



Cyber Brief (September 2026)

October 2, 2026 - Version 1

TLP:CLEAR

Executive summary

- We analysed 358 open source reports for this Cyber Brief¹.
- Relating to **cyber policy and law enforcement**, Dutch police arrested a man as part of an ongoing investigation into the ShinyHunters cybercrime group, while leaders from the United States' (US) top artificial intelligence (AI) companies briefed the UN Security Council on advanced AI risks and malicious misuse.
- Regarding **cyberespionage**, Serbian pro-democracy activists and dissidents were reportedly targeted with spyware tools. Globally, we observed several China-linked campaigns. Notably, US authorities reported industrial distillation campaigns against leading US models by Chinese AI companies.
- On the **cybercrime** front, Google reported a surge in LLM-jacking, with cybercrime actors stealing and reselling access to premium AI tools and hijacking cloud servers for unauthorised workloads.
- Regarding **data exposure and leaks**, Berlin's city administration confirmed that cybercrime group Rhysida stole approximately 5.79TB of government data. Further, ShinyHunters claimed to have breached FBI systems by exploiting an alleged Oracle PeopleSoft zero-day.
- As for **artificial intelligence** activity, this month continued to see cyberattacks instigated by rogue AI agents from leading US companies. Notably, OpenAI alerted multiple organisations globally after rogue OpenAI agents had engaged with their websites beyond passive browsing, with some agents allegedly accessing both public and non-public files without altering the contents.

For more information regarding CERT-EU's analytical and operational standards to classify, assess, and prioritise malicious cyber activities, please review our Cyber Threat Intelligence Framework [here](#).

Europe

Cyber policy and law enforcement

Sweden's data privacy regulator fined Miljödata over August data breach

On September 22, Sweden's Authority for Privacy Protection (IMY) has imposed an administrative fine of 1.8 million Swedish kronor on IT provider Miljödata after finding that the company failed to maintain adequate security for personal data, following an August cyberattack

that caused disruptions nationally and impacted the personal data of 2.2. million people. [technology](#) [link](#)

Dutch citizen arrested in ShinyHunters investigation

On September 28, Dutch police confirmed the arrest of a 24-year-old man in an investigation into the ShinyHunters cybercrime group. A Dutch cybersecurity company identified the individual as their offensive security lead, who was previously convicted of data theft and extortion. His employer said it had found no evidence that he targeted the company or its clients. [law enforcement](#) [link](#)

Russian-owned Oxygen Forensics concealed control in EU law enforcement use

On September 27, journalistic outlets reported that Oxygen Forensics, a US-based digital forensics vendor allegedly owned and operated from Russia, concealed Russian control after 2022 sanctions. Prosecutors link its Russian sister firm MKO Systems to sales to the FSB. Oxygen tools were used in EU-backed projects and bought by multiple European police agencies. [law enforcement](#) [spyware](#) [link](#)

Cyberespionage & prepositioning

Serbian pro-democracy movement members targeted with Pegasus and NoviSpy spyware

On September 2, Citizen Lab and the SHARE Foundation reported that Serbian pro-democracy movement members were targeted with Pegasus and NoviSpy spyware ahead of key 2026 election cycles. Cellebrite was used to plant the NoviSpy spyware. The SHARE Foundation reported at least 14 cases of individuals in Serbia's student movement and civil society, as well as an opposition member of parliament, who recently received Apple Threat Notifications. [civil society](#) [link](#) [link](#)

Iranian-linked social-engineering campaign targeted dissidents, activists and journalists

On September 15, United Kingdom (UK) National Cyber Security Centre reported that an Iran-linked threat actor engaged in a social engineering campaign across messaging platforms targeting dissidents, activists and journalists globally, including in the UK, US, and the Netherlands. Messages aimed to trick victims into installing malicious files, enabling surveillance and theft of contacts, e-mail and social media messages. Stolen personal data may be leaked, increasing personal safety risks. [civil society](#) [iran](#) [link](#)

Israeli intelligence targeted French advisor for North Africa and the Middle East with phishing e-mail

On September 16, French magazine Challenges reported that Israel's intelligence agency targeted the phone of the advisor for North Africa and the Middle East at the Élysée's diplomatic unit. The diplomat received phishing e-mails posing as the Palestinian authority. The French Ministry of Foreign Affairs didn't provide any comments on this incident. The Élysée's diplomatic unit stated that the information contains inaccuracies. [diplomacy](#) [israel](#) [link](#)

Disruption & destruction

A Portuguese telecommunications company suffered a disruptive cyberattack

On September 14, Portuguese telecommunications operator MEO reportedly experienced a distributed denial-of-service attack by an unnamed threat actor, resulting in disruption of its services. The company did not report a data leak nor other types of intrusion. [telecommunications](#) [link](#)

Data exposure and leaks

Rhysida ransomware hack-and-leak attack against Berlin city administration

On August 31, Berlin's city administration confirmed that the Rhysida ransomware group conducted a hack-and-leak operation against two Senate departments within Berlin's state administrative network. The threat actor exfiltrated approximately 5.79TB of data, including personnel records, financial documents, credentials, and sensitive government material. The attackers demanded 30 Bitcoin, threatening to publish the stolen data. Berlin's city administration stated it would not pay the ransom. [public administration](#) [link](#)

Second medical data breach targeting Polish healthcare

On September 25, TVP World reported a second cyberattack in weeks affecting Poland's healthcare sector, involving unauthorised access to medical data by an unnamed threat actor. The incident follows a recent similar breach, indicating continued targeting of health-related organisations in Poland. The impact included potential exposure of sensitive patient information and disruption risks for affected medical services. [health](#) [link](#)

World

Cyber policy and law enforcement

An undercover Google analyst infiltrated cybercrime group TeamPCP

On September 18, WIRED reported that Google infiltrated TeamPCP, a cybercrime group responsible for widespread open-source supply-chain attacks affecting over 1,000 companies. The undercover analyst monitored stolen credentials, helped revoke them, exposed an AI-developed zero-day, and provided evidence to law enforcement. Google's intelligence, alongside disclosures by ShinyHunters, reportedly contributed to the arrests of two Australian citizens. [law enforcement](#) [link](#)

Microsoft dismantles EvilTokens infrastructure and operators arrested

On September 22, Microsoft reported that its Digital Crimes Unit, acting with court authorisation from the US District Court for the Eastern District of Virginia, seized 50 websites used to operate phishing-as-a-service platform EvilTokens and disabled more than 150 additional domains tied to its infrastructure. The action was carried out in coordination with other industry partners. Concurrently, the Metropolitan Police Service arrested two men in connection with the operation. [law enforcement](#) [link](#)

Anthropic withheld latest AI model from UK testing agency

On September 9, the Financial Times reported that Anthropic withheld its latest AI model, Claude Mythos 5.1, from the UK's AI Security Institute (AISI) prior to release, making it available only to vetted US organisations. AISI was granted access to OpenAI's latest model, Astra, ahead of its release. [artificial intelligence](#) [link](#)

US AI leaders briefed UN Security Council on global security risks

On September 24, leaders from prominent US-based AI companies, including OpenAI and Anthropic, briefed the United Nations Security Council. Organised by France, the meeting addressed international security risks of advanced AI and malicious misuse. This briefing follows unprecedented warnings from industry leaders who are now calling for a coordinated slowdown to reportedly ensure adequate global human control and safety standards. [artificial intelligence](#) [link](#)

Cyberespionage & prepositioning

China-linked industrial-scale AI model distillation against US frontier models

On September 8, US authorities reported China-based AI companies, including DeepSeek, Moonshot AI, Alibaba, among others, conducted knowledge distillation against US frontier AI models since at least late 2024 at an industrial scale to extract proprietary functionalities and capabilities. The companies used multiple obfuscated access routes to avoid single-point detection. They claimed that companies conducting industrial-scale distillation see shortened AI development timelines and reduced costs in training frontier models. [artificial intelligence](#)

[technology](#) [china](#) [link](#)

China-linked fake news sites used to deliver Chrome and Windows zero-day chain

On September 21, Volexity reported that China-linked threat actor UTA0565 used spoofed websites and phishing e-mails to lure victims into visiting attacker-controlled pages that delivered a chained Chrome and Windows zero-day compromise. The activity targeted organisations globally, including government and policy-focused entities. Multiple lookalike domains were assessed as part of the same campaign infrastructure. [technology](#) [public](#)

[administration](#) [china](#) [link](#)

China-linked WordPress global exploitation campaign stole government records

On September 21, security researchers reported a suspected Chinese-speaking threat actor conducting a global exploitation campaign. The actor compromised WordPress sites and other internet-facing systems, stole credentials, and exfiltrated sensitive data. Impact included theft of over 18,000 records from a Western government and data collection from 996 ZyXEL switches across 48 countries, affecting government and small business entities. [public administration](#)

[china](#) [link](#)

China-linked Fire Ant expanded operations targeting trusted network infrastructure

On August 27, Sygnia reported that a China-linked threat actor dubbed Fire Ant expanded its operations beyond hypervisors into trusted network infrastructure. Active since 2025 and continuing into 2026, the actor compromised edge routers, TACACS servers, and Linux hosts globally. Fire Ant harvested credentials, captured network traffic, deployed persistent backdoors, and manipulated telemetry to obscure activity, positioning itself to reach connected high-value environments, including critical infrastructure. [technology](#) [china](#) [link](#)

Cybercrime

Surge in LLM-jacking attacks throughout 2026

On September 28, Google Threat Intelligence Group reported an increase in “LLM-jacking”, where cybercrime actors steal and sell access to premium AI tools and hijack cloud servers to run unauthorised AI workloads. Darkweb markets offered heavily discounted access and “guaranteed” replacement credentials. Google also noted at least one very active China-linked threat actor abusing victim infrastructure, increasing financial and operational risk for organisations adopting AI. [technology](#) [link](#)

Passkey-themed phishing campaign targeted Microsoft 365 accounts

On September 9, Microsoft reported a global campaign, active since May 2026, in which threat actors linked to ShinyHunters, Helix, and other cybercrime groups conduct passkey and SSO-themed social engineering attacks against corporate employees. Victims were directed to adversary-in-the-middle phishing sites or manipulated into device-code authentication flows, enabling account takeover and large-scale data theft from Microsoft 365 services including SharePoint, OneDrive, and Exchange. [technology](#) [link](#)

API key theft and AI credit abuse targeting METR, non-profit evaluating frontier AI models

On September 1, METR, a non-profit evaluating frontier AI models, disclosed two security incidents. In March 2026, attackers exploited a fail-open authentication vulnerability in a researcher's publicly exposed EC2 instance, extracting an API key via agent prompting and consuming approximately 600,000 US dollars worth of AI credits over three weeks. In May 2026, a likely financially motivated threat actor systematically probed METR's infrastructure using automated agents, apparently seeking unauthorised access to frontier AI models.

artificial intelligence technology [link](#)

Data exposure and leaks

Revolut data leak via government agency impersonation

On September 12, Revolut disclosed a breach where a threat actor impersonated a government agency via e-mail to obtain customer personal data. The shared information included identity documents, account statements with IBANs, and full transaction history including Bitcoin transactions. Revolut said a limited number of customers were affected, reportedly targeting high net worth users.

law enforcement finance [link](#)

Exploitation of GitLab repository API flaw to steal secrets

On September 14, US Cybersecurity and Infrastructure Security Agency (CISA) reported that threat actors were exploiting a critical GitLab vulnerability. Unnamed threat actors were observed probing and attempting to access sensitive information from exposed GitLab servers, potentially including credentials and other secrets. CISA urged rapid remediation due to the broad enterprise use of GitLab and the risk of widespread compromise.

technology [link](#)

ShinyHunters claimed FBI breach via PeopleSoft zero-day

On September 22, ShinyHunters claimed to have breached FBI systems by exploiting an alleged Oracle PeopleSoft zero-day, then accessing additional internal services and cloud-hosted environments. The group alleged theft of 2–3TB of data, including personal data relating to current and former employees and job applicants, sharing samples and a defacement screenshot. The FBI reported their investigation is ongoing.

law enforcement [link](#)

Artificial intelligence

Autonomous Gemini AI intrusions during security testing

On September 21, the BBC reported that Google's Gemini AI model autonomously compromised three companies during a May 2026 security evaluation conducted by independent company Irregular. Gemini exploited publicly available information and guessed credentials to gain unauthorised access. All affected entities were notified in July.

artificial intelligence [link](#)

Unauthorised OpenAI agentic activity discovered on over 20 websites globally

On September 9, a security researcher reported OpenAI-linked agents repurposed multiple third-party websites, including university YOURLS link shorteners, as unauthorised communication channels. The activity allegedly created at least 170 short links, generated tens of thousands of log entries, and leveraged public web services to relay queries, including to the FBI Crime Data Explorer using exposed API keys. OpenAI agents were reportedly active on 21 websites. Some operators restricted access afterwards.

artificial intelligence law enforcement public administration education research [link](#)

Rogue OpenAI agents bypassed web restrictions to access public and staging data

On September 26, OpenAI alerted multiple organisations globally after rogue OpenAI agents

had engaged with their websites beyond passive browsing. The agents were designed to identify authoritative sources of public information, but reportedly bypassed security measures on websites in some cases to achieve full access. In other instances, AI agents further accidentally uploaded user-provided images to third-party image-hosting services. OpenAI identified 53 such cases, with links not publicly listed, and worked with providers to remove most content. Most users were not affected, and opted out, enterprise, business, and API data was excluded.

artificial intelligence

technology

public administration

[link](#)

Spanish Data Protection Agency reported first AI agent-linked personal data breach

On September 14, the Spanish Data Protection Agency (AEPD) reported on an incident by an unknown threat actor involving an AI agent powered by a known large language model. The agent reportedly searched for vulnerabilities, logged onto systems, and autonomously probed applications for further weaknesses. In the final stage, it modified personal data and accessed invoices and other financial documents. This is the AEPD's first AI agent-enabled incident notification.

artificial intelligence

[link](#)

AI-built Carbonato botnet targeted exposed Docker APIs and steals AI API keys

On September 22, ThreatDown reported a publicly exposed, unauthenticated Docker registry leaking an operation's toolchain. The activity links a Docker-worm botnet dubbed Carbonato using Hermes Agent via Telegram tasking with a focus on stealing AI API keys and other credentials. Compromised hosts were used for persistence and lateral spread, with much of the supporting infrastructure still online.

artificial intelligence

technology

[link](#)

Researchers leveraged Claude to exploit OpenAI forum flaw to access employee ChatGPT account

On September 18, security researchers from a company called Hacktron AI reportedly exploited a misconfiguration in OpenAI's third-party hosted community forum to gain access to internal sign-ons and an OpenAI employee's ChatGPT account. The access enabled viewing of private software information and suggesting changes, with the account reportedly linked to internal code access. OpenAI said the issues were fixed and paid a 6.500 US dollar bug bounty.

artificial intelligence

technology

[link](#)

Opportunistic

Active exploitation of SonicWall SMA1000 zero-day chain

On September 1, SonicWall reported that threat actors are actively exploiting a chained set of SonicWall SMA1000 zero-day flaws, CVE-2026-83548 and CVE-2026-83549, to run commands on exposed remote access appliances. The activity targets SMA1000 6210, 7210, and 8200v models and may enable full device compromise, affecting enterprises, government, and critical infrastructure. SonicWall urged rapid patching and credential resets where compromise is suspected.

technology

[link](#)

Active exploitation of critical JFrog Artifactory authentication bypass CVE-2026-82329

On September 2, SOC Prime reported active exploitation of CVE-2026-82329, a critical authentication bypass flaw in self-hosted JFrog Artifactory instances. Threat actors are abusing the vulnerability to gain unauthenticated administrative access, generate admin tokens, enumerate users and credentials, and in some cases creating backdoor accounts.

technology

telecommunications

[link](#)

Exploitation attempts against Citrix NetScaler CVE-2026-19490

On September 3, vulnerability intelligence company Previdian reported unnamed threat actors targeting a critical Citrix NetScaler flaw (CVE-2026-19490) in the wild following publication of a credible proof-of-concept. Sensors observed exploitation-like requests from multiple IPs across

Australia, the US and Germany. Successful compromise was not confirmed, but exposed NetScaler ADC and Gateway deployments may face unauthorised access risks. [technology](#) [link](#)

Langflow's unauthenticated RCE flaw CVE-2026-0768 exploited in the wild

On August 31, VulnCheck reported active exploitation of Langflow's unauthenticated remote code execution flaw CVE-2026-0768. The vulnerability affects the code validator in Langflow's custom component editor. Malicious actors are conducting reconnaissance, checking the SSH directory, and harvesting credentials such as OpenAI and AWS keys. VulnCheck stated that the source traffic primarily originated from Russia and has thus far exclusively hit Canaries in the UK. The vulnerability was patched in January. [technology](#) [link](#)

SharePoint exploitation attempts chained anonymous access and RCE

On September 25, Previdian reported two-stage exploitation attempts by unnamed threat actors against Microsoft SharePoint, observed on a honeypot. The activity appeared to chain an anonymous access weakness with CVE-2026-65660 to achieve pre-authentication code execution and establish persistence via a webshell. The scope indicates opportunistic probing that could lead to server compromise where vulnerable SharePoint sites permit anonymous viewing. [technology](#) [link](#) [link](#)

Chrome zero-day CVE-2026-87491 exploited in attacks

On September 8, Google patched 230 vulnerabilities including an actively exploited Chrome zero-day (CVE-2026-87491). Unnamed threat actors were exploiting the flaw in the wild to run unauthorised code within Chrome by luring users to crafted web pages. Impact may include compromise of affected Chrome users and potential exposure of sensitive data. [technology](#) [link](#)

1. Conclusions or attributions made in this document merely reflect what publicly available sources report. They do not reflect our stance.

TLP definition

TLP	Disclosure	Message
RED	Not for disclosure, restricted to participants only.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed.
AMBER+STRICT	Limited disclosure, restricted to participants' organisations.	Recipients may share TLP:AMBER+STRICT information only with members of their own organisation.
AMBER	Limited disclosure, restricted to participants' organisations and their clients.	Recipients may share TLP:AMBER information only with members of their own organisation and its clients.
GREEN	Limited disclosure, restricted to the community.	Subject to standard copyright rules, TLP:GREEN information may be distributed with peers and partner organisations within their sector or community, but not via publicly accessible channels.
CLEAR	Disclosure is not limited.	TLP:CLEAR information may be distributed freely.