GitProtect

DevOps Threats Unwrapped Report

2026 EDITION

Introduction

For the third year in a row, here at GitProtect, we have been analyzing the threats that DevOps face on a daily basis. Malware and ransomware attacks, outages, and infrastructure downtime - all of that can lead to data loss as well as higher financial costs. Thus, analyzing threats of the previous year can give us a better picture of what to expect in the upcoming year and what security measures to take to foresee them.

What could we remember 2025 for when it comes to DevOps threats? It was a year when trusted development platforms, automation pipelines, and cloud identities became a playground for cyber criminals. Attackers leveraged platforms such as GitHub, GitLab, Atlassian, and Microsoft as part of their malware campaigns - they used trusted DevOps platforms as malware distribution channels, command-and-control infrastructure, and credential harvesting pipelines.

Campaigns such as GhostAction, Shai-Hulud, the nx supply-chain breach, GPUGate, and GitVenom demonstrated how abused automation and stolen tokens could ripple through hundreds, and in some cases even thousands, of repositories.

Malware families, including PyStoreRAT, SmartLoader, Lumma Stealer, Amadey, Lumma Stealer, RedLine Stealer, AsyncRAT (and its variants like DCRat and Venom RAT), ZeroCrumb, and remote access trojans were distributed via fake libraries, branded impostor repos, and poisoned packages. Even AI-generated repositories and dormant accounts became vectors for credential theft and covert reconnaissance.

However, attackers weren't only limited to distributing the code on DevOps platforms. Identity was another attack direction. Hackers abused OAuth flows, long-lived Personal Access Tokens (PATs), and MFA-bypassing phishing kits to bypass defenses on Microsoft 365, GitHub, and collaboration tools at scale. Let's just name a few Whisper 2FA, Salty 2FA, Tycoon2FA, and RaccoonO365, for example.

On the operational side, compared to 2024, the number of incidents has grown - in 2024 GitHub, GitLab, Azure DevOps, and Jira experienced 364 in total (according to their official status pages) (editor.: The information for January 2024 for Azure DevOps is not included to this figure), while in 2025 there were recorded 607 incidents recorded across these platforms' public incident reporting channels. So, we can see about a 40% rise. Of these, 156 were critical or major events, consuming over 1,750 hours of downtime and representing a 69% year-over-year increase in high-severity disruptions. In 2024, there were only 48 events of critical and major importance.

Vulnerability management reflected growing pressure in 2025 as well. With 236 patched vulnerabilities reported in vendors' public security advisories across DevOps services, 14 of which were of critical CVSS 9.0 and higher, and 126 of which were of high severity, we can see an 30% increase in patched flaws comparing H1 to H2.

In GitProtect's 2024 CISO's guide to DevOps Threats, Technology and Software were the most targeted industries. In 2025, the trend remained the same. Ransomware groups and extortion-driven actors, including Hellcat and Crimson Collective, were behind a significant share of data breaches in 2025, targeting globally recognized brands, including Red Hat, Nissan, Europcar, and Ascom, among others, across multiple industries, like telecommunications, automotive, and education.

As attackers blend trusted platforms, hardware-aware evasion, malicious AI-generated code, and phishing-as-a-service into their arsenals, the 2025 threat landscape makes one thing clear: traditional perimeter defenses and reactive monitoring are no longer enough. Organizations need to ensure the resilience of their environment.

This report explores the chronology of techniques, incidents, and lessons that defined DevOps and SaaS threats in 2025, and what teams must do now to build a more resilient future.

Daria Kulikova
HEAD OF GITPROTECT LAB

Disclaimer: This report is based exclusively on publicly available information published by vendors on their official status pages, security advisories, databases, and publicly reported incidents documented in industry media. All data and statistics have been compiled from these publicly accessible sources. The analysis presented herein represents GitProtect's independent assessment and is intended for educational and informational purposes only.

Table of contents

01

AI in DevOps

Key takeaways

1 Hidden risks

AI in DevOps expanded attack surfaces, enabling supply-chain attacks and code abuse at scale.

2 Human authority

The steady rise in AI-related incidents makes human oversight and strict access controls essential.

3 Security gaps

AI-powered platforms can be vulnerable to prompt injection, remote code execution, and credential exfiltration.

01

AI in DevOps

Our recommendation

The year 2025 shows that organizations should treat AI tools as untrusted actors by default. They should enforce strict input sanitization, least-privilege access for AI agents, short-lived credentials, and human-in-the-loop reviews, as well as verify repository provenance, sandbox AI outputs, harden CI/CD workflows, and continuously monitor AI-driven actions across dev environments.

Vulnerabilities, prompt injections, and AI risks

AI coding tools have been on the rise in recent years, and in 2025, the tendency didn't change. AI tools help developers write code faster and automate repetitive tasks. However, it's important to remember that AI-generated code can accelerate defects rather than reduce them. Moreover, they may introduce additional challenges related to security, logic, and long-term maintainability.

According to a CodeRabbit study, AI-generated pull requests can contain significantly more issues than human-authored ones, averaging about 1.7 times more problems.

Major DevOps platforms, including GitHub, Atlassian, and GitLab, have already integrated AI tools to assist users in software development - a lot of developers use GitHub Copilot, GitLab Duo, or Rovo on a daily basis at work. In our analysis of these platforms' status pages, we identified 68 AI-related issues of different impact in 2025.

FIGURE 1

Monthly number of incidents (2025)

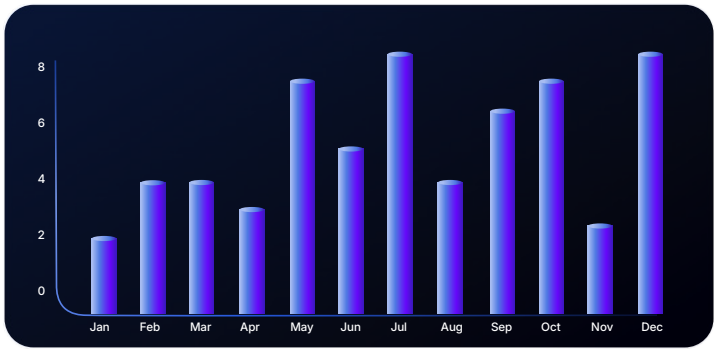

It's worth mentioning that the number of issues related to AI increased in the second half of the year - 40 cases versus 28 in the first half, which represents an 16% growth.

Quarterly analysis also shows a rising trend: 10 AI-related cases in Q1, 18 in Q2, and 20 each in Q3 and Q4.

FIGURE 2

Quarterly distribution (2025)

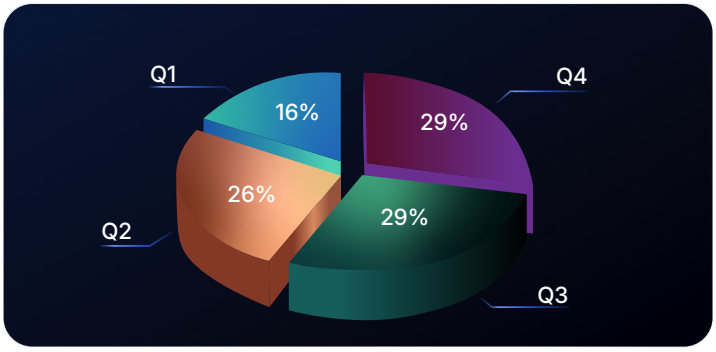

In this chapter, we have covered the most significant AI-related incidents that shook the DevOps world in 2025.

DECEMBER

AI-powered supply chain attack targets GitHub developers

SOURCES:
gbhackers
Menafn

The final AI-assisted supply chain attack of 2025 was aimed at researchers, developers, and security professionals through compromised GitHub repositories. The attackers used dormant accounts to distribute a backdoor called PyStoreRAT. The campaign leveraged polished AI-generated repositories that mimicked legitimate tools and utilities. It helped hackers to persuade victims to clone or run compromised projects.

Once executed, PyStoreRAT gathered system intelligence, exfiltrated files, harvested credentials, and evaded detection by endpoint security tools, while maintaining persistence through rotating command-and-control infrastructure. By focusing on niche developer communities, attackers exploited trust in open-source repositories. That, in turn, allowed them to bypass traditional malware detection and social vetting mechanisms.

> **The campaign demonstrates how AI can lower barriers for sophisticated attacks, producing realistic commit histories, documentation, and code that appear legitimate.**

What to do to mitigate such threats? Teams are required to have strict repository hygiene, sandbox experimental projects, verify provenance, and evolve community-level trust systems beyond simple popularity metrics.

Millions of users are affected by critical vulnerabilities found in GitHub Copilot and other AI tools

SOURCES:

Cyber Security News

gbhackers

Security vulnerabilities, dubbed "IDEsaster," were discovered across major AI-powered integrated development environments (IDEs) and coding assistants at the end of 2025. These IDEs and coding assistants included GitHub Copilot, Cursor, Zed.dev, Roo Code, JetBrains Junie, Gemini CLI, and Claude Code. Unlike traditional exploits, IDEsaster leveraged the base IDE's features, such as configuration files, workspace settings, and multi-root environments, rather than targeting a single application component.

The attack chain typically follows three stages:

1. prompt injection,
2. misuse of AI agent tools,
3. and exploitation of underlying IDE features, enabling remote code execution and data exfiltration.

Critical vulnerabilities included IDE Settings Overwrite, Remote JSON Schema manipulation, and Multi-Root Workspace Settings, with at least 24 CVEs assigned and over 30 vulnerabilities identified.

In proof-of-concept attacks, malicious prompts tricked AI agents into modifying files like .vscode/settings.json or .idea/workspace.xml. In turn, it allowed attackers to execute arbitrary commands or steal sensitive information.

> **All of these highlight that legacy IDE features, safe for human users, can become dangerous when autonomous AI agents operate with elevated privileges under adversarial influence.**

GitHub, Cursor, and Roo Code instantly released patches. However, developers should also remember to adopt "Secure for AI" principles, including human-in-the-loop verification, capability-scoped tools, input sanitization, and continuous monitoring of IDE features to mitigate risks.

Critical AI flaw puts GitHub and GitLab workflows at risk

SOURCES:

Cyber Security News

Red Hot Cyber

Cyber Press

Cyber Security News

A critical vulnerability class called "PromptPwnd" was discovered in AI-integrated CI/CD pipelines, including GitHub Actions and GitLab workflows. The flaw allowed attackers to inject malicious prompts via untrusted user inputs, like issue titles or pull request bodies, tricking AI agents such as Google Gemini CLI, Anthropic Claude

Code, and OpenAI Codex into executing privileged commands.

This could lead to leakage of sensitive credentials, API keys, and cloud tokens, or even unauthorized workflow modifications.

The attack exploits the common practice of embedding raw user content into AI prompts without proper sanitization or validation, and misconfigured workflows that grant AI models elevated access amplify the risk. Mitigation requires sanitizing inputs, treating AI outputs as untrusted code, restricting AI toolsets from making repository changes, and limiting token scopes by IP.

NOVEMBER

CrewAI token exposure shows hidden risks in AI development

SOURCES:

Security Boulevard

CSO

A critical vulnerability in the open-source CrewAI platform exposed an internal GitHub token with administrative access to all private repositories, highlighting the risks of long-lived credentials in AI development environments.

The flaw, discovered by Noma Security, occurred when an exception response inadvertently returned the token during a provisioning failure, allowing potential access to source code, workflows, and additional secrets.

This incident reflects a broader industry pattern, as research from Wiz found that **65% of top AI companies had inadvertently leaked secrets in GitHub repositories**, exposing credentials tied to AI firms valued at more than $400 billion and including model-access and workflow tokens. Long-lived tokens, often used for automation and continuous deployment, create high-risk non-human identities that can outlast their intended scope and amplify security fragility.

> **Thus, it's better to adopt short-lived, context-specific credentials, restricting scope, rotating secrets regularly, and auditing historical repositories to limit exposure.**

Moreover, the exposed secrets could have granted access to private AI models, training data, and internal systems, highlighting how speed-to-market pressures are outpacing secure development practices in the AI sector.

Well, cloud misconfigurations are not new, but the stakes are higher as AI leaks can enable model theft, data poisoning, regulatory violations, and supply chain attacks.

Critical AI flaws in GitHub Copilot and Visual Studio expose developers to code and data risks

SOURCES:

Cyber Security News

gbhackers

Microsoft disclosed two important security flaws in GitHub Copilot and Visual Studio/ Visual Studio Code that exposed a growing risk in AI-assisted development.

One of the vulnerabilities could allow path traversal attacks on developer machines, potentially exposing source code, configs, and secrets, while the other allowed attackers to manipulate Copilot's AI-generated code to bypass security checks.

Together, the flaws show how trusted AI coding assistants can become a direct attack vector when output validation and local protections fail. The risk is amplified by how widely these tools are used and how often developers accept AI suggestions with minimal review.

Microsoft quickly released patches, but the disclosures highlight a deeper issue: AI speed is still outrunning secure-by-design development practices. As generative AI becomes embedded in coding workflows, validating AI output and hardening DevSecOps pipelines is no longer optional; it's critical.

OCTOBER

Copilot chat vulnerability lets attackers hijack AI responses

SOURCES:
Security Week
SC Media
Bank Info Security

Security researchers disclosed a now-patched vulnerability in GitHub Copilot Chat that could have allowed attackers to leak sensitive data and hijack the AI assistant's responses using hidden prompt injection techniques.

Legit Security demonstrated that malicious instructions embedded in concealed pull request comments could be silently ingested by Copilot, influencing its behavior for other users who viewed the same repository context. By chaining this prompt injection with a creative bypass of GitHub's Content Security Policy, the researcher was able to exfiltrate secrets such as AWS keys and zero-day vulnerabilities from private repositories.

The attack abused the Camo image proxy by pre-generating signed image URLs for individual characters, effectively turning image fetch requests into a covert data exfiltration channel. Copilot, which has the same access privileges as the user invoking it, could be manipulated to search private code, encode the results, and leak them through these proxy requests without visible indicators.

The flaw highlighted how Copilot's strength, such as deep contextual awareness across repositories, can become a liability when untrusted input is not properly isolated.

AUGUST

GitHub Copilot 'YOLO mode' vulnerability allows attackers hijack developer systems

SOURCES:
Cyber Security News
gbhackers

A critical vulnerability in GitHub Copilot and Visual Studio Code, tracked as CVE-2025-53773 (with a CVSS score of 7.8) was discovered in August. It could allow attackers to achieve remote code execution via prompt injection.

The flaw exploited Copilot's ability to modify configuration files, particularly .vscode/settings.json, enabling attackers to activate "YOLO mode," which removed all user confirmations and granted the AI unrestricted access to execute shell commands and perform privileged operations. By embedding malicious prompts in source code, web pages, or GitHub issues, attackers could compromise Windows, macOS, and Linux systems, recruit developer machines into botnets called "ZombAIs," and propagate self-spreading AI viruses through Git repositories.

Researchers demonstrated that the exploit can also manipulate .vscode/tasks.json and inject malicious MCP servers, expanding the potential attack surface.

Microsoft addressed the vulnerability in the August 2025 Patch Tuesday release, requiring user approval for security-relevant configuration changes to prevent unauthorized AI actions.

JULY

Hacker slips malicious code into Amazon Q AI extension via GitHub pull request

SOURCES:

Bleeping Computer
WPN

Amazon moved to contain a supply-chain security incident after a hacker successfully injected malicious code into the Amazon Q Developer Extension for Visual Studio Code via a compromised GitHub workflow. The attacker, operating under the alias "lkmanka58," submitted a pull request that embedded a prompt instructing the AI assistant to wipe local systems and cloud resources, which was approved and merged without detection. The compromised version, Q 1.84.0, was published to Microsoft's VS Code Marketplace on July 17, exposing nearly one million users before the issue was flagged by security researchers days later.

AWS pulled the affected release and issued version 1.85.0, removing the unapproved code and revoking exposed credentials, while maintaining that the injected command was improperly formatted and posed no real risk. Some researchers, however, disputed this assessment, saying the code did execute in limited cases, underscoring the seriousness of the breach.

> **The incident highlights how AI-powered developer tools can amplify supply-chain risks when open-source contributions are insufficiently vetted.**

This case might prompt broader scrutiny of AI-assisted coding platforms and stricter controls over how AI prompts and code changes are reviewed before reaching production.

JUNE

‘Living off AI’ attack allows hackers steal Jira data via support tickets

SOURCES:

Infosecurity Magazine
SC Media
gbhackers

Researchers from Cato Networks have demonstrated a “Living off AI” attack vector that abused Atlassian’s Model Context Protocol (MCP) by delivering prompt injection payloads through Jira Service Management support tickets.

In the proof-of-concept attack, an external, often unauthenticated user might submit a crafted ticket that is later processed by an AI-powered workflow used by internal support staff. When the AI agent, connected to internal Jira or Confluence resources via MCP, executed the injected instructions, it did so with the support engineer’s privileges.

FIGURE 3

Prompt injection via Jira Service Management

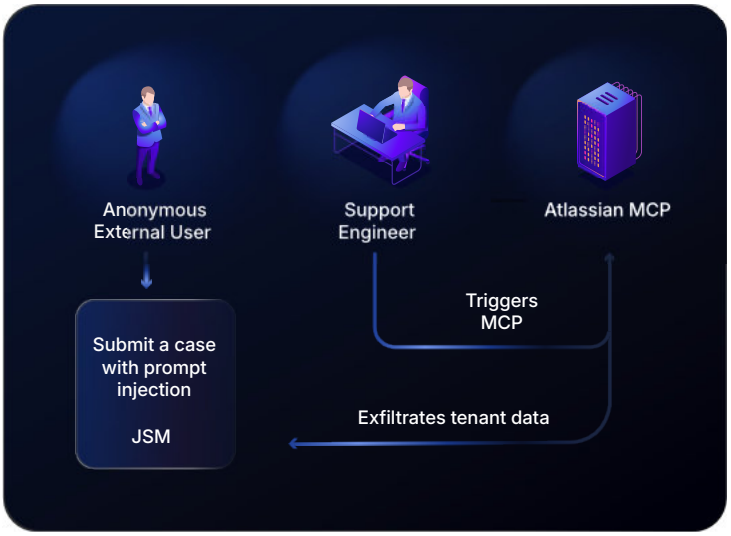

This could allow attackers to extract sensitive internal data, such as other Jira tasks or tenant information, and have it written back into the attacker’s own ticket. Crucially, the threat actor never directly accesses internal systems, instead using the AI and the employee as an unwitting proxy.

> **This attack highlights a broader systemic risk wherever AI agents act on untrusted external input without proper prompt isolation or access controls.**

Microsoft fixes zero-click AI flaw exposing corporate data

SOURCES:

Bleeping Computer
The Hacker News

Researchers from Aim Labs uncovered a critical “zero-click” AI vulnerability in Microsoft 365 Copilot, dubbed EchoLeak, that allowed attackers to exfiltrate sensitive corporate data without any user interaction.

The flaw, assigned CVE-2025-32711 with the severity score of 9.3 and rated critical, exploited a class of vulnerabilities called LLM Scope Violations, tricking Copilot’s retrieval-augmented generation (RAG) engine into mixing untrusted content with privileged internal data.

Attackers embedded malicious prompt injections in innocuous-looking emails or markdown content, which Copilot automatically incorporated into its AI context, silently leaking data via trusted channels like Teams or SharePoint URLs.

SOURCES:

Cyber Security News
gbhackers

While Microsoft fixed the vulnerability server-side by May 2025 and confirmed no real-world exploitation, EchoLeak demonstrates the risks posed by deep AI integration into enterprise workflows. This attack can highlight a broader threat landscape, where advanced tool poisoning attacks (e.g., Full-Schema Poisoning) and flaws in Model Context Protocol (MCP) could allow LLMs to access sensitive resources or execute malicious actions.

Additional attack vectors include DNS rebinding via MCP’s Server-Sent Events, enabling adversaries to pivot from external websites to internal enterprise systems.

The incident shows that organizations should implement granular input scoping, prompt injection filters, and strict permission controls for AI agents interacting with internal tools.

As AI assistants grow more autonomous, such vulnerabilities reveal fundamental architectural risks that cannot be fully mitigated without redesigning trust boundaries and monitoring LLM interactions with corporate data sources.

Cybercriminals exploit AI and GitHub trust to steal sensitive data

SOURCES:

Tend Micro
Cyber Security News
gbhackers.

Cybersecurity researchers found out about a cyber campaign that was based on AI-generated fake GitHub repositories to distribute SmartLoader and Lumma Stealer malware.

In this campaign, the malicious repositories mimicked legitimate projects with polished AI-written documentation, emojis, and logos to bypass human scrutiny and automated security checks.

Users who were downloading ZIP files from the Releases section triggered a multi-stage attack chain, where SmartLoader executed obfuscated Lua scripts that connected to command-and-control servers.

> **The malware establishes persistence, evades security software, and ultimately deploys Lumma Stealer, which targets browser credentials, cryptocurrency wallets, two-factor authentication data, and system metadata.**

Operators used modular payloads with anti-analysis checks, making detection and mitigation difficult while stolen data is sold for up to $5,000 per digital identity package.

FIGURE 3

Attack chain

The campaign exploited GitHub's trusted reputation to trick users seeking gaming mods, cracked software, or crypto tools into executing malicious code.

It proves that AI-assisted attacks can evolve supply chain threats, combining generative AI with legitimate development platforms to craft highly convincing attack vectors. Defenses include application allowlisting, behavior-based detection, repository authenticity verification, sandboxing downloads, and stricter monitoring of scripting activity.

FEBRUARY

Hackers exploit AI art tool to steal 1+ TB of Disney's internal data

SOURCES:

Security Week

ITC

A major data breach at Disney was triggered when the company's employee downloaded a seemingly free AI tool from GitHub, which secretly contained spyware. The malicious tool was disguised as friendly software for generating AI art.

The malware, deployed by a hacker who claimed to represent the hacktivist group NullBulge, managed to capture every keystroke on a Disney employee's device and harvested login credentials. Using these credentials, the threat actor accessed

thousands of Disney Slack channels, stealing approximately 1.1 TB of internal data over five months, including 44 million messages, unreleased project details, source code, employee salaries, and customer information.

The breach went largely unnoticed until a Disney employee reported suspicious activity after receiving a message on Discord, but Disney initially underestimated the severity of the incident.

When the extortion attempt failed, the attacker published the stolen data online, prompting Disney to discontinue Slack for internal communications and terminate the employee who had downloaded the malware.

SOURCES:

C Tech

Tech Crunch

Private GitHub repositories exposed through Copilot, affecting 16K+ organizations

Security researchers at Lasso uncovered a major data exposure issue where private GitHub repositories were surfaced through Microsoft Copilot after being indexed and cached by Bing. Even repositories that were made public only briefly, or later deleted or set to private, remained accessible via Copilot prompts.

In its research, the company managed to identify over 20,000 exposed repositories affecting more than 16,000 organizations, including major enterprises such as Google, IBM, PayPal, Tencent, and Microsoft. The leaked data included intellectual property, internal packages vulnerable to dependency confusion, and hundreds of active credentials and API tokens.

In some cases, these secrets enabled unauthorized access to corporate environments, creating a serious security and compliance risk.

Although Microsoft restricted public access to Bing's cache, researchers found that Copilot could still retrieve the data, indicating incomplete remediation.

The findings highlight a new AI-driven attack surface where data exposure can persist long after content is removed, making prompt-based data leakage a critical enterprise risk.

JANUARY

SOURCES:

Cybersecurity News

Dark Reading

Attackers can turn GitHub Copilot into a malicious coding tool

Researchers from the Apex Security team identified two techniques that could allow attackers to manipulate GitHub Copilot and bypass its security and usage controls.

The first, known as an "affirmation jailbreak," exploits Copilot's tendency to complete helpful responses by replacing refusal messages. Using words like "Sure," the attacker can cause the model to generate malicious code or instructions.

This technique works inside normal code files, making it easy for developers to

coerce Copilot into producing malware, exploit instructions, or other harmful outputs.

The second attack uses a proxy hijack, where Copilot's traffic is rerouted through a malicious proxy by modifying client-side settings. With it, attackers could manage to intercept Copilot's authentication token and gain unrestricted access to underlying OpenAI models without paying or respecting usage limits. Moreover, the intercepted traffic could also expose system prompts and conversation history, enabling further manipulation of model behavior and data leakage.

This report is based exclusively on publicly available information published by vendors on their official status pages, security advisories, databases, and publicly reported incidents documented in industry media. All data and statistics have been compiled from these publicly accessible sources. The analysis presented herein represents GitProtect's independent assessment and is intended for educational and informational purposes only.

02

Software supply chain compromise and malicious packages

Key takeaways

1 Silent compromise

Supply chain attacks have become more targeted, abusing trusted dev tools to breach thousands via one flaw.

2 Credential fallout

CI/CD misconfigurations and long-lived tokens amplify risk, enabling large-scale cascading attacks.

3 Implicit trust

Blind trust in public code and tools is risky; DevOps teams need verification, monitoring, and secrets control.

Software supply chain compromise and malicious packages

Our recommendation

Analyzing all the supply chain incidents of 2025, we can conclude that organizations should harden their CI/CD and developer workflows by enforcing short-lived, least-privilege tokens; continuously monitor repositories, workflows, and dependencies for anomalies; and verify all third-party code, PoCs, and tools before using them. They need to treat automation, secrets, and open-source trust as critical security boundaries, not defaults.

The anatomy of supply chain attacks

Attackers used software supply chain attacks a lot in 2025. Who were their prime targets? After analyzing dozens of cases, we can assume that developers and security researchers were the main targets across DevOps platforms.

Multiple large-scale campaigns, including GhostAction, Shai-Hulud, the nx breach, and other software supply chain attacks, demonstrated how stolen tokens and malicious workflows could cascade into hundreds or thousands of affected repositories.

During 2025, attackers disguised malware as legitimate security tools, vulnerability scanners, and proof-of-concept exploits, tricking even experienced researchers into executing malicious code.

Several incidents revealed how misconfigured GitHub Actions workflows and long-lived npm or GitHub tokens enabled attackers to publish trojanized packages, exfiltrate secrets, and propagate malware automatically across dependency chains.

We could also observe how the hackers were abusing AI developer tools and automation to accelerate credential harvesting and reconnaissance. Despite rapid takedowns, many leaked credentials might still remain active, allowing follow-up attacks and prolonged exposure. Together, these events highlight a growing systemic risk in modern development pipelines and underscore the need for stronger controls over automation, secrets management, and open-source trust.

DECEMBER

Fake React2Shell scanner on GitHub exposed as malware

SOURCES:
Hackread

A GitHub repository claiming to be a vulnerability scanner for CVE-2025-55182, dubbed "React2Shell," was taken down after being exposed as malware.

The project, hosted on GitHub under the user niha0wa, was flagged after it was identified to contain a hidden malicious payload in the code. Analysis showed the script abused the Windows tool mshta.exe to execute a PowerShell command that downloaded a second-stage payload from an external server.

Showing up as a legitimate security utility, the tool specifically targeted security researchers investigating the React2Shell vulnerability.

GitHub removed the repository following reports, but experts warn that cached copies or forks may still circulate.

> **The incident highlights the growing trend of threat actors weaponizing fake security tools to compromise researchers themselves.**

SEPTEMBER

GhostAction campaign compromises over 80 repos

SOURCES:
gbhackers
Hackread
Security Boulvard

Security researchers identified a large-scale software supply chain attack, named GhostAction. In it, attackers managed to compromise secrets that belonged to 327 developers across 817 repositories on GitHub. The campaign was first detected after a malicious GitHub Actions workflow was injected into the popular FastUUID project, where it silently exfiltrated CI/CD secrets such as API keys and tokens to an attacker-controlled server.

Researchers at GitGuardian found that similar personalized workflows were planted across hundreds of public and private repositories, ultimately leaking over 3,300 secrets, including npm, PyPI, DockerHub, Cloudflare, and cloud credentials.

However, some of the repos were detected immediately and removed; there were still some of the stolen credentials that might have been abused, with attackers attempting unauthorized access to cloud and database environments.

> **The GhostAction incident highlights the growing danger of malicious CI/CD workflows and the need for tighter monitoring of automation pipelines and secrets management.**

Shai-Hulud and Shai-Hulud 2.0 compromises hundreds of npm packages

SOURCES:

Security Week
The Hacker News
Bleeping Computer
Cyber Security News

A major supply-chain attack, named Shai-Hulud, compromised hundreds of npm packages, making it one of the most severe JavaScript ecosystem incidents in 2025. Researchers report that attackers hijacked more than 40 developer accounts and injected self-replicating malware into over 180 packages initially, a figure that later climbed beyond 500 affected packages and more than 700 malicious versions. The malware abused post-install scripts to run TruffleHog, harvest secrets such as GitHub, npm, and cloud credentials, and then publish those secrets into newly created public repositories labeled "Shai-Hulud Migration" on GitHub.

Using stolen npm tokens, the worm automatically trojanized additional packages maintained by compromised developers, creating a cascading infection across the npm ecosystem.

Security firms, including Wiz, Socket, and ReversingLabs, warned that inter-dependencies between packages might amplify the blast radius, exposing developers, enterprises, nonprofits, and even security vendors. Although many leaked credentials were quickly revoked, dozens could still remain active, allowing the campaign to persist.

And it happened again... Later, in December 2025, there was a second wave of this attack, which exposed up to 400K dev secrets. According to the researchers from Wiz, the Shai-Hulud 2.0 attack could spread over 30K GitHub repos. Among the exposed secrets, 70% of repos had a contents.jason file (including GitHub usernames and tokens), 80% of the repos contained environment.jason file (including OS info, metadata, and GitHub credentials), around 50% had the truffleSecrets.jason file with TruffleHog scan results, and 400 repos contained actionSecrets.jason containing workflow secrets from GitHub Actions.

> **The best practices to protect yourself against such threats? DevOps should have immediate dependency audits, token rotation, and close monitoring of GitHub accounts for unauthorized repositories or workflows as the ecosystem works to contain the fallout.**

AUGUST

A stolen token allows hackers poison Nx packages and abuse AI tools

SOURCES:

The Hacker News
Bleeping Computer
The Register

Maintainers of the widely used nx build system warned its users in late August of a supply chain attack. This nx abuse allowed threat actors to publish malicious versions of the nx npm package and related plugins, which impacted millions of developers. The compromised packages, which were available only briefly on August 26, 2025, contained post-install scripts that scanned file systems, harvested credentials, and exfiltrated data to attacker-controlled GitHub repositories.

The attack stemmed from a vulnerable GitHub Actions workflow that allowed command injection via a crafted pull request title and ran with elevated permissions, ultimately exposing an npm publishing token. Using this access, attackers released tainted nx versions that could steal secrets, modify shell configuration files, and even abuse installed AI developer tools to automate reconnaissance.

The nx team removed the malicious packages immediately, rotated credentials, enforced stronger publishing controls, and urged users to assume compromise, rotate secrets, and audit affected systems.

> **This incident highlights how the software supply chain attacks grow and raises concerns about the misuse of developer AI tools as part of malicious campaigns.**

JULY

Attackers use a phished npm account to distribute Scavenger malware

SOURCES:

Cyber Security News

A sophisticated supply chain attack struck the JavaScript ecosystem on July 18, 2025, after attackers compromised several popular npm packages to distribute a malware strain named Scavenger. The campaign centered on eslint-config-prettier, alongside related developer tools such as eslint-plugin-prettier, snyckit, @pkgr/core, and napi-postinstall, with malicious releases appearing on npm despite no matching changes in their GitHub repositories.

The package maintainer later confirmed their npm account was hijacked via a phishing email, enabling attackers to publish trojanized versions including 8.10.1, 9.1.1, 10.1.6, and 10.1.7.

Analysis by Humpty's RE showed that the malware targeted Windows systems and focused on Chromium-based browsers. It harvested data from extensions, caches, and visited links to collect developer credentials and session information. The infection chain relied on an obfuscated install script that covertly executed a bundled DLL via rundll32.exe, using string concatenation to evade detection.

Additional analysis measures, including virtual machine checks and security tool detection, highlighted a deliberate effort to stealthily compromise developer environments through trusted tooling.

JUNE

Banana Squad malware campaign hijacks GitHub repos to target developers

SOURCES:

Info Security Magazine
The Hacker News
SC Media

Cybersecurity researchers found out about a malware campaign, dubbed Banana Squad, in which threat actors published 67+ trojanized GitHub repositories

masquerading as Python-based hacking and utility tools.

According to ReversingLabs, the operation is a continuation of a rogue Python campaign first observed in 2023, when malicious PyPI packages infected tens of thousands of Windows systems.

The repositories impersonated legitimate projects and targeted users searching for tools such as Discord account cleaners, game cheats, and bulk account checkers, delivering hidden backdoors and information-stealing payloads instead.

All identified repositories were removed by GitHub, but researchers still warned that similar campaigns might continue to exploit trust in open-source platforms as a software supply chain attack vector. That's why it's a good practice for developers to verify repository authenticity, be cautious of inactive single-repository accounts, and closely inspect code before execution to reduce exposure to such threats.

Critical path traversal flaw found in over 1,700 GitHub projects

SOURCES:

Cyber Security News

gbhackers

A large-scale security study uncovered a critical path traversal vulnerability - CWE-22. It affected around 1,750 open-source projects hosted on GitHub, many of which are widely used across the software ecosystem.

The flaw originates from a common Node.js pattern in which developers use path.join() with user-controlled URL input without proper sanitization, allowing attackers to access files outside intended directories using sequences like ../.

Researchers found that many of the vulnerable projects carry CVSS scores above 9.0, as the issue can be exploited remotely without authentication and may lead to data exposure or denial-of-service conditions.

The insecure pattern had been circulating since around 2010, spreading through developer resources such as GitHub Gists and highly viewed posts on Stack Overflow, despite repeated warnings from security-conscious developers. To assess the scale of the problem, researchers built an automated pipeline to detect, test, score, patch, and responsibly disclose the vulnerabilities, with only 14% of reported issues being fixed by June 2025.

The study also found that the flawed pattern has contaminated large language models, with up to 95% of AI-generated static file server examples being vulnerable. This might reinforce a cycle of insecure code reuse in modern development.

MAY

Supply chain threat escalates as malicious Go modules destroy systems

SOURCES:

The Hacker News
gbhackers
Bleeping Computer

A destructive supply chain attack was uncovered in April 2025 targeting the Go programming language ecosystem, in which attackers published three malicious Go modules through public repositories on GitHub. The bad actors presented 3 Go-modules, including *github[.]com/truthfulpharm/prototransform*, *github[.]com/blankloggia/go-mcp*, and *github[.]com/steelpoor/tlsproxy*, as legitimate libraries while embedding obfuscated code.

When executed on Linux systems, the modules fetched remote shell scripts from attacker-controlled domains that used the dd utility to overwrite the primary disk (/dev/sda) with zeros, rendering systems permanently unbootable.

Researchers noted that the attack exploits Go's decentralized module system, where direct imports and namespace ambiguity make typosquatting and malicious lookalike packages difficult to detect.

Unlike typical supply chain malware focused on credential theft or espionage, this campaign prioritized irreversible destruction, eliminating any possibility of data recovery or forensic analysis.

> **This incident highlights an escalating trend in open-source supply chain abuse and underscores the need for rigorous dependency audits, automated scanning, and continuous runtime monitoring to prevent similar attacks.**

APRIL

GitHub patches critical CodeQL flaw with massive supply chain potential

SOURCES:

gbhackers
Dev Class
Cyber Security News

A vulnerability in GitHub's CodeQL actions revealed how a publicly exposed token in workflow artifacts could have enabled a large-scale supply chain attack across hundreds of thousands of repositories. The issue, tracked as CVE-2025-24362, stemmed from a GitHub App installation token with full write permissions that was inadvertently stored in an artifact generated by a GitHub Actions workflow in the github/codeql-action repository.

Although the token was valid for just over one second, researchers demonstrated via a race-condition exploit that it could be extracted and abused to create branches, push commits, and retag releases. Because CodeQL workflows rely on mutable tags, an attacker could have replaced trusted tags with malicious code, causing default CodeQL workflows to execute attacker-controlled logic in dependent repositories.

Potential impacts included source code exfiltration, theft of GitHub Actions secrets, execution of malicious code in CI environments, and cache poisoning to maintain persistence.

GitHub responded within hours by disabling the affected workflow, preventing debug artifact uploads, and later patching the issue in CodeQL Action version 3.28.3, stating that no evidence of exploitation was found.

The "CodeQLEAKED" incident shows how short-lived secrets, mutable tags, and artifact handling in CI/CD systems can combine into high-impact supply chain risks.

MARCH

Supply chain attack hits Coinbase and hundreds of GitHub repositories

SOURCES:

The Hacker News
Be(in)Crypto
Bleeping Computer
The Hacker News
Info World

A targeted supply chain attack against Coinbase's open-source project agentkit has revealed a broader compromise of popular GitHub Actions used across hundreds of repositories.

The attacker forked Coinbase repositories — 'agentkit' and 'onchainkit' — and abused GitHub Actions with "write-all" permissions to inject malicious code into CI/CD workflows, aiming to steal secrets and enable further compromise.

While Coinbase detected the activity in mid-March 2025 and prevented deeper access, the same technique was later used to compromise the widely used tj-actions/changed-files and its dependency reviewdog/action-setup.

Security researchers from Endor Labs estimated that at least 218 GitHub repositories leaked secrets, including short-lived GitHub tokens and credentials for AWS, DockerHub, and npm. Though it's important to keep in mind that even limited exposure can have serious downstream implications due to the popularity of affected repositories, some of which had over 350K stars and 63K forks.

Later, in an updated Palo Alto Networks Unit 42 investigation, it was stated that ultimately compromised users of the tj-actions/changed-files GitHub Action were traced back to the theft of a personal access token (PAT) connected to SpotBugs. The attackers initially exploited a vulnerable GitHub Actions workflow in the SpotBugs ecosystem in late November 2024, using the pull_request_target trigger to leak a maintainer's PAT through a malicious pull request. That same token later enabled lateral movement into reviewdog, where the attackers poisoned the reviewdog/action-setup dependency that eventually flowed into tj-actions/changed-files.

All of these show how a single leaked token, combined with permissive CI/CD configurations, enabled a months-long, multi-project compromise across the GitHub Actions supply chain.

FEBRUARY

Lazarus targets developers with stealthy npm and GitHub supply chain attacks

SOURCES:

Techradar

Cyberscoop

A supply chain campaign by the Lazarus Group that targeted software and Web3 developers through malicious npm packages and GitHub repositories was uncovered by the security researchers.

According to SecurityScorecard, the campaign, dubbed Marstech Mayhem, involved JavaScript implants hidden inside legitimate-looking code, some of it committed by a GitHub account previously linked to Lazarus. The malware, known as Marstech1, scanned infected systems for cryptocurrency wallets such as MetaMask, Exodus, and Atomic, modifying browser configuration files to silently intercept transactions.

Researchers from STRIKE confirmed at least 233 victims across the US, Europe, and Asia, underscoring the scale of the operation and its focus on cryptocurrency theft to fund North Korea's state activities.

Meanwhile, Socket identified six additional Lazarus-linked npm packages distributing BeaverTail malware via typosquatted library names, some mimicking widely trusted modules. These packages used layered obfuscation, multi-stage payload delivery, and GitHub-hosted repositories to appear legitimate before being removed by GitHub.

The findings highlight Lazarus' continued reliance on open-source ecosystems like npm to compromise developer supply chains and siphon cryptocurrency assets.

JANUARY

Fake LDAPNightmare PoC on GitHub spreads infostealer malware

SOURCES:

Bleeping Computer

Techzine

Security researchers at Trend Micro uncovered a deceptive proof-of-concept (PoC) exploit for CVE-2024-49113, also known as LDAPNightmare, hosted on GitHub, that infected users with infostealer malware.

The malicious repository appears to be forked from a legitimate PoC published by SafeBreach Labs, exploiting confusion caused by an early misreference between CVE-2024-49113 and the more severe CVE-2024-49112.

Although CVE-2024-49113 is only a denial-of-service flaw affecting Windows LDAP and was patched by Microsoft in December 2024, the mix-up generated heightened interest that threat actors leveraged. Users who downloaded and executed the fake PoC received a UPX-packed executable that deployed a multi-stage PowerShell infection chain.

The final payload gathered extensive system and network information and uploaded it as a compressed archive to an external FTP server using hardcoded credentials.

Thus, developers and researchers should treat public exploit code with caution, always verify repository authenticity, and avoid executing obfuscated binaries without thorough inspection.

This report is based exclusively on publicly available information published by vendors on their official status pages, security advisories, databases, and publicly reported incidents documented in industry media. All data and statistics have been compiled from these publicly accessible sources. The analysis presented herein represents GitProtect's independent assessment and is intended for educational and informational purposes only.

Secret leaks, token exposure, and identity risks

Key takeaways

1 Credential abuse

Exposed tokens and OAuth abuse enable silent, high-impact compromises without exploits.

2 Hidden threats

Attackers hide malicious activity inside legitimate DevOps workflows to evade detection.

3 Continuous security

MFA and manual controls are insufficient; modern attacks require short-lived creds and continuous monitoring.

Secret leaks, token exposure, and identity risks

Our recommendation

The year 2025 shows that organizations should protect their data by enforcing least privilege access, rotating credentials frequently, and using short-lived tokens. It's critical to continuously monitor CI/CD workflows, repositories, and cloud accounts for anomalies. Adopt phishing-resistant MFA, secret management, and automated dependency scanning to reduce supply chain and identity-based risks.

What's worth knowing about secret leaks

Over 2025, attackers systematically shifted their focus from traditional perimeter breaches to the software supply chain and cloud identity layer, exploiting the trust embedded in developer tools, CI/CD pipelines, and authentication workflows.

From leaked GitHub tokens and exposed secrets to MFA-bypassing phishing and malicious dependencies, these incidents reveal how small misconfigurations or credential leaks can quietly escalate into full cloud and enterprise compromise. What makes these attacks particularly dangerous? Their stealth, as adversaries increasingly abuse legitimate platforms, OAuth flows, and developer infrastructure to blend into normal activity.

The timeline of this chapter highlights how this trend unfolded month by month, underscoring the urgent need for stronger controls across code, cloud, and identity ecosystems.

DECEMBER

Exposed GitHub PATs enable silent Cloud takeovers

SOURCES:
SC Media

Attackers are increasingly abusing exposed GitHub Personal Access Tokens (PATs) to steal GitHub Actions secrets and pivot into cloud environments, according to research from Wiz.

The company found that 73% of organizations store cloud credentials in GitHub Actions secrets, meaning a compromised PAT can let attackers impersonate developers, scan repositories, and access platforms like Amazon Web Services, Microsoft Azure, and Google Cloud Platform.

> **Info-Tech Research Group also said that a valid PAT can act like a "backstage pass," enabling data theft, malware deployment, and persistent access.**

Researchers urged organizations to treat PATs as privileged credentials, enforce least privilege, rotate tokens, and strengthen monitoring to reduce detection blind spots.

NOVEMBER

17,000 secrets found leaking from public GitLab repositories

SOURCES:
SC Media
Techradar
Bleeping Computer

Security researcher Luke Marshall uncovered more than 17,000 exposed secrets in public GitLab Cloud repositories, nearly three times the number found on Bitbucket, highlighting a widespread software supply chain risk.

Most of the leaked credentials were generated after 2018, but some, dating back to 2009, were still valid, indicating long-term exposure. Google Cloud Platform credentials made up the majority of the findings, followed by MongoDB keys, Telegram bot tokens, OpenAI keys, and GitLab secrets.

Marshall used GitLab's public API and automated TruffleHog scans to analyze 5.6 million repositories in just over 24 hours, notifying affected organizations with the help of AI tools and earning about $9,000 in bug bounties.

While many developers remediated the issues, some projects might remain exposed, leaving them vulnerable to data theft, cryptomining, and deeper infrastructure compromise.

OCTOBER

PhantomRaven campaign hides malware in npm dependencies to steal GitHub tokens from devs

SOURCES:

The Hacker News

Security researchers uncovered an active software supply chain attack targeting the npm ecosystem, involving more than 100 malicious packages designed to steal authentication tokens, CI/CD secrets, and GitHub credentials from developer environments.

The campaign, dubbed PhantomRaven by Koi Security, began in August 2025 and spread to about 126 packages with over 86K installs. Instead of declaring normal dependencies, the attackers hid malicious code in remote dynamic dependencies fetched from an attacker-controlled website, allowing the packages to appear harmless to automated security scanners. Once installed, lifecycle scripts could silently execute malware that fingerprints systems, harvests environment data, and exfiltrates information to external servers.

This campaign exploits blind spots in dependency analysis and abuses "slopsquatting," where developers install AI-hallucinated package names that sound legitimate.

SEPTEMBER

VoidProxy uses AitM tactics to evade MFA and steal Cloud accounts

SOURCES:

gbhackers
Cyber Security News
Cyber Press

A Phishing-as-a-Service platform called VoidProxy was identified to drive a surge in highly evasive phishing campaigns against Microsoft and Google accounts. First observed in August 2025, the operation used adversary-in-the-middle techniques to intercept credentials, MFA codes, and session tokens in real time, effectively bypassing common MFA protections.

According to analysis by Okta, attacks began with phishing emails sent from compromised legitimate email services, followed by multi-stage redirects, CAPTCHA challenges, and Cloudflare Workers that blocked security scanners while allowing real users through.

VoidProxy proxies victims to real login services, captures the session cookies returned after authentication, and immediately hands attackers access via an admin panel. Its use of disposable domains, dynamic DNS, and Cloudflare masking makes detection and takedown difficult.

Researchers warn that the platform represents a major escalation in phishing sophistication and stress the need for phishing-resistant authentication and layered defenses.

AUGUST

Attackers abuse Microsoft ADFS to steal Office 365 credentials

SOURCES:

Cyber Security News

A novel at that time phishing campaign was reported to actively steal Microsoft 365 credentials by abusing Active Directory Federation Services to redirect users from legitimate office.com links to malicious login pages.

Researchers at Push Security said the attack represented a major evolution in phishing, as it leverages Microsoft's own infrastructure to evade user suspicion and traditional security filters. Instead of phishing emails, attackers used search engine malvertising, luring victims who clicked ads that initially led to genuine Microsoft URLs. By configuring their own Microsoft tenant's ADFS settings, threat actors forced trusted Microsoft servers to redirect users to pixel-perfect phishing pages under attacker control. The final site operated as an attacker-in-the-middle proxy, capturing credentials and session cookies to bypass MFA protections.

This "ADFSjacking" technique is difficult to detect. Thus, it's important to monitor unusual ADFS redirects and block malicious ads as key mitigations.

Fake Microsoft OAuth apps fuel high-success MFA bypass attacks

SOURCES:

The Hacker News

SC Media

Cyber Security News

Attackers managed to abuse fake Microsoft OAuth applications to steal credentials and bypass multi-factor authentication in Microsoft 365 environments. First observed in early 2025 and still ongoing in mid-2025, the activity involved more than 50 malicious apps impersonating trusted services such as RingCentral, SharePoint, Adobe, and DocuSign to lure victims.

According to Proofpoint, attackers used phishing emails from compromised accounts to redirect users to legitimate Microsoft OAuth consent pages, then funneled them to attacker-in-the-middle phishing sites powered largely by the Tycoon Phishing-as-a-Service platform. Notably, victims were redirected to fake Microsoft login pages whether they accepted or denied the requested app permissions, eliminating user choice as a defense.

Researchers estimate nearly 3,000 attempted account compromises across more than 900 Microsoft 365 environments, with confirmed success rates exceeding 50%.

The campaign highlights how cloud app abuse, combined with OAuth trust, can significantly undermine identity security controls.

JULY

GitHub-hosted Raven stealer hijacks chromium browsers for data theft

SOURCES:
Cyber Security News

With a commodity infostealer called Raven Stealer, attackers curated their attack. Using a compact Delphi/C++ binary, they could steal browser credentials and exfiltrate them via the Telegram bot API.

First observed in mid-July 2025 on a GitHub repository operated by the ZeroTrace Team, the malware was distributed as UPX-packed executables or disguised attachments such as renamed invoice files. Once executed, Raven runs silently, harvesting passwords, cookies, autofill data, and cryptocurrency wallet information from Chromium-based browsers within seconds.

Cyfirma analysts found that the stealer organized stolen data into structured directories and compressed it before exfiltration. The stolen archives were sent using Telegram's /sendDocument API, allowing attackers to bypass many corporate firewalls that permit Telegram traffic. Its use of in-memory DLL injection into a suspended chrome.exe process further concealed activity and complicated detection.

Outlook-Based Malware is used for long-term email espionage

SOURCES:
Bleeping Computer
Techradar

The UK National Cyber Security Centre formally attributed the sophisticated "Authentic Antics" malware to APT28, also known as Fancy Bear, a cyber-espionage group. First observed in 2023, the malware targeted Microsoft Outlook, displaying fake Microsoft login prompts to steal credentials and OAuth 2.0 tokens that could grant access to Microsoft 365 services such as Exchange Online, SharePoint, and OneDrive.

Authentic Antics operates stealthily by running inside the Outlook process and exfiltrating stolen data through the victim's own email account, while disabling sent-mail records to avoid detection.

The NCSC noted the malware's high sophistication, including minimal disk footprint and reliance solely on legitimate services rather than traditional command-and-control infrastructure.

SharePoint online users hit by "Something Went Wrong" errors amid Microsoft incident

SOURCES:
Bleeping Computer

In July, it was reported that Microsoft was investigating an ongoing incident causing

intermittent access issues for SharePoint Online users, who were encountering “Something went wrong” and 503 errors when opening sites.

The company confirmed the problem began early July 4th and was linked to an authentication-related issue affecting some users on specific infrastructure. Initial findings pointed to a recurring authentication cookie, and Microsoft later said that a configuration change caused authentication request URLs to exceed character limits, leading to failures.

Engineers disabled the change for a subset of users to verify remediation before rolling it back more broadly. Until a permanent fix is deployed, Microsoft advised affected users to try accessing SharePoint Online via incognito or InPrivate browsing.

The incident was classified as a critical service issue, though Microsoft has not disclosed the full regional scope.

JUNE

Researchers warn of widespread GitHub Actions misconfigurations

SOURCES:
Dev Class
CSO

Security researchers warned that many open-source projects could expose themselves to takeover through insecurely configured GitHub Actions workflows that execute untrusted code.

An analysis by researchers at Sysdig found dozens of repositories using the risky pull_request_target trigger, which ran workflows with access to repository secrets and a write-privileged GITHUB_TOKEN. Because this trigger executed in the context of the base repository, attackers could submit malicious pull requests that exfiltrate credentials or modify workflows if safeguards are missing.

The issue was identified in high-profile projects, including repositories maintained by MITRE and Splunk, underscoring that even security-aware organizations are affected. Researchers demonstrated that exploited workflows could lead to credential theft, lateral movement, or full repository compromise. While some affected projects quickly fixed the flaws, Sysdig cautioned that insecure workflow configurations remain widespread across the open-source ecosystem.

PhaaS platforms drive global spike in MFA-bypassing phishing attacks

SOURCES:
gbhackers

A sharp rise in Adversary-in-the-Middle (AiTM) phishing attacks was observed globally in early 2025, driven by the rapid growth of Phishing-as-a-Service (PhaaS) platforms, according to researchers at Sekoia. These attacks used reverse proxy infrastructure to steal credentials and session cookies in real time, allowing attackers to bypass MFA protections on services like Microsoft 365 and Google. Turnkey PhaaS kits such as Tycoon 2FA, EvilProxy, and Sneaky 2FA have lowered the barrier to entry,

enabling even low-skill actors to launch large-scale campaigns that collectively generated over a million attacks in the first two months of 2025.

Security teams are warned that basic MFA and traditional email defenses are no longer sufficient, and that layered detection, behavioral monitoring, and user awareness are now critical.

APRIL

39 million secrets exposed: GitHub pushes new tools to tackle credential leaks

SOURCES:

Bleeping Computer

GitHub announced updates to its Advanced Security platform after detecting more than 39 million leaked secrets in repositories during 2024, including API keys and credentials that pose serious security risks.

According to GitHub, these exposures were identified through its secret scanning service, underscoring that secret leaks remain one of the most common and preventable causes of security incidents. Despite protections like Push Protection being enabled by default on public repositories, leaks persist due to developer convenience and secrets lingering in Git history.

To address this, GitHub made Secret Protection and Code Security available as standalone products, lowering the barrier for smaller organizations. New enhancements also included free organization-wide secret risk assessments, delegated bypass controls for push protection, AI-powered detection via Copilot, and improved scanners developed with cloud providers.

GitHub urged developers to eliminate hardcoded secrets, rely on secret managers, and integrate automated protections across CI/CD pipelines.

MARCH

Fake GitHub “Security Alerts” used to hijack thousands of developer accounts

SOURCES:

Bleeping Computer
Cyber Security News
SC Media

A large-scale phishing campaign targeted nearly 12K repositories on GitHub by posting fake “Security Alert: Unusual Access Attempt” issues that impersonated official account warnings.

The alerts falsely claimed suspicious logins from Reykjavik, Iceland, and directed developers to links that lead to authorization pages for a malicious OAuth app called gitsecurityapp. If approved, the app requests extensive permissions, including full access to public and private repositories, GitHub Actions workflows, gists, and even the ability to delete repositories.

Security researcher Luc4m first identified the campaign, noting that authorization

tokens were sent to attacker-controlled callback pages hosted on onrender.com. Once access is granted, attackers could steal code, modify workflows, or fully compromise developer accounts.

GitHub actively responded as the repository targeting numbers fluctuate, while affected users were urged to revoke suspicious OAuth apps and rotate credentials immediately.

FEBRUARY

Researchers expose massive GitHub abuse behind Redox stealer operation

SOURCES:

Cyber Security News
gbhackers

Security researchers uncovered a large-scale malware campaign abusing GitHub repositories to distribute the Redox information stealer through fake software cracks and gaming mods. More than 1,100 repositories were identified offering "free" versions of premium software and cheats for popular games, luring victims into downloading archives that deploy Python-based malware.

Once executed, Redox harvested browser credentials, cryptocurrency wallets, gaming accounts, and financial files, exfiltrating the stolen data through Discord webhooks, AnonFiles uploads, and Telegram channels.

The operation relied heavily on social engineering, including AI-generated README files, forged VirusTotal screenshots, and SEO-optimized repository topics to appear legitimate.

Researchers found that only a small fraction of confirmed malicious repositories carried visible user warnings, allowing most to evade detection. The campaign proves how open platforms can be systematically exploited at scale using automation, obfuscation, and trust abuse.

Infostealer infections expose millions of sensitive accounts worldwide

SOURCES:

Infosecurity Magazine

Cybersecurity researchers at Hudson Rock warned that sensitive US corporate, government, and military networks may be at risk after uncovering widespread credential theft linked to infostealer malware.

Their analysis of cybercrime marketplaces found stolen credentials allegedly tied to major defense contractors, federal agencies, and US military entities being sold for as little as $10 per log. These logs could provide access to corporate email, VPNs, internal development platforms (such as GitHub, Jira, and Confluence), and, in some cases, active session cookies that enable multi-factor authentication bypass. Hudson Rock cautioned that even organizations not directly infected could still be compromised through partners or suppliers whose credentials have been stolen.

JANUARY

Multiple GitHub tools patched after critical credential vulnerabilities

SOURCES:

Cyber Security News

Security researchers uncovered multiple critical vulnerabilities across Git-related projects, including GitHub Desktop, Git Credential Manager, Git LFS, GitHub CLI, and GitHub Codespaces, that could allow attackers to leak user credentials.

The flaws stemmed from improper handling of Git's text-based credential exchange protocol, where inconsistencies in parsing newline and carriage return characters enabled credential exfiltration through crafted URLs or configuration files. Several issues, tracked under CVEs such as CVE-2025-23040 and CVE-2024-50338, made it possible to redirect usernames, passwords, or access tokens to attacker-controlled hosts. A particularly severe flaw in GitHub Codespaces caused the credential helper to return the GITHUB_TOKEN regardless of the requested host.

The vulnerabilities were discovered through the GitHub Bug Bounty program. In response, Git and GitHub introduced stronger input validation, host checks, and new default protections to prevent carriage return smuggling and unauthorized credential disclosure.

04

DevOps platform outages and service reliability incidents

Key takeaways

1 Downtime landscape

Configuration and automation errors caused most major cloud and DevOps outages in 2025.

2 Centralized risk

Centralized SaaS and cloud platforms act as single points of failure at a global scale.

3 Resilience gap

Short outages still caused major disruption, exposing weak resilience and change controls.

04

DevOps platform outages and service reliability incidents

Our recommendation

To reduce outage impact, organizations should build a reliable backup and disaster recovery plan. The security strategy should be designed the way to avoid single-provider dependencies, implement multi-region and multi-service redundancy, and test their recovery regularly. Organizations should be able to maintain offline access, backup communication paths, and clear incident response plan to sustain operations during cloud disruptions.

The state of outages and deteriorated performance

In 2025, a total of 607 incidents were recorded across GitHub, GitLab, Azure DevOps, and Atlassian's Jira and Bitbucket, accumulating 9,255 hours and 26 minutes of impact. Of these, 156 incidents were classified as critical or major, lasting 1,769 hours and 43 minutes.

Year-over-year analysis shows a 69% increase in critical and major incidents compared to 2024, when only 48 such incidents were recorded, indicating a significant escalation in high-severity disruptions.

Incident distribution across the year was uneven. The first half of 2025 experienced more critical and major incidents - 81 incidents - than the second half - 75 incidents.

Q2 2025 recorded the highest quarterly total with 47 incidents, while June stood out as the single month with the most incidents overall - 21, marking a peak period of operational instability across the analyzed platforms.

Looking at platform-specific data, GitLab recorded the highest number of critical and major incidents - 62 incidents - with a cumulative duration of 754 hours and 49 minutes, followed by Jira with 44 incidents lasting 728 hours and 23 minutes, and GitHub with 35 incidents lasting 172 hours and 32 minutes. Bitbucket experienced 15 critical and major incidents totaling 113 hours and 59 minutes, while Azure DevOps recorded none during the analyzed period.

On GitHub, December had the highest number of critical and major incidents - 5 incidents, but the most significant impact occurred in March, when just two incidents resulted in 58 hours and 21 minutes of degraded performance. Between March 29 and March 31, GitHub experienced a prolonged service outage caused by expired credentials used for an internal service, which was mitigated by renewing the credentials and redeploying affected services.

GitLab's most severe month was July, with 11 critical and major incidents lasting a

combined 252 hours and 51 minutes. Among these was a single critical incident lasting 49 hours and 46 minutes, which caused service disruption due to OAuth refresh tokens being incorrectly deleted.

For Bitbucket, May was the most disruptive month, recording 5 critical and major incidents lasting 84 hours and 6 minutes in total. The longest of these was a major incident lasting 49 hours and 44 minutes, during which pipelines failed to trigger, impacting some customers' workflows.

In Jira, June saw the highest number of critical and major incidents - 9 incidents, lasting 47 hours and 42 minutes in total. The longest incident that month lasted 23 hours and 56 minutes and affected Jira Software and Jira Service Management due to Forge invocation errors impacting some instances in the Singapore region.

Across all incidents, degraded performance dominated both frequency and duration, accounting for 374 incidents, which makes 62% of all, and 3,059 hours and 32 minutes - 34% - of total time. Maintenance activities represented 4% of incidents but consumed 29% of total time, while no-impact incidents accounted for 8% of incidents and 16% of total duration.

FIGURE 5

Incident distribution by impact (total 374 incidents)

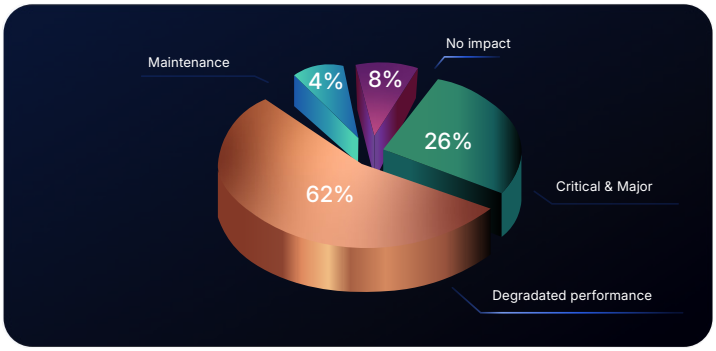

Overall, the data shows that while critical and major incidents represent a smaller share of total events, they consume a disproportionate amount of time and continue to increase year over year, signaling growing operational risk across developer and collaboration platforms.

FIGURE 6

Incident impact by total duration (9255hrs 26min total)

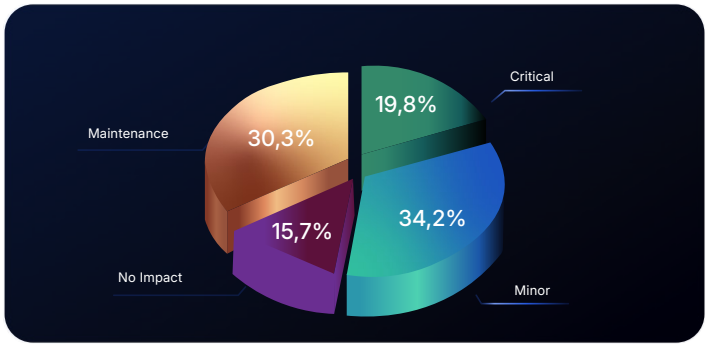

In this Chapter, we will remember the loudest outages of 2025. We will cover not only GitHub, Atlassian, and GitLab disruptions, but also Microsoft, AWS, and Cloudflare outages that were impacting developers all around the world in 2025.

DECEMBER

Microsoft 365 outage disrupts businesses across Australia

SOURCES:

Cyber Security News

In December, Microsoft 365 users across Australia were experiencing significant disruptions to Microsoft 365 services, with Microsoft Teams among the most severely affected tools. Microsoft confirmed the incident on December 10, 2025, classifying it as a "service degradation" that has led to latency, login failures, and service timeouts.

The issue began during local business hours and has disrupted access to core collaboration and cloud-based applications relied on by enterprises and individuals alike. Microsoft said the impact was limited to the Australian region.

Two Cloudflare incidents in weeks highlight risks of centralized infrastructure

SOURCES:

The Register
Forrester
SC Media
Fin Tech Weekly

Two significant outages linked to Cloudflare occurred within a short span, reinforcing concerns about the fragility of modern internet infrastructure.

The first incident caused multiple high-profile websites to become unreachable, with users encountering error messages, failed logins, and even outage-monitoring tools going offline.

Shortly afterward, a separate global outage temporarily disrupted services including X, ChatGPT, and numerous other platforms, with Cloudflare attributing the issue to an automatically generated configuration file intended for security management.

In another related incident, routine maintenance and a change to Web Application

Firewall request parsing led to dashboard, API, and site failures across regions, again affecting customers worldwide.

While Cloudflare restored services within hours and confirmed the incidents were not cyberattacks, the repeated failures highlighted how configuration changes and maintenance errors can cascade rapidly.

Together, the incidents intensified industry debate around cloud concentration risks, shared points of failure, and the need for stronger resilience planning across the digital ecosystem.

OCTOBER

AWS outage triggered by DNS automation bug causes global internet disruption

SOURCES:

The Guardian

Cyber Press

Amazon Web Services suffered a major outage on October 20, 2025, after a failure in its US-East-1 region triggered widespread disruptions across the internet.

The incident began with elevated error rates in DynamoDB caused by a latent defect in AWS's automated DNS management system, specifically an empty DNS record that failed to self-repair and required manual intervention. As the issue cascaded, core services such as compute and storage were impacted, knocking thousands of platforms offline, including messaging apps, streaming services, financial platforms, and internet-connected devices. Outage trackers recorded millions of user reports globally, highlighting the scale of disruption and AWS's central role in digital infrastructure.

Although AWS restored services within hours and later confirmed full recovery, the incident exposed the risks of cloud centralization and single points of failure. The company has since disabled parts of its DNS automation worldwide while implementing additional safeguards to prevent similar outages.

Microsoft 365 outage in North America caused by infrastructure misconfiguration

SOURCES:

Bleeping Computer

The Register

Microsoft experienced a Microsoft 365 outage in North America after a network infrastructure misconfiguration disrupted access to services, including Teams, Exchange Online, and the admin center.

The incident began around 18:10 UTC and was resolved roughly an hour later after Microsoft rerouted traffic to healthy infrastructure.

A separate but related disruption also caused authentication and MFA failures, triggered by imbalanced directory operations during a period of high traffic. Microsoft said it had rebalanced affected systems and was reviewing configuration and traffic management policies to prevent similar outages in the future.

AUGUST

Millions affected as GitHub suffers global service degradation

SOURCES:

Web Pro News

GitHub experienced a major service disruption in the early hours of August 12, 2025, affecting core features such as repositories, pull requests, issues, and API access. The incident, tracked as 9rfydl2xdqqj, began around 15:20 UTC and was marked by increased latency, stale data, and intermittent Git operation failures that disrupted workflows worldwide.

GitHub acknowledged the issue quickly, attributing the impact to problems linked to recent database infrastructure changes and initiating a rollback to stabilize services. Developers and enterprises reported blocked deployments and halted CI/CD pipelines, underscoring GitHub's critical role in the global software supply chain.

Partial recovery was observed within hours as engineers applied flow controls and rate limits to manage load, and later the incident was fully mitigated.

JULY

GitHub outage disrupts APIs and pull requests worldwide

SOURCES:

Cyber Press
Cyber Security News
gbhackers

GitHub suffered a widespread outage on July 28–29, 2025, disrupting core services, including API requests, issues, and pull requests for developers worldwide.

The incident began around 22:40 UTC and was traced to networking-related infrastructure problems, at its peak affecting roughly 4% of requests. While the disruption caused intermittent failures rather than a total shutdown, it still delayed deployments and workflows across global CI/CD pipelines.

GitHub deployed fixes in stages and fully restored all affected services by 02:06 UTC on July 29th.

Microsoft restores Teams after global outage

SOURCES:

Cyber Security News
Neowin

A significant outage affecting Microsoft Teams took place in July, temporarily preventing users from accessing core communication and collaboration features. The disruption impacted logins, chat, and video conferencing, affecting organizations and individuals who rely on Teams for remote work and virtual meetings.

Microsoft acknowledged the issue and directed administrators to incident reference TM1112332 in the admin center for ongoing updates. While automated recovery mechanisms were activated to restore service, the company stated that the underlying cause of the outage was still under investigation.

Microsoft emphasized that technical teams were prioritizing both immediate remediation and long-term analysis to prevent similar incidents. In a later update, the company reported that service telemetry indicated a full recovery, suggesting Teams access had largely been restored.

APRIL

Degraded Jira performance hits multiple Atlassian products

SOURCES:

Bleeping Computer

Cyber Security News

In April, Atlassian experienced a widespread service disruption that caused degraded performance across several Jira products, impacting users worldwide. The incident, which began at approximately 06:59 UTC on April 16, 2025, affected Jira, Jira Service Management, Jira Work Management, and Jira Product Discovery.

During the outage, users encountered error messages, slow or failed loading of the Jira user interface, and non-functional or timing-out dashboard widgets. These issues disrupted core workflows such as sprint planning, issue tracking, and service management for teams that rely heavily on Jira dashboards for real-time visibility.

Atlassian confirmed that the problem centered on certain Jira Dashboard widgets failing to load properly. After several hours of investigation, the company identified the root cause around 12:48 p.m. ET and began remediation efforts. Full recovery was completed later that day, with Atlassian confirming that services returned to normal operation by approximately 6:57 p.m. ET on April 16.

The company didn't disclose the specific technical cause or the total number of affected users, but stated that the issue was fully resolved and that all impacted Jira services were operating normally.

JANUARY

GitHub resolves major outage caused by caching infrastructure failure

SOURCES:

Bleeping Computer

GitHub resolved a major service outage that caused widespread disruptions across multiple core features, including pull requests, issues, repository access, commits, and GitHub Actions.

The incident was first acknowledged on GitHub's official status page after users reported degraded availability for Issues and Pull Requests, along with timeouts affecting various services. According to GitHub, the root cause was linked to problems within its caching infrastructure, which led to slow responses and request failures across the platform.

Thousands of users reported issues via DownDetector, citing website errors, server connection problems, and failures when using GitHub Actions. Affected users encountered error messages such as: "We couldn't respond to your request in time.

Sorry about that. Please try refreshing and contact us if the problem persists."

GitHub classified the disruption as a "major outage," though it did not disclose the total number of affected users or specific regions. As mitigation efforts progressed, the company warned of temporary additional disruptions while failing over one of its primary caching hosts.

By January 30 at 10:30 EST, GitHub confirmed that the failover process was complete and that services had returned to normal operation. The company stated it would continue monitoring systems to ensure stability.

Bitbucket is "Hard Down" in worldwide service disruption

SOURCES:
Bleeping Computer

In January, a major global outage took Atlassian Bitbucket Cloud services completely offline, with the company describing the platform as "hard down." The incident lasted for more than two hours and affected customers worldwide, according to hundreds of reports logged on DownDetector.

Atlassian confirmed that all Bitbucket Cloud services were impacted, including the website, API, Git over SSH and HTTPS, authentication, pipelines, webhooks, Git LFS, source downloads, and email delivery. In an update posted to its official status page, Bitbucket said it had identified a saturated Bitbucket database as the cause, which is disrupting all repository and Git operations.

GitHub's Git outage disrupts developers after a faulty configuration change

SOURCES:
Dev Class

GitHub experienced at least 49 minutes of Git downtime on January 13 due to a faulty internal load balancer configuration. Hovewer, some users were reporting outages closer to two hours.

While not all GitHub services were down, the disruption affected core Git operations and downstream systems such as CI/CD pipelines relying on GitHub repositories. GitHub said it will improve monitoring, deployment practices, and automated mitigation to reduce detection and recovery time for similar incidents.

Data source: Incident data compiled from official status pages of GitHub, GitLab, Azure DevOps, and Atlassian/Jira covering the period January-December 2025. Incident classifications and severity levels as reported by the respective vendors.

05

Critical vulnerabilities, exploits, and security patches

Key takeaways

1 High-severity flaws

High- and critical-severity flaws dominated 2025, accounting for 59% of all patched DevOps vulnerabilities.

2 Vulnerability rise

Vulnerability volume and severity both accelerated sharply in H2, with Q4 emerging as the most active and riskiest quarter.

3 Flaws severity

Core DevOps platforms could expose enterprise risk through authentication, CI/CD, and third-party dependency flaws.

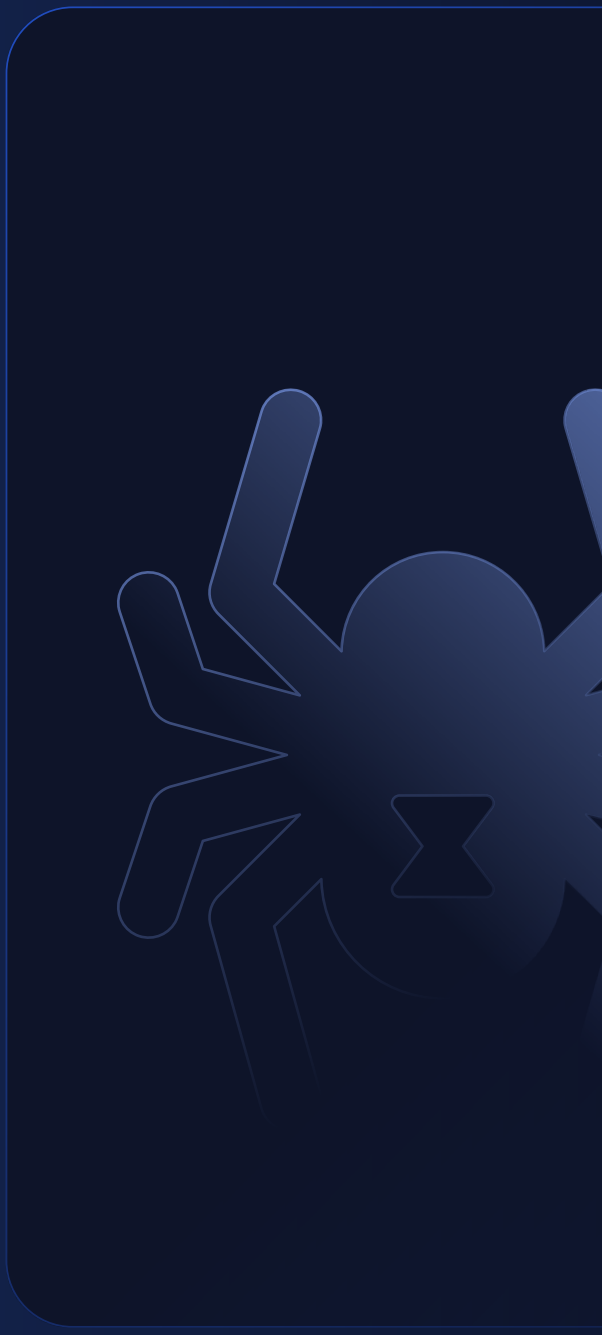

05

Critical vulnerabilities, exploits, and security patches

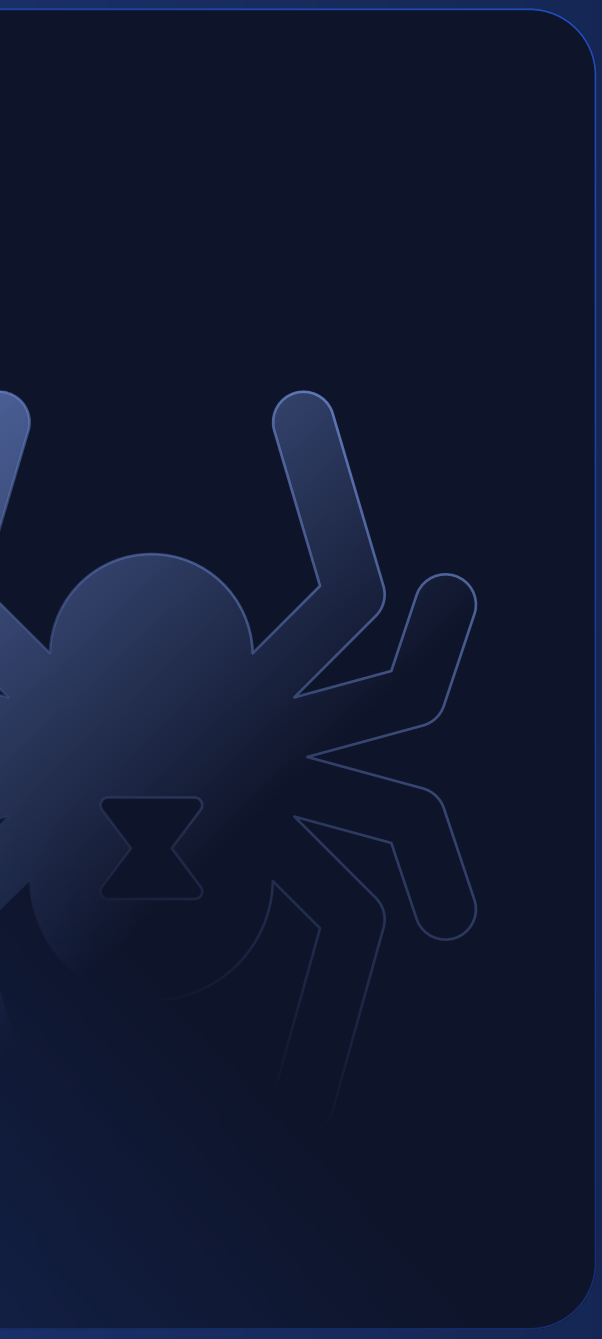

Our recommendation

The analysis of 2025 shows that organizations should keep an eye on vulnerability updates and recommendations DevOps platforms release. Implementing on-time patches, auditing third-party dependencies, restricting guest access, and continuous monitoring for anomalies to reduce the risk from high-severity vulnerabilities and supply-chain exploits are also necessary measures.

Patched vulnerabilities across DevOps platforms

In 2025, major DevOps platforms, such as GitHub, Atlassian's Bitbucket and Jira, GitLab, and Microsoft's Azure DevOps, patched a total of 236 vulnerabilities.

Of these, 14 were rated critical and 126 high severity, meaning 59% of all patched vulnerabilities could be exploited to cause serious damage, including unauthorized access to sensitive data, privilege escalation, or partial system compromise. The remaining vulnerabilities were classified as medium severity - 75, and low severity - 21.

FIGURE 7

Severity breakdown of patched vulnerabilities in 2025

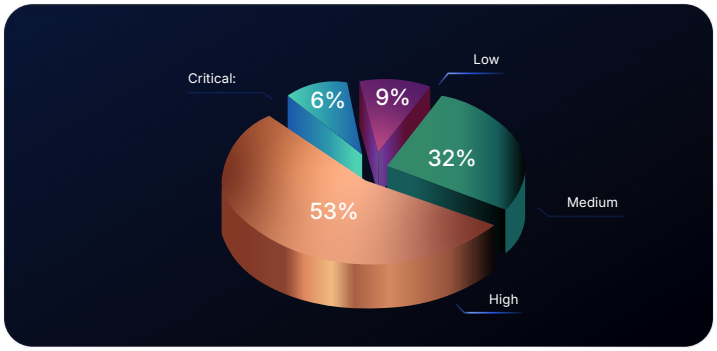

Comparing the first and second halves of 2025 shows a clear upward trend. Between January and July, platforms patched 97 vulnerabilities, while August to December saw 139 patched vulnerabilities. Thus, we can see a 30% increase in the second half of the year.

Quarterly data confirms steady growth:

- Q1: 45 vulnerabilities
- Q2: 52 vulnerabilities (+16%)
- Q3: 60 vulnerabilities (+15%)
- Q4: 79 vulnerabilities (+32%)

Q4 alone accounted for 34% of all patched vulnerabilities, representing a 76% increase from Q1 to Q4, making it the most active quarter of the year.

FIGURE 8

Patched vulnerabilities by quarter (2025)

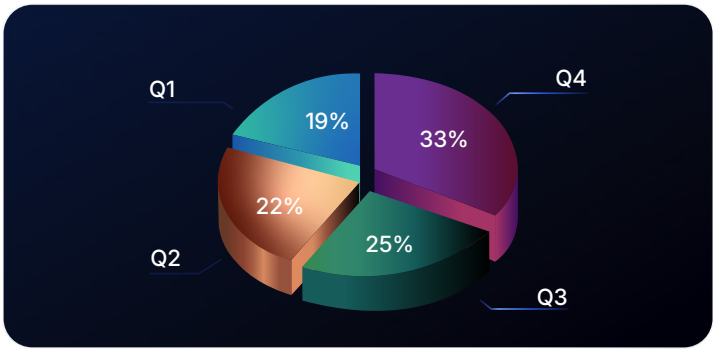

The increase was not limited to volume; it also applied to severity. While Q1 and Q2 each saw only two critical vulnerabilities, Q3 recorded four, and Q4 six. Overall, the number of patched critical vulnerabilities rose from 4 in H1 to 10 in H2.

High-severity vulnerabilities followed a similar pattern, increasing from 39 in the first half of the year to 87 in the second half, a 123% increase from H1 to H2.

November 2025 saw the highest number of patched vulnerabilities - there were a total of 36, accounting for 15% of the annual total.

FIGURE 8

Month-by-Month breakdown of patched vulnerabilities (2025)

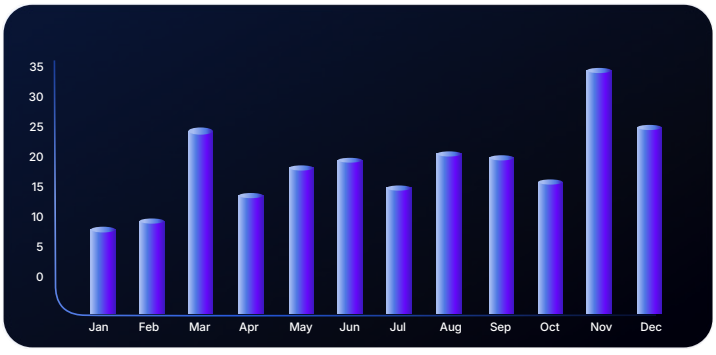

When it comes to platforms, GitHub patched 18 vulnerabilities in 2025. Five affected GitHub Enterprise Server versions and all the patches were released before June 19, 2025, after which GitHub announced that no further patches would be issued for GitHub Enterprise Server versions.

For GitHub Cloud, 13 vulnerabilities were patched, most rated critical or high severity. Four were critical, including CVE-2025-178, a composite GitHub Action vulnerability with a CVSS score of 10.0, allowing arbitrary code execution and patched in version 1.0.6.

Microsoft patched a large number of vulnerabilities across its ecosystem in 2025. Within Azure DevOps, two critical vulnerabilities were fixed. One of them, CVE-2025-47158, was an authentication bypass vulnerability that could allow an unauthenticated attacker to manipulate assumed-immutable data, enabling network-based privilege escalation.

GitLab patched 129 vulnerabilities in 2025, down from 153 in 2024, representing a 16% decrease year over year. The highest activity occurred in August and September, with 16 vulnerabilities patched each month. Only two vulnerabilities were identified as critical — CVE-2025-25291 and CVE-2025-25292 — both related to the ruby-saml library used when SAML SSO was enabled. These flaws, patched in March, could allow an attacker with a legitimately signed SAML response to authenticate as another valid user within the same SAML environment without knowing their credentials.

Atlassian patched 48 vulnerabilities in Bitbucket and 39 in Jira in 2025. All were classified as critical or high severity.

- Bitbucket: 4 critical vulnerabilities
- Jira: 2 critical vulnerabilities

Two of those vulnerabilities received the highest possible CVSS score of 10.0:

- CVE-2024-38999 (Bitbucket DC & Server): remote code execution via a third-party dependency
- CVE-2025-66516 (Jira Software DC & Server): XML external entity injection in a third-party dependency, allowing unauthenticated attackers to significantly impact confidentiality, integrity, and availability

Comparing Bitbucket's patched vulnerabilities in 2024 and 2025 reveals a 58% increase year over year. In 2024, Bitbucket patched 20 vulnerabilities.

Let's see how patched vulnerabilities across Atlassian, GitHub, GitLab, and Microsoft were covered in the media throughout 2025.

DECEMBER

Atlassian addresses dozens of high-risk vulnerabilities in Confluence and Jira

SOURCES:

Security Week

SC Media

Atlassian released fixes for nearly 30 third-party security vulnerabilities across multiple products, including a maximum-severity XML external entity (XXE) injection flaw in Apache Tika tracked as CVE-2025-66516. This vulnerability affects the toolkit's core, PDF, and parser modules and impacts several Atlassian products that use Tika, including Confluence, Bamboo, Crowd, Fisheye/Crucible, Jira, and Jira Service Management. Successful exploitation could lead to information disclosure, server-side request forgery (SSRF), denial-of-service (DoS), and, in some scenarios, remote code execution.

In addition, Atlassian patched other critical third-party issues. Confluence received updates addressing a prototype pollution flaw in webpack loader-utils, tracked as CVE-2022-37601, while Jira and Jira Service Management were updated to fix a critical prototype pollution vulnerability in the ZRender library, tracked as CVE-2021-39227.

Beyond these, Atlassian resolved more than two dozen additional high-severity vulnerabilities spanning categories such as XXE, DoS, SSRF, RCE, prototype pollution, file inclusion, and improper authorization, reducing the overall third-party risk exposure across its product portfolio.

NOVEMBER

GitLab patches high-risk flaws enabling credential theft and DoS

SOURCES:

Cyber Press

In November, GitLab released critical security patches to address six vulnerabilities affecting Community and Enterprise Edition versions 18.6.1, 18.5.3, and 18.4.5, urging all self-managed users to upgrade immediately.

The most severe issues included CVE-2024-9183, a high-risk race condition in the CI/CD cache that could allow authenticated attackers to steal credentials and escalate privileges, and CVE-2025-12571, an unauthenticated denial-of-service flaw in JSON validation that could crash GitLab instances with crafted payloads.

Additional flaws included an authentication bypass in account registration that could let unauthenticated users join arbitrary organizations, as well as further DoS and information disclosure vulnerabilities.

Microsoft Teams guest access vulnerability lets attackers bypass defender protections

SOURCES:

Hackread

Cyber Security News

A security gap in Microsoft Teams' B2B guest access was identified by Ontinue. This vulnerability allows attackers to bypass Defender for Office 365 protections, leaving users exposed to phishing and malware. The issue stems from Teams' "Chat with Anyone" feature, which enables cross-tenant chats by default, allowing malicious tenants to deliver unsafe links and files without triggering Safe Links, Safe Attachments, or Zero-hour Auto Purge protections.

Attackers can exploit this by creating basic or trial tenants, sending pretextual invites via email or LinkedIn, and luring victims into unsafe environments where malware can deploy, and data can be exfiltrated undetected.

The problem affects all Teams licenses, including low-cost SMB plans, and disabling outbound invites does not protect users from inbound threats.

Security experts recommend restricting guest invitations to allowlisted domains, deploying cross-tenant access policies, and training users to reject unsolicited invites to mitigate this exposure.

OCTOBER

Atlassian patches high-severity Jira flaw allowing arbitrary file writes

SOURCES:

Cyber Press

gbhackers

Cyber Security News

In November, Atlassian disclosed a high-severity path traversal vulnerability in Jira Software Data Center and Server that allowed authenticated attackers to arbitrarily write files accessible to the Jira Java Virtual Machine (JVM) process.

Tracked as CVE-2025-22167 with a CVSS score of 8.7, the flaw can be exploited remotely by low-privileged users to modify configuration files, application data, or other sensitive resources, potentially leading to service disruption or system compromise.

The vulnerability affected versions introduced from Jira 9.12.0 onward across multiple release branches, making it a widespread risk for enterprises relying on self-managed Jira deployments.

Atlassian urged its customers to upgrade immediately to fixed versions such as 9.12.28, 10.3.12, or 11.1.0, depending on their branch.

Security teams are advised to prioritize remediation, as the flaw poses high confidentiality, integrity, and availability risks, particularly in shared or multi-tenant environments.

GitLab fixes critical runner hijack and denial-of-service vulnerabilities

SOURCES:

Cyber Press
Cyber Security News

GitLab released urgent security updates for both Community Edition and Enterprise Edition in November to patch multiple high-risk vulnerabilities, including several denial-of-service (DoS) flaws that could allow attackers to crash GitLab instances remotely.

The patched versions — 18.5.1, 18.4.3, and 18.3.5 — addressed critical DoS issues in event collection, JSON validation, and upload endpoints, some of which could be exploited without authentication using specially crafted payloads.

One of the most severe fixes also resolved an improper access control flaw in the runner API that could allow authenticated users to hijack CI/CD runners across projects in GitLab EE.

Additional patches corrected authorization bypasses in pipeline builds, business logic errors in group memberships, and missing authorization checks in quick actions.

AUGUST

GitLab users urged to update after critical XSS and account takeover vulnerabilities

SOURCES:

Cyber Security News
Cyber Press
gbhackers

Multiple high-severity vulnerabilities in both GitLab Community and Enterprise editions were addressed in August. Those vulnerabilities could allow account takeovers and stored cross-site scripting (XSS) attacks. The flaws, some affecting versions as far back as 14.2, include CVE-2025-6186, CVE-2025-7734, and CVE-2025-7739, which enabled attackers to inject malicious HTML into work items, labels, or the blob viewer to execute unauthorized actions.

Another high-severity issue, CVE-2025-8094, allowed maintainers to misuse project API permissions, potentially causing denial of service to other users' CI/CD pipelines.

GitLab urged all self-managed installations to update immediately to versions 18.0.6, 18.1.4, or 18.2.2, noting that patched GitLab.com instances and GitLab Dedicated customers were already protected.

JULY

Local file inclusion bug in Microsoft 365 PDF export fixed after disclosure

SOURCES:

gbhackers
Cyber Security News
Cyber Press

A critical Local File Inclusion (LFI) vulnerability in Microsoft 365 was patched in July. That vulnerability could allow attackers to access sensitive server-side files through the platform's Export to PDF functionality.

The flaw stemmed from an undocumented behavior in Microsoft Graph APIs that

permitted HTML-to-PDF conversion without proper input validation or file path restrictions. By embedding malicious HTML tags such as <embed>, <object>, or <iframe>, attackers could force the PDF conversion service to include local system files like configuration files, credentials, and potentially application source code. Exploitation required only uploading a crafted HTML file, requesting PDF conversion via the API, and downloading the resulting PDF containing the exposed data.

Microsoft classified the issue as “Important” severity and remediated it by tightening HTML processing and enforcing strict file path sanitization.

MAY

Prompt injection bug exposes new risks in GitHub AI coding agents

SOURCES:

Cyber Security News
Dev Class
Cyber News

A critical vulnerability was discovered in the widely used GitHub Model Context Protocol (MCP) server that allowed attackers to steal data from private repositories through malicious prompt injection.

Researchers found that attackers can plant hidden instructions inside seemingly legitimate issues in public repositories, which hijack AI coding agents when they review those issues. Once triggered, the compromised agent could be coerced into accessing and leaking sensitive information from private repositories without any direct breach of the underlying systems.

The vulnerability, identified by Invariantlabs, affected all agent systems using the GitHub MCP integration, which had more than 14,000 stars, and showed a major blind spot in current AI agent security models.

GitLab fixes 11 security flaws, including high-impact DoS and auth issues

SOURCES:

gbhackers

Critical security updates patching 11 vulnerabilities across its GitLab Community and Enterprise editions were released in May. Several flaws, which enabled denial-of-service (DoS) attacks, posed the greatest risk.

The most severe issue, CVE-2025-0993, with a CVSS score of 7.5, allowed authenticated attackers to exhaust server resources via an unprotected large blob endpoint, potentially causing prolonged outages. Additional DoS vulnerabilities affected Kubernetes integrations, note positioning, and Discord webhooks, expanding the attack surface across both core features and third-party integrations.

The update also fixed authentication and data exposure flaws, including a SAML weakness that could bypass two-factor authentication and bugs that might leak CI/CD secrets and user email addresses.

Atlassian fixes DoS and privilege escalation risks in

SOURCES:
gbhackers

Atlassian published its May 2025 Security Bulletin, detailing eight high-severity vulnerabilities affecting multiple Data Center and Server products across its enterprise portfolio.

The flaws, identified through bug bounty research, penetration testing, and third-party dependency analysis, include denial-of-service (DoS) conditions and a notable privilege escalation risk. Five of the vulnerabilities stem from weaknesses in third-party components such as Apache Tomcat, XStream, JSON Smart, and Netty, impacting products including Bamboo, Confluence, Fisheye, and Jira.

Also, Atlassian disclosed a high-severity privilege escalation vulnerability (CVE-2025-22157) affecting Jira Core and Jira Service Management Data Center deployments.

All issues were patched in updated releases published in May 2025, with Atlassian strongly urging customers to upgrade immediately. The bulletin highlighted growing supply-chain security challenges in enterprise software and reinforced the need for timely patch management in mission-critical environments.

Microsoft fixes maximum-severity Azure DevOps vulnerability

SOURCES:
The Hacker News
The Register
Bleeping Computer

On Microsoft's May 2025 Patch Tuesday, Microsoft addressed 78 flaws, 11 of which were rated critical. One of the most critical ones, CVE-2025-29813 with a CVSS score of 10.0 (maximum possible!), was patched in Azure DevOps.

This was a privilege escalation and authentication bypass vulnerability in Azure DevOps that could allow an unauthorized attacker to elevate privileges over a network.

APRIL

GitHub urges immediate updates after severe enterprise server vulnerabilities

SOURCES:
Cyber Security News

Urgent security updates for GitHub Enterprise Server were released in April, after the company identified multiple high-severity vulnerabilities, including a critical remote code execution flaw tracked as CVE-2025-3509.

The most serious issue allowed attackers with site administrator or repository modification privileges to exploit dynamically allocated ports during hot patch upgrades, potentially leading to full system compromise under specific conditions.

Additional flaws included a medium-severity authorization bug (CVE-2025-3124) that exposed private repository names in the Advanced Security Overview and a high-risk

cross-site scripting vulnerability (CVE-2025-3246) enabling malicious HTML or CSS injection through Markdown math blocks.

The vulnerabilities affected GitHub Enterprise Server versions 3.13.0 through 3.16.1 and where patched in releases 3.13.14, 3.14.11, 3.15.6, and 3.16.2.

MARCH

Critical ruby-saml vulnerabilities could allow user impersonation and code execution in GitLab and GitHub

SOURCES:

Cyber Security News
Bleeping Computer
The Hacker News

In March, GitLab released a few security patches addressing multiple high-impact vulnerabilities in GitLab Community Edition and Enterprise Edition that could allow attackers to bypass authentication or achieve remote code execution under certain conditions.

The most severe flaws, CVE-2025-25291 and CVE-2025-25292, affected the ruby-saml library used for SAML Single Sign-On and could let attackers authenticate as legitimate users if they obtain a valid signed SAML document from an identity provider.

The same vulnerabilities also affected GitHub, which discovered and reported those flaws in November 2024. At that time, GitHub said that those vulnerabilities attackers could use to perform takeover attacks.

GitLab also fixed a high-severity vulnerability, CVE-2025-27407, in the Ruby GraphQL library that could enable remote code execution when a malicious project is transferred via the Direct Transfer feature, which is disabled by default on self-managed instances.

To mitigate risk, GitLab urged all self-managed customers to upgrade immediately to versions 17.9.2, 17.8.5, or 17.7.7 for both Community and Enterprise editions.

FEBRUARY

GitLab patches high-risk XSS and authorization flaws across platform

SOURCES:

Cyber Security News
gbhackers

Urgent security updates to address multiple high- and medium-severity vulnerabilities that could allow attackers to execute malicious scripts, bypass access controls, and expose sensitive data were released by GitLab in February.

The patches, delivered in versions 17.9.1, 17.8.4, and 17.7.6 for both Community and Enterprise editions, were strongly recommended for all self-managed installations.

The most serious issues included two high-severity cross-site scripting flaws: CVE-2025-0475 in the Kubernetes proxy endpoint and CVE-2025-0555 in the Maven Dependency Proxy, both of which could enable session hijacking and arbitrary

script execution under specific conditions.

Additional medium-severity flaws allowed HTML injection leading to XSS, unauthorized access to security policy configuration files by Guest users, and improper data exposure for Planner-role users.

Atlassian urges immediate updates after multiple critical vulnerabilities found

SOURCES:

Security Week

Patches for 12 critical- and high-severity vulnerabilities affecting its Bamboo, Bitbucket, Confluence, Crowd, and Jira products were released by Atlassian in its February bulletin.

Five critical flaws fixed in Confluence Data Center and Server, as well as Crowd, stem from vulnerable third-party dependencies, including multiple issues in Apache Tomcat that could allow unauthenticated attackers to achieve remote code execution or bypass authentication.

The Tomcat vulnerabilities, tracked as CVE-2024-50379, CVE-2024-56337, and CVE-2024-52316, all carry CVSS scores of 9.8 and posed a severe risk if left unpatched.

Bitbucket Data Center and Server received patches for a high-severity remote code execution bug in the Apache Avro Java SDK, and the Protocol Buffers DoS issue was also remediated in Jira.

Although Atlassian reported no evidence of active exploitation, it strongly recommended customers patch all affected installations as soon as possible to reduce exposure.

JANUARY

Clone2Leak bugs allow attackers to steal Git authentication tokens

SOURCES:

The Hacker News

Security researchers disclosed multiple vulnerabilities across Git-related tools, including GitHub Desktop, that could allow attackers to leak users' Git credentials through malicious repositories.

The flaws, collectively dubbed Clone2Leak, stem from improper handling of messages in the Git Credential Protocol, enabling credential leakage via crafted URLs, carriage return smuggling, and recursive cloning behaviors. Affected components include Git Credential Manager, Git LFS, and GitHub CLI, each exposing credentials or access tokens under specific conditions.

In GitHub Desktop and Git Credential Manager, attackers could exploit carriage return or CRLF injection to trick the tools into sending credentials to attacker-controlled

hosts. The GitHub CLI issue was particularly risky in GitHub Codespaces, where environment variables caused access tokens to be sent to non-GitHub hosts when cloning malicious repositories. Successful exploitation could allow attackers to reuse leaked tokens to access privileged resources and repositories.

In response, the Git project and GitHub have released patches addressing these issues and urged users to update immediately or avoid cloning untrusted repositories with recursive submodules.

Researchers expose multiple high-risk vulnerabilities in Azure DevOps

SOURCES:

Cyber Security News

Security researchers from Binary Security uncovered multiple vulnerabilities in Azure DevOps that could allow attackers to carry out Server-Side Request Forgery (SSRF), CRLF injection, and DNS rebinding attacks.

The first issue affects the platform's endpointproxy functionality, where attackers could manipulate a URL parameter to force Azure DevOps to make unauthorized requests to internal services, including metadata endpoints.

A second flaw was found in Azure DevOps Service Hooks, enabling both SSRF and CRLF injection, which allowed malicious actors to inject arbitrary HTTP headers into outbound requests.

Researchers demonstrated that attackers could add headers such as Metadata: True, a requirement for interacting with Azure metadata APIs, significantly increasing the potential impact.

Together, these vulnerabilities highlighted how misconfigurations in request-handling features could expose internal cloud services and sensitive infrastructure data.

Data source: Vulnerability data compiled from vendors' official security advisories, CVE database entries, and public security bulletins covering the period January-December 2025. Severity ratings (CVSS scores) as published by the vendors and CVE program.

Phishing, social engineering, and identity attacks

Key takeaways

1 Identity abuse

Attackers bypass MFA via trusted identity flows, cloud services, and OAuth, not passwords.

2 Phishing risks

Phishing-as-a-service ecosystems are at the center of cybercrime and state-aligned attacks.

3 Platform exploit

Platforms like Microsoft 365 and GitHub can be abused for attacks that evade defenses.

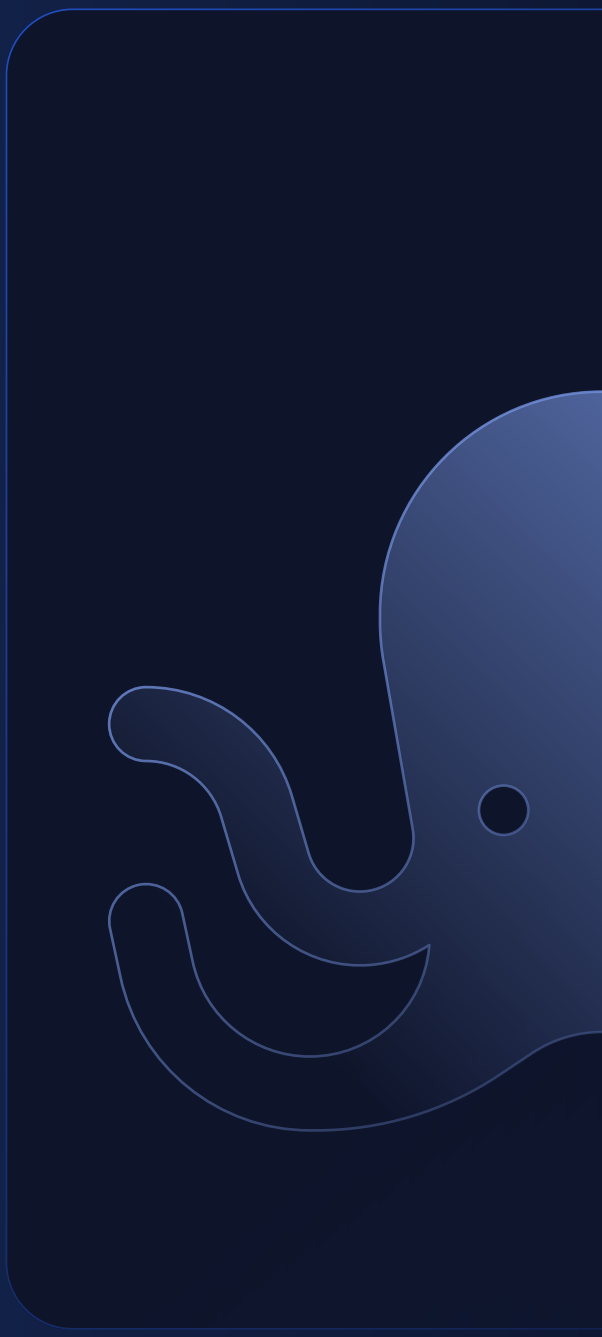

Phishing, social engineering, and identity attacks

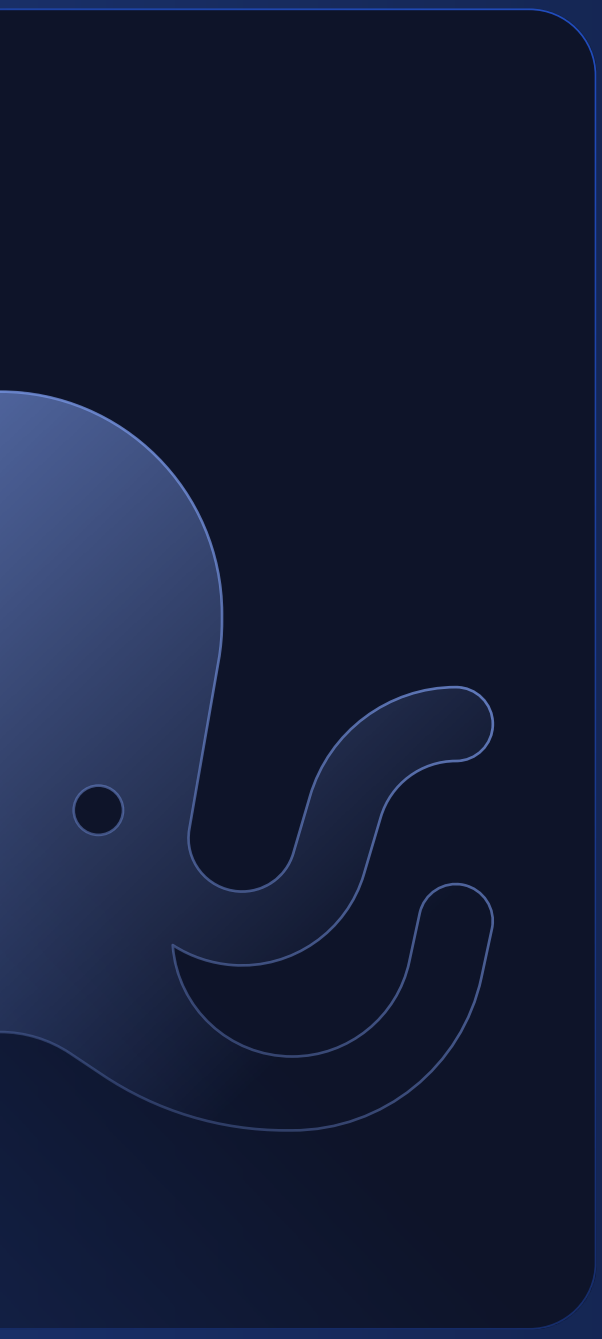

Our recommendation

The phishing and identity attacks of 2025 show that to resist such threats, organizations should restrict OAuth device code flows with Conditional Access, enforce app consent approvals, and continuously audit authorized applications. They should combine phishing-resistant MFA with behavior-based detection, user training focused on device codes, and trusted platform abuse.

Phishing attack surface

In 2025, identity became the primary attack surface for enterprises, as threat actors systematically abused trusted platforms, OAuth workflows, and phishing-as-a-service ecosystems to bypass MFA, persist inside cloud environments, and compromise high-value accounts.

Rather than exploiting software flaws alone, attackers increasingly targeted the assumptions behind identity, access, and user trust across Microsoft 365, GitHub, and collaboration tools, exposing critical gaps in detection, policy enforcement, and user-driven security controls that CISOs can no longer afford to treat as edge cases.

Among the most abused phishing-as-a-service campaigns were Whisper 2FA, VoidProxy, Salty 2FA, “ADFSjacking” technique, Tycoon, PyStoreRAT, SquarePhish2, Graphish kit, and RaccoonO365.

DECEMBER

Hackers exploit OAuth device codes to hijack Microsoft 365 accounts

SOURCES:

Bleeping Computer
Cyber Press
Cyber Security News

Microsoft 365 accounts were being targeted through phishing campaigns that abuse the OAuth 2.0 device code authorization flaw, according to the findings from Proofpoint.

From September 2025 to December 2025, threat actors used emails masquerading as document shares, security alerts, or salary notifications to trick victims into entering device codes on Microsoft's legitimate device login page, unknowingly granting attackers valid OAuth access tokens. This technique allowed attackers to bypass traditional credential theft and even multifactor authentication, enabling persistent access, data exfiltration, and account takeover once a malicious application is authorized.

Proofpoint observed both financially motivated groups, such as TA2723, and suspected state-aligned actors like UNK_AcademicFlare adopting the method in large-scale campaigns targeting government, education, transportation, and enterprise users. Attackers were leveraging automated phishing frameworks, including SquarePhish2 and the Graphish kit, which used QR codes, Azure app registrations, and reverse-proxy infrastructure to make the attacks appear fully legitimate.

Wider adoption of phishing-resistant MFA alone did not stop this abuse; organizations should restrict OAuth device code flows via Conditional Access policies and train users to never enter device codes from untrusted prompts.

RaccoonO365 crackdown signals rising pressure on PhaaS networks

SOURCES:

The Hacker News
Bleeping Computer

Authorities in Nigeria arrested three suspected cybercriminals linked to large-scale phishing operations, including the alleged developer behind the RaccoonO365 phishing-as-a-service scheme.

The Nigeria Police Force National Cybercrime Centre said the investigation, conducted with support from Microsoft and the Federal Bureau of Investigation, identified Okitipi Samuel, also known as Moses Felix, as the principal suspect responsible for building and selling phishing infrastructure via Telegram. Investigators alleged the operation sold phishing links for cryptocurrency and hosted fake Microsoft 365 login portals on Cloudflare using stolen or fraudulently obtained email accounts.

RaccoonO365, tracked by Microsoft as Storm-2246, was believed to have facilitated

the theft of at least 5,000 Microsoft credentials across 94 countries, leading to business email compromise, data breaches, and financial losses.

Searches in Lagos and Edo states resulted in the seizure of laptops, mobile devices, and other digital evidence, while police clarified that the two additional suspects were not involved in developing the phishing platform. The arrests followed earlier legal action by Microsoft and Health-ISAC and come amid a broader global crackdown on phishing-as-a-service operators, highlighting growing international cooperation against cybercrime.

PyStoreRAT malware campaign abuses GitHub to target security researchers

Cybersecurity researchers at Morphisec Threat Labs uncovered a previously undocumented malware campaign abusing GitHub to distribute a remote access trojan dubbed PyStoreRAT.

The operation specifically targeted IT professionals, cybersecurity specialists, and open-source intelligence researchers by masquerading as legitimate open-source projects, including OSINT tools, DeFi trading bots, and AI-driven applications. Attackers revived long-dormant GitHub accounts and populated them with polished, AI-assisted repositories that gained visibility and trust before quietly introducing malicious updates containing the PyStoreRAT backdoor.

Once deployed, the malware provided persistent remote access, profiled infected systems, executed commands, and could deliver additional payloads while adapting its behavior to evade endpoint security tools.

Morphisec also found that PyStoreRAT supports multiple propagation methods, including removable media, and relies on rotating command-and-control infrastructure to complicate detection and takedown efforts.

GitHub removed most of the identified malicious repositories fast. However, this campaign shows a growing trend of attackers blending AI-generated code, social engineering, and trusted cloud platforms to bypass traditional security defenses.

SOURCES:

IT Security News

Cyber Security News

OCTOBER

Whisper 2FA emerges as major threat to Microsoft 365 users

SOURCES:

ITPro

Security researchers at Barracuda warned of a rapidly growing phishing-as-a-service platform called Whisper 2FA that was targeting millions of Microsoft 365 accounts. Being active since July 2025, the tool had already powered close to 1M phishing attempts by October, making it one of the most prevalent PhaaS kits in use.

Whisper 2FA was designed to steal both login credentials and multi-factor authentication tokens in real time by abusing AJAX-based workflows that intercepted

sessions without page reloads. Researchers said its heavy obfuscation, anti-debugging defenses, and realistic brand impersonation made detection difficult and could allow it to bypass many traditional security controls.

The kit shows similarities to Salty 2FA, a related threat highlighted by AnyRun, and underscores a broader trend toward highly sophisticated phishing platforms.

SEPTEMBER

Phishers abuse GitHub notifications in Y Combinator crypto scam

SOURCES:

Bleeping Computer
gbhackers
Cyber Security News

Cybercriminals launched a large-scale phishing campaign that abused GitHub's issue notification system to impersonate Y Combinator and target developers' cryptocurrency wallets.

The attackers created multiple look-alike GitHub accounts and repositories, then opened hundreds of issues per repo while tagging random users, causing GitHub to send legitimate-looking notifications directly to victims. These messages falsely claimed recipients had been selected for Y Combinator's Winter 2026 funding round and urged them to verify a crypto wallet to proceed.

Victims were redirected to a typosquatted domain, y-comblnator.com, hosting a convincing fake YC application page that prompted wallet connections. Signing the requested "verification" transactions silently authorized malicious transfers, draining wallets of cryptocurrency.

GitHub removed the malicious accounts and repositories immediately. Though the campaign highlights how trusted collaboration platforms can be weaponized for high-impact phishing attacks.

AUGUST

Salty 2FA highlights the growing sophistication of phishing kits

SOURCES:

gbhackers

Phishing remains the dominant cause of global cyber incidents, fueled by low-cost Phishing-as-a-Service platforms such as Tycoon2FA, EvilProxy, and Sneaky2FA that lower the barrier for sophisticated attacks.

Researchers at ANY.RUN uncovered another, previously undocumented, PhaaS framework called Salty 2FA, primarily targeting Microsoft 365 credentials. The framework used a multi-stage execution chain with "salted" payloads, heavy obfuscation, and Cloudflare-based protections to evade static analysis and sandbox detection. Salty 2FA dynamically rendered fake Microsoft login pages, intercepted credentials and multiple 2FA methods in real time, and exfiltrated data through encrypted POST requests to attacker-controlled endpoints.

Activity surged in mid-2025 and affected organizations across finance, telecom, energy, education, and other sectors worldwide. Analysts warn that its constantly mutating infrastructure makes static indicators unreliable, urging defenders to rely on behavioral detection patterns instead.

Phishers exploit proofpoint and intermedia link wrapping

SOURCES:

Bleeping Computer

A threat actor was abusing link-wrapping services from Proofpoint and Intermedia to disguise malicious links leading to Microsoft 365 phishing pages. According to findings from Cloudflare's Email Security team, the attacker compromised email accounts protected by these services and used the built-in URL rewriting features to make phishing links appear trustworthy.

The campaign, active between June and July, relied on multi-stage redirects, often starting with URL shorteners before links were automatically wrapped by the security platforms. Victims were lured with fake voicemail alerts or Microsoft Teams document notifications that ultimately redirected to credential-harvesting Microsoft 365 login pages. By laundering malicious URLs through legitimate security infrastructure, the attacker significantly increased the likelihood of bypassing email defenses and user suspicion.

Researchers note that while abusing trusted services is not new, exploiting link-wrapping features for phishing represents an emerging tactic.

07

APT campaigns and targeted attacks

Key takeaways

1 DevOps platform abuse

Threat actors increasingly abuse DevOps platforms as trusted infrastructure for malware delivery, command-and-control, and data exfiltration.

2 Tools criminals exploit

Both financially motivated criminals and nation-state APT groups rely on open-source tools and cloud platforms to stay resilient and evade detection.

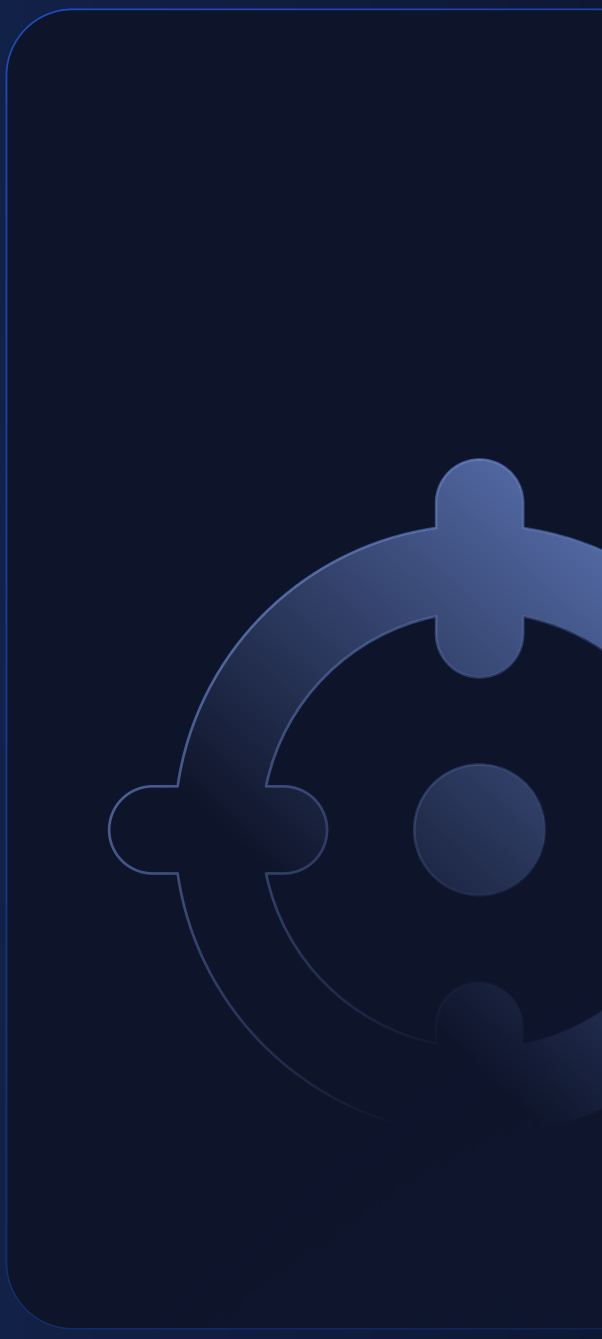

07

APT campaigns and targeted attacks

Our recommendation

Organizations should monitor and restrict the use of public repositories for internal workflows, enforce strict access controls on API tokens, validate code from third-party sources, and deploy behavioral detection to flag unusual repository activity. Thus, they can prevent the impact of malware, ransomware, and APT campaigns that abuse trusted platforms.

Advanced persistent threat

Throughout 2025, we could notice threat actors were abusing trusted platforms, like GitHub, as part of their sophisticated malware, ransomware, and espionage campaigns.

Multiple groups, including North Korean APTs, leveraged GitHub repositories as command-and-control infrastructure, malware hosting platforms, and data exfiltration channels to evade detection. All the threat campaigns ranged from banking credential theft and ransomware attacks to employment fraud and long-term espionage operations.

Let's look deeper at the loudest cases of the advanced persistent threat (APT) campaigns that shook the DevOps world in 2025.

OCTOBER

Banking malware turns to GitHub as resilient command channel

SOURCES:

The Hacker News

Cyber Press

In October, cybersecurity researchers identified an Astaroth banking trojan campaign that abused GitHub as a resilient fallback infrastructure to survive takedowns of traditional command-and-control servers.

The attack chain begins with DocuSign-themed phishing emails that deliver a ZIP file containing a malicious Windows shortcut (LNK). In its turn, this LNK file later executed obfuscated JavaScript to fetch additional payloads, including AutoIt scripts and shellcode. Once deployed, the Delphi-based malware injects itself into legitimate processes and actively monitors browser activity to steal credentials from banking and cryptocurrency websites, particularly targeting victims in Brazil and other Latin American countries.

To evade detection, Astaroth employs extensive anti-analysis checks, geofencing, and locale restrictions, shutting down if security tools, debuggers, or English-language systems are detected.

What sets this campaign apart is its use of GitHub repositories to host configuration data hidden inside image files using steganography, allowing the malware to retrieve fresh instructions whenever primary infrastructure is disrupted.

This campaign shows how threat actors can exploit trusted cloud platforms to maintain persistence and complicate disruption efforts.

JUNE

North Korean APT Abuses GitHub PATs in Targeted Spearphishing Campaign

SOURCES:

Cyber Security News

gbhackers

A highly targeted spearphishing campaign attributed to Kimsuky, the North Korean threat group, demonstrates how bad actors can abuse GitHub as a main point of destination for distributing malware.

Security researchers uncovered the operation after they analyzed a malicious PowerShell script shared on X, which revealed extensive use of GitHub and Dropbox as malware hosting and command-and-control infrastructure. The campaign began in March 2025 with carefully crafted spearphishing emails aimed primarily at South Korean targets, impersonating trusted entities such as law firms or financial authorities. Victims received password-protected archives that contained malicious attachments that could execute PowerShell scripts, which later could retrieve payloads from attacker-controlled GitHub repositories or Dropbox URLs.

A defining feature of this operation is the embedding of hardcoded GitHub Personal Access Tokens (PATs) with repository-level permissions directly into the malware. These tokens granted attackers persistent read and write access to private repositories used to store malware payloads, decoy documents, and exfiltrated victim data.

Researchers identified repositories such as “hole_311” and “star,” linked to attacker accounts “Dasi274” and “luckmask,” which contained infostealer scripts, downloader components, and scheduled-task logic that uploaded victim logs every 30 minutes. The malware payloads also included the open-source XenoRAT, deployed through obfuscated .NET executables featuring state-machine obfuscation and dynamically loaded encrypted strings.

Thus, once executed, the malware performed system reconnaissance, keylogging, and data exfiltration directly back to GitHub, blending malicious traffic with legitimate API activity to evade detection.

This campaign shows how advanced persistent threat actors can transform trusted developer platforms into resilient espionage infrastructure.

APRIL

How CrazyHunter turned GitHub tools into a ransomware arsenal

SOURCES:

gbhackers

Cyber Security News

A ransomware group, known as CrazyHunter, was identified as a growing threat to critical infrastructure organizations in Taiwan in April, though its attacks had intensified since early 2025.

Researchers from Trend Micro reported that the group was targeting healthcare providers, educational institutions, and industrial organizations using a highly structured attack chain built largely from open-source tools hosted on GitHub. Roughly 80% of CrazyHunter’s toolkit consisted of modified public tools, enabling the group to lower development costs while maintaining operational effectiveness. The attackers in their campaign relied heavily on the Bring Your Own Vulnerable Driver (BYOVD) technique, abusing legitimate but flawed drivers such as Zemana Anti-Malware’s zam64.sys to disable security controls and evade detection.

Once the defense of their victims’ was neutralized, CrazyHunter deployed a customized ransomware strain based on the Prince ransomware builder, encrypting files with the “.Hunter” extension and issuing ransom demands through notes and desktop wallpaper changes.

Analysts noted that the group’s use of redundant execution paths and failover mechanisms suggested a level of maturity uncommon in observed ransomware operations, that underscors the growing risk posed by the misuse of open-source tooling.

MARCH

GitHub portfolios used in sophisticated North Korean employment fraud

SOURCES:

gbhackers

Cyber Security News

Cybersecurity researchers from Nisos uncovered a coordinated network of suspected North Korea-linked IT workers who abused GitHub to build fake professional identities and secured remote jobs in Japan and the United States. The investigation showed that these operatives pose as developers from Vietnam, Japan, and Singapore, focusing on engineering and blockchain roles to generate foreign currency for Pyongyang's weapons programs.

The actors manipulated GitHub credibility by fabricating contribution histories, coordinating co-authored commits with other suspected DPRK-linked profiles, and avoiding broader social media exposure that could reveal inconsistencies. Researchers identified recurring technical markers, including shared email patterns and synchronized commit activity, linking multiple personas into a single organized network.

Experts warned that beyond financial fraud, these infiltrations pose serious cybersecurity risks, as embedded workers may gain access to sensitive systems, source code, or infrastructure.

Blind Eagle intensifies malware campaigns against Colombian institutions

SOURCES:

The Hacker News

Cyber Security News

gbhackers

The threat actor Blind Eagle, who is also tracked as APT-C-36 and AguilaCiega, was linked to cyber campaigns targeting Colombian judicial bodies, government agencies, and private organizations.

According to the findings from Check Point, one campaign alone infected more than 1,600 victims, highlighting an unusually high success rate for a typically targeted APT operation. The group relied heavily on spear-phishing emails and malicious URL files, rapidly incorporating a variant of the patched CVE-2024-43451. Once victims interacted with the lure, Blind Eagle deployed readily available remote access trojans such as Remcos RAT, protected by tools like HeartCrypt and PureCrypter, with payloads hosted on trusted platforms including GitHub and Bitbucket.

Researchers also uncovered an operational mistake in a GitHub repository that exposed around 1,630 credential records tied to Colombian individuals and institutions, reinforcing the group's regional focus.

Check Point noted that Blind Eagle's continued abuse of legitimate cloud services and underground crimeware tools which allowed it to evade traditional defenses while maintaining persistent access across compromised networks.

FEBRUARY

Lazarus group uses fake LinkedIn jobs to spread cross-platform malware

SOURCES:

The Hacker News

The North Korea-linked Lazarus Group was found to be tied to an active malware campaign that abused fake job offers on LinkedIn to target professionals in the cryptocurrency and travel sectors.

According to research from Bitdefender, attackers pose as recruiters offering flexible, well-paid remote roles before requesting CVs or links to personal repositories to build trust and harvest data. Victims were then directed to check the supposed project code from GitHub or Bitbucket, where obfuscated scripts were embedded in fake decentralized exchange projects. These scripts deployed a cross-platform JavaScript information stealer that targeted cryptocurrency wallet browser extensions and acted as a loader for a Python-based backdoor. The infection chain ultimately delivered additional .NET malware capable of enabling persistent access, exfiltrating system data, logging keystrokes, launching crypto miners, and routing traffic through Tor.

Bitdefender noted that the activity overlapped with the long-running “Contagious Interview” cluster but showed signs of ongoing evolution, with attackers refining their tooling and recruitment lures across multiple platforms.

Data breaches and enterprise DevOps compromises

Key takeaways

1 Targeted industries

The most targeted industry by attackers in 2025 was technology and software.

2 Detection gaps

Weak credential hygiene and delayed detection allowed many breaches to persist long enough to cause significant operational and financial impact.

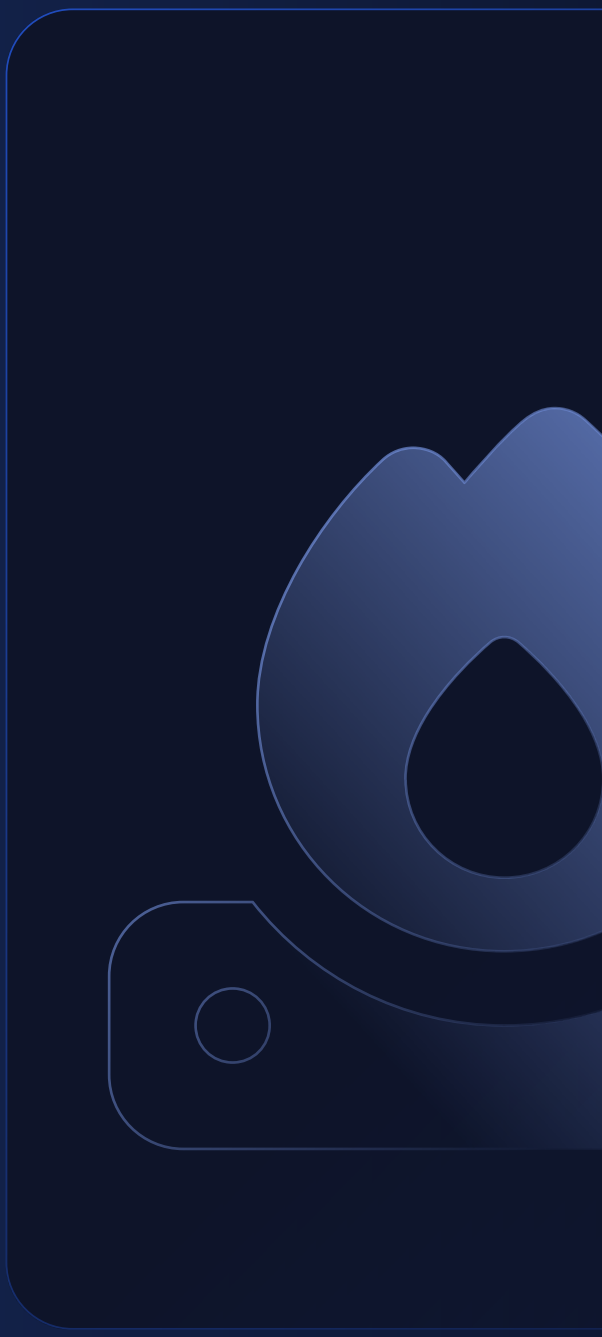

Data breaches and enterprise DevOps compromises

Our recommendation

Organizations should enforce strict credential hygiene, rotate and revoke old access tokens, and implement least-privilege access across DevOps platforms. Continuous monitoring, repository audits, and employee security training are essential to prevent unauthorized access, insider threats, and supply chain abuse.

Tech and IT remaine the most targeted sector

In 2025, cyberattacks against major organizations increasingly focused on the platforms used to build, manage, and operate digital services. Technology and software companies remained the most targeted sector for the second year in a row, while telecommunications and automotive firms saw a sharp rise in attacks, overtaking fintech and media compared to 2024. Retail and consumer businesses also continued to face sustained pressure as attackers sought access to customer data and operational systems.

A defining feature of the 2025 attacks was the growing reliance on compromised credentials, with attackers repeatedly abusing trusted tools such as GitHub, GitLab, Jira, and cloud identity services. Long-lived credentials stolen by infostealer malware enabled threat actors to quietly access internal systems months or even years after the initial compromise. In several cases, attackers moved laterally from developer environments into production systems, supply chains, and third-party integrations.

Ransomware groups and extortion-focused actors, including Hellcat and Crimson Collective, were responsible for many data breaches of 2025. Under attack, there were globally recognized brands across multiple industries, such as Red Hat, Nissan, Europcar, Ascom, etc.

Let's look at the most significant cases reported throughout 2025 and the techniques attackers used to carry them out.

DECEMBER

Nissan Confirms 21K Customers Affected by Red Hat Data Breach

SOURCES:

Bleeping Computer
Info Security Magazine
Security Affairs
Technadu

Nissan Motor Co. confirmed that a third-party data breach at Red Hat exposed personal information belonging to around 21,000 customers in Japan. The incident stemmed from unauthorized access to a Red Hat–managed GitLab environment used to develop a customer management system for Nissan Fukuoka Sales.

Nissan said the leaked data included names, addresses, phone numbers, partial email addresses, and sales-related customer information, but no credit card or financial details.

Red Hat notified Nissan of the breach on October 3, after which the automaker reported the incident to regulators and began contacting affected customers.

Nissan stated there was currently no evidence that the data had been misused and confirmed that no additional customer information remained at risk. The breach follows claims by the Crimson Collective hacking group, which alleged a large-scale theft of data from Red Hat's consulting repositories.

University of Sydney discloses major data breach linked to coding repository

SOURCES:

Web Pro News

The University of Sydney disclosed a significant cybersecurity breach after attackers accessed a self-hosted GitLab repository and exfiltrated historic personal data. The incident affected more than 13K individuals, including current and former staff, students, alumni, and donors, with exposed information such as names, contact details, resumes, and limited financial data.

University officials said the compromised data was not connected to active systems and that no ransomware was involved. The affected repository was taken offline, external forensic experts were engaged, and impacted individuals are being notified and offered credit monitoring.

The breach highlights ongoing cybersecurity risks facing universities that rely on internally managed development platforms.

Home Depot left GitHub token exposed for a year

SOURCES:

Techradar

Security researchers found that Home Depot left a sensitive GitHub access token exposed for more than a year, unintentionally granting broad access to its internal systems.

The token that the company revoked in December, tied to an employee account on GitHub, allowed access to hundreds of private repositories and enabled modifications to source code connected to cloud infrastructure, inventory management, and order fulfillment systems.

This incident highlights a persistent industry-wide risk, as hardcoded secrets and misconfigured repositories remain common entry points for attackers.

NOVEMBER

Apple accidentally leaks App Store front-end code after web redesign

SOURCES:

9to5Mac

Apple briefly exposed the complete front-end source code of its web App Store after accidentally shipping the site with production sourcemaps enabled. The mistake allowed a GitHub user to extract Svelte and TypeScript source files, UI components, routing logic, and API integration code directly from the live App Store website using standard browser tools.

The extracted code was published to GitHub for educational purposes before Apple moved quickly to have the repository and its forks taken down.

While the incident does not appear to pose an immediate security or privacy risk, it shows how even mature technology companies can inadvertently expose internal implementation details during major product launches.

OCTOBER

Red Hat confirms GitLab breach after hackers claim massive data theft

SOURCES:

Security Week
Bleeping Computer
The Register
SC Media

Red Hat confirmed that an unauthorized third party accessed and copied data from a GitLab instance used by its consulting team, validating claims made by an extortion group calling itself Crimson Collective.

The company said the incident was limited to a consulting-managed, self-hosted GitLab environment and did not affect its products, services, or software supply chain. Attackers claim to have stolen data from as many as 28,000 repositories, including customer engagement reports that may contain configuration details, architecture information, and credentials.

Red Hat stated it engaged external security experts, notified law enforcement, and isolated the affected instance while its investigation continues.

While the company said no sensitive personal data was identified so far, cybersecurity authorities in Belgium warned organizations to rotate any credentials or tokens shared with Red Hat's consulting team as a precaution.

SEPTEMBER

Attackers warn of SK Telecom data leak after alleged Bitbucket compromise

SOURCES:

Cybernews
SC Media

SK Telecom faced cyber extortion claims in September after a ransomware group, CoinbaseCartel, alleged it stole the company's source code and internal project files in an August breach.

According to Cybernews, the attackers claimed to have exfiltrated code repositories containing Docker files, build configurations, Python extensions, and even exposed AWS access keys, and have threatened to leak the data if there were no negotiations.

Cybernews researchers believed the intrusion may have originated from the compromise of an employee's Bitbucket account, granting access to multiple internal projects. While analysts noted the leaked source code didn't directly expose production secrets or customer data, it could still weaken SK Telecom's security posture by enabling vulnerability discovery or intellectual property theft.

The claims came months after SK Telecom disclosed a separate ransomware attack by the Qilin ransomware gang, which previously alleged the theft of 1TB of data.

JLR Confirms Data Breach as Cyberattack Triggers $890M Impact

SOURCES:

The Register
Cyfirma
Technadu

Jaguar Land Rover was hit by a major cyberattack that forced the shutdown of global IT systems, halted manufacturing and retail operations, and extended factory closures into late September. The incident forced JLR to halt production for more than a month, contributing to a financial shortfall exceeding $890 million and disrupting its supply chain.

A group calling itself Scattered Lapsus$ Hunters claimed responsibility, publishing screenshots allegedly showing access to JLR's internal systems.

The incident followed the earlier 2025 breaches by the Hellcat ransomware group, which exploited stolen Jira credentials to leak internal documents and employee data.

Jaguar Land Rover confirmed that a cyberattack compromised personal data belonging to current and former employees and contractors. The breach exposed employment-related information, including payroll, benefits, and staff scheme details, though the company said there was no evidence of misuse.

Salesloft Confirms Drift Breach Contained After GitHub Compromise

SOURCES:

Hackread

Cyber daily

Salesloft confirmed that a security breach involving its Drift application had been contained after an investigation led by Mandiant. The inquiry found that attackers gained access to a Salesloft GitHub account between March and June 2025, where they downloaded private repository content, added workflows, and conducted limited reconnaissance.

The threat actors later pivoted to Drift's AWS environment and stole customer OAuth tokens, which were used to access Salesforce data through third-party integrations.

Salesloft responded by taking Drift offline, rotating affected credentials, isolating infrastructure, and hardening its environments, with Mandiant confirming no deeper compromise of Salesloft's core systems. Investigators say the incident is part of a broader campaign targeting Salesforce integrations, which affected multiple high-profile companies, though attribution remains unconfirmed.

JULY

Toptal GitHub breach leads to malicious npm packages in supply chain attack

SOURCES:

Bleeping Computer

The Hacker News

SC Media

A software supply chain attack hit Toptal after attackers compromised the company's GitHub organization and published ten malicious packages to the npm registry.

According to security firm Socket, the packages were modified to steal GitHub authentication tokens during installation and then attempt to wipe victim systems on Windows and Linux. The attackers also made 73 Toptal repositories public, briefly exposing private projects and source code.

The malicious releases were downloaded roughly 5,000 times before removal, though Toptal later said most of that activity likely came from automated security scanners rather than real users.

In its statement, Toptal said the breach was traced back to credentials exposed in the historic LastPass breach and insisted no customers or partners were impacted.

The incident underscores how stolen credentials and trusted package ecosystems continue to be abused to target developers.

Hacker threatens to leak 106GB of Telefónica data as company denies new breach

SOURCES:

Bleeping Computer

In July, a hacker claiming ties to the Hellcat ransomware group threatened to release 106GB of data allegedly stolen from Telefónica, even as the company didn't acknowledge a second breach in 2025 (editor - the first attack took place in January

2025).

The threat actor, using the alias Rey, leaked a 2.6GB archive as proof and claimed the intrusion occurred on May 30, exploiting a Jira misconfiguration left after a previous incident. According to Bleeping Computer, the exfiltrated data included internal communications, invoices, logs, and customer and employee records spanning multiple countries.

JULY

Indian Grocery App KiranaPro Taken Offline in Targeted Insider Breach

SOURCES:

The Register
Technadu

KiranaPro suffered a severe cyberattack that wiped its GitHub repositories and AWS infrastructure, leaving the grocery delivery platform completely inoperable.

The breach occurred between May 24 and 25, when attackers gained access to the company's AWS root account and GitHub, deleting application code, EC2 instances, and core services. The incident was described by the company as a targeted and deliberate attack, attributing it to a malicious insider, potentially a former employee with lingering privileged credentials.

The attack exposed sensitive data and disrupted operations that supported more than 30,000 active users and over 2,000 daily orders across 50 cities.

Company executives discovered the breach after losing access to AWS services, despite multi-factor authentication being enabled, suggesting credential compromise and privilege escalation.

KiranaPro contacted GitHub for forensic support, initiated legal action against suspected former employees, and began rebuilding its systems with stronger security controls.

APRIL

Europcar confirms GitLab breach affecting up to 200K customers

SOURCES:

Bleeping Computer
Tech hq

Europcar Mobility Group confirmed a data breach after a hacker gained access to its GitLab repositories, resulting in the theft of source code for its Android and iOS applications and personal data belonging to up to 200,000 customers.

The threat actor attempted to extort the company by threatening to release 37GB of the stolen data, including more than 9,000 SQL backup files and at least 269 configuration (.ENV) files containing sensitive environment settings.

According to the company, the exposed customer data is limited to names and email addresses associated with its Goldcar and Ubeeqo brands, with no evidence that

passwords, payment card details, or banking information were compromised. Some of the affected records date back to 2017 and 2020, and the total number of impacted customers is still being assessed.

Europcar was notifying affected users as the investigation continued, while the method of initial access remains unclear.

MARCH

Hellcat hackers claim 44GB data theft from Ascom

SOURCES:

Bleeping Computer
Cyber Security News

Ascom confirmed a cyberattack on its IT infrastructure after hackers breached the company's technical ticketing system.

The intrusion was claimed by the Hellcat hacking group, which exfiltrated roughly 44GB of data, including source code, project information, invoices, and confidential documents.

Ascom stated that the breach affected a non-critical system, had no impact on business operations, and did not require customers or partners to take action.

The incident aligned with Hellcat's established pattern of targeting Atlassian Jira servers using credentials harvested by infostealer malware.

Security researchers warn that long-lived, unrotated credentials continue to make Jira systems a high-value target for attackers.

Jaguar Land Rover breached in Hellcat attack

SOURCES:

Cyber Security News
gbhackers

Started in mid-2024 and continued in 2025, the Hellcat ransomware group attacked a number of companies, including Jaguar Land Rover, which was hit by a major data breach, exposing large volumes of sensitive internal information.

The attackers gained access to JLR's Atlassian Jira environment using credentials stolen years earlier by infostealer malware such as Lumma. Those credentials reportedly originated from an infected third-party employee system and had remained valid since at least 2021, enabling attackers to infiltrate JLR's internal systems.

The breach escalated further when a second hacker, operating under the alias APTS, exploited additional infostealer-derived credentials to exfiltrate an estimated 350GB of data beyond the initial leak. Stolen information is said to include proprietary documents, source code, Jira issues, and employee data.

Security researchers warned that the incident highlighted the long-term risks posed by unrotated credentials and the growing abuse of Jira as a high-value entry point for enterprise breaches.

FEBRUARY

Hacker claims theft of thousands of fles from Orange systems

SOURCES:

Bleeping Computer

After a threat actor claimed to had stolen thousands of internal documents containing employee and user information, Orange Group confirmed a data breach.

In February, Orange said the incident affected a non-critical back-office application and that an investigation is ongoing to assess and limit the impact. The attacker, using the alias Rey and linked to the HellCat ransomware group, alleged they had access to Orange's systems for over a month and exfiltrated nearly 12,000 files totaling about 6.5GB.

According to the claims, the stolen data, largely tied to Orange's Romanian operations, included email addresses, internal documents, source code, contracts, invoices, and partial customer payment card details.

Rey stated that access was gained using compromised credentials and by exploiting vulnerabilities in the company's Jira software and internal portals, adding that a ransom note was left, but no negotiations followed

Orange said there has been no impact on customer operations and that it is cooperating with authorities while continuing its internal response.

JANUARY

HPE investigates hacker claims of source code and credential theft

SOURCES:

Bleeping Computer

In January, Hewlett Packard Enterprise was investigating claims of a security breach after a threat actor known as IntelBroker alleged they stole data from the company's developer environments.

In its statement to BleepingComputer, HPE said it found no evidence of a breach but activated its cyber response protocols and disabled related credentials as a precaution.

IntelBroker claimed they accessed HPE systems, including the company's API, WePay, and bot private and public repositories, for at least two days and exfiltrated items, including certificates, source code, Docker builds, and older personal data used for deliveries.

This is not the first time the actor made such allegations, having previously posted data said to be stolen from HPE in February 2024, which the company also said showed no confirmed breach.

HPE stated there is currently no operational impact and no indication that customer data was compromised.

Hackers breach Telefónica ticketing system using stolen credentials

Telefónica confirmed that its internal ticketing system was breached after stolen data was leaked on a hacking forum.

The company took all the necessary steps to block unauthorized access. According to information shared with the publication, attackers gained access to an internal Jira system using compromised employee credentials and were able to scrape approximately 2.3 GB of tickets and documents.

Separate reporting indicated the leaked data may include hundreds of thousands of Jira issues, employee email addresses, and internal documents containing sensitive operational details. The breach was claimed by four threat actors, three of whom are linked to the Hellcat ransomware operation.

Telefónica stated that password resets were performed on affected accounts as part of its response while the investigation was in progress.

SOURCES:

Bleeping Computer
Dark Reading
hackread

This report is based exclusively on publicly available information published by vendors on their official status pages, security advisories, databases, and publicly reported incidents documented in industry media. All data and statistics have been compiled from these publicly accessible sources. The analysis presented herein represents GitProtect's independent assessment and is intended for educational and informational purposes only.

09

Abuse of dev platforms for malware hosting and delivery

Key takeaways

1 Malware channel

In 2025, GitHub was actively exploited by cyber criminals as a trusted malware distribution channel through fake repositories, trojanized tools, and malvertising.

2 Impersonation techniques

Attackers combined brand impersonation and advanced evasion techniques to bypass security controls and infect victims at scale.

3 Developer threat

Trusted developer platforms can now be a major supply-chain risk in cybercriminal threat campaigns.

09

Abuse of dev platforms for malware hosting and delivery

Our recommendation

Organizations should enforce strict code vetting, limit access to repositories, and scan third-party projects for malware before use. Also, it's worth combining repository monitoring, behavioral analysis, and credential hygiene with employee awareness to prevent attacks exploiting fake or trojanized dev projects, malvertising, and supply-chain abuse.

Dev platforms growing abuse by bad actors

Analyzing the year 2025, we can see the surge in malware campaigns abusing GitHub to distribute malicious code, steal credentials, and compromise both enterprise and consumer systems at scale.

When it comes to abusing dev platforms, we can say that threat actors increasingly relied on fake repositories in 2025. They were impersonating trusted brands, trojanizing open-source tools, and running malvertising campaigns to lure victims into executing malware hosted on GitHub.

There were several threat campaigns that combined GitHub's trusted reputation and advanced evasion techniques, including hardware-based execution checks, sandbox avoidance, and multi-stage payload delivery. Malware families distributed through these campaigns ranged from infostealers and remote access trojans to ransomware loaders and cookie hijackers.

Some security researchers also warned that the open-source nature of popular malware frameworks continues to lower the barrier to cybercrime and fuel malware-as-a-service ecosystems. These findings highlight a growing supply-chain and trust-abuse problem affecting developer platforms worldwide.

SEPTEMBER

Fake GitHub repositories impersonate big brands to spread malware

SOURCES:

Cyber Press
gbhackers
Cyber Security News

In September, cybersecurity researchers noticed a large-scale campaign that aimed at abusing GitHub to distribute malware through repositories impersonating well-known brands such as Malwarebytes, Citibank, LastPass, and SentinelOne.

Being active since late August 2025, the campaign used cloned logos, polished README files, and fake download buttons to lure users into running trojanized installers and scripts. On Windows systems, malicious PowerShell-based build scripts deployed stealthy loaders that injected payloads into legitimate processes, established registry-based persistence, and quietly exfiltrated system data and credentials.

GPUGate campaign exploits trusted dev platforms and hardware evasion

SOURCES:

Cyber Security News
Cyber Press

Cybersecurity researchers at Arctic Wolf identified a sophisticated malware campaign dubbed GPUGate that abused Google Ads and GitHub to distribute malware.

The campaign used sponsored search results for terms like “GitHub Desktop” to lure victims to a malicious GitHub commit containing a concealed redirect to a trojanized installer hosted on a lookalike domain. The downloaded installer mimicked the legitimate GitHub Desktop setup and was intentionally bloated with dummy executables to evade sandbox analysis.

GPUGate’s defining feature is its hardware-gated execution, using the OpenCL API to decrypt its payload only on systems with genuine GPUs and leaving it dormant in most virtualized analysis environments. Once activated, the malware establishes persistence, weakens Microsoft Defender protections, and can deploy additional payloads such as ransomware and credential stealers.

The malware campaign shows that combining advertising abuse, trusted developer platforms, and hardware-based evasion has become a trend among attackers to compromise high-value IT targets.

JULY

Attackers spread Amadey malware via public GitHub repositories

SOURCES:

The Hacker News
Hackread

Cybersecurity researchers at Cisco Talos uncovered a malware-as-a-service (MaaS) campaign abusing public GitHub repositories to distribute multiple infostealer families.

Being active since at least February 2025 and observed in April, the operation relies on the Amadey bot to download malicious payloads directly from attacker-controlled GitHub accounts.

Investigators found that fake accounts such as “Legendary99999,” “Milidmdds,” and “DFfe9ewf” hosted hundreds of repositories containing single malicious files, leveraging GitHub’s trusted status to bypass web filtering.

The attack chains also reused the Emmenhtal loader, previously linked to SmokeLoader phishing campaigns targeting Ukrainian organizations, with only minor changes to deliver different payloads.

Malware families distributed through this infrastructure included Amadey, Lumma Stealer, RedLine Stealer, AsyncRAT, and related tools.

Although GitHub removed the identified accounts, researchers warned that the campaign demonstrates how MaaS operators can efficiently weaponize legitimate platforms to distribute malware at scale while evading detection.

AsyncRAT forks highlight risks of open-source malware framework

SOURCES:

The Hacker News

There has been seen the evolution of AsyncRAT, a remote access trojan first released on GitHub in January 2019, which has spawned numerous forks and variants since then. In its report, ESET explained that AsyncRAT’s open-source, plug-in-based architecture is central to its widespread adoption and long-term threat relevance. Although its core capabilities are relatively basic, the malware is commonly distributed through phishing campaigns and bundled with loaders such as GuLoader and SmokeLoader, enabling information theft and remote system control.

Also, the researchers noted that AsyncRAT builds on the groundwork laid by Quasar RAT but represents a major rewrite rather than a simple fork. Over time, it gave rise to more advanced variants, including DCRat and Venom RAT, which add stronger evasion, credential theft, and surveillance features.

This analysis notes how open-source malware frameworks continue to lower the barrier to cybercrime and fuel the growth of malware-as-a-service ecosystems.

JUNE

Trojanized dev tools fuel multi-stage malware campaign on GitHub

SOURCES:

The Hacker News
gbhackers

A financially motivated threat actor dubbed Water Curse was actively abusing GitHub to distribute weaponized repositories containing multi-stage malware.

According to the analysis from Trend Micro, the campaign was linked to at least 76 GitHub accounts and had been active since at least March 2023, with broader detection beginning in May 2025.

The attackers embedded malicious payloads into seemingly legitimate open-source projects, including penetration testing tools such as SMTP email bombers and Sakura-RAT, by hiding code within Visual Studio project configuration files. When compiled, the trojanized projects trigger obfuscated VBS and PowerShell scripts that download additional payloads, conduct system reconnaissance, disable defenses, and establish long-term persistence.

Researchers noted that the malware is capable of stealing credentials, browser data, and session tokens while using services like Telegram and public file-sharing platforms for command-and-control and data exfiltration.

The entire campaign proves once again a growing supply chain threat in which trusted developer platforms are leveraged to blur the line between legitimate security tooling and active malware distribution.

MAY

ZeroCrumb malware abuses browser cookies to hijack accounts

SOURCES:

Cyber Security News

An information-stealing malware dubbed ZeroCrumb was uncovered in the first half of 2025. The malware was distributed through repositories hosted on GitHub. Specifically targeting browser cookies from Chrome, Brave, and Edge, the malware helped attackers to extract authentication data without triggering User Account Control alerts.

According to analysis by KrakenLabs, ZeroCrumb exploits browser security mechanisms to decrypt protected cookie data, allowing threat actors to hijack active web sessions. Unlike traditional cookie stealers, the malware bypasses standard privilege requirements, making it particularly dangerous for both enterprise and individual environments.

Researchers observed the threat after detecting unusual data exfiltration patterns across multiple corporate networks. The findings showed a significant evolution in

cookie-stealing techniques at a time when browser-based authentication underpins access to critical personal and business services.

MARCH

Nearly 1M devices hit in GitHub-linked malvertising attack

SOURCES:

Bleeping Computer
The Register
Cyber Security News

An undisclosed number of malicious repositories were removed from GitHub by Microsoft, after they uncovered a large-scale malvertising campaign that affected nearly one million devices worldwide.

The activity was detected in early December 2024 when Microsoft threat analysts observed malware being downloaded from GitHub repositories and later used to deploy additional payloads. According to the investigation, attackers injected malicious ads into videos on illegal pirated streaming websites, redirecting viewers through multiple layers of malicious sites before ultimately leading them to GitHub-hosted malware. The initial payload performed system discovery and data exfiltration, while subsequent stages deployed remote access and information-stealing malware, including NetSupport RAT, Lumma stealer, and the Doenerium infostealer.

Microsoft also observed alternate payload hosting on Dropbox and Discord during the campaign. Tracked under the name Storm-0408, the campaign impacted both consumer and enterprise environments, underscoring its broad and indiscriminate reach.

FEBRUARY

GitVenom campaign turns GitHub repos into malware traps

SOURCES:

Bleeping Computer
SC Media

A large-scale malware operation named GitVenom was uncovered by security researchers. In their campaign, threat actors abused fake repositories on GitHub to distribute malicious payloads disguised as legitimate software projects.

The fraudulent repositories impersonated tools such as Telegram Bitcoin wallet bots, Instagram automation software, and even game-related utilities, complete with polished README files and artificially inflated commit histories to appear credible. Once executed, the embedded malware delivered information stealers, remote access trojans, and clipboard hijackers capable of exfiltrating credentials, cryptocurrency wallets, and browsing data.

The scale and persistence of GitVenom highlight how trusted development platforms continue to be exploited, reinforcing the need for strict vetting of third-party code before use.

Security experts warn of misuse after Wi-Fi Password Stealer appears online

SOURCES:

Cyber Security News

A GitHub repository titled Windows-WiFi-Password-Stealer drew attention from cybersecurity experts due to its potential misuse despite being labeled as an educational project. The repository contained a Python script that leverages legitimate Windows netsh commands to extract saved Wi-Fi network credentials and store them in plaintext. By exporting wireless profiles into temporary XML files, the tool is able to parse and retrieve pre-shared keys before deleting traces of its activity, making detection more difficult.

Security researchers warn that the script's simplicity and open-source availability significantly lower the barrier for malicious actors. The inclusion of instructions to compile the script into a standalone executable further increases its accessibility for misuse.

Although not inherently malicious, the project presents ongoing risks associated with how operating systems manage and expose sensitive credential data.

10

Compliance failures, privacy incidents, and governance issues

Key takeaways

1 Regulatory risk

Weak governance, unclear data practices, and poor compliance can trigger hefty fines.

2 Shared accountability

Using third parties doesn't remove accountability; organizations remain legally responsible.

3 Trust threats

Cybersecurity failures now threaten trust and services, even without confirmed breaches.

10

Compliance failures, privacy incidents, and governance issues

Our recommendation

Organizations should enforce strict compliance requirements, ensuring transparency, consent, and clear data handling with cloud providers. Regular vulnerability management, rapid incident response, and continuous monitoring are essential to prevent data exposure, protect sensitive information, and assign accountability across all parties.

Organizations still fail to meet compliance requirements

As regulatory scrutiny intensifies, more organizations are actively investing in meeting compliance requirements, particularly around data protection and cybersecurity. However, compliance failures continue to occur, and in some areas, they are becoming more visible.

In 2025, according to the Enforcement Tracker, regulators recorded 335 GDPR-related compliance incidents, up from 297 in 2024, highlighting a year-over-year increase in enforcement activity. At the same time, the longer-term trend shows progress: compared to over 520 incidents in 2023, GDPR failures in 2025 have declined by 36%, indicating gradual improvement despite ongoing challenges.

GDPR-compliance gaps remain concentrated in specific sectors. Industry and commerce lead with 560 recorded incidents, followed by media, telecommunications, and broadcasting with 352 cases, and individuals and private organizations with 327 incidents, according to data from the Enforcement Tracker.

In the final chapter of this report, we examine the evolving compliance landscape and highlight the most significant cases where organizations failed to meet security and cybersecurity compliance requirements, offering critical lessons for risk, legal, and security leaders navigating today's regulatory environment.

OCTOBER

Austrian regulator finds Microsoft 365 education breached student privacy

SOURCES:

Cyber Insider

Techzine

The Austrian Data Protection Authority (DSB) ruled that Microsoft 365 Education unlawfully tracked students and violated core GDPR principles, including transparency and the right of access to personal data. The case arose after a student's data access request revealed that neither Microsoft, the school, nor the Ministry of Education could clearly explain what student data was collected, how it was used, or whether it was shared with third parties.

The DSB found that Microsoft set non-essential tracking cookies without consent and used student data for vague "business purposes," such as business modeling and internal reporting, without providing sufficient explanations, an especially serious issue given that the data subjects were minors. Crucially, the authority rejected Microsoft's attempt to shift responsibility to its Irish subsidiary, determining that key data-processing decisions were made by Microsoft Corporation in the US.

The ruling holds Microsoft, the Ministry, and the school jointly responsible, ordering the deletion of unlawfully collected data and requiring clear disclosures about data use and potential sharing with entities like LinkedIn, OpenAI, or Xandr.

The decision underscores a structural problem: schools cannot fulfill GDPR obligations when cloud providers retain control and transparency over data processing. With Microsoft 365 Education used by millions across Europe, the ruling may have far-reaching implications for both educational institutions and the broader use of Microsoft 365 under EU data protection law.

Capita fined £14m for failing data protection requirements

SOURCES:

The Guardian

BBC

The UK Information Commissioner's Office has fined outsourcing giant Capita £14 million after a 2023 cyber-attack exposed the personal data of 6.6 million people, including pension members, employees, and clients' customers.

Hackers stole nearly one terabyte of data, including financial details, criminal records, and other sensitive information, after Capita failed to patch known vulnerabilities and left systems inadequately monitored.

Although the company detected the attack within minutes, it did not isolate the compromised device for more than two days, giving attackers time to deploy ransomware, reset passwords, and exfiltrate data.

The ICO concluded that Capita's security operations were understaffed and its defensive testing insufficient, despite the firm handling millions of sensitive records across public and private sector services.

An initial £45 million fine was reduced after Capita demonstrated post-incident security improvements, cooperation with regulators, and support for affected individuals. The ruling comes amid a surge in major cyber-attacks across the UK, including incidents affecting retailers and automotive firms, prompting warnings from the National Cyber Security Centre about rising national cyber risk.

Regulators said the case underscores that even large, critical service providers will be held accountable for data protection failures, as cyber resilience becomes a matter of public trust and national economic stability.

AUGUST

Illumina pays $9.8M over false cybersecurity claims to federal agencies

SOURCES:
White & Case
Crowel

The U.S. Department of Justice has reached a $9.8 million False Claims Act settlement with Illumina, marking the first case where a medical technology company was penalized for selling products with inadequate software cybersecurity to federal agencies.

The case, brought by a former Illumina employee under whistleblower provisions, alleged that for years the company failed to embed cybersecurity into the design, development, and post-market monitoring of its genomic sequencing systems while falsely certifying compliance with NIST and ISO standards.

Critically, the DOJ argued that Illumina's claims were false regardless of whether an actual cyber breach occurred, signaling that misrepresentation and weak internal controls alone can trigger liability. Regulators found that Illumina under-resourced product security, failed to remediate known vulnerabilities, and lacked effective processes to detect and manage software risks tied to sensitive genetic data.

The settlement reflects a broader DOJ strategy to use the FCA to enforce cybersecurity obligations across regulated industries, extending beyond traditional defense contractors into healthcare and life sciences.

Recent enforcement actions show a consistent pattern: companies are being penalized for inaccurate certifications and unverified security claims, not just for breaches. The Illumina case underscores that secure-by-design software, verifiable controls, and accurate cybersecurity representations are now legal requirements, not just optional best practices, for federal contractors and medical device manufacturers.

JULY

Bank of England fines Vocalink £11.9M for risk management and governance failures

SOURCES:
Fintech Global
Reuters

The Bank of England fined Vocalink Limited £11.9 million for failing to meet compliance obligations under the Banking Act 2009, marking the first fine for a financial market infrastructure firm.

Despite implementing a remediation programme following a 2021 Direction, Vocalink's ineffective risk management framework and weak governance, controls, and escalation processes prevented full compliance by the February 2022 deadline. The firm processes the majority of salaries, household bills, and state benefits in the UK, making robust risk management critical to national payment systems.

The fine was reduced from £20 million due to Vocalink's early admission of fault and cooperation during the investigation. This enforcement highlights the increasing rigor of UK regulators and underscores the need for integrated risk frameworks, clear escalation paths, and proactive governance in the financial services sector.

MAY

Raytheon pays $8.4M over cybersecurity failures on DoD contracts

SOURCES:
Arnold & Porter
Compliance Week

The U.S. Department of Justice announced that Raytheon, RTX Corporation, and Nightwing Group would pay $8.4 million to resolve whistleblower allegations that Raytheon violated the False Claims Act by failing to meet DoD cybersecurity requirements on certain contracts.

The lawsuit, filed by a former Raytheon engineer, alleged that Raytheon and its subsidiary did not implement required cybersecurity controls or system security plans, in violation of DFARS and FAR regulations.

The settlement underscores the DOJ's increased focus on cybersecurity compliance under the Civil Cyber Fraud Initiative and highlights the role of qui tam whistleblowers, who receive a portion of government recoveries. This action reflects broader enforcement trends holding defense contractors accountable for protecting sensitive government data against evolving cyber threats.

MARCH

MORSE Corp fined $4.6M for cybersecurity compliance failures on DoD contracts

SOURCES:
SC Media
Compliance Week

Massachusetts-based defense contractor MORSE Corp agreed to pay $4.6 million to

settle allegations under the False Claims Act that it failed to meet cybersecurity requirements in contracts with the U.S. Army and Air Force.

The DOJ alleged that MORSE did not properly disclose its third-party email provider's compliance, failed to provide system security plans from 2018 to early 2021, and submitted falsified Supplier Performance Risk System scores regarding its NIST SP 800-171 controls. Although the company emphasized that the settlement addressed historic allegations, the case underscores the government's enforcement focus on accurate reporting and adherence to cybersecurity standards for contractors handling sensitive data.

This settlement signals that misrepresenting cybersecurity compliance, even without evidence of a breach, can carry significant financial and regulatory consequences.

FEBRUARY

Health Net and Centene settle $11 M false claims act case over cybersecurity compliance failures

SOURCES:

SC Media
The National Law Review
The HIPAA Journal

Health Net Federal Services (HNFS) and its parent company, Centene Corporation, agreed to pay $11.25 million to settle allegations that they falsely certified compliance with cybersecurity requirements while managing the TRICARE healthcare program for U.S. service members and their families.

Between 2015 and 2018, HNFS allegedly failed to implement critical cybersecurity controls, including vulnerability scanning, patch management, access controls, and firewall configurations, yet submitted annual reports asserting full compliance.

The DOJ investigation found that HNFS ignored internal and third-party audit warnings, potentially putting sensitive military health data at risk, although no breaches were reported.

The settlement resolves the False Claims Act allegations without admission of liability and underscores the government's enforcement focus on cybersecurity compliance for contractors handling sensitive federal data. Whistleblowers played a key role in bringing the case forward, highlighting the importance of the FCA in protecting taxpayer-funded programs and national security. This action serves as a warning to contractors that false certifications of cybersecurity compliance can carry substantial financial and reputational consequences.

Conclusion

2025 IN REVIEW

The DevOps threat landscape of 2025 made one reality unmistakably clear: modern development environments have become primary attack surfaces. AI-assisted attacks, supply chain compromises, identity abuse, outages, and large-scale data breaches showed that security gaps in DevOps environments are no longer isolated technical issues but systemic business risks.

To withstand this reality, organizations should move decisively toward a DevSecOps-first model, where security is embedded into every stage of the SDLC and shared across development, operations, and security teams. Shifting security left, automating continuous testing and monitoring, and treating AI tools, prompts, and outputs as untrusted by default are no longer optional practices. They should become baseline requirements.

Equally critical is strong identity and access governance. Least-privilege access, short-lived and scoped credentials, continuous secret rotation, and strict token hygiene are essential to limit blast radius when accounts or automation pipelines are compromised. Clear security policies, compliance alignment, and ongoing employee awareness must reinforce these controls across the organization.

Finally, resilience must be designed, not assumed. As outages, ransomware, human error, and SaaS failures continue to rise, comprehensive DevOps backup and disaster recovery strategies become the last line of defense. Immutable, encrypted, multi-location backups; point-in-time restore; cross-platform recovery; and centralized monitoring are essential to guarantee business continuity under any failure scenario.

The lesson of 2025 is clear: protecting DevOps is building a resilient, governed, and security-aware development ecosystem. Organizations that invest now in DevSecOps culture, automation, identity controls, and robust backup and recovery will be far better positioned to face the next wave of threats, whatever form they take.

So, let's sum up the lessons learned from the attacks and incidents in 2025:

01 Organizations should treat AI tools as untrusted actors by default. They should enforce strict input sanitization, least-privilege access for AI agents, short-lived credentials, and human-in-the-loop reviews, as well as verify repository provenance, sandbox AI outputs, harden CI/CD workflows, and continuously monitor AI-driven actions across dev environments.

02 Companies should harden their CI/CD and developer workflows by enforcing short-lived, least-privilege tokens; continuously monitor repositories, workflows, and dependencies for anomalies; and verify all third-party code, PoCs, and tools before using them. They need to treat automation, secrets, and open-source trust as critical security boundaries, not defaults.

03 To protect their data, organizations should enforce least privilege access, rotate credentials frequently, and use short-lived tokens. It's critical to continuously monitor CI/CD workflows, repositories, and cloud accounts for anomalies. Adopt phishing-resistant MFA, secret management, and automated dependency scanning to reduce supply chain and identity-based risks.

04 To reduce outage impact, organizations should build a reliable backup and disaster recovery plan. The security strategy should be designed the way to avoid single-provider dependencies, implement multi-region and multi-service redundancy, and test their recovery regularly. Organizations should be able to maintain offline access, backup communication paths, and a clear incident response plan to sustain operations during cloud disruptions.

05 Companies should keep an eye on vulnerability updates and recommendations that DevOps platforms release. Implementing on-time patches, auditing third-party dependencies, restricting guest access, and continuous monitoring for anomalies to reduce the risk from high-severity vulnerabilities and supply-chain exploits are also necessary measures.

06 The phishing and identity attacks of 2025 show that to resist such threats, organizations should restrict OAuth device code flows with Conditional Access, enforce app consent approvals, and continuously audit authorized applications. They should combine phishing-resistant MFA with behavior-based detection, user training focused on device codes, and trusted platform abuse.

07 Organizations should monitor and restrict the use of public repositories for internal workflows, enforce strict access controls on API tokens, validate code from third-party sources, and deploy behavioral detection to flag unusual

repository activity. Thus, they can prevent the impact of malware, ransomware, and APT campaigns that abuse trusted platforms.

08 Enforcing strict credential hygiene, rotating and revoking old access tokens, and implementing least-privilege access across DevOps platforms should become the core of security practices as well. Continuous monitoring, repository audits, and employee security training are essential to prevent unauthorized access, insider threats, and supply chain abuse.

09 Organizations should enforce strict code vetting, limit access to repositories, and scan third-party projects for malware before use. Also, it's worth combining repository monitoring, behavioral analysis, and credential hygiene with employee awareness to prevent attacks exploiting fake or trojanized dev projects, malvertising, and supply-chain abuse.

10 Companies should enforce strict compliance requirements, ensuring transparency, consent, and clear data handling with cloud providers. Regular vulnerability management, rapid incident response, and continuous monitoring are essential to prevent data exposure, protect sensitive information, and assign accountability across all parties.

GitProtect