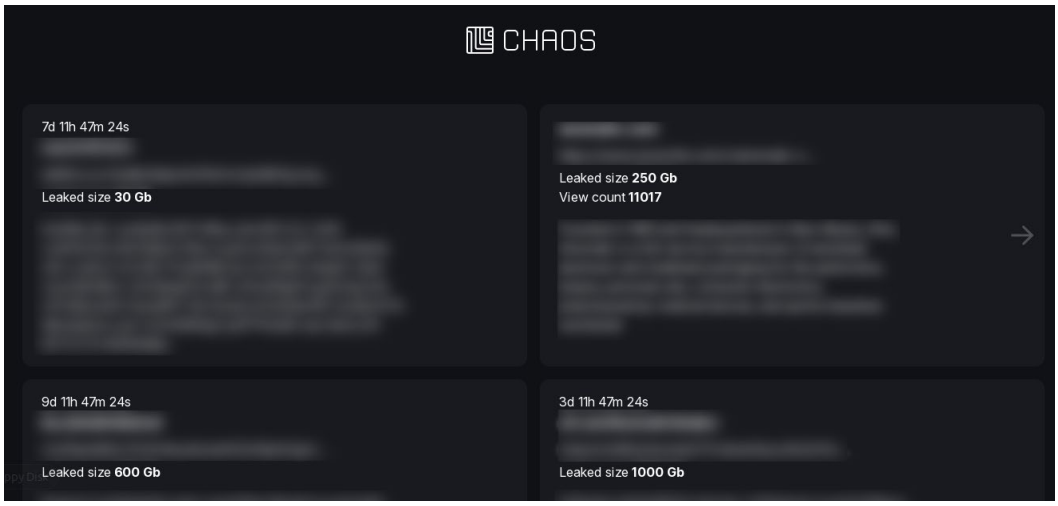

Detail: CHAOS leak site.

Trusted-relationship abuse emerged as a critical ransomware-adjacent access vector. ShinyHunters, while historically associated with data theft rather than ransomware, conducted high-impact breaches affecting organisations including Qantas, Allianz Life, and LVMH by compromising a Salesforce environment belonging to a third-party service provider. These incidents reinforced the systemic risk posed by identity compromise within SaaS platforms and supplier ecosystems, which increasingly serve as initial access points for downstream extortion and ransomware activity.

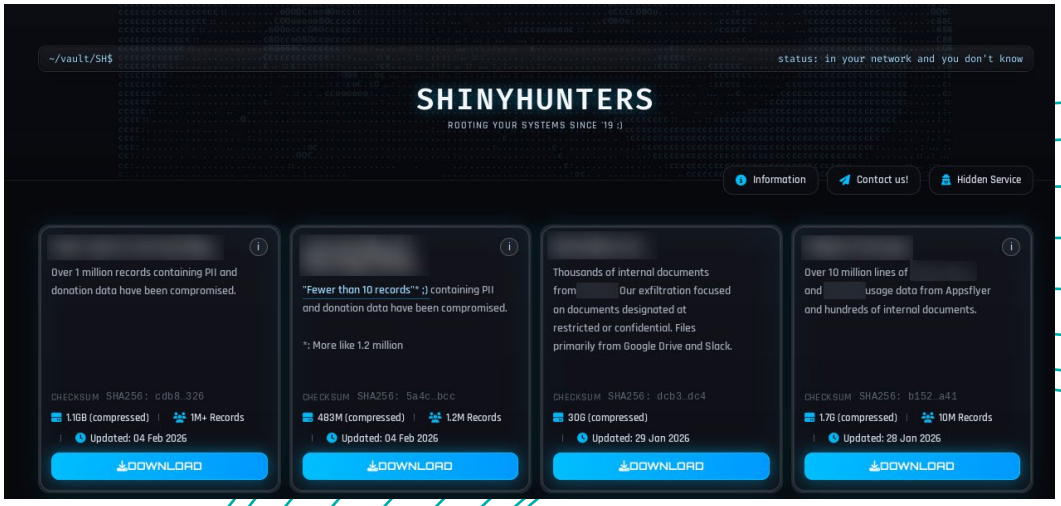

Detail: ShinyHunters leak site.

Multi-platform encryption

Ransomware families increasingly targeted Windows, Linux, and VMware ESXi environments in parallel, reflecting a strategic focus on virtualised infrastructure, backups, and hypervisors to maximise operational disruption and recovery pressure on enterprises.

Ecosystem churn masked continuity

Group shutdowns and rebrands (e.g., Hunters International to World Leaks, SatanLock) did not reduce overall ransomware capacity; instead, tooling, affiliates, and codebases rapidly re-emerged through new brands, recycled malware (e.g., Babuk-derived variants), and alternative extortion models.

New and evolving ransomware families continued to appear throughout H2 2025. Devman ransomware emerged as a Go-based family targeting both Windows and Linux systems, supporting online and offline encryption modes through embedded cryptographic material. Bert ransomware similarly targeted Windows, Linux, and ESXi environments across multiple regions. Ailock employed intermittent or selective encryption to accelerate execution and reduce detection, combined with multi-stage loading chains designed to hinder analysis. SafePay gained visibility as a ransomware operation focused on Managed Service Providers, using MSP compromise to amplify downstream impact. NailaoLocker, while technically unsophisticated and using weak or hard-coded passwords in its encryption routine, nonetheless demonstrated that low-quality implementations can still pose operational risk. Kawa4096 leveraged the leaked Babuk source code, underscoring the continued downstream impact of historic ransomware leaks.

The Chaos ransomware operation—distinct from earlier Chaos malware builders—was assessed as a rebrand or evolution of the BlackSuit/Royal ecosystem. Its campaigns followed a big-game-hunting model, frequently initiating access via voice phishing to deploy legitimate remote access tools, followed by lateral movement and deployment of multi-platform ransomware using selective encryption to optimise speed and impact.

Ransomware delivery and extortion tactics continued to diversify. Social-engineering techniques originally associated with loaders and stealers were adapted for ransomware delivery. ClickFix-style lures, which coerce users into executing attacker-supplied commands, were observed in campaigns leading to ransomware deployment, including activity associated with Epsilon Red.

The Malware-as-a-Service ecosystem remained a key enabler: Matanbuchus 3.0 was actively marketed as a loader capable of delivering ransomware and post-exploitation tooling, lowering barriers for affiliates and smaller crews. Extortion strategies also evolved; the Qilin ransomware group publicly advertised a so-called "legal department," attempting to increase pressure on victims by invoking regulatory exposure, including GDPR-related consequences.

Finally, unusual developments were observed even within ransomware tooling itself. A security researcher identified a remote code execution flaw within a Cl0p-associated encryption utility, theoretically enabling defensive or law-enforcement intervention against the ransomware process. While not operationalised at scale, the discovery highlighted that attacker tooling can itself introduce exploitable weaknesses, particularly in rapidly iterated ransomware codebases.

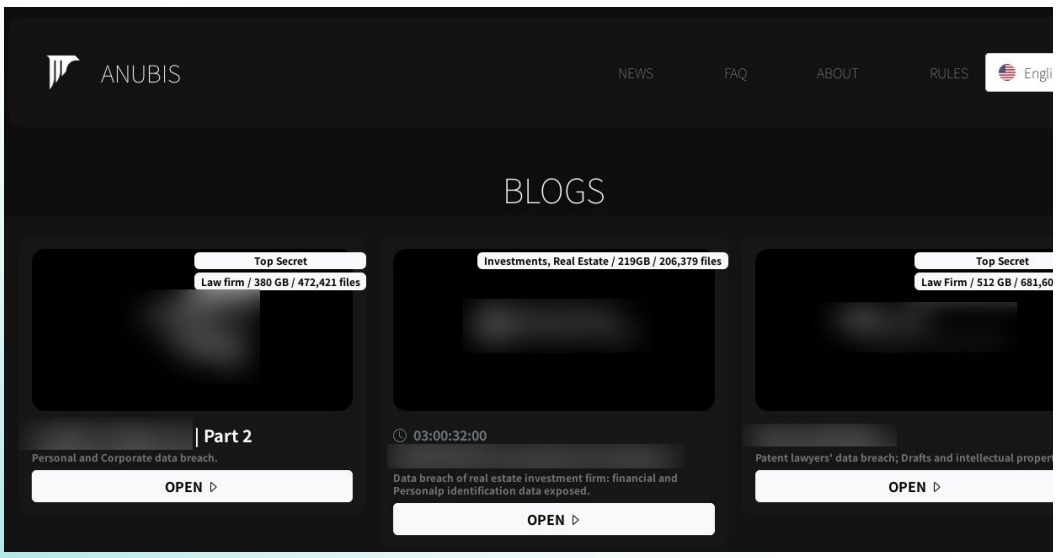

Detail: Anubis leak site.


EILEENA

Lumma Stealer resurfaced after law enforcement disruption

Denmark targeted with ClickFix and steganographic in-memory injection

A Denmark-targeted incident observed in late 2025 shows a mature infostealer ecosystem adapting delivery and evasion rather than changing objectives: rapidly stealing credentials, browser data, and authenticated sessions to enable financial fraud and downstream intrusion.

The campaign's defining feature is the pairing of ClickFix—a social-engineering technique that convinces users to manually execute a malicious command—with a steganography-based loader that hides later-stage payloads inside legitimate-looking images, then executes them in memory using injection into a trusted Windows process.

This matters because ClickFix shifts the security boundary: instead of exploiting a software vulnerability, it exploits a workflow and human behaviour, creating an attack path that can evade "automatic" controls. Microsoft notes that the required user interaction (the user running the command) is specifically what helps ClickFix slip past conventional automated defences, and that final payloads are frequently loaded in memory via living-off-the-land execution paths rather than dropped as obvious executables.

Denmark targeting

Based on the incident analysed by CSIS, the observed delivery chain in Denmark followed a well-established ClickFix "user-run command" pattern and then transitioned into a staged, memory-heavy loader sequence: ClickFix delivery → mshta/PowerShell → reflective .NET loader → PNG steganography → Donut shellcode → explorer.exe injection → Lumma Stealer execution. The use of mshta to retrieve remote script content and PowerShell to decode and load .NET payloads aligns closely with public analyses of ClickFix chains delivering infostealers.

At a high level, the campaign used script-based stages to load a .NET component that extracted and decrypted a PNG-embedded payload, ultimately producing shellcode that was injected into explorer.exe (a ubiquitous Windows user process) and then used to run the infostealer. This core pattern has been documented multiple times: image-based concealment, extraction into a byte array, and subsequent injection into explorer.exe, showing how the approach blends execution into common system behaviour while limiting obvious binaries on disk.

A notable operational detail in this campaign is that this represents the first confirmed case of this steganographic loader chain being used in a geographically specific Denmark-focused incident rather than broadly opportunistic distribution. Public reporting from Denmark is limited on this exact victimology, but Microsoft's broader Lumma analysis emphasises the malware's global reach and flexible distribution infrastructure—making regional "focus" a meaningful signal when it appears in reliable incident telemetry.

Why ClickFix works (and why it evades controls)

ClickFix is best understood as "social engineering as an execution primitive." A typical ClickFix chain begins with users being routed (via phishing, malvertising, or compromised sites) to a convincing visual lure and then being tricked into executing a malicious command themselves—an explicit design choice that can bypass conventional automated security solutions because the user becomes the initiator of execution.

LummaC2 has perfected this technique: victims are presented with a fake CAPTCHA and instructed to open the Windows Run dialogue and paste clipboard contents; pressing Enter launches a subsequent Base64-encoded PowerShell process. User-invoked execution is a repeatable, scalable initial access mechanism for LummaC2, not a one-off trick.

Operationally, ClickFix also degrades the value of several traditional controls. First, it can reduce reliance on malicious attachments (which email gateways are optimised to detect) by moving the decisive step into a browser-based or site-based lure and the user's Run prompt. Second, ClickFix-delivered payloads are often "fileless" or memory-resident, loaded via trusted binaries and .NET/CLR execution paths, complicating purely file-based detection strategies.

Steganography and Donut injection at a high level

Steganography in these campaigns is not "hiding an executable at the end of a file"; it is often encoding payload material directly into PNG pixel data so that the malware can reconstruct and decrypt it only at runtime. Malicious code can be embedded within specific colour channels of PNG images and then reconstructed in memory, an approach that helps evade signature-based scanning and complicates triage because the file looks like a benign image until processed by the loader.

The in-memory "bridge" between a concealed payload and execution is frequently shellcode plus process injection. In ClickFix clusters, extracted shellcode is injected into a target process (notably explorer.exe) using memory allocation and remote-thread style techniques, with observed Donut-packed shellcode as part of this chain. Injection into common Windows processes increases stealth and complicates attribution because malicious activity executes inside a legitimate host process.

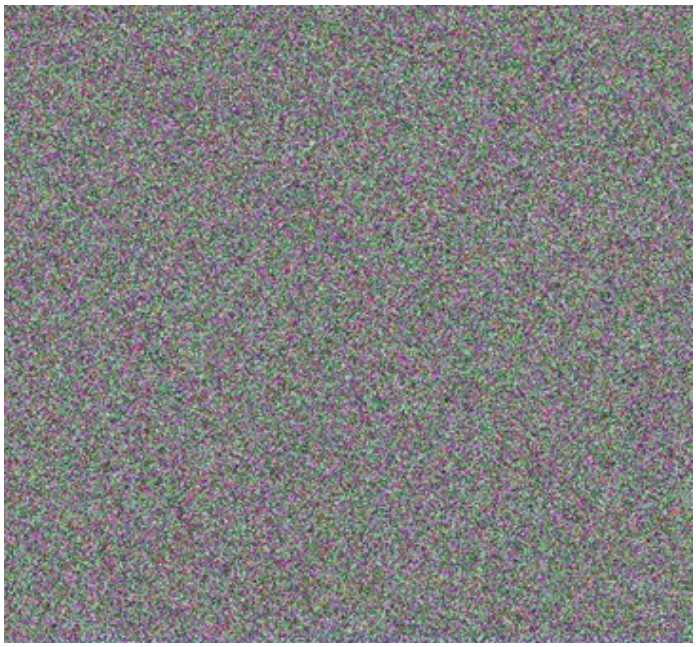

Detail: Encoded PNG Image.

Donut itself is widely documented as an offensive and dual-use tool that converts .NET assemblies (and other payload types) into position-independent shellcode designed for in-memory execution, optionally supporting staged retrieval and encryption. The official Donut project describes its purpose as enabling in-memory execution of .NET assemblies and other modules, illustrating why it is attractive to infostealer operators who want to minimise disk artefacts and hinder straightforward scanning.

Lumma takedown and resurgence

In May 2025, coordinated public- and private-sector actions significantly disrupted Lumma's infrastructure. The U.S. Department of Justice announced court-authorised seizure of multiple domains used to operate the LummaC2 information-stealing malware service, explicitly linking Lumma to large-scale credential theft used to facilitate crimes such as fraudulent bank transfers and cryptocurrency theft.

Industry participation underscored the operation's breadth. ESET stated it collaborated with Microsoft and multiple partners to target Lumma infrastructure, aiming to make the botnet "in large part, inoperative," while also describing Lumma as one of the most prevalent infostealers over a multi-year period and emphasising that harvested credentials are commonly sold to other criminals, including ransomware affiliates.

However, telemetry indicates disruption did not end operations. Lumma re-emerged shortly after the takedown and that activity began resurging in June–July

2025, with operators shifting to more discreet channels and stealthier evasion tactics while continuing to expand reach. This "regroup and adapt" behaviour is consistent with malware-as-a-service economics: infrastructure can be rebuilt, affiliates can rotate delivery, and tooling can be repackaged.

Lumma and ClickFix timeline (non-technical)

2025/05 Global disruption actions and domain seizures (public + private partners)

2025/06 Activity begins rebounding; operators shift tactics

2025/10 ClickFix lures use fake verification/update themes with image-based concealment

2025/11 Denmark-targeted incident observed using stego loader + in-memory injection

Business impact and enterprise risk

The central business impact is credential and session theft. Lumma is designed to exfiltrate sensitive information that enables account takeover and fraud. Lumma can steal browser data, financial credentials, cryptocurrency wallets, browser extensions, and even MFA-related details, creating a multi-dimensional identity exposure issue rather than just a workstation malware problem.

Session theft is particularly consequential because it can weaken MFA's practical protections. MITRE ATT&CK documents the technique "Steal Web Session Cookie," explaining that adversaries can steal session cookies and use them as authentication tokens to access web applications as an authenticated user without needing the original credentials.

The Federal Bureau of Investigation has specifically warned that criminals steal "remember-me" cookies to gain access to accounts, describing how these cookies can enable access without repeated logins and can persist for extended periods (commonly around 30 days). This creates a window in which remediation must include not just password resets, but also session invalidation, device sign-out, and stronger conditional access controls.

Finally, the downstream risk includes BEC and ransomware enablement pathways. Attackers can use stolen session cookies and credentials to access mailboxes and conduct follow-on business email compromise activity, reinforcing how identity compromise can quickly become financial fraud and reputational risk even without deep network intrusion.

Mitigations and detection priorities

The practical mitigation strategy should prioritise breaking the ClickFix step, containing identity fallout, and improving behaviour-based detection for file-light execution. Microsoft explicitly frames ClickFix as depending on user-executed commands and notes that attackers refine obfuscation and in-memory loading; this supports an executive mandate to invest in user-resistant controls (policy hardening, restricted script execution paths) alongside awareness training that focuses on the specific ClickFix behaviour pattern (Run box + paste + execute).

Because infostealers often lead to session/token abuse, executive readiness should include a playbook for rapid identity containment: revoke sessions, rotate credentials, prioritise privileged and finance accounts, and hunt for anomalous sign-ins and mailbox activity. Microsoft's guidance on token theft emphasises that as MFA coverage increases, adversaries increasingly pursue token/session theft, which requires prevention and response measures beyond password resets alone.

Behaviour-based detection should focus on the chain's "tells": unusual mshta and PowerShell execution relationships, reflective .NET loading patterns, suspicious image processing followed by memory allocation/injection behaviours, and explorer.exe being used as an execution host. Huntress specifically recommends monitoring suspicious process lineage such as explorer spawning mshta or PowerShell and highlights that the attack's "simplicity" is the user executing a malicious command, making high-fidelity endpoint telemetry and prevention of Run-box abuse valuable controls.

Control	Purpose	Effort	Risk reduction
Restrict or disable Windows Run for standard users (where feasible)	Breaks the ClickFix "Win+R paste" execution step	Medium	High
Constrain scripting + enforce robust PowerShell logging (including script block telemetry)	Surfaces and limits in-memory script staging	Medium	High
Block or tightly control mshta and other high-risk LOLBins via application control	Reduces living-off-the-land launcher abuse	Medium	Medium-High
Identity hardening (conditional access, device compliance, session lifetime tuning)	Limits usefulness of stolen cookies/tokens	Medium-High	High
Rapid session revocation + credential rotation playbook	Reduces attacker dwell time after theft	Medium	High
Security awareness focused on "never paste commands to fix updates/CAPTCHAs"	Targets the core deception vector	Low	Medium

Executive takeaway

- This is an identity-first campaign: the likely "blast radius" is stolen credentials and session tokens, not just a single endpoint infection.
- ClickFix is engineered to bypass automation: by making a user perform the decisive step, the chain can evade many pre-execution controls.
- Steganography + in-memory injection raises dwell-time risk: hiding payloads in images and injecting into explorer.exe reduces obvious on-disk indicators and blends into normal processes.
- Takedown ≠ eradication: Lumma's May 2025 disruption was significant, but reputable telemetry shows activity rebounded shortly after as operators shifted tactics.

Geolocation

Lumma Stealer C2

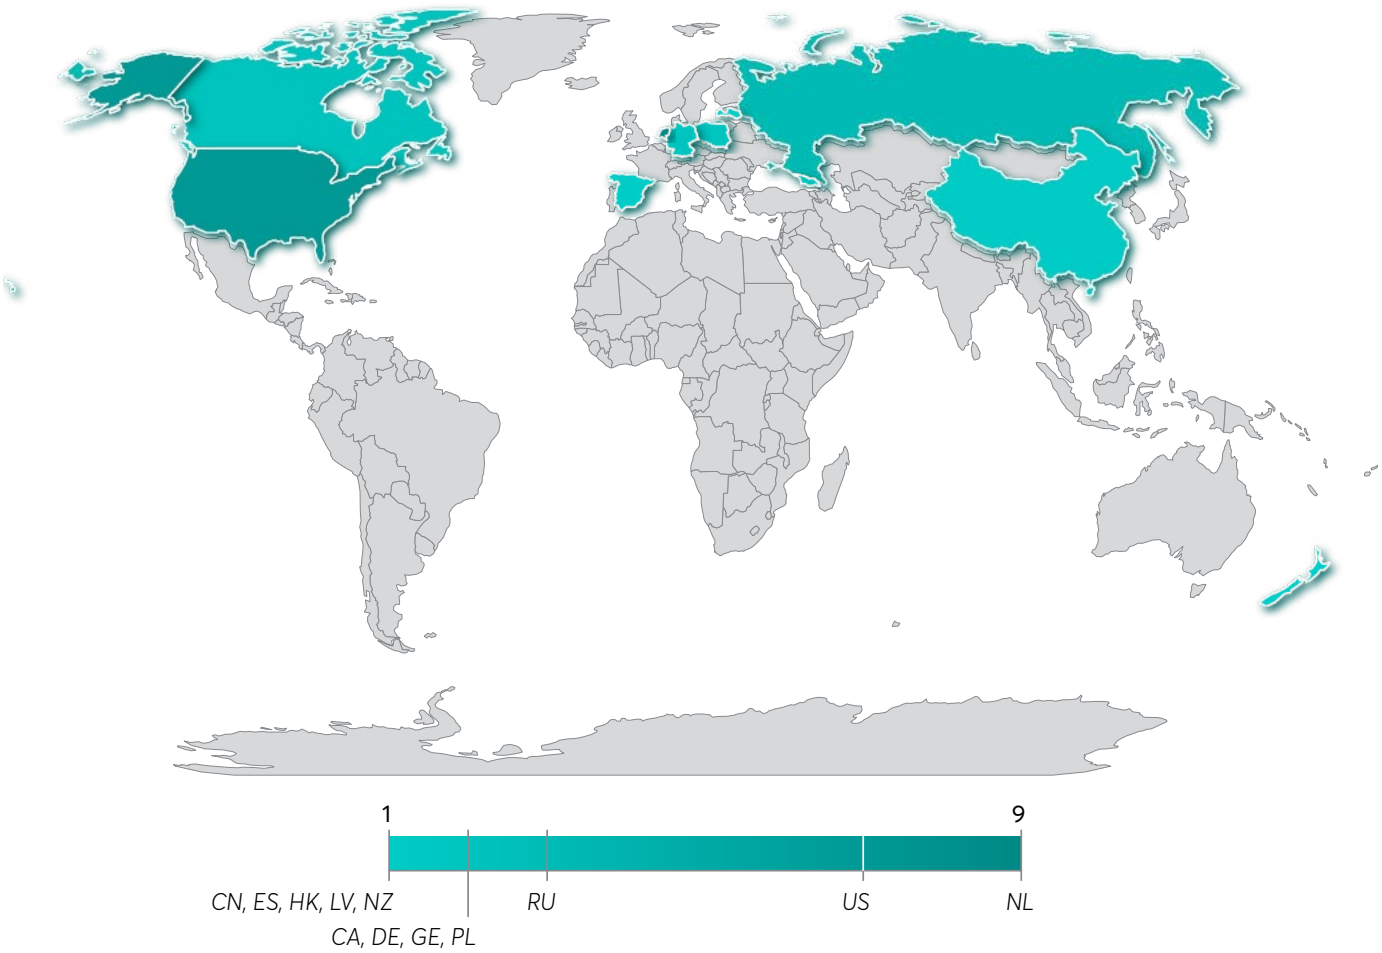

Lumma Stealer C2

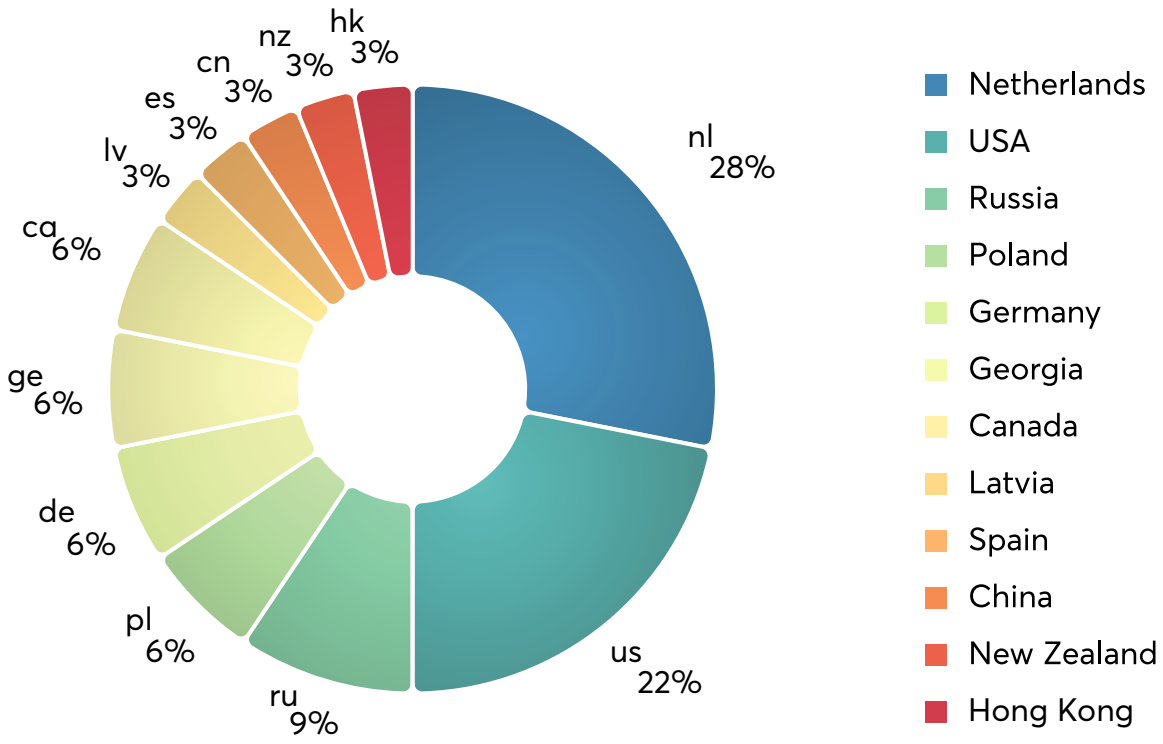

Control by consent

Accessibility abuse and the industrialisation of Android malware

The mobile threat landscape in H2 2025 was characterised by a sustained increase in the scale and sophistication of Android malware, with banking trojans remaining the dominant operational threat.

While Android continued to be the primary target, the defining shift was not novelty of malware families but the industrialisation of delivery, permission abuse, and post-installation control. Banking trojans were observed operating across multiple regions, with several campaigns expanding beyond their original geographic focus to target users in North America, Europe, and parts of Asia.

A consistent and highly effective attack model underpinned these campaigns. Threat actors repeatedly abused the official Android app ecosystem by publishing seemingly benign dropper applications on the Google Play Store. These applications—often masquerading as PDF readers, document scanners, system utilities, or productivity tools—contained little or no malicious code at the time of review, allowing them to bypass automated and manual vetting. Malicious functionality was subsequently introduced via staged updates or dynamic payload retrieval, after which users were socially engineered into enabling Android Accessibility Services. Once granted, this permission enabled comprehensive device control, including UI interaction, keystroke capture, screen content access, interception of one-time passwords, and fully automated on-device fraud.

Abuse of Accessibility Services remained the single most important enabler for modern Android banking malware. Families including Anatsa, DoubleTrouble, and RedHook all relied on carefully crafted in-app prompts, overlays, and persistence mechanisms to coerce users into granting these permissions. Following activation, the malware was able to operate with minimal further user interaction, performing credential theft, session hijacking, and fraudulent transactions directly on the device, thereby bypassing many server-side fraud detection controls.

Accessibility services abused

The most impactful Android banking trojans rely less on exploit development and more on coercing users to grant Accessibility permissions, which effectively provide full UI-level control of the device and enable on-device fraud, credential theft, and OTP interception at scale. Exploitation of 2FA mechanisms, mnemonic recovery phrases, and transaction interfaces.

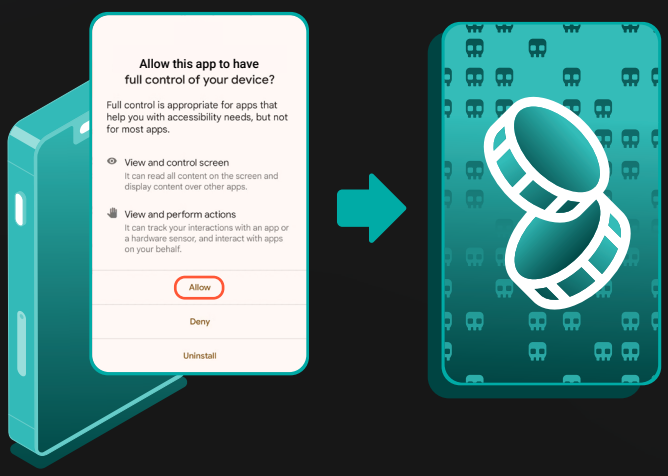

Google Play operationally exploited

Threat actors repeatedly publish benign dropper applications that pass initial vetting and later activate malicious functionality through staged updates or dynamic payload delivery, allowing sustained campaigns despite periodic takedowns.

Threat actors demonstrated growing proficiency in maintaining long-term presence on the Play Store. Dropper applications were routinely disguised as legitimate consumer utilities and achieved significant download volumes before removal. This approach was notably used by Anatsa operators and by the Konfety malware family, which historically functioned as adware but evolved to include more robust evasion and loader capabilities. In these cases, the Play Store served not as a one-time delivery channel but as a recurring access vector, reinforcing its attractiveness despite takedown efforts.

While earlier Android banking campaigns were often regionally constrained, recent activity showed increasing geographic diversification. Anatsa, also known as TeaBot, expanded from a primarily European focus to actively target financial institutions and users in the United States and Canada. RedHook focused on users in Vietnam, employing lures and overlays tailored to local banking and e-commerce applications. Separately, an ongoing campaign targeted customers of more than 20 Indian banks via smishing, distributing malware disguised as bank rewards or service applications and abusing Accessibility Services to harvest credentials, card data, and personal information.

Malware developers also invested heavily in evasion and loader sophistication. The Ducex packer exemplified this trend, using a multi-stage architecture that embedded encrypted payloads across multiple DEX files and relied on reflection-based loading to execute code entirely in memory. This approach significantly degraded the effectiveness of static analysis and signature-based detection. Konfety, which resurfaced during the period, incorporated improved sandbox and emulator detection, allowing it to evade automated analysis while continuing to function primarily as an ad-fraud and unwanted-app distribution platform.

Beyond financially motivated malware, mobile surveillance remained a persistent concern. Researchers identified DCHSpy, an Android surveillance tool attributed to an Iranian threat actor, distributed through targeted social-engineering campaigns. Masquerading as legitimate services such as VPN clients or messaging application updates, DCHSpy enabled extensive data exfiltration, including contacts, call logs, SMS messages, device location, photos, and other sensitive artefacts, highlighting the continued use of mobile platforms for targeted espionage.

Additional delivery mechanisms diversified the mobile malware ecosystem. The QwizzSerial campaign used malicious online quizzes promoted via social media to lure victims into downloading Android installers disguised as "results" files. While primarily associated with Android distribution, this campaign also delivered information-stealing malware, including commodity stealers, reinforcing the convergence between mobile and desktop credential-theft ecosystems. Separately, trojanised Android applications impersonating Telegram were observed delivering SpyNote RAT variants capable of credential theft, audio recording, and location tracking.

From regional to global targeting

Families such as Anatsa expanded beyond their original geographies to target users in North America and Europe, while parallel campaigns targeted Southeast Asia and India using locally tailored lures, indicating increasing campaign maturity and scalability.

Finally, hybrid monetisation models gained traction. Security researchers identified Android applications that combined background ad-click fraud with credential-harvesting functionality via phishing overlays for popular social media platforms. This blending of low-risk monetisation with higher-value credential theft reflects a broader trend toward maximising return on infection rather than specialising in a single fraud outcome.

Malware timeline

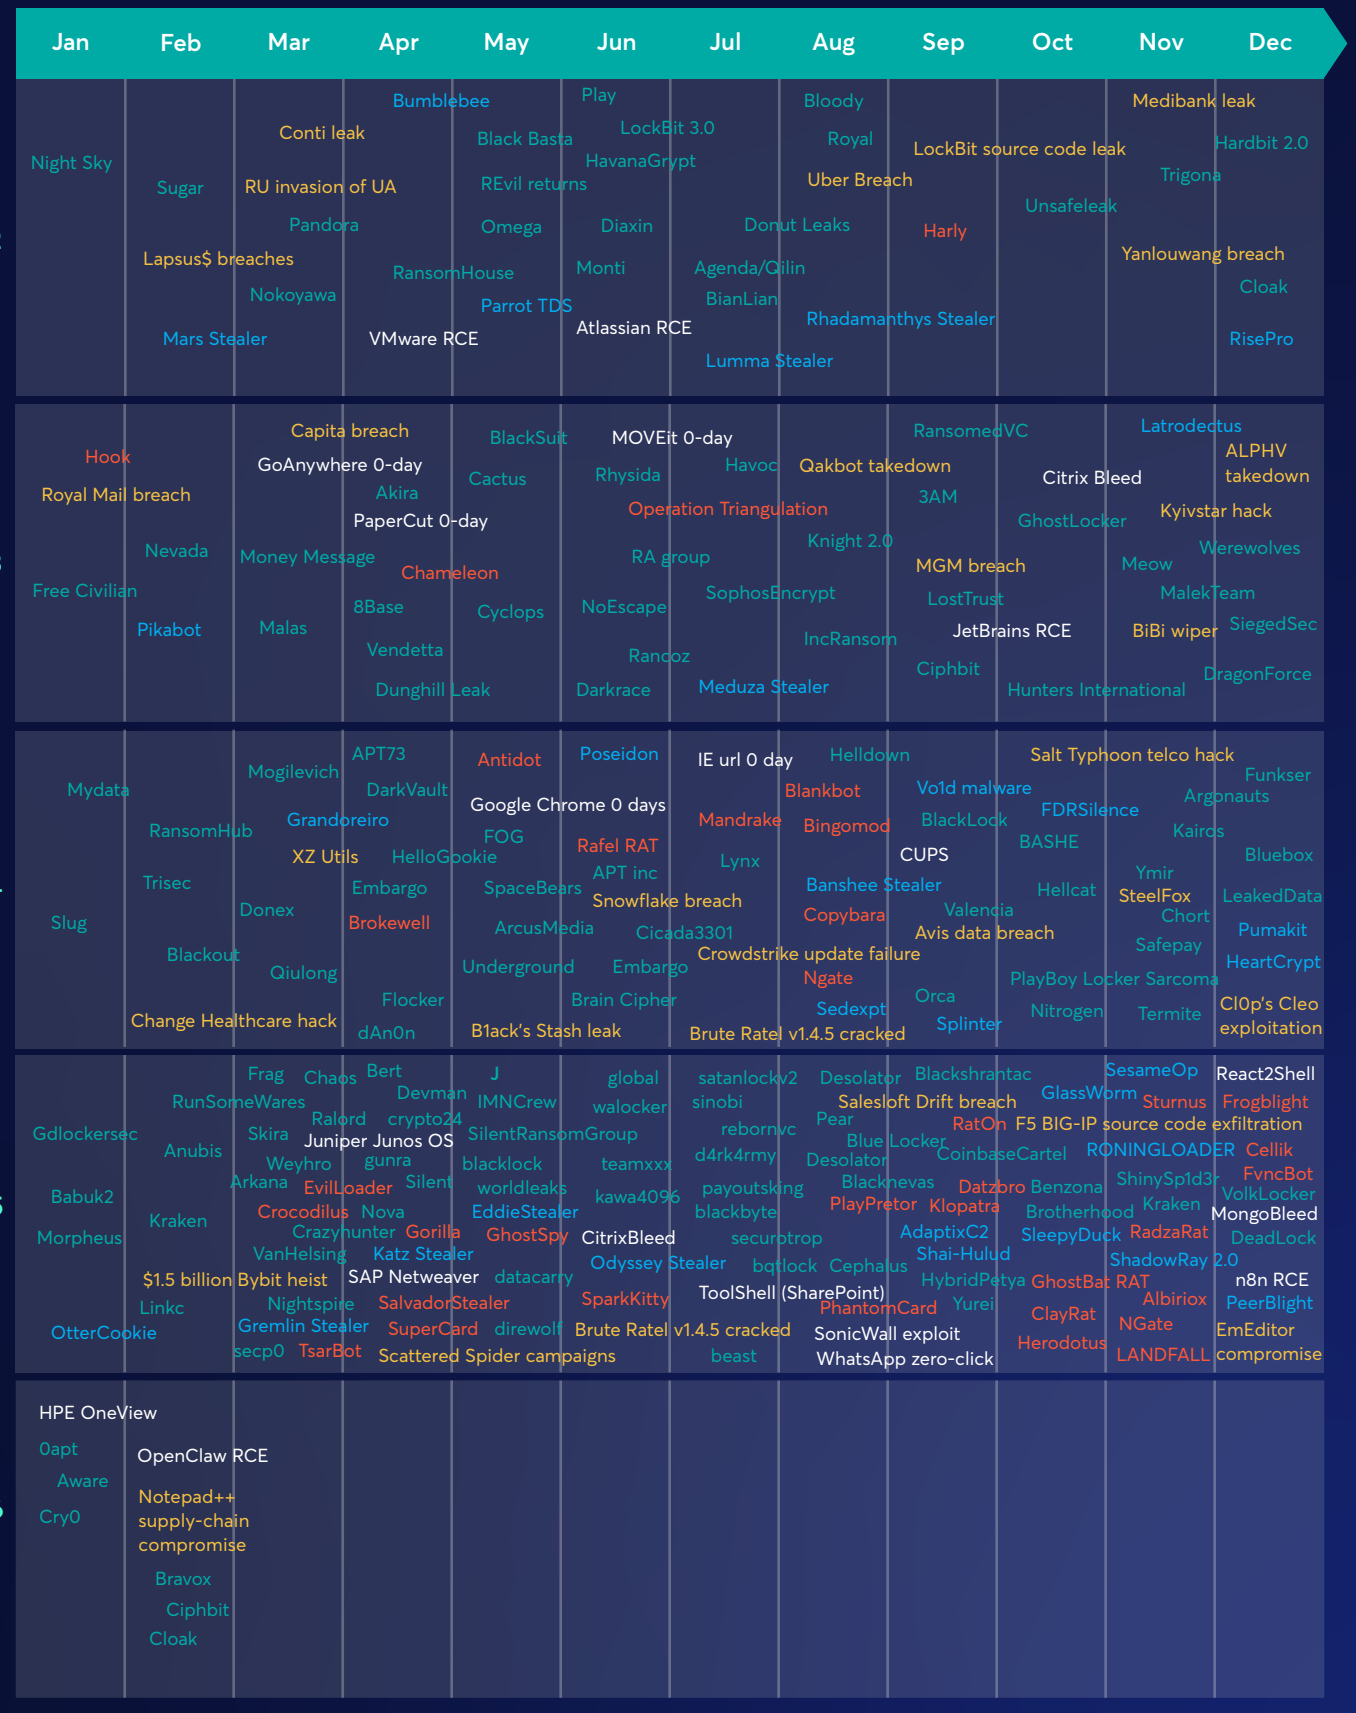

News from the phish pond

How Chinese New Year reveals the logic of the Chinese phishing underground

Chinese New Year, also known as Spring Festival, is usually associated with family, renewal, and commerce. In the Chinese phishing underground, it also marks something else: a seasonal sales cycle.

Holiday promotions, discounted platform access, bundled services, and "new year" relaunches show an ecosystem that increasingly behaves like a real digital industry—competitive, product-driven, and built for growth.

> ***"What makes this ecosystem dangerous is not just its scale. It is the fact that it now thinks like a business."***

Chinese New Year

Over the past two years, Chinese-speaking phishing actors have evolved from operators running delivery-themed SMS scams into a broader underground economy built around phishing platforms, fake online shops, payment-card theft, digital wallet abuse, and account takeover. The tools have matured, the workflows have become more efficient, and the market now shows many of the traits of a legitimate software sector: subscriptions, upgrades, customer support, training, bundled add-ons, and aggressive competition.

Chinese New Year offers an unusually clear view into that transformation. The holiday period has long been one of the biggest commercial moments in the Chinese calendar—a time for promotions, customer acquisition, relaunches, and symbolic fresh starts. In the underground market, the same pattern now appears to apply. Seasonal offers tied to the new year suggest that phishing platforms are no longer marketed only as criminal tools. They are being sold as business opportunities.

Ecosystem evolution

That shift matters because it helps explain why the ecosystem has scaled so quickly. What was once a narrower phishing problem now resembles a service economy. Operators are not only buying kits; they are buying access to infrastructure, onboarding help, storefront support, traffic generation, and updated versions of an evolving product line. The language of the market is no longer just technical. It is commercial.

At the centre of this ecosystem is a simple but powerful idea: remove friction for the customer. That means lowering the technical barrier for new entrants, packaging fraud capabilities into easy-to-buy tiers, and providing enough support that even less sophisticated actors can participate. The result is an underground market that can continuously recruit, train, and expand.

The new year theme sharpens that picture. Seasonal promotions are not important because they are festive. They are important because they reveal timing, intent, and confidence. An actor offering a limited-time holiday deal is behaving less like a lone scammer and more like a vendor trying to grow market share. Discounted access, free components, bundled features, and launch messaging all point to the same conclusion: these groups see themselves as operating inside a competitive market, and they plan around commercial cycles.

From phishing to Phishing-as-a-Service

This is why the Chinese ecosystem stands apart. It has moved well beyond one-off phishing pages and entered a more mature phase of platformisation. Under the surface, many of these operations are tied to highly organised back-end systems, multilingual admin panels, role-based access, reusable brand templates, payment workflows, and integrated support for different monetisation paths. The phishing page is no longer the product. It is only the front door.

The real value lies further downstream. These actors do not just steal credentials or card details; they increasingly convert stolen payment data into digital wallet tokens on attacker-controlled devices, allowing fraudulent

purchases to continue with much less friction after the initial compromise. Others push into fake e-commerce, building convincing storefronts that look like ordinary online retail sites while quietly harvesting payment data and authentication codes during checkout. Still others are branching into financial accounts, adapting the same logic to higher-value targets.

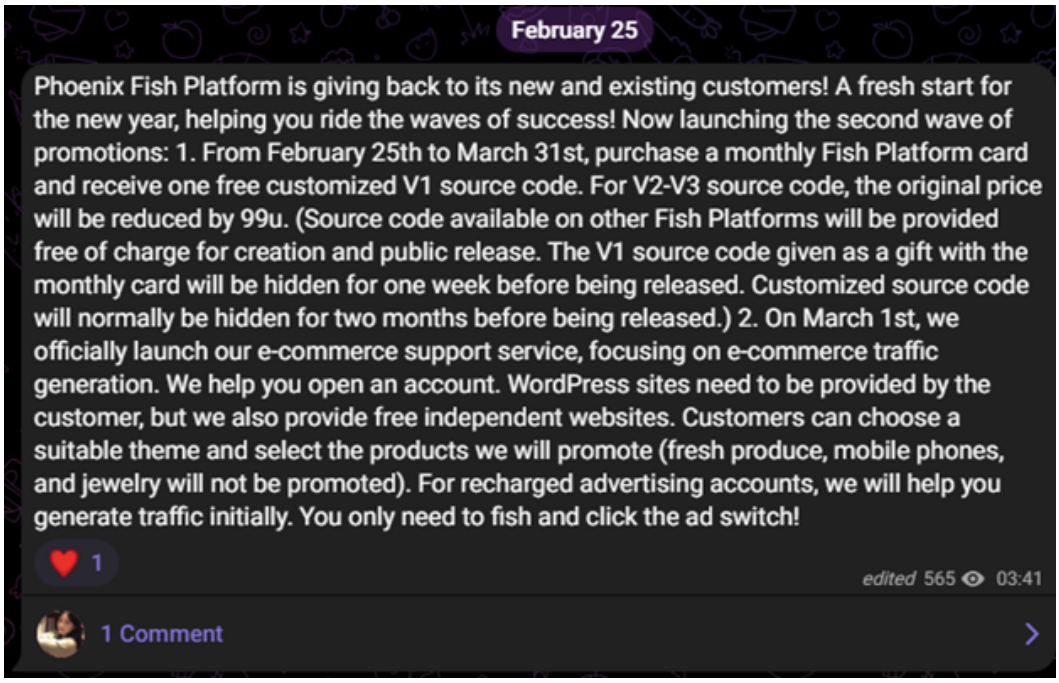

Detail: New Year greeting post framing the underground platform as a professional brand

Complete fraud chain for sale

That expansion helps explain the tone of new year offers inside the underground market. The promotions are not only trying to sell access to a phishing panel. They are often selling participation in a wider fraud economy: tools, websites, support, traffic, and ways to turn stolen data into cash. In that sense, Chinese New Year becomes more than a holiday reference. It becomes a business season.

There is a strategic lesson here for threat managers. Defenders often focus on the phishing lure itself: the message, the URL, the domain, the delivery method. But the bigger story is the market that sits behind it. A service-driven ecosystem can absorb takedowns, swap infrastructure, refresh branding, and onboard new customers with remarkable speed. If one storefront disappears, another appears. If one kit is exposed, a newer version replaces it. The danger comes from the resilience of the commercial model as much as from the individual scam.

Why seasonal promotions matter

That is why seasonal underground promotions deserve attention. They may offer an early signal of platform refreshes, recruitment drives, service expansion, or shifts in monetisation. They show when actors are investing in growth, and they reveal how openly the ecosystem now borrows from the rhythms of legitimate commerce.

Chinese New Year, in this light, is not just a cultural backdrop. It is a mirror. It reflects an underground market that has become organised enough to think in campaigns, sell in tiers, and grow by design.

DDoSIA attacks against Denmark

Election-centred DDoS activity

Over the past six months, Denmark has faced a sustained series of distributed denial-of-service (DDoS) campaigns linked to the pro-Russian hacktivist ecosystem around NoName057(16) and its DDoSIA project.

While the attacks have generally been moderate in technical sophistication, they have been persistent, highly visible, and strategically timed around politically sensitive events.

The clearest pattern is that DDoSIA activity against Denmark has clustered around elections and election-related political tension. The first major concentration occurred around the November 2025 Danish municipal and regional elections, when political party websites, public-sector portals, and high-visibility national services were targeted. Activity then intensified again in late February and early March 2026, following the announcement of Denmark's snap general election on 24 March 2026, with attacks expanding to include ministries, municipalities, transport, media, universities, and Greenlandic entities.

From a management perspective, these incidents are best understood not as attempts to cause lasting technical damage, but as hybrid influence operations designed to create disruption, attract media attention, and undermine public confidence at moments of democratic sensitivity.

The operational impact has so far remained limited to temporary website outages and service degradation, and there is no indication that voting processes themselves were affected. However, the repeated timing of attacks around elections materially increases their strategic significance.

The overall risk profile is therefore moderate from an operational standpoint, but high from a reputational, political, and resilience perspective. For organisations operating in Denmark, especially those with public-facing services, public-sector relationships or critical infrastructure exposure, DDoSIA should be treated as a persistent threat actor whose campaigns are likely to continue whenever Denmark is involved in high-profile geopolitical events.

Strategic context

DDoSIA is associated with the pro-Russian hacktivist collective NoName057(16), which has built a volunteer-driven model for carrying out disruptive attacks against countries perceived as supporting Ukraine. Its campaigns rely on broad participation, public target lists, and repeated claims of responsibility through Telegram and related channels.

This is important because it places the Denmark activity in a wider pattern of hybrid coercion, where technical effects are often less important than symbolic and informational effects. In this model, the desired outcome is not necessarily deep technical compromise. Instead, the aim is to embarrass public institutions, show that disruption is possible, generate media coverage, reinforce geopolitical messaging, and create a perception of instability during politically sensitive moments.

Denmark's profile makes it a natural target in this context. It is a strong supporter of Ukraine, a NATO member, and a state with both domestic political visibility and international significance due to Greenland and the Arctic.

Focus on elections

The strongest conclusion from the past six months of monitoring DDoSIA is that elections have acted as focal points for DDoSIA targeting.

1. November 2025 municipal and regional elections

The first major concentration of activity occurred around the November 2025 Danish municipal and regional elections. CSIS data and OSINT reporting indicates that DDoSIA waves in November targeted political party websites, municipal websites, public-service platforms,

Parliament-related sites, transport services and other nationally visible institutions.

This timing matters. Even though Danish voting is conducted manually and was not disrupted, the attacks were well suited to generating attention on election eve and during a period of heightened public scrutiny. The target mix suggests an effort to affect the wider information environment around the election rather than the ballot process itself.

The November wave also established an early pattern: political parties were included for symbolic impact, government and municipal portals were included for public visibility, and transport and public-service sites were included to maximise inconvenience and media interest.

Campaign overview

- DDoSIA attacks against Denmark in the past six months have been strongly centred on elections and election-related political narratives.
- The attacks are not random nuisance activity. They are part of a broader pattern of Russian-aligned hybrid campaign.
- Threat actors are applying pressure that seeks to exploit moments of public attention, especially elections, to maximise symbolic effect.

2. February–March 2026 snap general election period

A second and more sustained phase emerged in late February 2026, just after Prime Minister Mette Frederiksen announced a snap general election for 24 March 2026. CSIS data from end of February shows renewed targeting of Danish institutions and political entities, indicating a major surge in volume from 23 February to 1 March 2026.

This wave broadened beyond political party sites to include ministries, municipalities, courts and prosecution-related services, transport and maritime entities, media outlets, universities, energy providers, and Greenlandic government and commercial websites.

This expansion is consistent with a campaign designed to shape the broader climate around an election rather than to focus only on electoral institutions. In effect, the attackers appear to be targeting the wider ecosystem of state visibility, public trust, and national resilience during an election period.

Greenland as pressure point

The early 2026 wave also stands out because of the heavy targeting of Greenland alongside Denmark. This appears strategically relevant rather than incidental. Greenland had become central to heightened geopolitical debate and domestic political positioning, and it featured prominently in the broader context around the March 2026 election.

The inclusion of Greenlandic public and commercial domains suggests that DDoSIA was not only exploiting the election window, but also aligning its activity with the Arctic sovereignty narrative. That makes the campaign more than an anti-Denmark nuisance effort; it becomes part of a wider attempt to pressure Denmark at a moment when domestic politics and international geopolitical symbolism overlapped.

Threat actors' objectives

Political signalling

The actor is using cyber disruption to reinforce a pro-Russian geopolitical message: Denmark is being punished for support to Ukraine and for alignment with Western policies.

Election-period amplification

By concentrating activity around elections, the attacker increases the chance of media coverage, public discussion, and political sensitivity. The impact of even minor outages becomes greater when institutions are under election scrutiny.

Erosion of trust

Temporary outages of government, transport, or political websites can create a cumulative impression that public institutions are vulnerable or unreliable, even when real-world services continue functioning.

Demonstration of persistence

Repeated waves against overlapping targets show that the actor can return repeatedly. That persistence is itself part of the message: disruption can be sustained at low cost over a prolonged period.

Targeted sectors

- Government and citizen services: ministries, municipalities, courts, prosecution services, and citizen-service portals.
- Political parties: party sites were repeatedly included, especially around election-linked periods.
- Transport and logistics: rail, ports, aviation-related websites, road tolling, and postal/logistics assets.
- Media & education: attractive because disruption affects public information flows and can generate further coverage of the attacks themselves.
- Greenland: locally relevant public and commercial services, transportation and critical infrastructure.

Impact assessment

To date, the practical operational impact appears to have been limited. The attacks have primarily caused temporary website outages, intermittent service degradation and reputational inconvenience. There is no evidence that DDoSIA compromised systems or altered election processes. Voting in Denmark remained unaffected.

The strategic impact is more significant. Because the campaigns are timed around elections and public political tension, they draw disproportionate attention, create avoidable pressure on public institutions, raise public concern about foreign interference and contribute to an atmosphere of contested democratic legitimacy.

This is why the risk should not be dismissed simply because the underlying technique is "just DDoS." In the context of elections, low-complexity disruption can still produce meaningful strategic effects.

Finally, organisations hit during these periods face a heightened reputational burden. A short outage to a municipal website, ministry page, broadcaster, or party domain may be technically minor, but it can still be interpreted publicly as evidence of weak preparedness or institutional fragility.

Outlook and recommendations

The past six months of DDoSIA activity against Denmark show a clear and meaningful pattern: the attacks have been centred on elections and election-adjacent political moments. The November 2025 municipal and regional elections marked the first major concentration.

The near-term outlook remains elevated, especially through the 24 March 2026 general election period. Based on the pattern observed, further DDoSIA activity is likely because Denmark and Greenland have become a priority target in recent DDoS attack distribution.

Implications for management

- Election periods create a significantly heightened risk window for public-facing disruption.
- Website outages may have limited direct operational impact, but can still damage trust, create media amplification, and support hostile narratives.
- Political parties, ministries, municipalities, media, transport, and citizen-service platforms are especially exposed.
- Greenland has emerged as an additional pressure point in early 2026, linking cyber activity to broader Arctic sovereignty and election narratives.

The actor has demonstrated persistence and willingness to update target lists frequently around elections, aiming to create conditions for symbolic disruption. Even if attack sophistication does not increase, the timing alone can preserve strategic effect. Organisations should therefore plan on the assumption that additional waves may occur before, during, and shortly after the election period.

Any organisation with public-facing services should move to an elevated monitoring and response posture during elections and major geopolitical events, both internal and external. From a technical perspective, priority should be given to protect high-visibility web services and any citizen-facing services likely to attract public attention.

Because reputational effect is central to these campaigns, organisations should prepare not only traffic mitigation but also clear external messaging for public outreach.

DDoSIA itself has mainly produced website disruption, but the broader Russia-linked threat picture includes other actors and potentially more damaging forms of cyber activity. DDoS campaigns should be treated as part of a larger hostile operating environment.

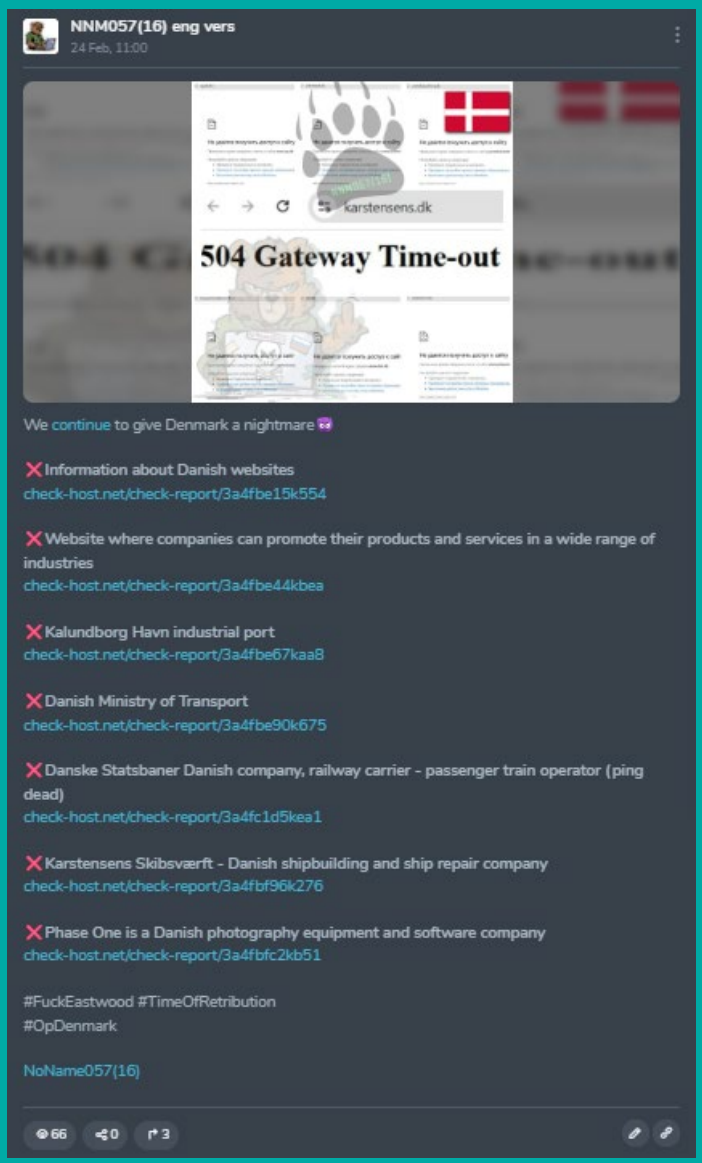

Detail: Hacktivist post from NoName057(16) showing claims of disruption activity targeting Denmark

Hybrid escalation, espionage push, and strained alliances driving cyber risk

From July to December 2025, Europe's security environment was dominated by sustained Russian hybrid pressure and persistent Chinese espionage risk. Russia combined cyberespionage, disinformation, hacktivism, and deniable drone incursions to probe red lines, disrupt infrastructure, and erode public trust. High-level warnings that Russia could threaten NATO militarily within five years, alongside major EU financial support for Ukraine, are likely to sustain Russian initial-access, espionage, and influence operations against European governments, critical infrastructure, and financial nodes.

From a Chinese perspective, European "tech sovereignty" measures—such as seizing a Chinese-owned semiconductor firm in the Netherlands and scrutinising Huawei's access to sensitive data in Spain—have raised the likelihood of Chinese retaliation and increased the likelihood of cyberespionage for IP. This likelihood has already manifested throughout the period with China observed targeting government and advanced technology sectors.

In the Americas, US domestic policy volatility and an increasingly aggressive counter-narcotics posture were the main drivers of cyber and geopolitical risk in the hemisphere, with knock-on effects for alliances. The 'Big Beautiful Bill' accelerated AI, automation, and offensive cyber capability but expanded the attack surface and supply-chain dependencies.

Diplomatic frictions—visible in the inconclusive Trump–Putin summit and US–India disputes over tariffs and Russian oil—strained trust, threatened defence and cyber cooperation, and offered Russia and China openings to widen rifts. In the US tighter skilled-immigration measures and an October government shutdown weakened cyber capacity and vulnerability management, just as US–Mexico tensions and intensified kinetic action against drug traffickers and Venezuela incentivised cartels to adopt GNSS spoofing, cyber surveillance, and potentially direct attacks on US agencies.

By year end, allied concerns over the legality of US counter-narcotics strikes led to reduced intelligence sharing, whilst Washington's mixed isolationist rhetoric and ongoing commitments encouraged greater European "digital sovereignty" moves.

Across MEA the threat environment was defined by a steady baseline of Israel–Hamas-related espionage and hacktivism, intermittent but unpredictable Iran-linked retaliation risk, and increasingly sophisticated influence operations across Africa. Exposure for private-sector and "military-adjacent" technology targets rose following controversy over cloud support to Israeli military surveillance, raising the likelihood that Iranian or proxy actors would treat major tech providers and host countries as legitimate targets.

Regional security shifts—including a Saudi–Pakistan defence pact and a US security guarantee for Qatar—had longer-term implications for offensive capabilities and espionage focused on Gulf diplomacy, AI, and technology investment. Conflict theatres in Gaza, Sudan, and Yemen remained focal points for hacktivism and proxy competition, with risks of spillover to UAE-linked entities, shipping routes, global supply chains, and Gulf digital infrastructure.

Asia-Pacific risk was dominated by China's persistent cyber-espionage, tightening information controls, and strategic competition over AI, rare earths, and critical supply chains, alongside mounting maritime and political tensions. Beijing's actions—from the BRICS AI governance push, temporary blocking of outbound internet traffic, and expanded rare earth export controls to the "Created in China" innovation drive—reinforced a trajectory toward tech nationalisation, fragmented standards, and greater espionage against advanced STEM sectors and semiconductor ecosystems.

Regional flashpoints, including South China Sea coercion, Japan's political instability, nuclear-policy debates in Japan and South Korea, and Taiwan-related frictions, were assessed as triggers for intensified Chinese (and North Korean) cyber operations and influence activity, while measures such as an exit ban on a Wells Fargo employee heightened corporate concern over travel, device seizure, and data exposure.

India's new data protection regime tightened privacy obligations whilst its growing role as a technology hub drew increased Russian and Chinese intelligence interest. Australia, meanwhile, both bolstered regional resilience through Indo-Pacific cyber capacity-building and highlighted the challenges of emerging online safety regulation via its under-16 social media ban, exposing enforcement and circumvention issues which are likely to occur elsewhere.

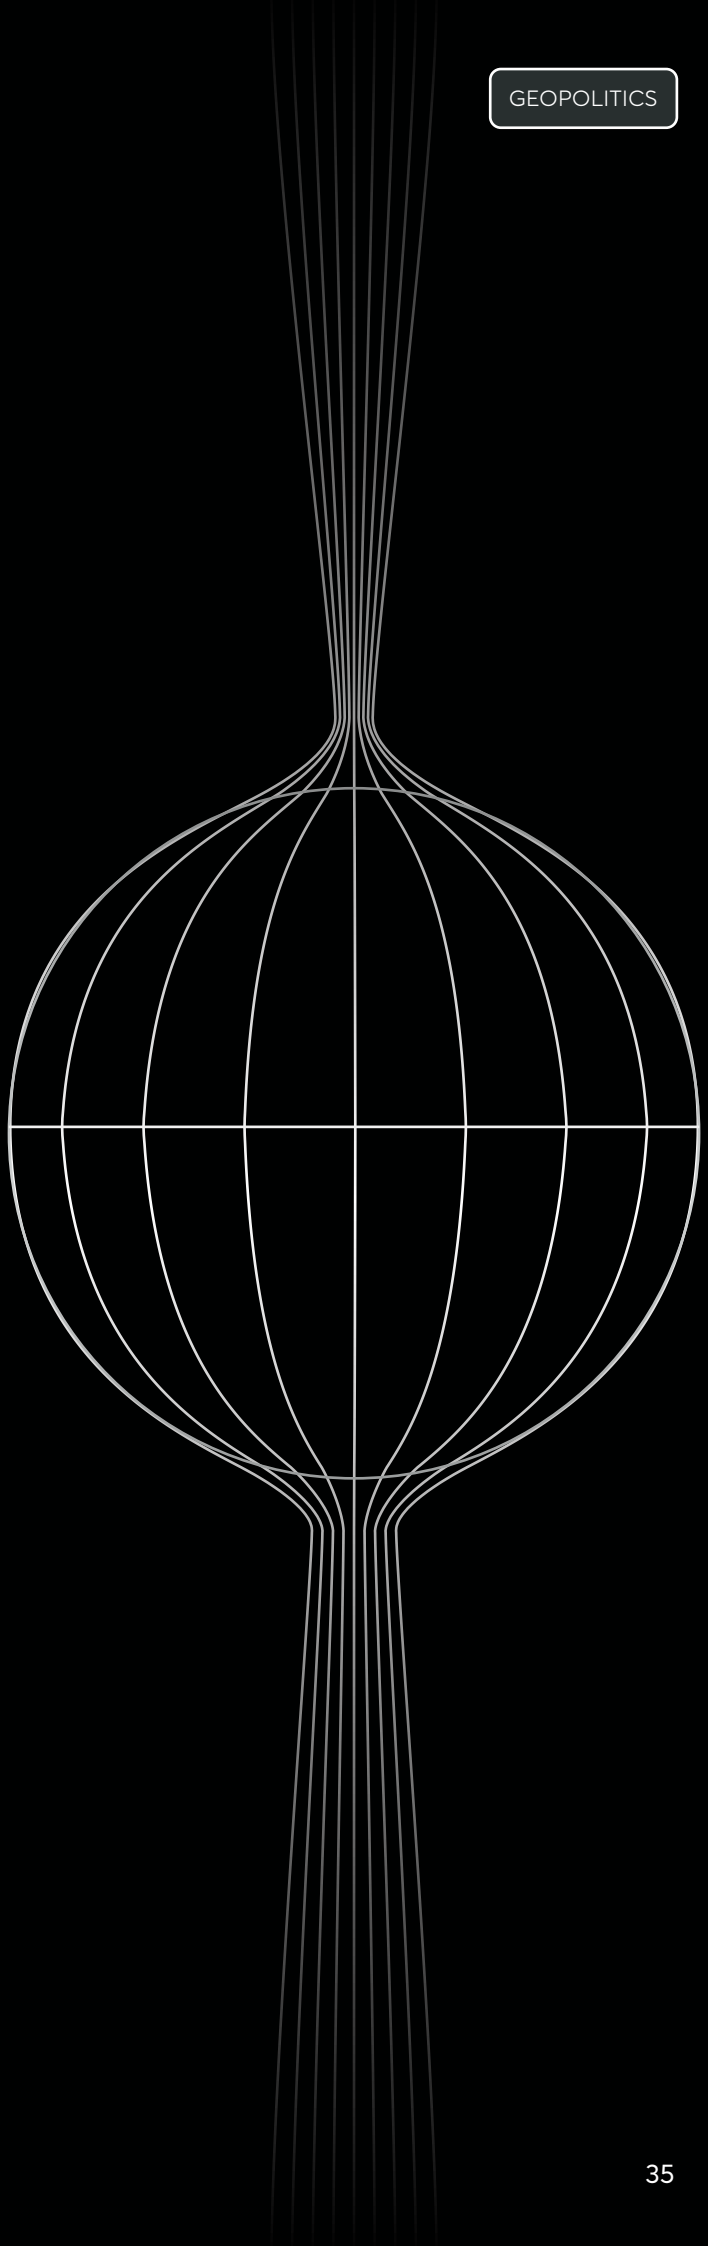

Europe

From July to December 2025, Europe's security environment was dominated by sustained Russian multi-domain pressure and parallel Chinese espionage risks. In July, the UK–Germany friendship treaty formalised deeper cooperation on hybrid resilience, intelligence coordination, and cyber/emerging tech; UK sanctions against GRU units underscored that Russian cyberespionage, information operations, and sabotage-linked activity were continuing and unlikely to be deterred. Also in July, reporting that Spain contracted Huawei to manage and store wiretap data raised allied concern about sensitive-access risk and the realistic possibility of intelligence-sharing friction, even if extensive exploitation via the contract was assessed as comparatively unlikely given safeguards.

In August–September, attention centred on Russia's interference and escalation dynamics on Europe's eastern periphery. Ahead of Moldova's 28 September elections, reported influence activity followed an established playbook (disinformation, political financing/proxies, disruption); despite PAS retaining power, the period featured pro-Russian hacktivist DDoS and reinforced expectations of replication elsewhere. Repeated Russian drone incursions across NATO's eastern flank in September signalled ongoing testing of red lines; deniable drone tactics were assessed as likely to persist and expand, complementing cyber-enabled influence, espionage, and CNI prepositioning. The September UK–US tech prosperity deal promised major AI/quantum investment but increased the attractiveness of UK ecosystems for state-led industrial espionage, particularly by China, and raised concerns that lighter-touch regulation could elevate security risk.

October clarified China-related tensions. In the UK, the collapse of a prosecution involving alleged China-linked spying highlighted enduring Chinese interest in political/policy intelligence and reinforced expectations that HUMINT is paired with cyber-enabled espionage and prepositioning across government and policy-adjacent sectors. In the Netherlands, seizure of control of Nexperia intensified "tech sovereignty," introducing short-term disruption risks and increasing the likelihood of Chinese retaliatory measures, while reinforcing that Beijing would almost certainly lean more heavily on cyberespionage for semiconductor IP if legitimate routes narrowed.

In November–December, Russia remained the primary driver of cyber/hybrid risk. A US-led draft framework for a Ukraine ceasefire did not materially reduce escalation expectations; Russia was assessed as highly

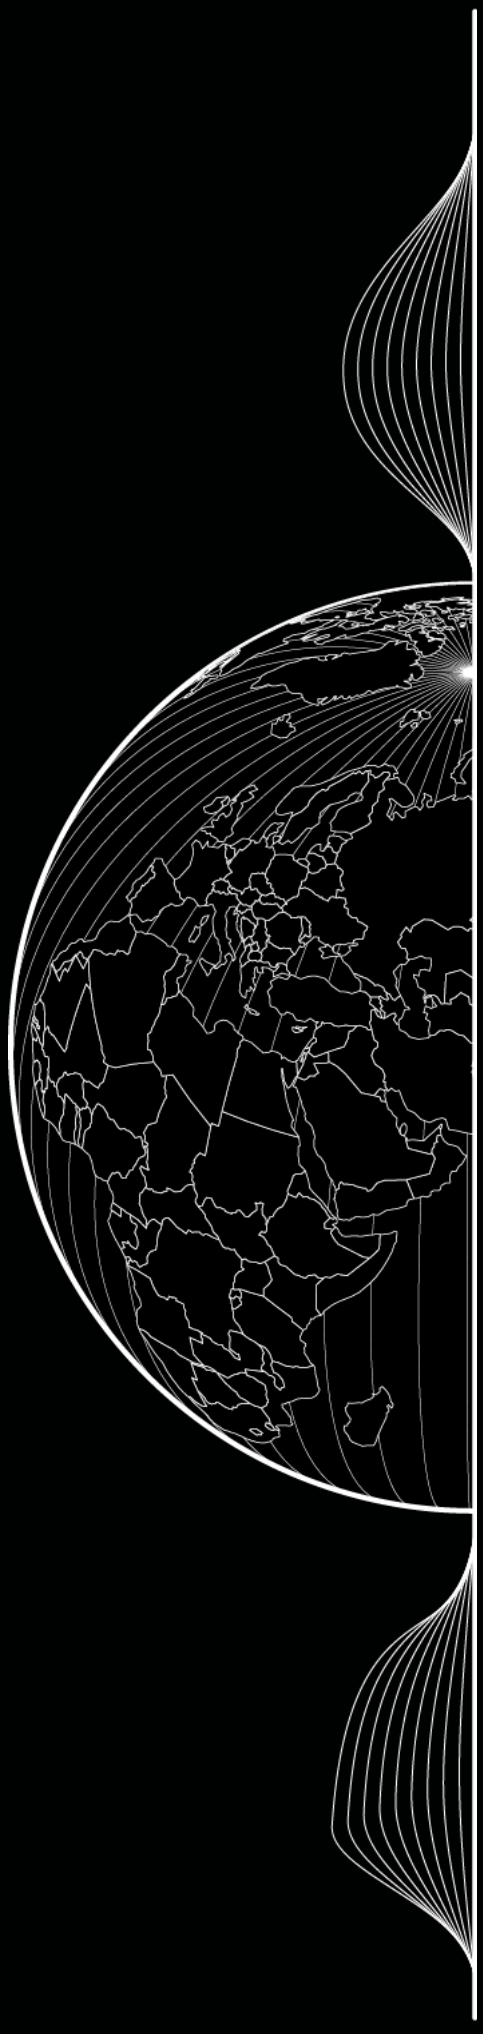

unlikely to accept constraining terms and almost certain to continue cyber operations against Ukraine while pressuring Western partners via espionage, influence, and sabotage. Russia-linked drone overflights disrupted airports and appeared near sensitive sites, assessed as serving intelligence collection and intimidation—likely to continue. The Commission's Digital Omnibus proposal raised the realistic possibility of increased personal-data exposure and additional vulnerabilities if rapid AI deployment outpaced secure-by-design controls. In December, leaders warned Russia could be positioned to use force against NATO within five years; this aligned with assessments that Russian initial-access, espionage, and CNI foothold-building are highly likely to persist, while destructive cyber activity remains unlikely outside direct conflict. The EU's EUR 90bn Ukraine loan was assessed as likely to drive Russian initial access and espionage against EU bodies and financial-sector nodes (including Euroclear), alongside continued influence pressure and intimidation.

Russia as hybrid cyber driver

Sustained cyberespionage, disinformation, hacktivism, and "deniable" actions (e.g., drone incursions) aimed at probing NATO/EU red lines, undermining trust, and prepositioning in critical infrastructure and government networks—likely to persist as Ukraine support continues.

The Americas

Across H2 2025, US domestic policy volatility and an expanding counter-narcotics posture were central drivers of cyber and geopolitical risk, with secondary effects on alliance cohesion. July's "One Big Beautiful Bill" accelerated AI/automation across government and military functions and expanded funding for offensive cyber operations; while potentially improving capability, it was assessed as expanding the targetable attack surface and supply-chain dependency risks, especially if disruption elsewhere reduced resilience. The July decision to supply weapons to Ukraine via NATO increased expectations of pro Russian hacktivist targeting of US government, defence, manufacturing, logistics, and CNI, alongside likely Russian espionage.

The August Trump–Putin summit in Alaska ended without concrete agreements, reinforcing allied concerns about US concessionary signalling and straining transatlantic trust, increasing the likelihood of allied caution in intelligence sharing and alignment. US–India tensions over tariffs and Russian oil purchases were assessed as creating a realistic possibility of undermining defence and cyber cooperation, offering Russia and China opportunities to widen the rift. September moves affecting skilled immigration (notably the H-1B fee) were assessed as likely to exacerbate specialist staffing constraints. The October US government shutdown materially increased cyber risk by reducing federal capacity (including CISA), slowing vulnerability handling/enforcement, and undermining information-sharing, creating greater opportunity for cybercriminal and nation-state intrusion.

US–Mexico tensions remained relevant: while both likely used cyber-enabled measures in counter-cartel activity, friction risked limiting joint effectiveness, and cartels were assessed as likely to continue leveraging cyber capabilities for surveillance and intelligence on law enforcement. In September–October, intensified US kinetic action against suspected drug trafficking vessels introduced the realistic possibility of cartel adaptation via GNSS spoofing/jamming and even direct cyber targeting of US agencies; US cyber activity against Venezuela was assessed as likely to feature in escalation scenarios. In November, concerns that US strikes violated international law drove the UK, Colombia, and others to limit intelligence-sharing with Washington, raising the realistic possibility of reduced US visibility and greater unilateral collection, while creating opportunities for traffickers to exploit gaps.

Counter-narcotics escalation

Intensified US action against cartels and Venezuela is driving adversary adaptation (GNSS spoofing/jamming, cyber surveillance) and raises the possibility of direct cyber targeting of US agencies and broader offensive cyber use alongside kinetic measures..

In December, alliance dynamics were complicated by the new National Security Strategy's isolationist themes and criticism of Europe, contrasting with NDAA commitments (including support for Ukraine and a sustained troop posture in Europe); this inconsistency was assessed as likely to accelerate European "digital sovereignty" and may possibly limit future cooperation. US escalation against Venezuela broadened through naval blockade activity, tanker seizures, and expanded legal framing (including fentanyl as a WMD and cartel terror designations), increasing the realistic possibility of offensive cyber operations alongside kinetic pressure. Russia and China were assessed as likely to expand cyberespionage in the region to support decision-making and potentially undermine the US, while overt disruptive operations remained assessed as unlikely.

Middle East and Africa

Throughout the period, MEA risk featured a steady baseline of Israel–Hamas-related espionage and hacktivism, intermittent Iran-linked retaliation risk, and growing influence-operations sophistication across Africa. In July, pressure on Israel intensified and domestic politics became more fragile; while patterns did not fundamentally change, any move toward Israeli elections was assessed as almost certain to attract heightened influence operations and intelligence activity by adversaries such as Iran. The UK assessment of Iran—emphasising cyber operations and response unpredictability—reinforced that shocks could trigger disproportionate disruptive activity, particularly against Western states and Iran-adversary-linked entities.

August developments increased exposure of private sector and diaspora-linked targets. Revelations about Microsoft's customised Azure services for Israeli military surveillance triggered protests and raised the realistic possibility Iranian actors/proxies could increase targeting of large tech providers and host countries, with concern "military-adjacent" infrastructure could be reframed as legitimate targets. Australia's attribution of antisemitic attacks to Iran's IRGC via proxies and organised crime underlined Tehran's reliance on deniable networks abroad and suggested continued targeting of dissidents and Jewish/Israeli-linked entities outside the region.

September's snapback reimposition of UN nuclear-related sanctions on Iran introduced a key escalation driver; analysts assessed a realistic possibility of Iranian retaliatory cyber-attacks—potentially disruptive—against European and US entities and possibly critical infrastructure. A Saudi Arabia–Pakistan mutual defence pact, including technology transfer, introduced longer-horizon cyber implications, creating a realistic possibility of improved offensive cyber capabilities for both states and elevating risk for rivals. October's Israel–Hamas ceasefire paused hostilities but was assessed as unlikely to produce durable settlement, so cyberespionage and hacktivism were expected to continue largely unchanged. A separate US security guarantee for Qatar increased its intelligence value, with Iran, Russia, and China assessed as highly likely to prioritise collection on negotiations and investment flows, particularly given Gulf AI ambitions.

In November, Sudan's civil war remained a locus for hacktivism focusing on UAE-linked entities, with limited spillover (e.g., low-level DDoS against UAE partners) a realistic possibility. The fragile Gaza ceasefire and UN-backed plan did not shift baseline expectations, though Western policy flashpoints were assessed as capable of

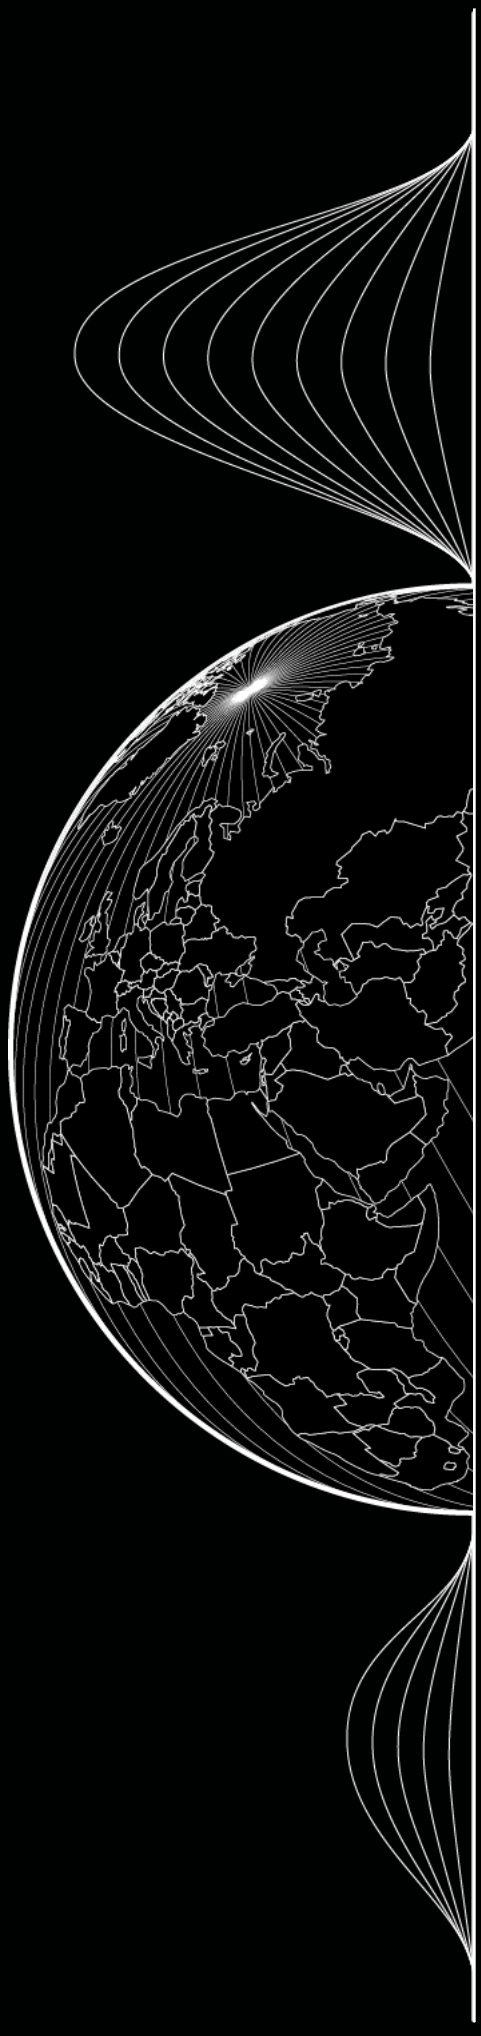

prompting temporary surges in disruptive hacktivism. In December, Yemen re-emerged as an escalation concern after UAE-backed STC advances; while immediate cyber effects were expected to remain regionally bounded, a wider proxy dynamic could affect shipping routes and supply chains, with downstream impacts on Gulf digital infrastructure. Allegations of Israeli surveillance at the US-led Gaza coordination mechanism reinforced assessments Israel would continue collecting even against partners, creating risk for reconstruction organisations and financial intermediaries. Meta reported pro-Russian influence operations in sub-Saharan Africa using legitimate freelancers, likely increasing reach and credibility over time.

Risks are defined by espionage

Israel–Hamas and Iran-linked dynamics elevate targeting of "military-adjacent" tech providers and critical services; in Asia-Pacific, China's tightening controls, supply-chain leverage (rare earths), and regional flashpoints (South China Sea, Taiwan) sustain high espionage pressure and operational risk for multinationals.

Asia and Australasia

Across July–December 2025, Asia-Pacific risk was characterised by China's sustained cyber-espionage posture, technology and information control measures, and intensifying competition over AI and strategic supply chains, alongside political/regulatory shifts; Australia's October investment in Indo-Pacific cyber capacity contributed a resilience trend via deeper regional cooperation and intelligence sharing. Japan's political instability after an electoral setback raised the realistic possibility of Chinese influence exploiting nationalist narratives and elevated espionage risk against political and defence institutions. Corporate risk sharpened as China imposed an exit ban on a Wells Fargo employee, increasing concern about travel, device seizure, and sensitive-data exposure for multinationals, especially defence/advanced tech-linked.

South China Sea tensions—especially coercive activity near the Philippines—were assessed as likely to be mirrored in cyberspace via increased Chinese intelligence operations against the Philippines and partners. Uncertainty over US guarantees fuelled nuclear-policy debate in Japan and South Korea; any movement was assessed as almost certain to trigger intensified Chinese and North Korean cyberespionage against government, defence, and research networks. The BRICS AI governance declaration in July amplified contest over AI standards and access and increased the likelihood of IP theft against advanced tech and manufacturing across the US, Europe, and Asian partners. In August, China temporarily blocked outbound global internet access via TCP port 443, highlighting operational risk for firms reliant on offshore connectivity and signalling drift toward more segmented controls.

In September, the TikTok ownership/operations transfer in the US reduced some immediate concerns but was assessed as part of a wider "tech nationalisation" trajectory that could accelerate decoupling and fragment standards. The SCO summit's AI emphasis and a leak of Great Firewall internal data highlighted China's role in shaping surveillance-enabled governance models abroad, with the leak suggesting export and potential embedding of access mechanisms in partner infrastructure. October's Fourth Plenum discussions on the 15th Five-Year Plan reinforced expectations of heightened cyber/intelligence operations against key STEM sectors (AI, quantum, robotics, renewables) across allied Asian markets. China expanded rare earth export controls in October, underscoring enduring leverage; it was assessed as highly likely to continue intelligence-gathering against REE supply chains and adjacent semiconductor ecosystems.

China's espionage risk is rising

European restrictions/scrutiny (semiconductors, Huawei and sensitive data access) increase the likelihood of Chinese retaliation and intensified cyber-enabled IP theft, especially against advanced tech/STEM and semiconductor ecosystems.

Chinese espionage against Western targets remained persistent through November, combining cyber, HUMINT, and investment-linked pathways; LinkedIn-enabled targeting of UK parliamentary staff and intimidation of UK academic researchers reflected broad collection and influence objectives. China–Japan tensions in November over Taiwan-related statements increased expectations of cyber targeting of Japanese government/defence, alongside a realistic possibility of reciprocal activity and expanded Chinese influence to shape narratives, even as high-impact sabotage remained unlikely short term. In India, rollout of the DPDP Act tightened privacy and breach notification, creating meaningful compliance impacts and a potential model for others. Further Russia–India engagement in December and major Western investment flows elevated India's strategic value as a technology hub, increasing likelihood of Russian intelligence collection and Chinese cyber-enabled targeting for advantage. Australia's under 16 social media ban highlighted enforcement/circumvention dynamics, with migration to less-regulated platforms and "logged-out" access increasing exposure to harmful content and complicating monitoring.

The numbers behind the story

Malware trends

Infostealer top 10 C2 detection 2025

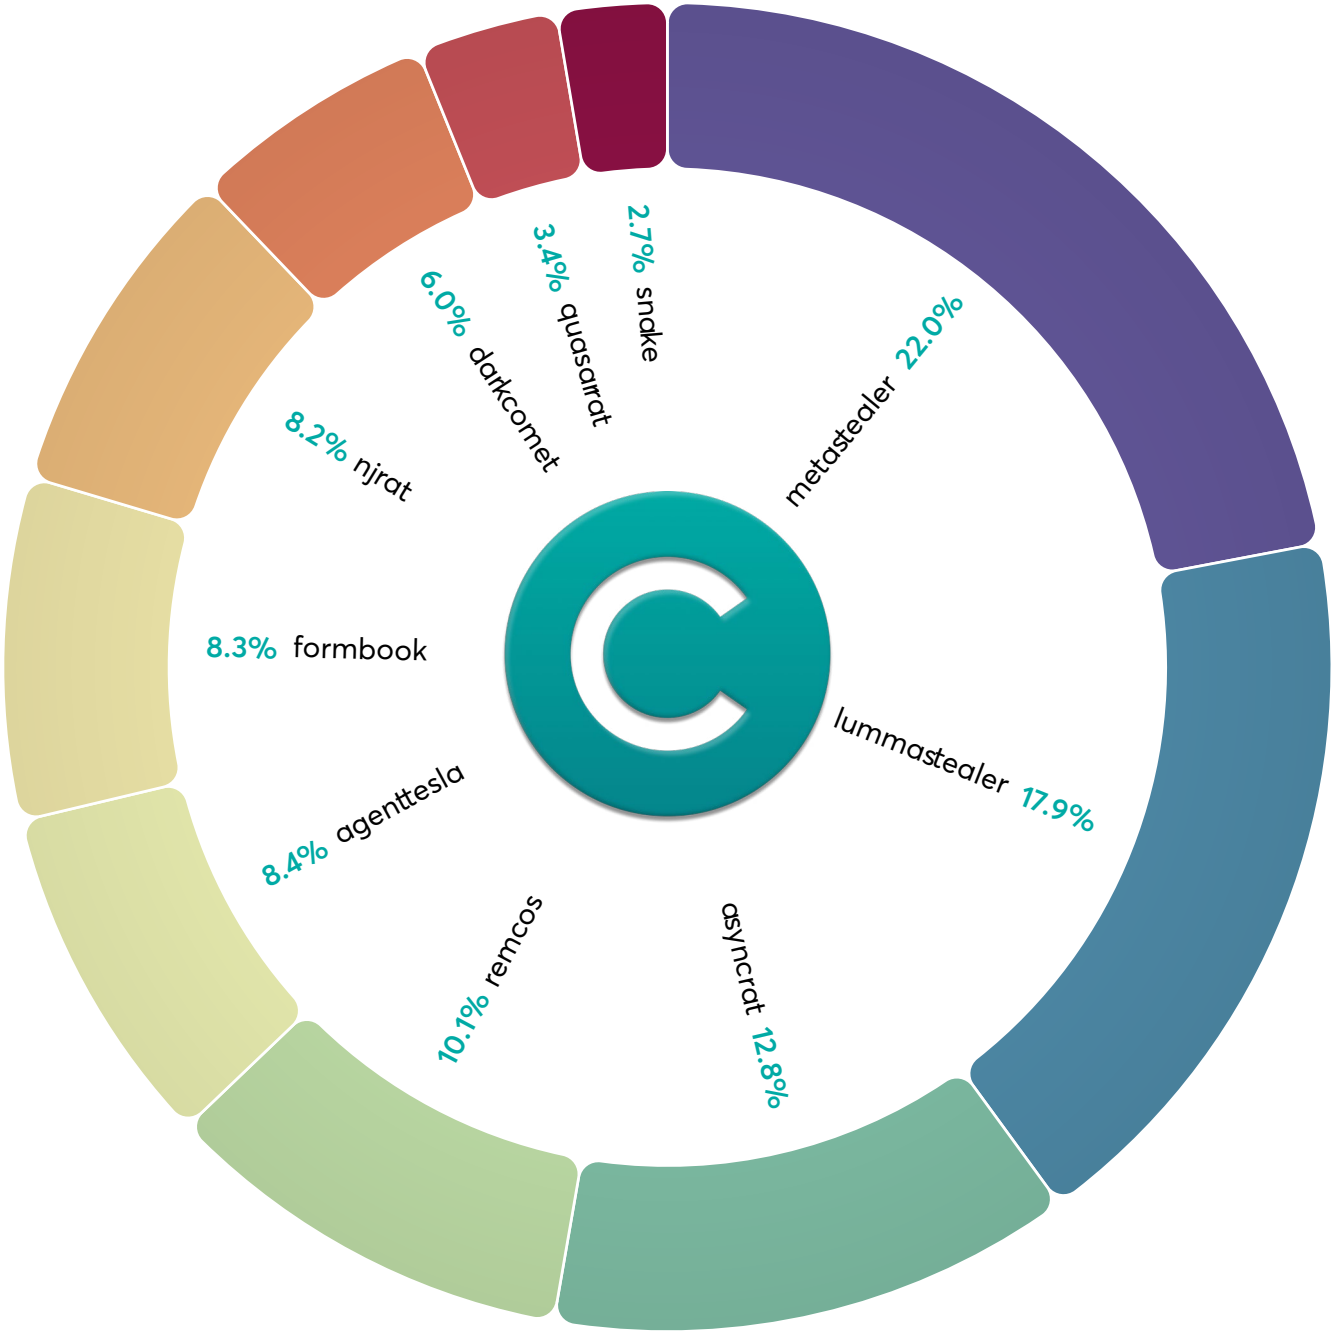

Infostealers are not merely low-level commodity threats, but a material business risk and a frequent precursor to more serious compromise. Their primary value to adversaries lies in stealing credentials, session cookies, and other access data that can be reused for account takeover, fraud, data theft, and onward intrusion activity, including ransomware.

From a threat intelligence perspective, monitoring dominant infostealer families provides an efficient way to prioritise detection, response, and control validation against threats that are most likely to translate into operational and commercial impact.

Post-exploitation C2 detection 2025

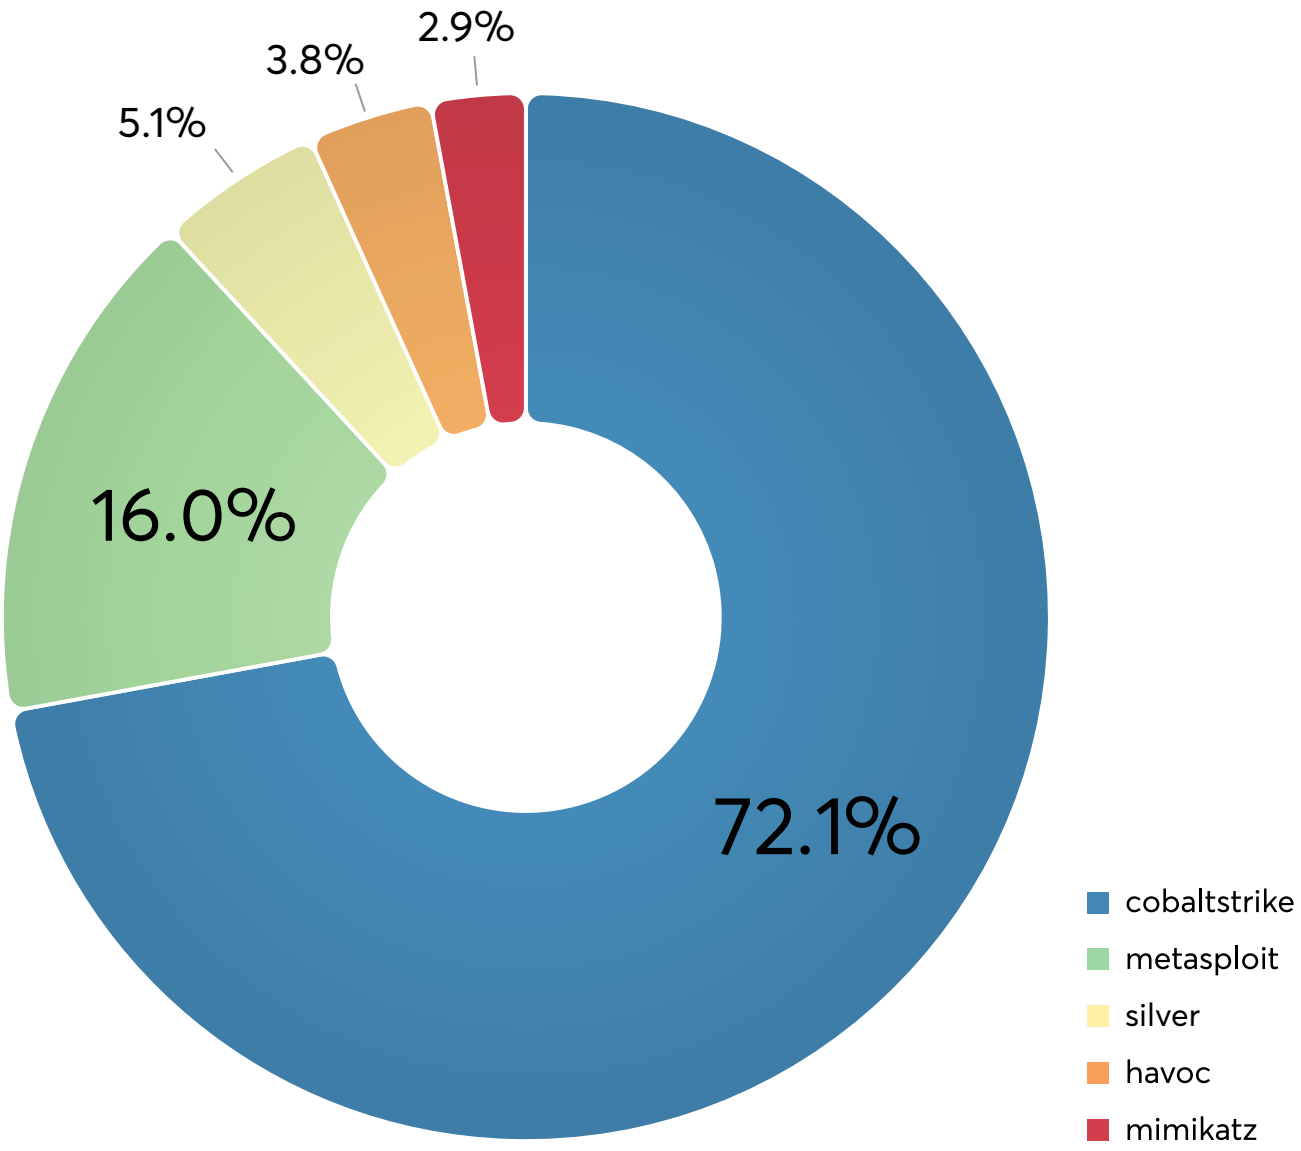

Post-compromise activity remains highly standardised: once attackers gain access, many continue to rely on a narrow set of mature, scalable tooling to maintain control, move laterally, and progress towards objectives such as data theft, extortion, or ransomware deployment. From a threat intelligence and defensive perspective, this concentration creates a practical opportunity to focus detection engineering, threat hunting, and control validation on the toolsets most likely to be used during the critical post-intrusion phase, thereby improving the organisation's ability to identify and contain serious incidents before business impact escalates.

New ransomware groups 2025

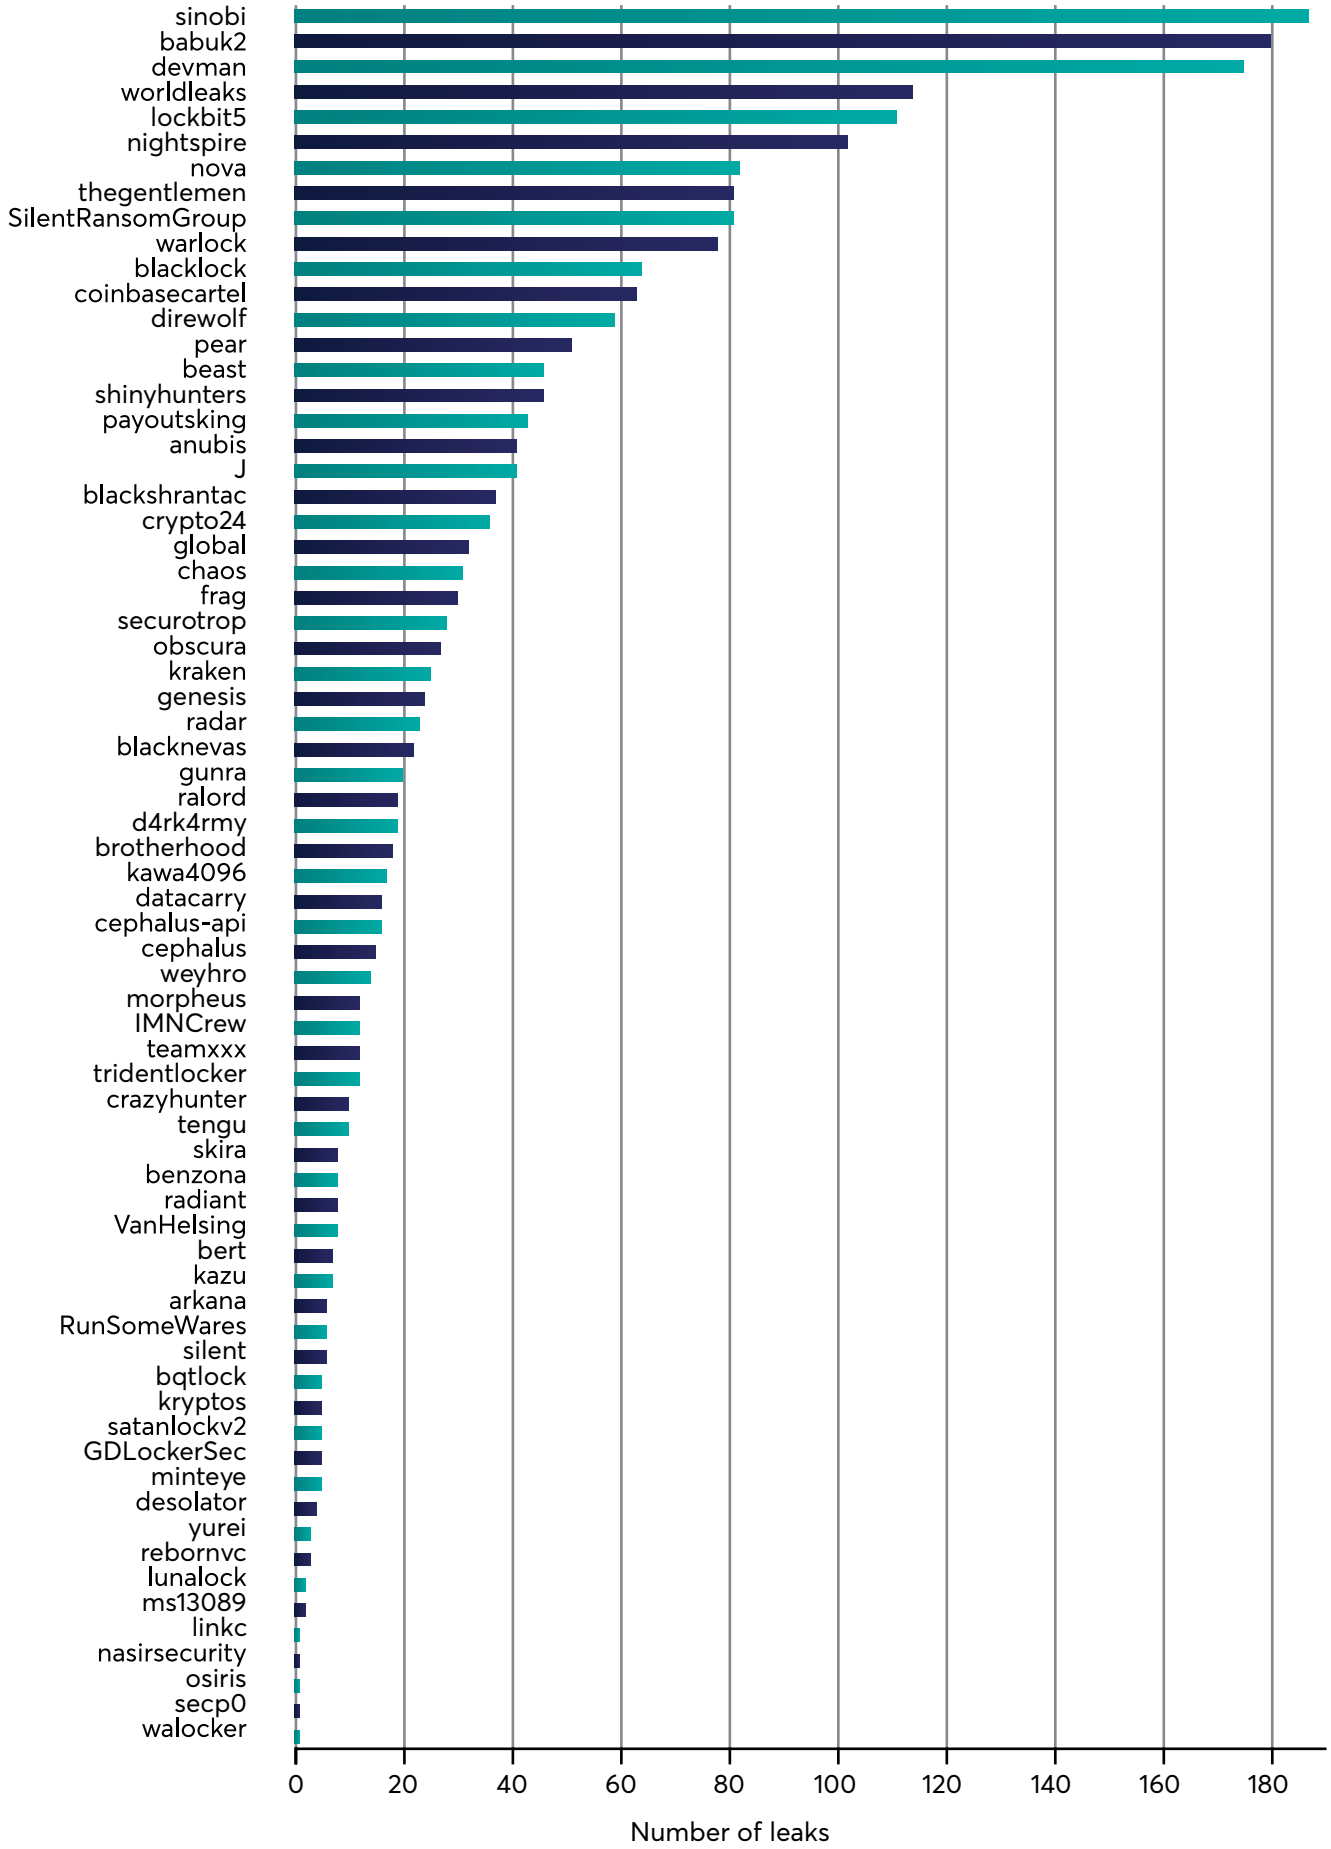

Top 10 ransomware groups 2025

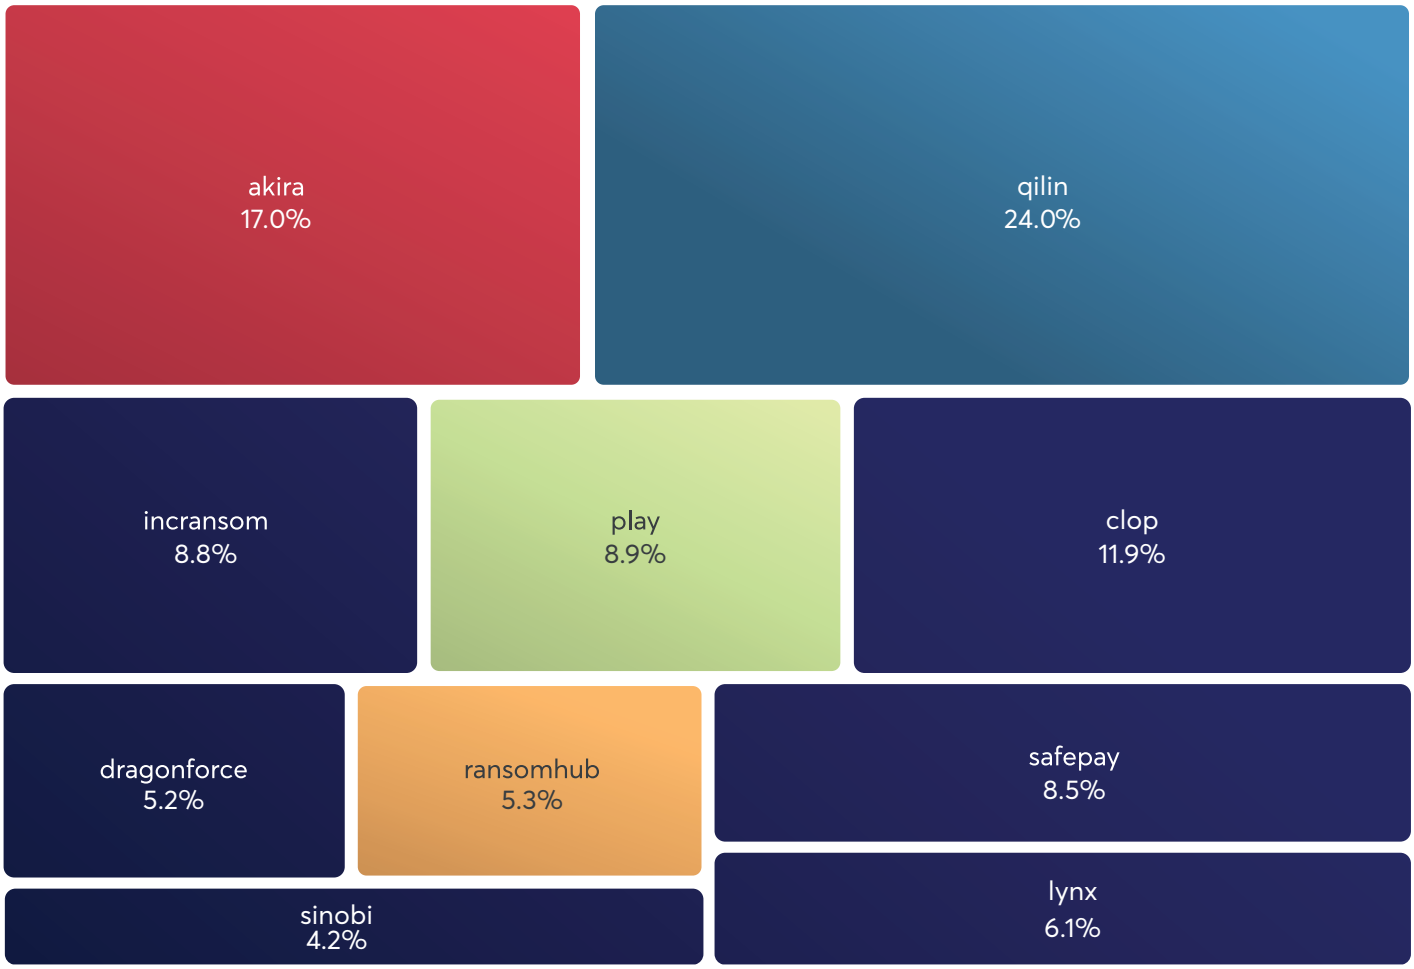

Top 10 ransomware groups 2024

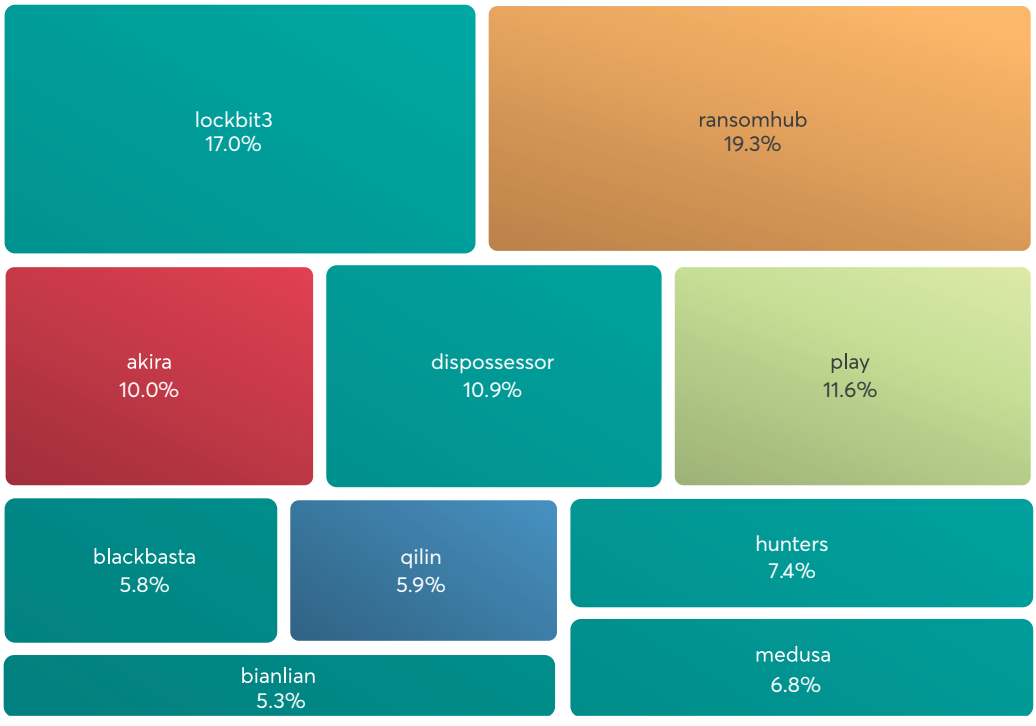

Top 10 groups by number of leaks on public leak sites. Percentages indicate how leaks are split among these 10, not the total volume of leaks.

Compromised data

One of the major threats coming from malware infections, phishing, and data breaches was the risk of sensitive information, such as payment card data and login credentials, being stolen and/or leaked by threat actors. CSIS was constantly monitoring multiple online sources to identify compromised data belonging to our customers to help them mitigate this threat.

Compromised payment cards H1 2025

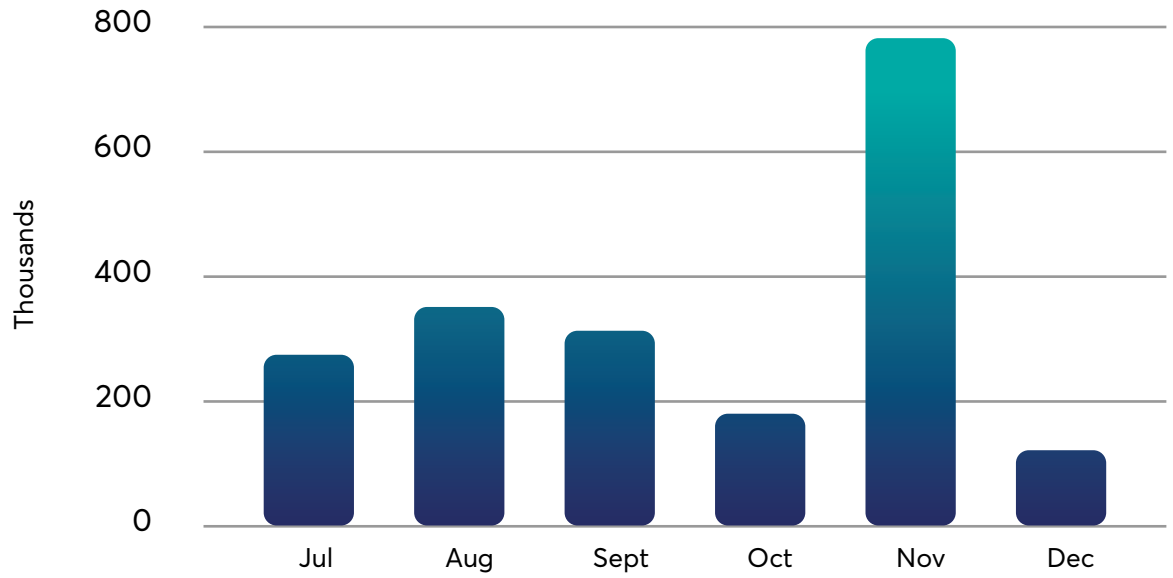

In terms of compromised credentials, the numbers were in the billions because much of the data was being reshared in so-called combo-leaks, which repackaged leaks from other sources thereby exposing the data even further to potential abuse.

Compromised credentials H1 2025

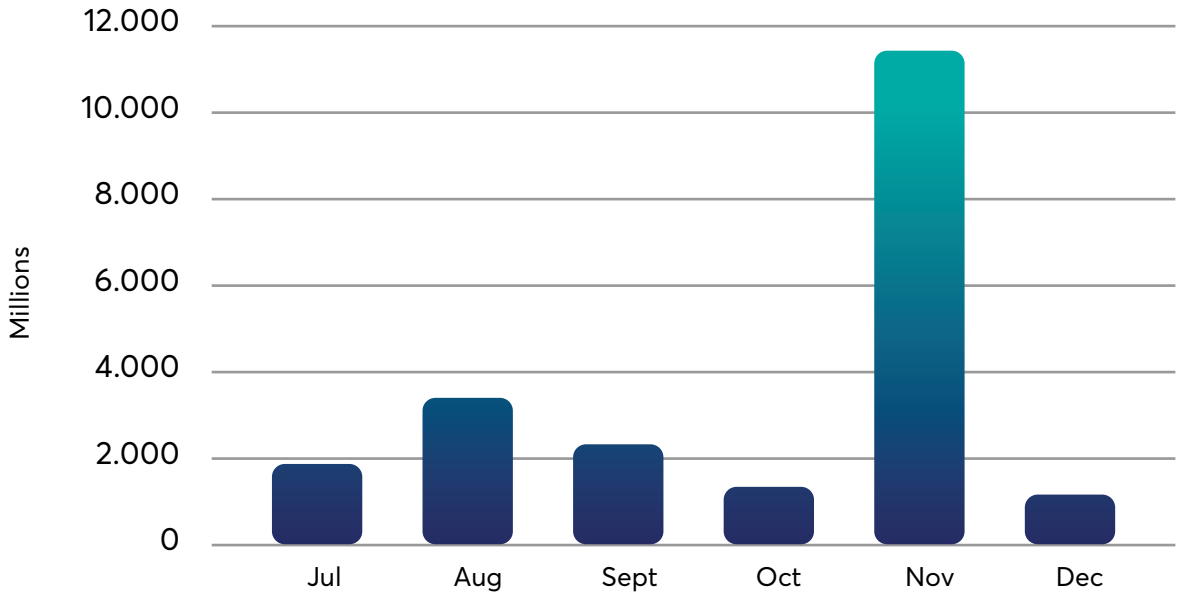

The November spike represented the adding of new major sources of compromised data.

Phishing by the numbers

Top 10 targeted phishing brands H2 2025

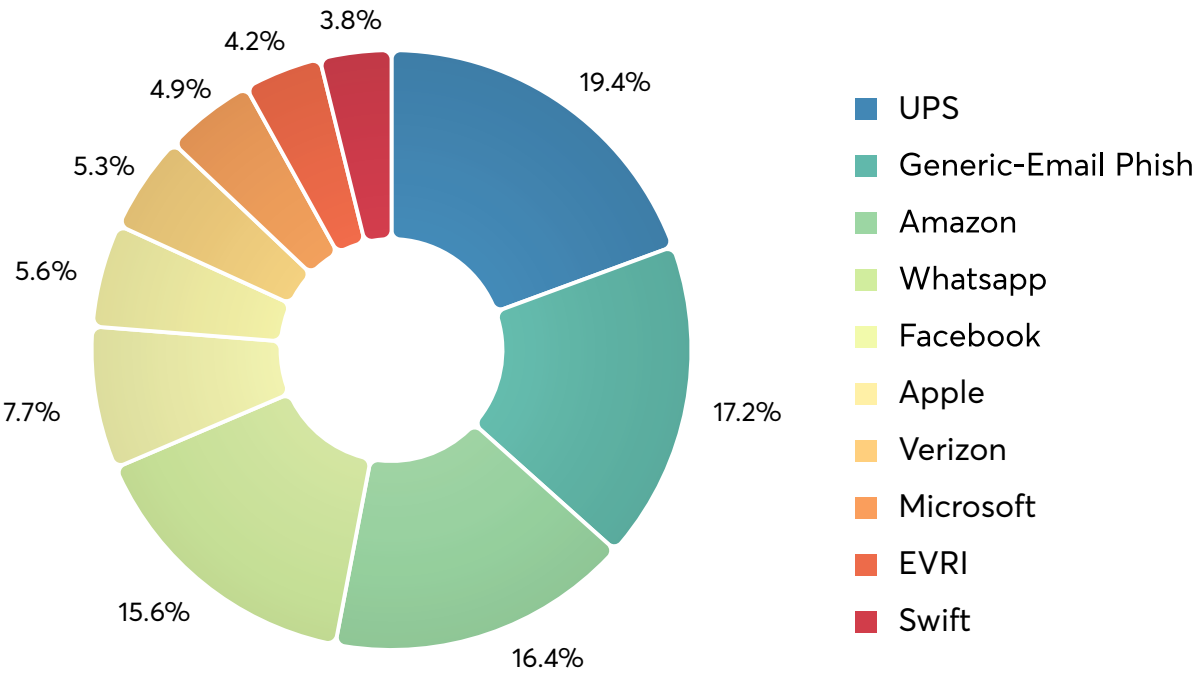

Top 10 phishing hosts H2 2025

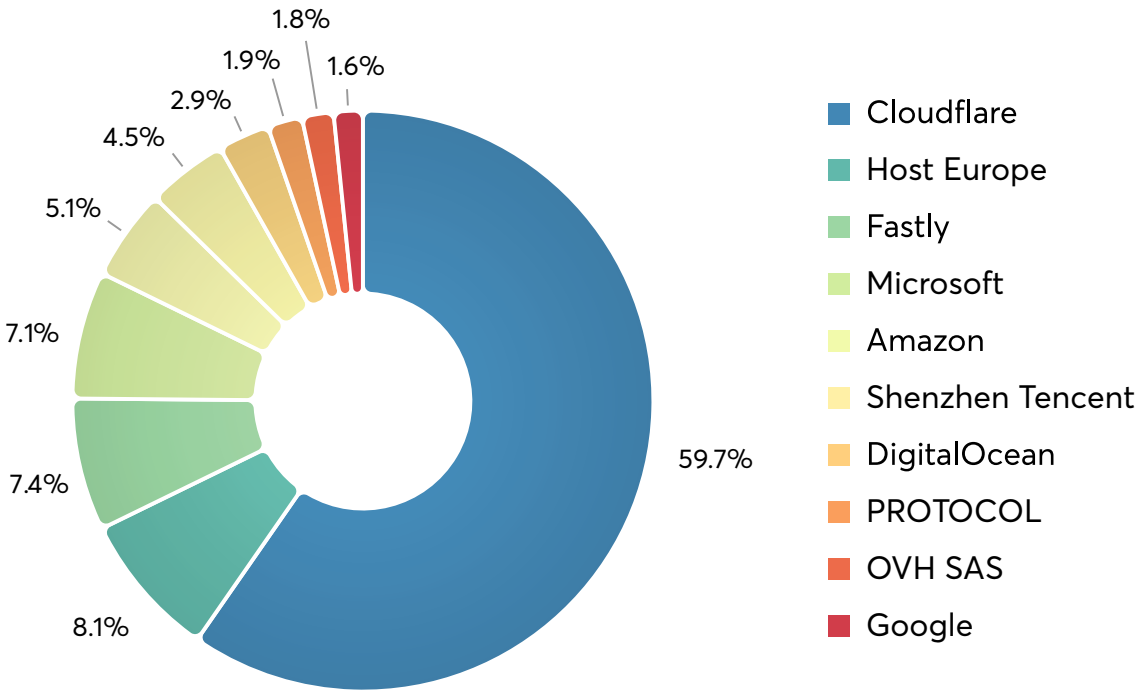

News from CSIS

General updates

Northern Cyber Alliance launched to strengthen executive resilience

The **Northern Cyber Alliance** was launched as a new Nordic partnership bringing together expertise across cybersecurity, law, crisis communication, negotiation, and strategic risk management.

The alliance is designed to support organisations and executive leadership in preparing for and navigating complex crises, including cyber incidents and hybrid threats. By combining complementary disciplines, the initiative aims to strengthen decision-making, preparedness, and resilience before, during, and after critical events.

Learn more here.

CSIS achieves ISAE 3000 Type I certification

CSIS has achieved **ISAE 3000 Type I certification**, an internationally recognised assurance standard validating the design and implementation of organisational controls.

The certification confirms that CSIS maintains structured risk management practices covering security, availability, and confidentiality. It reflects the organisation's ongoing commitment to strong governance, transparency, and trusted service delivery for customers and partners.

We regard this as a defining moment in our growth – underscoring our commitment to excellence in cybersecurity operations and enhancing stakeholder confidence in our response capabilities.

Learn more about our certifications here.

New publication explores ransomware negotiations

A new book titled "***En usynlig fjende***" (***The Invisible Enemy***) has been released, providing insight into the realities of ransomware incidents and crisis negotiations with cybercriminals.

The publication draws on real-world experiences from **Jan Kaastrup**, cybersecurity technical expert at CSIS Security Group, and **Michael Sjøberg**, crisis management expert and negotiator at Delta Crisis Management. The book examines how ransomware attacks unfold behind the scenes and highlights key considerations for organisations preparing for such incidents.

Learn more about the book here.

Cyber predictions for 2026 published

CSIS has published a new blog post featuring **Cyber Predictions for 2026**, with insights from Stefan Tanase, Jan Kaastrup, and Dean Cowlishaw.

The post explores how AI-driven automation is accelerating cyber threats and why identity has become the primary attack surface. It also examines the growing need for machine-to-machine defence and the role of threat intelligence as a strategic capability rather than a purely technical function.

Read the full blog post here.

Copyright

www.csis.com

CSIS
REST ASSURED

CSIS IN BRIEF

CSIS Security Group A/S is the leading European pure-play provider of tech-enabled cybersecurity and intelligence services. With a fully-fledged 24/7 capability, we deliver Managed Detection & Response, Incident Response, and Security Consulting services to customers across all sectors. Our acquisition of SecAlliance boosted our presence with a powerful world-class intelligence offering. Accredited and certified by various organisations, including CREST, we actively contribute to global security initiatives, ensuring that we have a positive impact on the cyber community.

CSIS IS IN THE ALLURITY FAMILY

Allurity is a group of tech-enabled cybersecurity service providers, comprised of best-in-class specialists with a common purpose to enable a safe digital world. We are the chosen growth partner for progressive entrepreneurs and talent. Together, we aim to develop the industry and support companies in reaching their full potential.

Allurity is backed by Trill Impact,
the pioneering Impact House.

OUR OFFERING

- SOC Services
- Emergency Response Services
- Digital Risk Protection
- Cyber Intelligence
- Consulting
- Specialised Services

CSIS Security Group A/S

DENMARK

Head office

Lindevangs Allé 8-12
2000 Frederiksberg

Regional office

Adelgade 115, st. tv.
8660 Skanderborg

+45 88 13 60 30
info@csis.com

UNITED KINGDOM

Head office

9 New Square
Lincoln's Inn,
London, WC2A 3QN