Intigriti Bug Bytes #240 - September 2026

BY AYOUB · SEPTEMBER 25, 2026

Want to read the latest version? Visit the link below:
https://www.intigriti.com/researchers/blog/bug-bytes/intigriti-bug-bytes-240-september-2026

Hi hackers,

Welcome to the latest edition of Bug Bytes! In this month's issue, we'll be featuring:

- Compromising OpenAI, Slack, Meta, and more via a vulnerable image library
- Hacking OpenAI employee accounts in under 72 hours
- Breaking into Google's GFile for $100K
- Hacking AI CX agents
- Turbo Intruder 2 surpassing 100K requests per second over HTTP/3

And so much more! Let's dive in!

Reconnaissance unleashed: meet CrowdRecon

CrowdRecon is officially here! CrowdRecon brings crowd-led context to the moments before and between vulnerability reports. It helps customers quickly identify where attention is needed across exposure, scope, coverage, investigation, and follow-up. For researchers, it introduces new ways to collaborate, contribute, and earn recognition for the valuable recon work that usually disappears.

The public beta is open! Register today and be among the first to try it.

Reconnaissance unleashed: meet CrowdRecon

How AI changed my day as a QA Engineer

In our latest post, Senior QA Engineer Martin Klimovski shares how he is leveraging AI for repetitive tasks to free up time for deep exploratory work and security testing. The piece covers how he removes friction while keeping judgment, and his views on AI security and safety.

From sceptic to supercharged. How AI changed my day as a QA Engineer

How AI has changed the way I think, build, and work

Senior Software Engineer Koen Van Hauwe shares an honest look at a day in his life and his shift from writing code to steering AI. The post covers the Orchestrator shift and the critical role of human oversight, the risks of speed and subtle mistakes that build up over time, and how AI introduces new attack surfaces alongside opportunities for defenders.

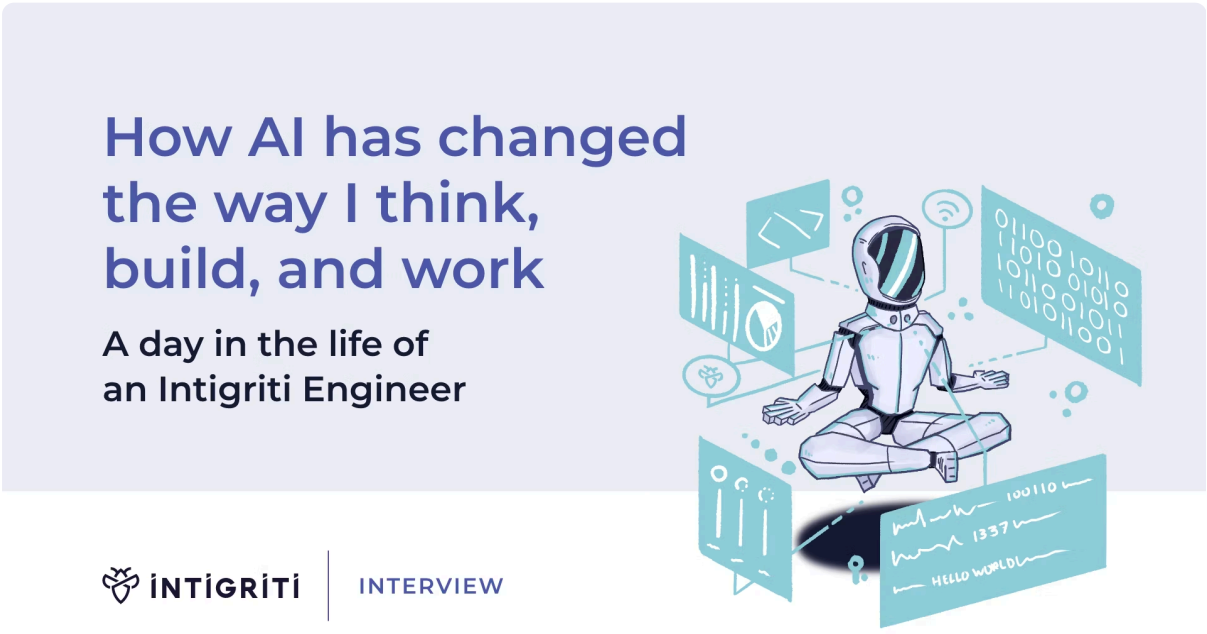

How AI has changed the way I think, build, and work. A day in the life of an Intigriti Engineer

Quick! Intigriti 0926 Critter Gallery Challenge is still live

Intigriti's 0926 Challenge, Critter Gallery by @khanhdlq, is still ongoing. Capture the flag before Monday the 28th of September for a chance to win €400 in swag prizes.

Take on the challenge!

Intigriti Challenge 0926

Intigriti 0826 Bad Reception CTF results are in

August's CTF challenge, Bad Reception, featured a broken TV that revealed the flag once fixed. This challenge definitely sent lots of participants down rabbit holes and dead ends, which made it more engaging than ever.

Quick recap:

- First blood was captured by martijnperdaan
- 79 hackers found the correct solution
- 21 hackers wrote a cool write-up

If you want to put your hacking skills to the test, be sure to give Bad Reception 0826 a go before heading over to Bugology, where you can find all the researchers' submitted solutions.

Intigriti Challenge 0826

Blogs & videos

Hacking AI customer service agents

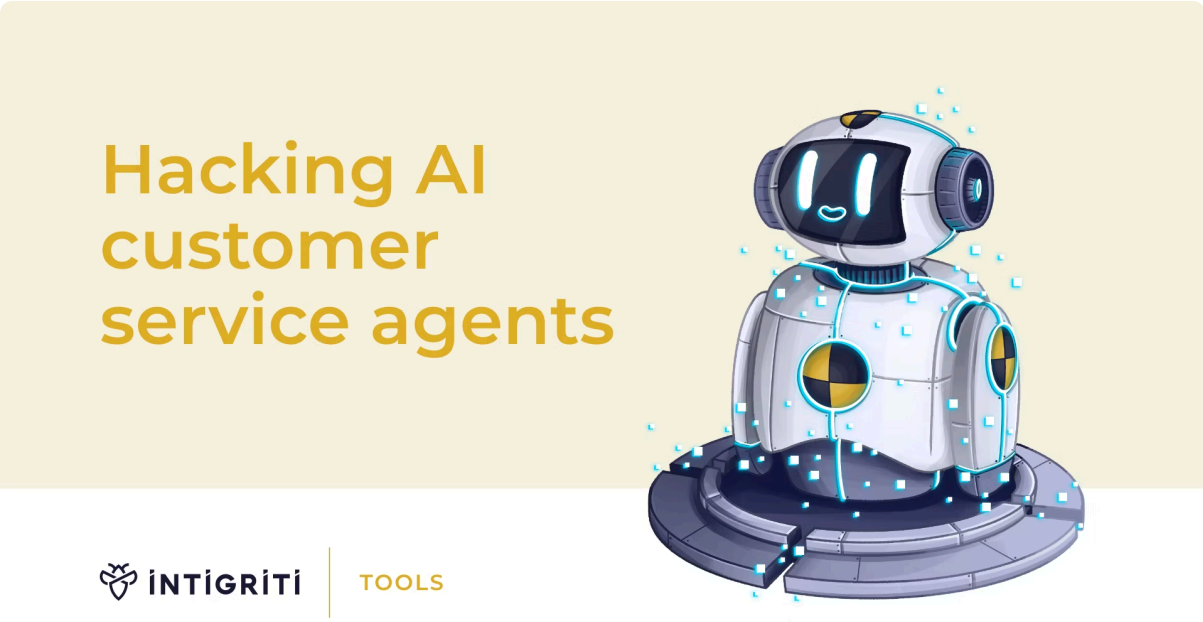

Hacking AI customer service agents Cover Image

Most support chatbots have evolved into fully autonomous agents that can help with almost anything, and that comes with additional risk. Our latest article, based on @intidc's talk at Bug Bounty Village during DEF CON 34, breaks down the full attack surface in AI customer service agents: from tricking them into leaking sensitive data via email spoofing to invoking unauthorized tool calls on behalf of the victim. All without ever touching an automated scanner or proxy interceptor.

- Looking for a complete guide on file upload vulnerabilities? Our full guide on exploiting insecure file uploads covers everything from basic bypass techniques to advanced exploitation scenarios.

- Reports sometimes get closed incorrectly, downgraded in severity, or left pending for longer than expected. When that happens, knowing how to respond through the right channels matters. Our comprehensive guide walks you through Intigriti's mediation process, the most common scenarios researchers face, and the mistakes to avoid during the process, such as unauthorized public disclosure, which can work against you rather than help your case.

- Going from zero to your first valid bug report? In case you missed it, we recently launched the Bug Bounty Starter Kit, a free guide covering everything from recon and tooling to the exploitation of SQLi, XSS, and BAC vulnerabilities, to finally how to learn to write a compelling vulnerability report that gets triaged faster. Get your copy now.

Tools & resources

Tools

GeminiHunter

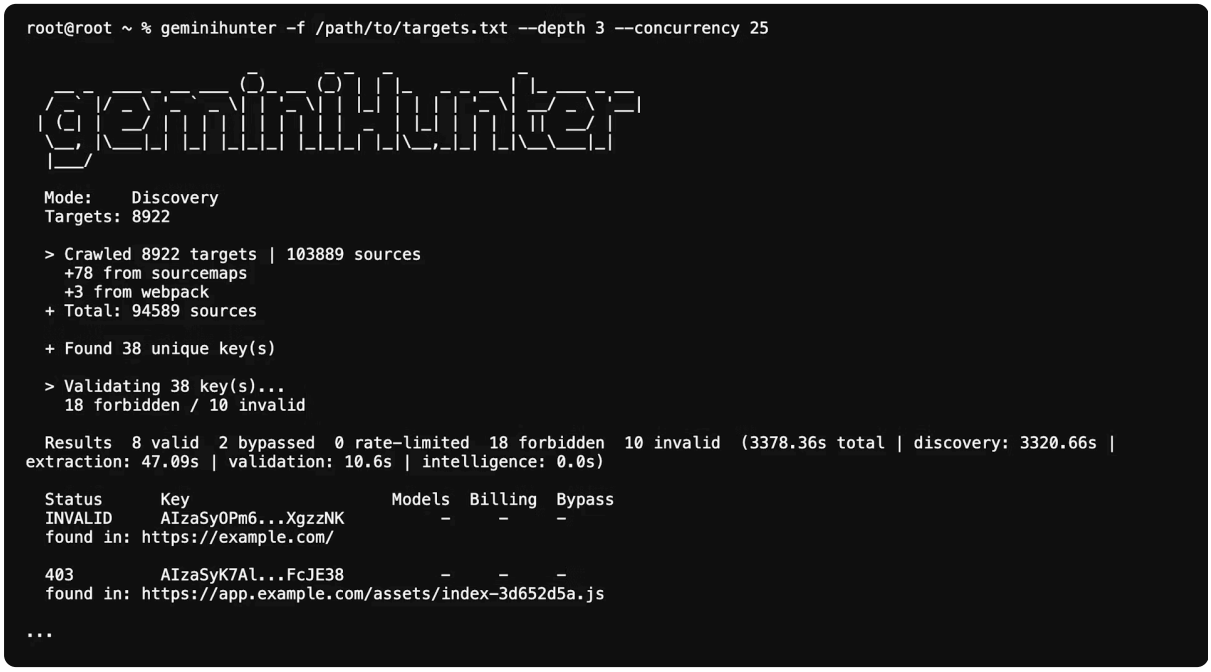

GeminiHunter

Did you know that some exposed Google Maps API keys also have access to Gemini models? GeminiHunter by @devploit finds and validates exposed Google Gemini API keys across web assets, source maps, Wayback snapshots, and Android APKs. It deduplicates findings, checks if keys are actually valid, and attempts bypass variations on restricted ones. If you are hunting Google API key exposures, this adds a new dimension to check.

- Prompt injection is the top risk in the OWASP LLM Top 10. Awesome Prompt Injection is a curated collection of resources covering everything from research papers and real-world attack chains to hands-on CTFs and detection tools for testing LLM and agentic applications.

- Want your AI agent to interact directly with Chrome? Chrome DevTools MCP lets AI agents use Chrome's debugging protocol to inspect network requests, DOM elements, execute JavaScript in the console, and capture screenshots, all programmatically.

Resources

HEIF Heist

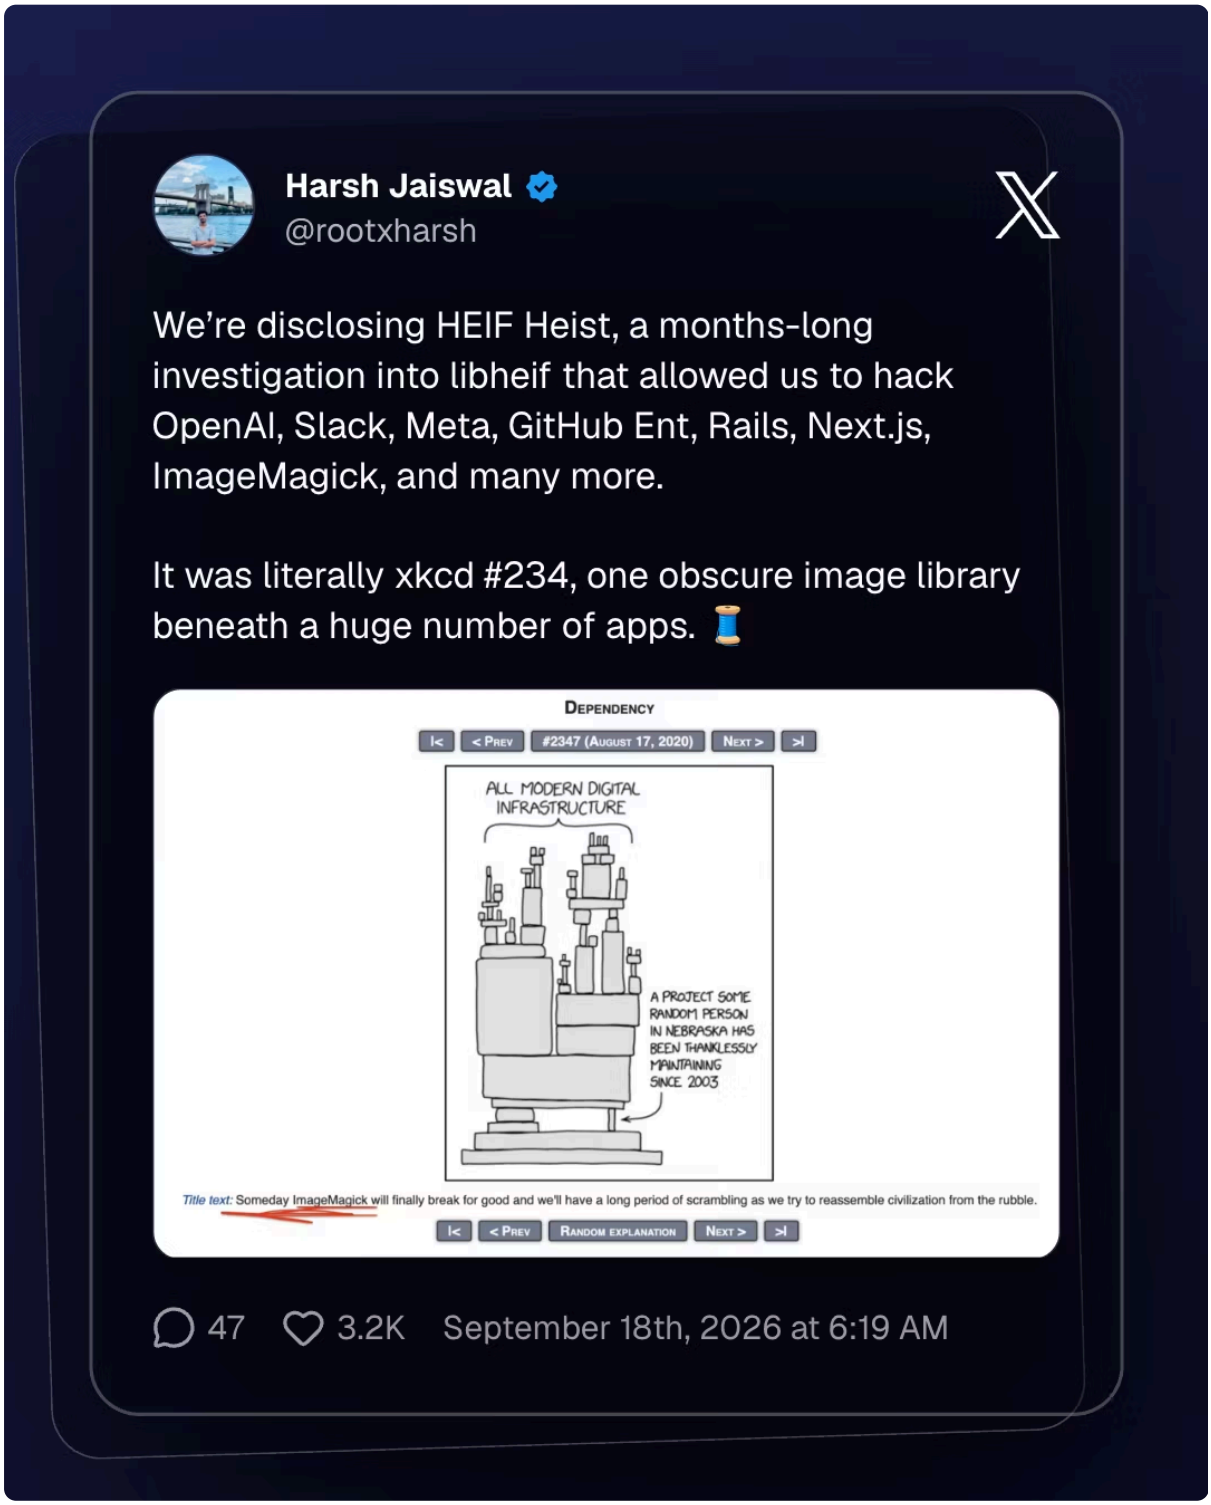

HEIF Heist research

@rootxharsh and team published HEIF Heist, a months-long investigation into libheif that compromised OpenAI, Slack, Meta, GitHub Enterprise, Rails, Next.js, ImageMagick, and more. The attack surface was a single obscure image parsing library sitting beneath a huge number of applications. Some of the RCE and information leak attempts only landed after thousands of image uploads, and only one organization caught the exploitation in progress.

- **Two bugs, 72 hours, and a PR in OpenAI's internal codebase.** @S1r1u5_ and team published the details on how they took over ChatGPT and Codex accounts of OpenAI employees, reaching connected services including Outlook, Slack, and GitHub.

- **Breaking into Google's GFile for $100K.** Brutecat documents how he leveraged AI to scan discovery documents at scale across Google's infrastructure, then found a new approach to exploit the GFile library for access to internal filesystems and storage.

- **A valid URL with a Command in its Shadow.** @saur1n found that a connectivity check was passing user input to cURL inside a shell, including a second feature, which ultimately led to a command injection vulnerability.

- **Turbo Intruder 2 has landed.** @t0xodile from PortSwigger released a new version of Turbo Intruder with a native HTTP/3 engine, surpassing 100,000 requests per second over WiFi. It also includes a Burp adapter for HTTP/3 exclusive targets and new research-grade race condition techniques.

- **Full control over 3,000 companies and all their registered branches.** @stuipds published a detailed write-up covering three weeks of studying a target to find a chain of vulnerabilities that led to complete company takeover via secondary context.

- **A remote kernel vulnerability in macOS SMBFS, $20,000 bounty from Apple.** @slinafirinne published the write-up for CVE-2026-84543, his second $20K Apple bounty.

- **Executing an ELF without touching the filesystem.** This research article demonstrates using the Linux kernel keyring to stage an ELF in slab memory and execute it via userland exec, skipping execve and the filesystem entirely.

- **Exploiting tool and function calling in LLM agents.** Sentry Security published a post covering the attack surface in LLM agents when they are given access to external tools.

- **Path traversal plus arbitrary file write turned into RCE.** @MrTuxRacer published the details on CVE-2026-28373 in the Stackfield desktop app, where an encrypted backup was enough to achieve remote code execution.

- **LFI filter bypass via emojis.** @0xConda shared a practical path traversal bug bounty tip with emojis to bypass strict filters.

- @0xacb shares a practical bug bounty tip on the fundamental concept of parser differentials. Understanding the concept in depth can help you land more vulnerabilities, consistently.

- **To read private chat rooms in Better Messages, just tell it you are the AI bot.** @_CryptoCat published the analysis of CVE-2026-89093: the identity was an IP-prefix check, and the IP came from a request header.

- **Guest checkout coordinates flowing straight into a SQL query.** @_CryptoCat also published another write-up on CVE-2026-18442 in WCFM Marketplace, where a store-distance shipping calculation was vulnerable to SQL injection.

- **How close is "one click away from shutting down every restaurant in Europe"?** This write-up covers a high-severity issue in a restaurant platform that allowed malicious users to disrupt the delivery service.

- **Two researchers bought cheap domains, including noreply.net and deleteduser.com, then set up email listeners.** Hundreds of companies have been caught sending corporate secrets to those emails. Read the full story covered by Wired.

- **$76,000 from a single program on Bugcrowd.** @4non_Hunter published a breakdown of how sustained focus on one target led to a significant payout.

- **A year of hacking with LLMs.** @edwardzpeng shared his keynote from the Offbyone security conference, covering a full year of integrating LLMs into vulnerability research.

- **Rickrolling the entire FIFA World Cup.** This write-up details how a single ID was enough to gain unauthorized control over an administrative panel that could, in theory, have let him rickroll the entire FIFA World Cup.

- **Escalating a blind SSRF to internal cloud metadata, $4,200 bounty.** This researcher documented how he was able to escalate a blind SSRF to read cloud metadata.

- **An IDOR in an invite cancellation flow exposed personal data.** The write-up covers how scheduled, recurring actions led to unintended PII exposure.

- **A default password let one researcher log into almost anyone's account.** The write-up covers a default-credential issue leading to mass account access.

- **11,000+ bug bounty reports in one searchable library.** @_Shark_byte compiled an open library of disclosed reports across platforms. Whether you're new to web hacking or a well-seasoned researcher, you can always learn from other researchers' findings and experiences.

- **Mutating Safari.** @garethheyes published a brief post on a Safari-specific behavior found by Shazzer.

- **Pwnfox, but for Brave.** Now that Brave supports containers, @leo__rac built pwnbrave to replicate the Pwnfox workflow without being tied to Firefox.

- **Built to pass OSCP, now used by the community.** @DotNetRussell built BloodBash for Active Directory practice and shared it at DEF CON 34, where others told him they used it to pass their OSCP as well.

- **James Kettle's (@albinowax) 'Can AI do novel security research?' talk is now available on YouTube.**

- **Did OpenClaw really hack a gym?** @PinkDraconian published a video investigating whether OpenClaw really hacked a gym to claim an already claimed spot.

Company news

Wild West Hackin' Fest

Intigriti is a Gold Sponsor of Wild West Hackin Fest in Deadwood, South Dakota. If you are attending, be sure to come say hi!

Wild West Hackin' Fest

Feedback & suggestions

Before you click away: Do you have feedback, or would you like your technical content to get featured in the next Bug Bytes issue? We want to hear from you. Feel free to send us an email at community@intigriti.com or DM us on X/Twitter, and we'll take it from there.

Did you like this Bug Bytes issue? Consider sharing it with your friends and tagging us along on X/Twitter, Instagram, or LinkedIn.

Wishing you a bountiful month ahead,

Keep on rocking!

AUTHOR

Ayoub

Senior security content developer

REQUEST A DEMO
intigriti.com/demo

VISIT THE WEBSITE
intigriti.com

GET IN TOUCH
hello@intigriti.com