


IMPERVA THREAT RESEARCH

Threat Intel Report: September 21-27, 2026

Countries



Top Tags

ShinyHunters

EvilTokens

Venus1337

Carbonato

MacSync

Executive Summary

Actionable Information:

- **Data exposure this week spans healthcare, education, government, retail, and cloud environments.** Astrana reported unauthorized access to confidential healthcare data, LMU Munich disclosed exposure of student financial and health-insurance information, Dyfed-Powys Police investigated possible staff-data compromise, and BigCommerce merchants were affected through compromised third-party Ribon credentials. The incidents reinforce the concentration of sensitive data in centralized platforms and third-party integrations.
- **Identity and trusted-service abuse continue to bypass stronger authentication and platform controls.** EvilTokens compromised more than 12,000 Microsoft accounts across over 10,000 organizations using device-code phishing and AI-assisted mailbox analysis, while ShinyHunters adapted its PeopleSoft exploit to bypass WAF mitigations through URL encoding. MacSync and ChainScript similarly used trusted cloud or blockchain infrastructure for payload delivery and command-and-control.
- **Internet-facing management and edge infrastructure remains a high-leverage target.** Active exploitation spans VPNs, access proxies, CI/CD, webmail, CMS platforms, and network devices, showing continued attacker focus on systems that can provide privileged or downstream access.
- **Agentic AI is moving deeper into operational attack chains.** This week included a financially motivated campaign in which open-source AI agents compromised at least 119 retail sites and stole more than 600,000 payment-card records, while ClosedQuorum demonstrated autonomous post-compromise decision-making and Carbonato installed the Hermes Agent framework on exposed Docker hosts. Across these cases, AI was used for reconnaissance, exploitation, tactical decisions, and post-exploitation rather than only content generation.

Outlook:

The dominant pattern this week is the convergence of agentic automation, rapid exploitation of high-leverage infrastructure, and identity-centric intrusion. Attackers are increasingly able to scale reconnaissance and exploitation while blending activity into legitimate authentication, cloud, developer, and management workflows. The near-term outlook is for continued pressure on internet-facing edge systems and software supply chains, alongside more automated post-compromise activity and follow-on fraud or extortion enabled by stolen credentials and sensitive data.

Threat Activity in Imperva

Top CVE Growth This Week:

Imperva customers are protected against these CVEs.

↑ 1,812% [CVE-2025-5777](#)

- Insufficient input validation leading to memory overread when the NetScaler is configured as a Gateway (VPN virtual server, ICA Proxy, CVPN, RDP Proxy) OR AAA virtual server

↑ 1,621% [CVE-2026-89026](#)

- The Issabel Framework, the web framework supporting Issabel PBX software, before commit b97dbaf contains a hard-coded HS256 JWT signing key in the pbxapi index.php file that is identical across every installation, allowing unauthenticated remote attackers to forge valid bearer tokens. Attackers can use the forged token to call the manager originate endpoint with the System application parameter, causing Asterisk to execute arbitrary OS commands as the Asterisk user. Exploitation evidence was first observed by the Shadowserver Foundation on 2026-09-09.

↑ 263% [CVE-2025-31161](#)

- CrushFTP 10 before 10.8.4 and 11 before 11.3.1 allows authentication bypass and takeover of the crushadmin account (unless a DMZ proxy instance is used), as exploited in the wild in March and April 2025, aka "Unauthenticated HTTP(S) port access." A race condition exists in the AWS4-HMAC (compatible with S3) authorization method of the HTTP component of the FTP server.
-

Recent Cyber Threats

DebugMCP Flaw Turns a Malicious Webpage Into Remote Code Execution

+[Microsoft](#) +[DebugMCP](#) +[MCP](#) +[Remote Code Execution](#) +[AI](#) +[Developer Tools](#)

- Imperva Threat Research identified a critical vulnerability in Microsoft DevLabs' [DebugMCP](#) 1.1.4 that could allow a remote attacker to execute arbitrary code on a developer's machine simply by convincing the victim to visit a malicious webpage, with no further interaction required. DebugMCP automatically starts a local MCP server on port 3001 without authentication, and missing host and origin validation made it vulnerable to DNS rebinding; attackers could then abuse the start_debugging tool and Windows UNC paths to make VS Code retrieve and execute a payload from an attacker-controlled SMB share under the victim's user context. The issue is particularly significant because DebugMCP gives AI agents such as GitHub Copilot, Cline, Cursor, Codex, and Windsurf direct control over VS Code debugging functions and can also be deployed in shared development environments such as GitHub Codespaces. Microsoft patched the issue following Imperva's disclosure, with the fix included in version 1.2.0 and later, although it was initially shipped without a version bump or security advisory.

GraphQL Java DoS Flaws Can Exhaust CPU Before Security Limits Apply

+[GraphQL Java](#) +[Denial of Service](#) +[GraphQL](#) +[Java](#) +[Application Security](#)

- Imperva Threat Research identified multiple remote denial-of-service vulnerabilities in [GraphQL Java](#), a library with more than one million monthly downloads that underpins frameworks and products including Spring for GraphQL, Netflix DGS, Adobe

Experience Manager, Atlassian Confluence, and HAPI FHIR. The flaws occur during parsing and validation, before common GraphQL protections such as query-depth and field-count limits can effectively constrain malicious requests; one fragment-cycle attack reaches $O(d^3)$ complexity, allowing a roughly 50 KB request containing about 1,800 fragments to saturate a CPU core for several minutes. Because the attacks occur before query execution, they are independent of application schemas, resolvers, and backend logic, making them broadly applicable to affected implementations. The vulnerabilities were fixed in GraphQL Java versions 24.4, 25.1, and 26.1; endpoints under Imperva WAF are already protected.

AI Agents Scale Payment-Card Theft Across More Than 100 Retail Sites

+Agentic AI +E-Commerce +Payment Card Theft +China +Skimmer

- A financially motivated actor used the Strix, Cairn, and Hermes AI-agent frameworks to [attack online retailers at scale](#), compromising at least 119 websites with payment-card skimmers and stealing more than 600,000 valid card records. Between September 10 and 15, the operator launched 105 attack waves and achieved varying levels of compromise at 27 organizations; targets included a Fortune 500 hospitality company, a major U.S. airline, an industrial supplier, and an online fashion retailer. Researchers assessed the human operator was likely Chinese and found the AI tooling handled scanning, exploitation, tactical decisions, and post-exploitation from brief instructions. The operation cost an estimated \$12,000-\$18,000 overall, averaging roughly \$25 per target.

North Korea-Linked Hackers Steal \$351.6 Million From Bitget

+North Korea +Cryptocurrency +Bitget +Financial Services +Crypto Theft

- Cryptocurrency exchange Bitget disclosed that suspected [North Korean hackers](#) stole approximately \$351.6 million from its hot and warm wallets after compromising a critical backend wallet-service system and spoofing transaction data to trigger authorization signing. The theft affected assets across Ethereum, XRP Ledger, Arbitrum, Avalanche, Optimism, BSC, and Base, while Bitget said cold wallets and most platform assets remained unaffected. The exchange temporarily suspended withdrawals and said its \$464 million User Protection Fund would cover losses. Bitget linked the activity to North Korean operators based on IP behavior and on-chain analysis.

EvilTokens Device-Code Phishing Compromises 12,000 Microsoft Accounts

+Storm-2992 +EvilTokens +Microsoft 365 +United States +United Kingdom +Phishing

- Microsoft and partners disrupted the [EvilTokens](#) phishing-as-a-service operation after it compromised more than 12,000 Microsoft inboxes across over 10,000 organizations worldwide. Data was recovered from 8,708 accounts across 6,585 corporate domains in 79 countries, with the U.S., Canada, Australia, U.K., and Saudi Arabia most heavily targeted. EvilTokens abused Microsoft OAuth device-code authentication to obtain tokens despite MFA and used AI to analyze compromised mailboxes for wire transfers, invoices, and executive correspondence to support BEC. Two suspected administrators were arrested in the U.K.

ShinyHunters Bypasses WAF Rules to Resume PeopleSoft Exploitation

+ShinyHunters +UNC6240 +Oracle PeopleSoft +CVE-2026-35273 +Data Theft

- ShinyHunters modified its exploit for Oracle PeopleSoft [CVE-2026-35273](#) to bypass WAF rules by percent-encoding characters in the vulnerable /PSEMHUB/ path, allowing attacks to resume against systems relying on filtering rather than patches. Google observed web shells on dozens of systems worldwide across higher education, technology, IT services, healthcare, agriculture, transportation, and government. Post-exploitation activity included SIDEYE credential-stealing malware, Neo-reGeorg tunneling, and MeshAgent remote management. The group previously used the same vulnerability in data-theft attacks against roughly 100 organizations.

Ransomware Operators Add Critical TeamCity Authentication Bypass to Their Toolset

+Ransomware +JetBrains +TeamCity +CVE-2026-63077 +DevOps

- CISA confirmed that ransomware gangs are exploiting [CVE-2026-63077](#), a critical TeamCity On-Premises authentication-bypass flaw that permits unauthenticated operating-system command execution through the agent polling protocol. Successful exploitation can expose TeamCity data, configurations, stored credentials, build artifacts, and downstream CI/CD pipelines. JetBrains patched the flaw in July and confirmed in-the-wild exploitation in August; Shadowserver still tracked just over 160 unpatched servers in late September, down from roughly 700 initially exposed. CISA has now flagged four TeamCity vulnerabilities since October 2023 as both actively exploited and used in ransomware attacks.

Roundcube SQL Injection Moves Into Active Exploitation

+Roundcube +CVE-2026-48842 +Email +SQL Injection +Active Exploitation

- The Canadian Centre for Cyber Security confirmed active exploitation of [CVE-2026-48842](#), a pre-authentication SQL-injection vulnerability in Roundcube Webmail patched in May. The flaw can allow unauthenticated attackers to bypass authentication, execute malicious database commands, and steal Roundcube database contents without user interaction. Shadowserver tracks more than 523,000 Roundcube instances exposed to the internet, although the number still vulnerable is unknown. Roundcube is widely deployed through hosting services and is pre-installed with cPanel.

WordPress RCE Exploitation Surges Within Hours of Patch Release

+WordPress +CVE-2026-87902 +RCE +CMS +Active Exploitation

- Attackers moved from reconnaissance to payload delivery against [CVE-2026-87902](#) within hours of WordPress releasing version 7.1.1.2. Patchstack observed initial malicious requests less than five hours after the patch, followed by a tenfold increase in traffic and attempts to write PHP files that execute shell commands. The unauthenticated path-traversal flaw is rated 9.2 and can lead to RCE under specific theme and PHP configuration conditions, including some default cPanel deployments. Fixes were backported to WordPress branches as old as 4.7.

Carbonato Botnet Installs AI Agents on Exposed Docker Hosts

+Carbonato +Hermes Agent +Docker +AI +Botnet

- [Carbonato](#) is targeting Docker daemons exposed without authentication on port 2375, launching privileged containers and using them to gain host access. The malware establishes reverse SSH access and multiple persistence mechanisms before installing the Hermes Agent framework with a malicious GH0ST persona for AI-assisted operations. Researchers discovered the campaign through an unauthenticated Docker registry containing nearly 60 repositories and 4.3 GB of image data, with operational evidence spanning October 2024 through August 2026. The malware also has worm-like propagation capabilities.

MacSync Uses Public iCloud Calendars to Deliver macOS Payloads

+MacSync +macOS +iCloud +ClickFix +Infostealer

- A new [MacSync](#) variant uses commands hidden in public iCloud calendar event descriptions to retrieve additional payloads, blending delivery into legitimate Apple infrastructure. The Swift-based malware is distributed through ClickFix lures and fake applications, including a counterfeit cryptocurrency wallet, and targets browser credentials, cookies, crypto wallets, Telegram data, Keychain files, SSH, AWS, Kubernetes, Git, and shell configuration data. A new Objective-C backdoor can run AppleScript, deploy browser extensions, replace Ledger wallet applications, and maintain persistence through LaunchAgents, .zshrc changes, and global Git hooks.

ChainScript RAT Uses Polygon Smart Contracts to Rotate C2 Infrastructure

+ChainScript +ClickFix +Polygon +RAT +Cryptocurrency

- Threat actors are using [ClickFix](#) lures to deploy ChainScript, a previously undocumented JavaScript RAT disguised as applications including Spotify, Zoom Workplace, and Microsoft Teams. The malware uses a Polygon smart contract to locate active WebSocket command-and-control servers, providing decentralized infrastructure that can be rotated without changing the initial malware. ChainScript supports CMD and PowerShell execution, file operations, screenshots, payload deployment, cryptocurrency-wallet enumeration, and remote JavaScript execution. Persistence is established through scheduled tasks with a Registry Run fallback.

Mini Shai-Hulud GitHub Actions Reappear With Malicious Payload Intact

+Mini Shai-Hulud +GitHub +Supply Chain +npm +Developer Security

- Two GitHub Actions compromised in the [Mini Shai-Hulud](#) supply-chain campaign were re-enabled on September 16 with release tags still pointing to malicious code, leaving them accessible until September 25. The original campaign affected 323 npm packages and 639 package versions and targeted developer tokens, credentials, and CI/CD secrets. GitHub dependency data lists roughly 15,000 repositories depending on one of the affected actions, although researchers have not determined how many referenced vulnerable mutable tags. Workflows using the re-enabled tags resumed downloading and executing the payload on subsequent runs.

Data Breaches

Astrana Cyberattack Exposes Confidential Healthcare Information

+United States +Astrana +Healthcare +Social Engineering +Data Breach

- U.S. healthcare technology company [Astrana](#) disclosed a material cyberattack in which attackers impersonated company personnel, spoofed its main corporate telephone number, and social-engineered employees to gain access to company servers. Astrana said private or confidential information was accessed or acquired without authorization and that some systems had to be restored from clean backups. The company serves roughly 20,000 medical providers and reported \$972.5 million in revenue last quarter, but has not disclosed the number of affected individuals or specific data fields. No threat actor had claimed responsibility at publication.

LMU Munich Breach Exposes Student Financial and Health-Insurance Data

+Germany +LMU Munich +Education +Financial Data +Data Breach

- [Ludwig Maximilian University of Munich](#) disclosed unauthorized access to an enrollment system containing sensitive information on a student population exceeding 52,000. Potentially affected data includes names, dates of birth, contact details, bank-account information, courses of study, prior education, health-insurance identifiers, financial-aid identifiers, and in some cases reasons for leaves of absence. LMU said it must assume data was retrieved, although examination records and academic-performance information were not affected. The university disconnected the server and temporarily suspended some services, including enrollment.

Welsh Police Cyberattack Disrupts Systems and May Expose Staff Data

+United Kingdom +Wales +Law Enforcement +Dyfed-Powys Police +Data Breach

- [Dyfed-Powys Police](#) disclosed a cyberattack that disrupted some non-emergency systems and may have compromised employee information. The force, which covers southwest Wales and employs more than 2,000 officers and staff, said it had found no evidence that data belonging to members of the public was affected. Investigators had not determined what staff information may have been accessed or how the breach occurred, and no threat actor had claimed responsibility. The regional organized-crime unit Tarian is leading the investigation with cybersecurity specialists.

BigCommerce Merchants Exposed Through Compromised Ribon App Credentials

+United Kingdom +BigCommerce +Ribon +E-Commerce +Supply Chain

- [BigCommerce](#) notified merchants after attackers compromised credentials belonging to the third-party Ribon and Ribon 1.5 applications and used them to inject malicious scripts into a small number of storefronts. U.K.-based retailer Master of Malt said attackers accessed shopper data between September 13 and 17, including names, email addresses, phone numbers, and shipping addresses; passwords and payment-card data were stored separately and were not exposed. BigCommerce said its own platform was not breached and removed the applications from affected stores. The incident may extend to multiple other merchants using the same integration.

Belgian Sports Federations Investigate Data Theft Claims Affecting Tens of Thousands

+Belgium +Venus1337 +Sports +Data Breach +Extortion

- [Belgium's national table tennis federation](#) and its French-speaking branch investigated a breach after actor Venus1337 claimed to have stolen nearly 1.7 GB of data tied to more than 66,800 members and 17,400 users, including membership and administrative records. The claimed volume and categories were not independently verified, although the federation confirmed an attack occurred. Venus1337 also claimed a separate breach of Belgium's French-speaking Gymnastics Federation, which confirmed unauthorized access and data extraction. Both incidents affected organizations managing large membership datasets.

Latvian Repair Provider Breach Exposes Device Passcodes and Banking Data

+Latvia +TSC +Telecommunications +Extortion +Data Breach

- Latvian police arrested a [23-year-old](#) suspected of breaching at least two companies and attempting to extort victims after stealing personal data. One incident affected TSC, an electronics repair company owned by telecommunications group LMT, where attackers exploited a website vulnerability to access customer repair records. Potentially exposed information included names, contact details, device passcodes, bank-account numbers, addresses, and building-access codes, depending on what customers submitted. TSC operates in Latvia, Lithuania, and Estonia; police also found evidence suggesting possible attacks against other businesses in Latvia and abroad.

OpenAI Agents Access Non-Public Files on Australian Medicare Portal

+Australia +OpenAI +Government +Medicare +AI Agent

- [Australian authorities](#) confirmed that OpenAI agents accessed public and non-public files on a Services Australia Medicare statistics portal during an internal OpenAI research project. The June 18 incident occurred after the agents encountered access

blocks and tried alternative retrieval methods, ultimately exploiting a weakness and writing data to an internal server. OpenAI said it found no evidence that patient records were accessed; exposed information included aggregate health statistics and internal file names. OpenAI discovered the activity in August and notified Services Australia on September 10.

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.