


IMPERVA THREAT RESEARCH

Threat Intel Report: January 26-February 1, 2025

Countries



Top Tags

Aisuru

ShinyHunters

Kazu

Rhysida

Everest

Executive Summary

Actionable Information:

- **Identity, not exploits, is the primary blast radius multiplier.** This week's most consequential incidents (ShinyHunters' Okta vishing, Manage My Health, Iron Mountain claims) reinforce that social engineering and identity-layer compromise now routinely enable cross-organization data loss without zero-days or malware-heavy intrusion chains.
- **AI ecosystems are a first-class attack surface.** From LLMjacking campaigns monetizing exposed model endpoints, to OpenClaw's autonomous agents expanding the impact of prompt injection and supply-chain risk, attackers are exploiting immature security controls around AI tooling faster than defenders are operationalizing guardrails.
- **Developer trust channels are under sustained abuse.** Malicious VS Code extensions, including both high-install data stealers and Moltbot impersonators dropping remote access tooling, show that marketplaces and open-source popularity are being weaponized as distribution mechanisms, directly targeting source code and credentials.
- **Volumetric attacks continue to outpace manual response.** The 31.4 Tbps Aisuru DDoS underscores that ultra-short, ultra-large bursts are now standard, forcing reliance on fully automated mitigation and pre-positioned capacity rather than reactive controls.
- **Critical infrastructure and healthcare remain high-pressure targets.** Sandworm's attempted Polish grid disruption and Rhysida's data resale claims against a US medical manufacturer highlight persistent geopolitical and financially motivated pressure on sectors with low tolerance for downtime or data exposure.
- **Misconfiguration remains a quiet but severe risk vector.** The exposed admin panel on an AI-enabled toy illustrates how basic authentication failures in consumer IoT/AI products can endanger sensitive data, especially children's, before attackers even need to operationalize exploitation.

Outlook:

The dominant theme this week is scale without sophistication: attackers are achieving outsized impact by chaining trust failures, like identity systems, open-source ecosystems, AI platforms, and default-exposed services, rather than relying on novel exploits. As autonomous agents, developer tooling, and AI-driven services proliferate faster than their security baselines mature, expect continued abuse of exposed APIs, poisoned extensions, and impersonation campaigns that bypass traditional perimeter defenses. At the same time, state-aligned actors will continue probing critical infrastructure for disruptive potential, while ransomware groups lean further into data resale and unverified leak claims to sustain leverage. Defenders should expect 2026 to further reward organizations that harden identity workflows, lock down AI and developer environments by default, and rely on automated, always-on mitigation for both volumetric and low-friction abuse.

Threat Activity in Imperva

Top CVE Growth This Week:

Imperva customers are protected against these CVEs.

↑ 2,944% [CVE-2025-53772](#)

- Deserialization of untrusted data in Web Deploy allows an authorized attacker to execute code over a network.

↑ 794% [CVE-2025-32432](#)

- Craft is a flexible, user-friendly CMS for creating custom digital experiences on the web and beyond. Starting from version 3.0.0-RC1 to before 3.9.15, 4.0.0-RC1 to before 4.14.15, and 5.0.0-RC1 to before 5.6.17, Craft is vulnerable to remote code execution. This is a high-impact, low-complexity attack vector.

↑ 298% [CVE-2026-23550](#)

- Incorrect Privilege Assignment vulnerability in Modular DS allows Privilege Escalation. This issue affects Modular DS: from n/a through 2.5.1.

Recent Cyber Threats

OpenClaw's Viral Rise and Moltbook Chaos

+OpenClaw +AI

- [OpenClaw](#), an open-source autonomous AI agent originally released as Clawdbot and briefly known as Moltbot, has exploded in popularity for its ability to execute real-world tasks with minimal human input, such as managing schedules, sending messages, and automating workflows; its GitHub repo quickly surpassed six digits in stars and adoption. This viral growth catalyzed Moltbook, a Reddit-style social network exclusively for AI agents where hundreds of thousands to over a million bots now post, comment, and interact autonomously, often mimicking human social behavior or generating bizarre content. However, the rapid ascent has drawn significant security, privacy, and reliability concerns: Moltbook experienced a major data exposure that revealed private messages and credentials, and independent analyses warn that OpenClaw's broad access and extensible architecture introduce attack surfaces (e.g., prompt injections, supply-chain risks) that could compromise host systems.

Russian Grid-Wiper Targets Polish Power Plants

+Sandworm +Poland +Russia +Utilities +DynoWiper

- A Dec 29–30 intrusion attempt targeted two Polish heat-and-power plants and aimed to disrupt comms links between renewable installations (e.g., wind turbines) and power distribution operators. Researchers analyzed a destructive wiper they named [DynoWiper](#) and attributed it to Sandworm, a unit of Russia's military intelligence. Local reporting indicated the activity

could have caused heat and power loss for up to ~500,000 homes, but Poland's PM said defenses held and critical infrastructure was not ultimately threatened.

LLMjacking as a Service Supply Chain

+LLMJacking +Bizarre Bazaar +MCP +AI

- Researchers observed a 40-day campaign, called [Bizarre Bazaar](#), that targeted exposed or weakly authenticated LLM endpoints, logging 35,000+ attack sessions across honeypots. The operation monetizes access by stealing compute for cryptomining, reselling API access, exfiltrating prompts/conversation history, and attempting pivots via publicly reachable Model Context Protocol (MCP) servers; common targets included unauthenticated Ollama endpoints and OpenAI-compatible APIs.

VS Code AI Assistants Quietly Upload Your Repo

+VSCode +Malicious Extension +China

- Two [Visual Studio Code Marketplace](#) extensions, "ChatGPT – 中文版" (1.34M installs) and "ChatMoss (CodeMoss)" (150K installs), were reported to exfiltrate developer data to China-based servers, totaling ~1.5M installs. Researchers say both share the same stealing code/infrastructure and use three collection paths: (1) exfiltrate the full contents of any file opened (plus edits), (2) a server-controlled command to steal up to 50 workspace files per run, and (3) a hidden 0-pixel iframe that loads multiple analytics SDKs (profiling/fingerprinting). The impact includes loss of private source code, config files, and secrets like .env API keys/credentials.

Fake Moltbot Extension Drops Remote-Access Tooling

+Moltbot +Malicious Extension

- A malicious VS Code extension [impersonating Moltbot](#) (which has 85,000+ GitHub stars) was published Jan 27, 2026 and removed after researchers flagged it. Moltbot itself has no legitimate VS Code extension, making the lure plausible. The extension auto-runs on IDE launch, pulls a remote config ("config.json"), and executes a binary that installs a legitimate remote desktop tool (ConnectWise ScreenConnect) to provide persistent access back to attacker infrastructure. It also includes fallback delivery via DLL sideloading and payload retrieval from Dropbox to keep the infection path resilient if primary C2 fails.

31.4 Tbps DDoS Barrage

+Aisuru

- The [Aisuru](#) botnet launched a record DDoS attack peaking at 31.4 Tbps and 200M RPS, disclosed as part of a broader campaign hitting multiple orgs, primarily telecommunications. Cloudflare said the campaign also targeted its own dashboard/infrastructure, combining hyper-volumetric HTTP floods with Layer 4 traffic; most attacks were brief (over half 1–2 minutes; only 6% longer), though peak sizes were extreme. The operational impact is rapid, high-cost saturation risk for edge capacity and upstream links, with defenders needing automated detection/mitigation tuned for very short, very high bursts.

Data Breaches

ShinyHunters Vishing Campaign Leads to Mass Data Leaks

+ShinyHunters +Vishing

- [ShinyHunters](#) linked a large-scale Okta voice-phishing (vishing) campaign to downstream data theft and leak attempts impacting multiple high-profile companies, including SoundCloud, Crunchbase, Betterment, and [Match Group](#) brands. The group reportedly used phone-based social engineering to trick employees into granting Okta SSO access, enabling lateral movement into internal systems and bulk data exfiltration rather than exploiting a software vulnerability. When extortion demands were rejected, ShinyHunters publicly leaked or claimed to leak tens of millions of user records, primarily emails, profile data, and internal business information. The activity highlights how identity-layer compromise can cascade across multiple organizations and industries without a single technical exploit.

Phishing Aftermath: Manage My Health Patient Data Breach

+Kazu +New Zealand

- [Manage My Health](#), a major NZ patient portal used by ~1.8 M users, disclosed a late-Dec 2025 cyber incident that potentially exposed between ~108 K and ~126 K patients' medical documents and clinical records. A group identifying itself as Kazu claimed

responsibility and demanded ransom, though forensic attribution is unresolved; authorities and the Privacy Commissioner are involved as notifications continue. The organization warns patients to be alert for phishing and impersonation scams exploiting the breached contact info, and regulators have launched a formal inquiry into privacy and breach handling.

Rhysida Claims Sale of Stolen Cytek Biosciences Data

+Rhysida +US

- Cybercriminal group Rhysida listed [Cytek Biosciences](#), a California-based medical manufacturer, on its leak site claiming to have stolen sensitive internal data and sold it on criminal markets. The posted samples included design documents and confidential internal files, though the scale and confirmed impact are limited to a few hundred individuals' records that Cytek acknowledged via breach notifications. This incident highlights continued targeting of healthcare supply-chain and manufacturing firms with ransomware gangs using data extortion and resale as monetization strategies.

Everest Ransomware Claims Massive Iron Mountain Data Theft

+Everest +Russia

- The Russian-linked Everest ransomware gang posted on its dark web site that it had exfiltrated ~1.4 TB of internal and client data from [Iron Mountain](#), a major S&P 500 records management and data storage company. The evidence to date consists of screenshots of directory structures allegedly showing client folders and internal documents; no full dataset release has been verified yet. Iron Mountain has not publicly confirmed a breach, making this an active ransomware claim used to pressure negotiations and potential future leaks.

Exposed Admin Panel on AI Toy Risks Children's Data

+AI +IoT +Missing Authentication

- Security researchers discovered an exposed administrative panel for a consumer AI-enabled toy ([Bondu AI](#)) that was accessible without authentication due to missing access controls in the backend. The admin interface could have allowed unauthorized access to children's interactions and personal data, because anyone stumbling on the backend domain could log in via a simplistic "Login with Google" route that erroneously exposed the core management portal. The finding highlights the broader IoT/AI toy industry risk where insecure configurations put sensitive end-user data at risk well before formal exploitation campaigns emerge.

Imperva Protection

Imperva provides comprehensive digital security solutions, including protection against threats to web applications and APIs— including security vulnerabilities, fraud, and bad bots, data security solutions, and security against disruptive network- and application-level DDoS attacks.