


IMPERVA THREAT RESEARCH

Threat Intel Report: September 28-October 4, 2026

Countries



Top Tags

Star Blizzard

Antino

Warlock

KillSec

Zammad

Executive Summary

Actionable Information:

- **Attackers are increasingly targeting trusted platforms and privileged infrastructure rather than relying on standalone malware delivery.** Warlock used SharePoint vulnerabilities to reach critical infrastructure and public-sector organizations, JADEPUFFER abused privileged Azure service principals for destructive cloud activity, and more than 543,000 still-valid credentials were identified in public GitHub repositories. The common thread is access through legitimate enterprise services, identities, and developer ecosystems that can provide substantial downstream reach once compromised.
- **Automation is reducing the interaction required across both state-sponsored and emerging AI-driven intrusion chains.** Star Blizzard's RedFlick technique streamlines delivery of the CosmicPulse backdoor, while the DIVD incident demonstrated an autonomous agent independently chaining two Zammad zero-days to achieve a successful network breach. At the same time, malicious custom ChatGPTs combined AI branding, paid search results, and ClickFix techniques to move victims into attacker-controlled execution workflows, showing automation influencing both technical exploitation and social engineering.
- **Sensitive identity and credential data remains highly concentrated across government, education, cloud, and third-party platforms.** The Pentagon breach affected more than three million people, Frontline Education exposed Social Security numbers and other employee information across the U.S. K-12 ecosystem, and DTU's identity-system compromise potentially affected up to 200,000 users. Separately, more than 16,000 misconfigured Supabase databases exposed PII, passwords, and authentication tokens, illustrating how both direct compromise and configuration failures can create large pools of data suitable for follow-on identity abuse.

Outlook:

This week's activity points toward continued convergence between identity compromise, trusted-platform abuse, rapid exploitation, and automation. State-linked operations are refining established intrusion chains while agentic systems are beginning to demonstrate autonomous exploitation of previously unknown vulnerabilities, and cloud identities, developer platforms, collaboration software, and third-party services increasingly provide attackers with access that extends beyond individual endpoints. The Bitget theft further demonstrates the potential impact of vulnerabilities in trusted third-party security products, while campaigns across Asia and attacks against critical infrastructure show that both espionage and financially motivated operations are targeting environments with broad organizational or strategic value. Near-term activity is therefore likely to remain concentrated around privileged identities, internet-facing and third-party infrastructure, and increasingly automated attack chains capable of moving from discovery to exploitation with less direct operator involvement.

Threat Activity in Imperva

Top CVE Growth This Week:

Imperva customers are protected against these CVEs.

↑ 4,914% [CVE-2026-57827](#)

- Joomla Extension - rsjoomla.com - Unauthenticated file upload in RSFiles component < 1.17.12 - The Joomla extension RSFiles is vulnerable to an unauthenticated arbitrary file upload that allows uploading executable files and leads to full RCE.

↑ 1,856% [CVE-2025-11833](#)

- The Post SMTP – Complete SMTP Solution with Logs, Alerts, Backup SMTP & Mobile App plugin for WordPress is vulnerable to unauthorized access of data due to a missing capability check on the __construct function in all versions up to, and including, 3.6.0. This makes it possible for unauthenticated attackers to read arbitrary logged emails sent through the Post SMTP plugin, including password reset emails containing password reset links, which can lead to account takeover.

↑ 1,359% [CVE-2025-54068](#)

- Livewire is a full-stack framework for Laravel. In Livewire v3 up to and including v3.6.3, a vulnerability allows unauthenticated attackers to achieve remote command execution in specific scenarios. The issue stems from how certain component property updates are hydrated. This vulnerability is unique to Livewire v3 and does not affect prior major versions. Exploitation requires a component to be mounted and configured in a particular way, but does not require authentication or user interaction. This issue has been patched in Livewire v3.6.4. All users are strongly encouraged to upgrade to this version or later as soon as possible. No known workarounds are available.

Recent Cyber Threats

Star Blizzard Automates Malware Delivery With New RedFlick Technique

+Russia +Star Blizzard +CosmicPulse +Cyber Espionage +Malware

- Researchers reported that Russian state actor Star Blizzard is using a new delivery technique dubbed [RedFlick](#) to deploy its CosmicPulse backdoor. The approach further automates malware installation and reduces the amount of victim interaction required, extending the group's established espionage tradecraft rather than introducing a fundamentally new exploitation primitive. Star Blizzard is a long-running Russian state-linked operation known for targeting government, diplomatic, defense, and policy communities. The technique illustrates continued investment by state actors in making established intrusion chains faster and less dependent on manual victim actions.

China-Nexus UAT-11587 Expands Antino Espionage Across Asia

+China +UAT-11587 +Antino +Cyber Espionage +Taiwan +India +Philippines +Cambodia +Pakistan +Thailand +Myanmar

- Research linked a China-nexus cluster tracked as UAT-11587 to an espionage campaign targeting government and policy organizations across Asia with the previously undocumented [Antino](#) backdoor. Activity expanded from Taiwan-focused spear phishing to 16 entities across eight Asian countries, including Taiwan, India, the Philippines, Cambodia, Pakistan, Thailand, and Myanmar. The Rust-compiled Windows backdoor supports reconnaissance, shell and PowerShell execution, file transfer, in-memory shellcode loading, and persistence, while its command-and-control can abuse Outlook and OneDrive. The campaign has been active since at least September 2025 and focuses on organizations with government, academic, think-tank, and policy relevance.

Warlock Uses SharePoint Exploits Against Water, Telecom, Government, and Education

+China +Warlock +Ransomware +SharePoint +Critical Infrastructure +Telecommunications

- China-linked [Warlock](#) ransomware operators exploited Microsoft SharePoint vulnerabilities for initial access against a water utility, telecommunications provider, regional government body, and university. The campaign demonstrates ransomware operators using vulnerabilities in widely deployed collaboration infrastructure to reach organizations across critical and public-sector environments. SharePoint exploitation provided an entry point for subsequent ransomware activity rather than remaining limited to opportunistic web compromise. The victim mix spans operationally sensitive sectors where disruption can extend beyond ordinary enterprise IT.

Apple Patches CoreGraphics Zero-Day Used in Sophisticated Targeted Attacks

+Apple +iOS +CoreGraphics +Zero-Day +Mobile Security

- Apple released emergency security updates for a [CoreGraphics](#) zero-day exploited in what the company described as extremely sophisticated targeted attacks against iOS users. The flaw was already being used in the wild before patches became available, placing it in the class of vulnerabilities typically associated with highly targeted surveillance or intrusion activity. Apple did not publicly identify the attackers or disclose the targeted individuals. The disclosure continues a pattern of zero-day exploitation against mobile-device components used to process attacker-controlled content.

Malicious Custom ChatGPTs Funnel Users Into ClickFix RAT Infections

+AI +ChatGPT +ClickFix +RAT +Malvertising +Google Ads

- Researchers found custom [ChatGPT](#) variants promoted through sponsored Google search results redirecting users to malicious sites that use ClickFix social engineering to deploy remote-access malware. The campaign combines trusted AI branding, paid search placement, and clipboard-based execution instructions to persuade victims to run attacker-supplied commands themselves. This approach avoids the need for a browser exploit or conventional malicious attachment and shifts execution into a user-authorized workflow. The activity reflects continued convergence between search-engine malvertising, AI-themed lures, and ClickFix delivery.

JADEPUFFER Uses Compromised Azure Service Principals for Destructive Cloud Attacks

+JADEPUFFER +Microsoft Azure +Cloud +Identity +Destructive Attack

- [JADEPUFFER](#)-linked attackers used compromised Microsoft Azure service principals to conduct destructive actions inside a victim cloud environment, demonstrating the impact of compromised non-human identities with privileged access. Rather than relying on endpoint malware, the operation abused legitimate cloud authentication and management functionality to manipulate or delete Azure resources. The activity highlights service principals as a high-value identity target because they can retain persistent permissions without the interactive controls applied to human accounts. The incident expands the threat model for cloud identity compromise from data access to direct infrastructure destruction.

KillSec Ransomware Infrastructure Seized in International Operation

+KillSec +Ransomware +Spain +Germany +United Kingdom +Romania +Law Enforcement

- An international law-enforcement operation seized [KillSec's](#) ransomware leak site and servers and resulted in three arrests, including a 16-year-old in Spain whom investigators identified as the group's suspected administrator and main operator. German authorities led the investigation, while additional suspects in the U.K. and Romania were arrested as part of the coordinated action. KillSec is accused of stealing organizational data and using publication threats to extort victims. The takedown removed core extortion infrastructure while giving investigators access to servers associated with the operation.

GitLab AI Gateway RCE Exposes Self-Managed AI Infrastructure

+GitLab +AI Gateway +RCE +DevOps +AI Security

- [GitLab](#) warned self-managed customers to immediately patch a critical vulnerability in its AI Gateway service that can allow attackers to execute arbitrary commands on vulnerable instances. AI Gateway brokers AI-related requests and responses, placing it in a sensitive position between developer workflows and model services. GitLab strongly recommended immediate updates for customers hosting their own gateway, particularly where the service is used to keep AI traffic inside organizational environments. The flaw adds AI-enablement infrastructure itself to the expanding set of enterprise systems exposed to high-impact code-execution vulnerabilities.

Hundreds of Thousands of Live Secrets Remain Exposed on GitHub

[+GitHub](#) [+Credential Exposure](#) [+Cloud Security](#) [+Developer Security](#) [+Secrets](#)

- Researchers found more than 543,000 credentials exposed in public [GitHub](#) repositories that were still valid in July despite platform controls intended to detect and prevent secret leakage. The exposed material represents usable authentication data rather than merely historical or revoked strings, creating opportunities for cloud compromise, API abuse, software-supply-chain access, and lateral movement. Public repositories remain a persistent source of long-lived credentials because secrets can survive in code, commit history, forks, and copied projects. The scale of still-valid credentials shows that detection without rapid revocation leaves a substantial residual attack surface.

Data Breaches

Pentagon HR Breach Exposes Records of More Than 3 Million People

[+United States](#) [+Pentagon](#) [+Defense](#) [+Government](#) [+Data Breach](#)

- The [U.S. Defense Manpower Data Center](#) began notifying more than three million military personnel and other individuals that attackers stole their information after breaching a Pentagon human-resources management system in October 2025. DMDC maintains personnel information supporting the Department of Defense workforce and benefits ecosystem, making the affected dataset unusually sensitive and concentrated. The October 2026 notification substantially clarifies the scale of an incident that occurred roughly a year earlier. The breach represents one of the largest recently disclosed exposures involving U.S. defense personnel data.

AI Agent Chains Zammad Zero-Days to Breach DIVD

[+Agentic AI](#) [+DIVD](#) [+Zammad](#) [+Zero-Day](#) [+Netherlands](#)

- The [Dutch Institute for Vulnerability Disclosure](#) said an autonomous AI-driven attacker breached its network by chaining two zero-day vulnerabilities in the open-source Zammad ticketing platform. DIVD described the intrusion as highly automated and noisy, with the agent leaving extensive evidence while independently conducting exploitation activity; the organization's initial investigation said the ultimate purpose and impact remained unclear. The incident is notable because the attacker used an agent to move beyond reconnaissance into successful exploitation of previously unknown flaws. DIVD disclosed the Zammad zero-day chain as it continued investigating the compromise.

Frontline Education Breach Exposes U.S. School Employee SSNs

[+United States](#) [+Frontline Education](#) [+Education](#) [+K-12](#) [+Data Breach](#)

- US-based [Frontline Education](#) notified school districts that attackers exploited a vulnerability in third-party software to access its environment and steal employee information. Exposed data includes Social Security numbers, email addresses, and physical addresses belonging to school-district personnel. Frontline provides administration and workforce-management software across the U.S. K-12 sector, giving the incident potential downstream impact across multiple customer organizations rather than a single district. Notifications to affected districts and individuals began in early October.

DTU Identity-System Breach May Affect Up to 200,000 People

[+Denmark](#) [+DTU](#) [+Education](#) [+Identity](#) [+Data Breach](#)

- The [Technical University of Denmark](#) disclosed that attackers accessed its identity and access management system and downloaded a large volume of data, potentially exposing information belonging to as many as 200,000 users. Compromise of an identity-management environment is particularly significant because it can concentrate account and personal information across students, staff, alumni, and other affiliated users. DTU began investigating the unauthorized access and assessing the affected population and data categories. The potential victim count makes the incident a major higher-education breach in Northern Europe.

More Than 16,000 Misconfigured Supabase Databases Expose Sensitive Data

+Supabase +Cloud +Data Exposure +PII +Credentials +Authentication Tokens

- Researchers identified more than 16,000 misconfigured [Supabase](#) databases exposing readable tables containing personally identifiable information, passwords, authentication tokens, and other sensitive application data. The exposure stems from insecure application and database configurations rather than a compromise of Supabase itself, but the shared backend platform makes the issue broadly repeatable across independently developed applications. Exposed credentials and tokens can create secondary compromise paths into user accounts or connected services. The findings demonstrate how cloud backend misconfiguration can aggregate into large-scale, distributed data exposure.

Bitget Says \$387.5 Million Theft Began With Third-Party Security Zero-Day

+Bitget +Cryptocurrency +North Korea +Zero-Day +Supply Chain +Financial Services

- Cryptocurrency exchange [Bitget](#) said attackers stole approximately \$387.5 million after exploiting a zero-day vulnerability in third-party security products used in its environment. The September 24-25 theft affected a limited number of hot and warm wallets, while the exchange said cold wallets and the overwhelming majority of platform assets remained secure; Bitget subsequently resumed withdrawals. Earlier reporting linked the operation to suspected North Korean hackers based on infrastructure and blockchain analysis. The incident combines third-party security-product exploitation with one of the year's largest cryptocurrency thefts.

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.