


IMPERVA THREAT RESEARCH

Threat Intel Report: July 20-26, 2026

Countries



Top Tags

CI0p

Laundry Bear

HollowGraph

BlueNoroff

APT28

Executive Summary

Actionable Information:

- **Enterprise applications and internet-facing business platforms remain the highest-value targets for initial access.** This week featured active exploitation of critical vulnerabilities affecting ServiceNow and Fastjson, alongside CI0p's continued targeting of PTC Windchill and FlexPLM deployments. Rather than opportunistic exploitation, attackers are focusing on widely deployed enterprise platforms that provide immediate access to sensitive business data, engineering intellectual property, and cloud-connected environments. *Organizations should prioritize rapid patching of externally accessible enterprise applications, continuously inventory vulnerable software dependencies, and monitor for post-exploitation activity rather than relying solely on perimeter defenses.*
- **AI is increasingly being integrated throughout the attack lifecycle, from intrusion to victim prioritization.** This week's reporting demonstrates that AI is no longer limited to phishing content generation. Threat actors are using autonomous AI agents to automate post-exploitation activity against government networks, while malware families such as Dolphin X advertise AI-powered victim scoring to identify the most valuable compromises. These developments suggest attackers are beginning to use AI to improve operational efficiency rather than simply increase attack volume, enabling smaller groups to execute more sophisticated campaigns with less manual effort. *Organizations should expect AI to become an operational force multiplier across reconnaissance, exploitation, and post-compromise decision making.*
- **Credential theft continues to evolve beyond traditional phishing into infrastructure and trust-based attacks.** Russian operators abused hotel Wi-Fi DNS settings to redirect travelers to Microsoft 365 phishing pages, BlueNoroff selectively profiled cryptocurrency wallets before deploying malware, and HollowGraph demonstrated how legitimate Microsoft 365 services can be repurposed as covert command-and-control infrastructure. Together, these campaigns illustrate a broader shift toward abusing trusted services, collaboration platforms, and existing enterprise infrastructure to evade detection while stealing credentials and

maintaining persistence. *Organizations should strengthen identity monitoring, secure DNS infrastructure, implement phishing-resistant authentication, and closely monitor abnormal use of legitimate cloud services.*

- **Third-party providers and customer-facing platforms remain attractive breach targets because they enable downstream attacks.** Recent breaches affected organizations across the energy, healthcare, logistics, retail, and government sectors, exposing sensitive identity, employment, financial, and customer data. Several incidents including Origin Energy, Craneware, OnTrac, and Estée Lauder demonstrate how attackers increasingly target organizations that serve as intermediaries for customers or business partners, while credential stuffing against Chick-fil-A reinforces the persistent downstream impact of previously breached credentials. *Organizations should expect stolen customer data to fuel follow-on phishing, fraud, and account takeover campaigns long after the initial compromise.*

Outlook:

The defining trend this week is the continued convergence of enterprise software exploitation, identity compromise, and AI-assisted operations. Threat actors are moving away from indiscriminate attacks and instead targeting business-critical platforms, trusted cloud services, and high-value enterprise applications that offer immediate access to sensitive data and operational systems. At the same time, AI is beginning to reduce the expertise and time required to conduct post-exploitation activities, while legitimate infrastructure, from Microsoft 365 calendars to hotel DNS services, is increasingly being repurposed to hide malicious activity. Looking ahead, organizations should expect faster exploitation of newly disclosed enterprise vulnerabilities, broader adoption of AI-enabled offensive tooling, and continued growth in data theft operations that leverage trusted platforms and third-party providers to maximize both operational impact and long-term monetization of stolen information.

Threat Activity in Imperva

Top CVE Growth This Week:

Imperva customers are protected against these CVEs.

↑ 1,858% [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. This issue has been patched in versions 3.9.15, 4.14.15, and 5.6.17, and is an additional fix for CVE-2023-41892.

↑ 770% [CVE-2026-6875](#)

- ServiceNow has addressed a remote code execution vulnerability that was identified in the ServiceNow AI platform. This vulnerability could enable an unauthenticated user, in certain circumstances, to execute code within the ServiceNow platform. ServiceNow addressed this vulnerability by deploying a security update to hosted instances. Relevant security updates have also been provided to ServiceNow self-hosted customers and partners. Further, the vulnerability is addressed in the listed patches and family releases, which have been made available to hosted and self-hosted customers, as well as partners. We are not currently aware of exploitation against ServiceNow instances. We recommend customers promptly apply appropriate updates or upgrade to a patched release if they have not already done so.

↑ 504% [CVE-2026-49049](#)

- The Helix3 plugin for Joomla exposes an ajax handler task, that allows unauthenticated attackers to delete arbitrary files, write arbitrary JSON files and update template parameters.

Recent Cyber Threats

ServiceNow Remote-Code-Execution Flaw Exploited in Active Attacks

+AI +CVE

- Attackers began exploiting [CVE-2026-6875](#), a critical pre-authentication sandbox-escape vulnerability in the ServiceNow AI Platform that can allow unauthenticated remote code execution against self-hosted instances. Researchers observed exploitation attempts targeting the /assessment_thanks.do endpoint through a code-execution route that differed from the publicly documented proof of concept. ServiceNow began addressing hosted instances in April and released patches for self-hosted systems on July 13, although its advisory had not initially classified the flaw as actively exploited. ServiceNow says its platform processes more than 100 billion workflows annually and powers over 100,000 enterprise AI applications across 85% of Fortune 500 companies, giving the vulnerability broad potential enterprise impact.

Clop Targets PTC Windchill and FlexPLM in Global Data-Theft Campaign

+Clop +Manufacturing +Automotive +Aerospace +Retail

- [Clop-linked operators](#) are targeting internet-exposed PTC Windchill and FlexPLM environments used by manufacturing, automotive, aerospace, and retail organizations. Attackers reportedly chain a pre-authentication information-disclosure flaw in the FlexPLM WSDL endpoint with a Windchill server-side vulnerability to obtain unauthenticated remote code execution and deploy hex-named JSP webshells. The campaign has involved file-system enumeration, staging of engineering and design information, and theft intended for double-extortion operations rather than immediate encryption.

Imperva Blocks Active Exploitation of Critical Fastjson Zero-Day

+FastJson +CVE +RCE

- Active exploitation of [CVE-2026-16723](#) has been reported, a critical CVSS 9.0 remote-code-execution vulnerability affecting Fastjson 1.2.68 through 1.2.83 in certain Spring Boot fat-JAR deployments. The flaw is exploitable under Fastjson's default configuration without authentication, enabled AutoType, user interaction, or an external gadget class, allowing crafted JSON containing malicious @type values and nested JAR URLs to bypass type restrictions and execute code with the application's privileges. Imperva observed attacks across financial services, healthcare, computing, retail, and other sectors, with activity concentrated primarily against U.S. organizations and smaller volumes targeting Singapore and Canada; browser impersonators generated most requests, while Ruby- and Go-based tools accounted for roughly 30% collectively. No patched Fastjson 1.x release is available, but Imperva Cloud WAF and WAF Gateway customers are protected against associated exploit attempts; organizations should immediately enable SafeMode or use a noneautotype build, inventory transitive Fastjson dependencies, review logs for suspicious @type, jar:http, and jar:file patterns, and prioritize migration to Fastjson 2.x.

Russian Espionage Group Exploits Zimbra Zero-Click Flaw

+US +Russia +Laundry Bear +Phishing

- The US Government warned that the Russian state-sponsored group [Laundry Bear](#), also tracked as Void Blizzard, is compromising organizations through phishing and exploitation of a patched Zimbra Collaboration vulnerability. The campaign uses a zero-click attack path to access email data, increasing risk for organizations operating exposed or insufficiently patched Zimbra infrastructure. Email servers remain particularly valuable espionage targets because a single compromise can expose diplomatic communications, authentication-reset messages, internal documents, and long-term organizational relationships. The activity is relevant globally, particularly to government, defense, foreign-policy, and critical-infrastructure organizations that operate self-hosted collaboration platforms.

HollowGraph Hides Espionage Traffic in Microsoft 365 Calendars

+HollowGraph +Malware +Israel +Iran

- [HollowGraph](#), a newly identified .NET espionage implant, uses compromised Microsoft 365 calendars as a bidirectional command-and-control and exfiltration channel. The malware retrieves instructions from attachments placed in calendar events dated May 13, 2050, and exfiltrates RSA- and AES-256-encrypted files through new far-future events, causing the traffic to resemble legitimate Microsoft Graph API activity. Researchers found the implant on at least 12 machines, with approximately three actively communicating during analysis, and observed victim traffic from June 3 through July 9; the compromised exfiltration mailbox belonged to an Israeli organization. The malware shares technical characteristics with Cavern, previously associated with an Iranian intelligence-linked cluster overlapping MuddyWater and Lyceum, although the researchers did not make a definitive actor attribution.

Hermes AI Agent Automates Attack on Thailand's Finance Ministry

+Thailand +Government +AI

- A threat actor reportedly used the open-source [Hermes AI agent](#) in unattended “YOLO” mode to automate post-exploitation activity during an alleged compromise of Thailand’s Ministry of Finance. Allowing the agent to operate without interactive approval enabled it to conduct malicious actions autonomously after initial access, demonstrating how agentic AI can reduce the personnel and time required to navigate a breached environment. The case follows other 2026 incidents in which general-purpose AI tools were used to generate attack infrastructure, process stolen information, and recommend operational improvements. Although the underlying ministry breach remained alleged at publication, the activity represents a significant Asia-Pacific example of autonomous AI being applied directly to government intrusion operations.

BlueNoroff Profiles Cryptocurrency Wallets Before Delivering Malware

+North Korea +BlueNoroff +Malware +Crypto +Typosquatting

- North Korea-linked [BlueNoroff](#) operators are using a phishing platform that impersonates Zoom and Microsoft Teams, profiles victims’ cryptocurrency wallets, and selectively delivers malware to higher-value targets. The operation uses compromised trusted contacts and Telegram communications to create a self-propagating social-engineering chain, increasing credibility and allowing attackers to approach victims through existing industry relationships. Typosquatted collaboration domains and ClickFix-style instructions are then used to move targets toward malware execution. The addition of wallet reconnaissance before payload delivery allows the group to allocate resources toward victims with greater potential financial returns, reflecting the increasingly commercialized targeting methods used in North Korean cryptocurrency operations.

SourTrade Malvertising Targets Crypto Users Across 12 Countries

+SourTrade +Malvertising +Crypto

- The [SourTrade](#) malvertising campaign impersonated TradingView, Solana, and Luno services to target retail traders and cryptocurrency investors across 12 countries and 25 languages. Instead of delivering a complete malware executable from a conventional download server, the operation uses JavaScript and a legitimate Bun runtime to make the victim’s browser assemble the Windows payload from separate components. The landing pages fingerprint visitors and display empty content to suspected researchers or automated scanners while presenting selected targets with convincing copies of legitimate trading platforms. The campaign has reportedly operated since late 2024 and illustrates how multilingual malvertising enables cybercriminals to scale financially motivated attacks across regional and linguistic boundaries.

Hotel Wi-Fi DNS Hijacking Redirects Guests to Microsoft 365 Phishing

+Hospitality +DNS Hijacking +Phishing +Russia +APT28

- Russian APT28-associated attackers are modifying DNS settings on Wi-Fi infrastructure at hotels and conference venues to redirect guests toward [counterfeit Microsoft 365](#) authentication pages. Compromising local network DNS allows the attacker to manipulate browsing sessions without first infecting each individual device, making business travelers and conference attendees particularly attractive targets. Stolen Microsoft 365 credentials or sessions can provide access to corporate email, documents, cloud applications, and password-reset workflows after the victim leaves the affected location. The campaign has global relevance for travel, hospitality, and multinational organizations whose employees frequently use unmanaged wireless networks.

Dolphin X RAT Uses AI to Prioritize High-Value Victims

+Dolphin X +RAT +AI

- The newly identified [Dolphin X](#) remote-access trojan claims to include AI-powered profiling that scores infected users and ranks them according to potential value. This capability could help operators prioritize systems containing financial data, enterprise credentials, cryptocurrency assets, or other information likely to produce higher returns. Automated victim ranking would allow commodity malware operators to manage large infection volumes while reserving manual exploitation for the most valuable targets. Even where claims of AI functionality are partly promotional, the model reflects a broader shift toward automated triage and data-driven victim selection within cybercrime operations.

Data Breaches

Origin Energy Breach May Affect Up to Two Million Australian Customers

+Australia +Energy

- [Origin Energy](#), Australia's largest energy retailer, confirmed that attackers accessed customer data potentially including names, physical addresses, dates of birth, phone numbers, account information, the last four credit-card digits, and the last three bank-account digits. Origin serves 4.8 million customers, while the threat actor claimed to possess records for two million people and threatened to publish the data unless the company entered negotiations. Origin said the incomplete banking and payment-card information could not independently be used to take over accounts or make unauthorized transactions. The company notified the Australian Federal Police, Australian Cyber Security Centre, and privacy regulator, while the exposed identity and account data creates substantial risk for targeted scams and impersonation.

South Korean Diplomatic Academy Breach Persists for Ten Months

+South Korea +Government +Education

- Attackers maintained access to the [South Korean National Diplomatic Academy's](#) online education environment from April 2025 through February 2026 after exploiting a server vulnerability. The breach affected at least 6,000 current and former Ministry of Foreign Affairs personnel, including approximately 350 government attachés stationed overseas, although domestic reporting placed the possible total as high as 10,000. Exposed information included account IDs, names, email addresses, encrypted passwords, and reportedly job titles and departmental affiliations. The compromised server was located inside the foreign ministry's headquarters but had been excluded from regular security scrutiny, contributing to the ten-month dwell time and creating espionage and spear-phishing risk for diplomats worldwide.

Craneware Cyberattack Exposes U.K. Healthcare Technology Data

+Healthcare +Technology +UK

- Edinburgh-based healthcare technology provider [Craneware](#) disclosed that attackers accessed part of its data environment and stole customer and employee information. Craneware's Trisus platform provides accounting and billing technology to approximately 40% of U.S. hospitals, giving the incident potential downstream relevance beyond the U.K. even though operations and customer services were not disrupted. The company engaged the FBI, U.K. regulators, external cybersecurity specialists, and other authorities while investigating the scope and sensitivity of the stolen data. Craneware said much of the affected information might be nonsensitive or publicly available, but the incident underscores the concentration of healthcare-sector data and relationships within specialist technology providers.

Estée Lauder HR Breach Exposes Financial, Health, and Employment Data

+Retail +CISO

- [Estée Lauder](#) notified employees that attackers accessed an Oracle E-Business Suite environment used for human-resources management and obtained highly sensitive personal information. Exposed fields included names, addresses, email addresses, dates of birth, Social Security and passport numbers, bank-account information, health data, payroll information, performance reports, and other employment records. The intrusion occurred around August 9, 2025, but the company did not determine that personal information had been compromised until June 19, 2026. The timing and technical circumstances align with CIOp's mass exploitation of CVE-2025-61882 against Oracle EBS versions 12.2.3 through 12.2.14, a campaign that affected multiple global companies and universities.

Upbound Breach Enables \$13 Million in Fraudulent Leases

+Financial Services

- [Upbound Group](#) disclosed that attackers used customer information and documents stolen from its systems to create approximately \$13 million in fraudulent Acima lease-to-own agreements. The threat actors used the compromised data to obtain merchandise through participating retailers; Acima paid those retailers, but the attackers retained the goods and failed to make lease payments. The company did not disclose the number of affected customers or precisely characterize the stolen information, describing some of it as nonsensitive. Upbound responded by strengthening authentication, fraud-detection controls, and monitoring, demonstrating how stolen customer information can generate direct operational losses even without public extortion or ransomware deployment.

OnTrac Network Intrusion Exposes Customer Information

+Delivery

- Parcel-delivery company [OnTrac](#) notified customers that attackers breached its corporate network and may have accessed personal information. Logistics providers hold data connecting names, physical addresses, delivery patterns, merchants, and order activity, which can support highly credible parcel-delivery phishing and impersonation campaigns. The incident is particularly relevant to retail and e-commerce organizations because compromise of a delivery partner can expose consumers

even when the merchant's own environment remains secure. At publication, the company had not publicly quantified the number of affected individuals in the available reporting.

Chick-fil-A Credential Stuffing Compromises More Than 13,000 Accounts

+Food & Beverage +Credential Stuffing

- [Chick-fil-A](#) confirmed that attackers compromised more than 13,000 customer accounts during credential-stuffing activity targeting its website and mobile application between June 17 and June 19. Credential stuffing relies on username and password combinations obtained from unrelated breaches, demonstrating how previously exposed credentials continue to create downstream risk across consumer platforms. Compromised restaurant accounts may expose contact details, loyalty information, saved preferences, transaction histories, or other account data and can also be used for fraudulent purchases where stored payment functionality is available. The incident reinforces the need for breached-password screening, rate limiting, bot detection, and multifactor authentication on consumer-facing account systems.

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.