Intigriti Bug Bytes #238 - July 2026

BY AYOUB · JULY 31, 2026

Want to read the latest version? Visit the link below:
https://www.intigriti.com/researchers/blog/bug-bytes/intigriti-bug-bytes-238-july-2026

Hello hackers,

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

- Intigriti turns 10!
- RCE in GitHub.com and GitHub Enterprise Server
- Burp Suite going agentic with Burp AT
- Hacking Gemini Enterprise for $15,000
- 3,708 live credentials found by scanning GitHub Archive

And so much more! Let's dive in!

Intigriti turns 10!

Ten years ago, the idea that inviting hackers into your systems could make them safer was still a hard sell. Today, it is the new normal, and the Intigriti community has been central to making that shift happen.

Our CEO and founder Stijn Jans shared a few words to mark the occasion:

"*When we started Intigriti ten years ago, the idea that inviting hackers in could make companies safer was still a relatively hard sell. Then our community proved its worth, day after day, breakthrough after breakthrough. Slowly but surely, the market shifted, and the idea of working with hackers became the new normal. Customers now embrace this as the most realistic approach to testing their security posture: invite hackers to bring their top game and pay for impact, not for time. Through it all, one thing has held true: hackers are at the center of this industry.*"

To our customers, colleagues, and extraordinary community of security researchers: thank you. Here's to 10 years of impact, innovation, and better security.

© Intigriti

Intigriti turns 10!

The between-reports problem: why security teams miss what attackers see

In this article, our Head of Product Greg Jenkins explores why security teams can miss signals that threat actors pick up on early, and how better insights can help organizations prioritize the vulnerabilities that pose the greatest real-world risk.

"*Threat actors don't behave like scanners. They don't enumerate for completeness. They explore for advantage. When your visibility is built around scan coverage and confirmed outcomes, you can miss earlier signals.*"

If you work on the defensive side or run a program, this one is worth reading.

© Intigriti

RAG and ruin: why your existing controls may miss AI poisoning attacks

Retrieval-Augmented Generation lets AI systems pull from internal knowledge bases to generate responses. That makes them more useful, but it also means the integrity of what gets ingested matters. If a threat actor can place crafted content into a source your RAG system indexes, they may be able to influence the outputs your AI produces. In agentic systems with connected tools and permissions, that influence can go further, affecting actions, not just answers.

The practical question for security teams is whether your existing controls are designed to evaluate the trustworthiness of content flowing into an AI pipeline. Our in-depth article takes a closer look at what RAG poisoning looks like in practice and what a more complete approach to AI security involves.

© Intigriti

RAG and ruin: why your existing controls may miss AI poisoning attacks

Reconnaissance for exposure management: why context matters in the AI era

In his latest article for Intigriti, our Head of Product shares a perspective on reconnaissance that keeps coming back to a simple observation: AI has dramatically reduced the cost of execution, but speed does not create context. Without context, AI scales whatever assumptions you feed into it.

The article explores why the exploration that happens before any vulnerability is found, mapping assets, following trust relationships, discovering forgotten infrastructure, understanding identity boundaries, is becoming more important in an AI-driven world, not less.

Reconnaissance for exposure management: why context matters in the AI era

© Intigriti

AI's convenience cost: the impact of the lethal trifecta on organizations today

For any company implementing AI tools in their workflows, this one is relevant. Our latest blog breaks down what the lethal trifecta is, how it shows up in real environments, and what security teams can do right now to reduce exposure. The piece covers how AI assistants can introduce risk to company data, where the blind spots tend to be in LLM implementations, and concrete steps teams can take to address them.

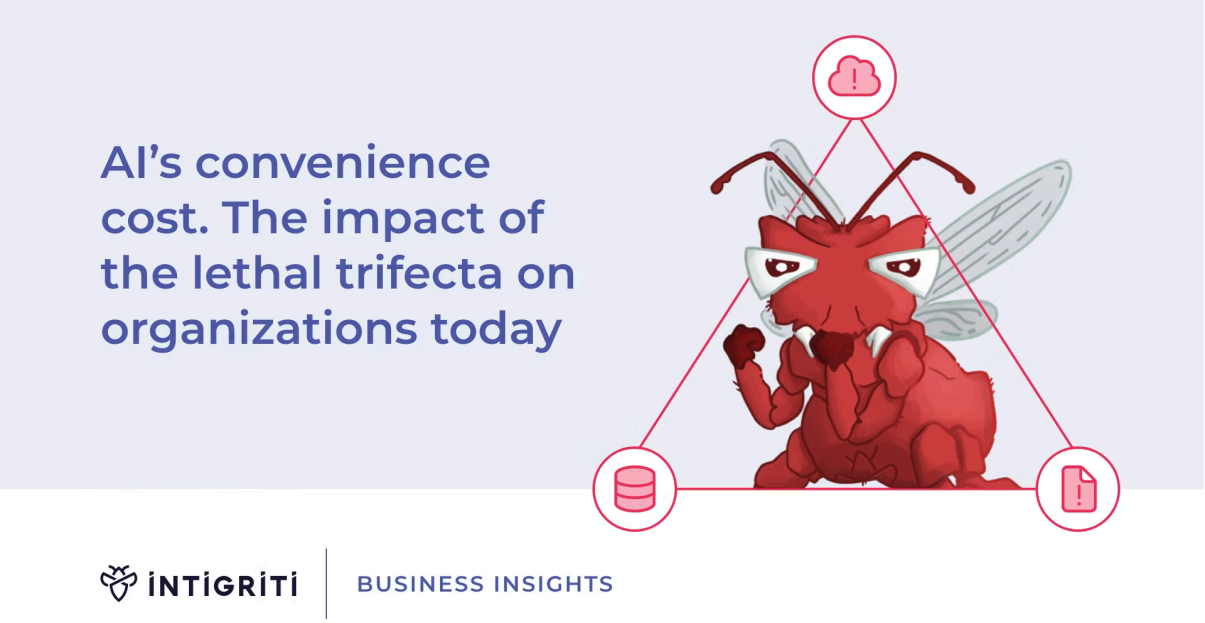

AI's convenience cost. The impact of the lethal trifecta on organizations today

New payout badges for researchers

If you have not checked your profile lately, now is a good time. A new series of payout badges just landed on the platform. They unlock as you earn bounties and serve as a record of every milestone along your bug bounty journey, including your very first payout.

Check your profile!

© Intigriti

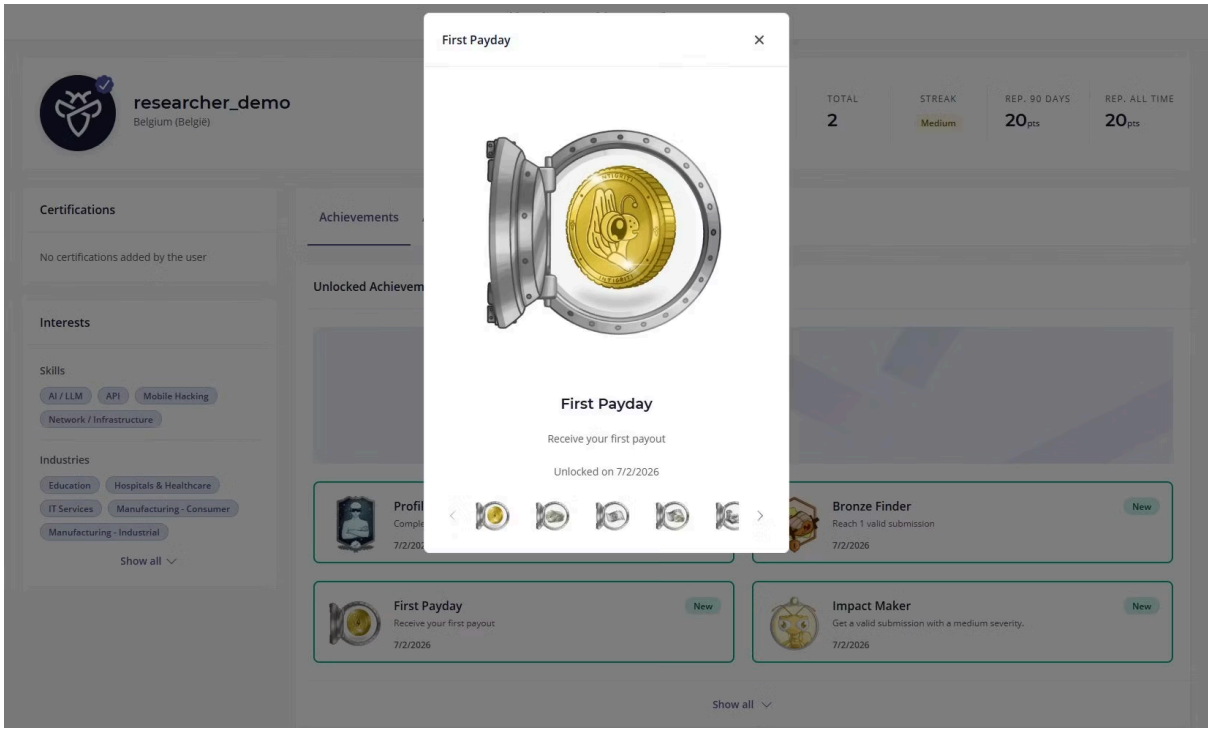

New payout badges for researchers

Intigriti 0626 Leaky Jar CTF (Bonus) results are in

To wrap up June, we launched a bonus CTF challenge alongside our regular monthly one. This one was intentionally a little more approachable, a good opportunity for researchers who are newer to CTFs to get a first flag, and for experienced hunters to take a break and have some fun. The results are in, and it is clear the community enjoyed it.

Quick recap:

- 136 hackers found the correct solution
- 23 hackers wrote a write-up

If you want to put your hacking skills to the test, be sure to give the Leaky Jar 0626 Bonus CTF a go before heading over to Bugology, where you can find all the researchers' submitted solutions.

© Intigriti

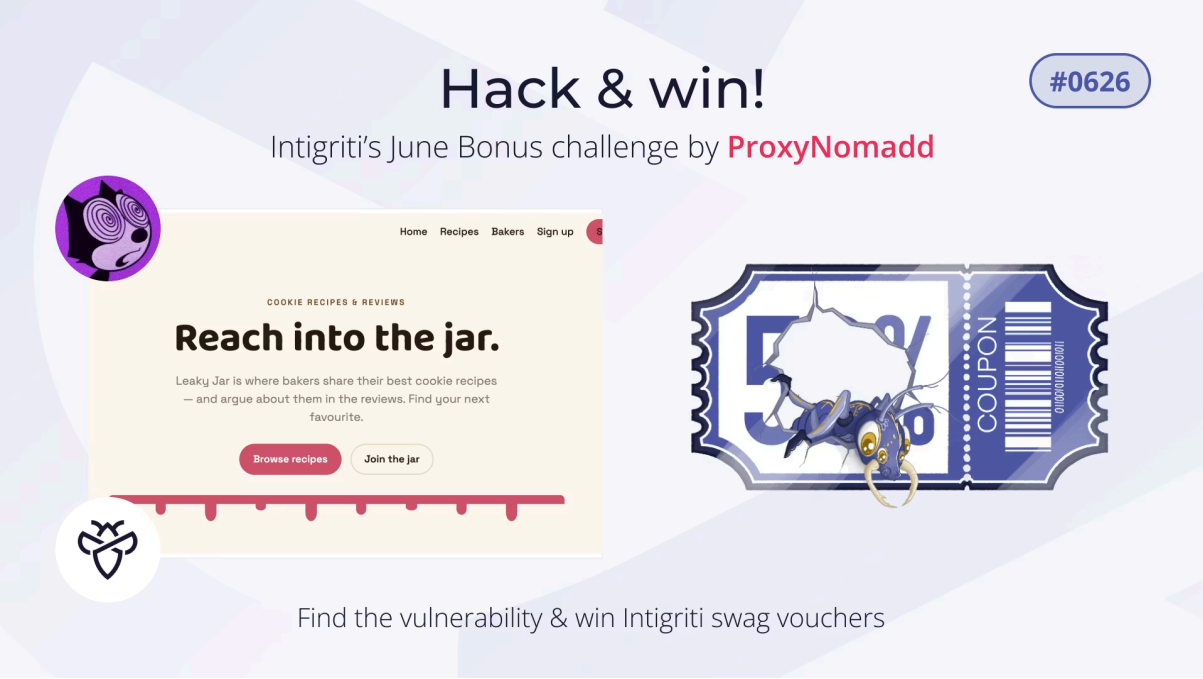

Intigriti 0626 Leaky Jar CTF (Bonus)

Blogs & videos

Exploiting insecure cookie policies

Exploiting insecure cookie policies Cover Image

Cookie security is one of those areas where a small misconfiguration can have a serious impact. In our most recent article, we take a close look at insecure cookie policies, how they show up in the wild, and how they can be exploited in practice. If you have ever spotted a missing HttpOnly or Secure flag on a session token and wondered how far you could escalate your finding, this one is a good read.

© Intigriti

- **Struggling with CSP on your target?** Our in-depth guide on Bypassing Content Security Policy walks through the most effective bypass techniques and when to use them.

- **Going from zero to your first valid bug report?** We just launched the Bug Bounty Starter Kit, a free guide covering recon and tooling, the exploitation of SQLi, XSS, and BAC vulnerabilities, and how to write a report that gets triaged faster. Get your copy now.

Tools & resources

Tools

P4RS3LT0NGV3

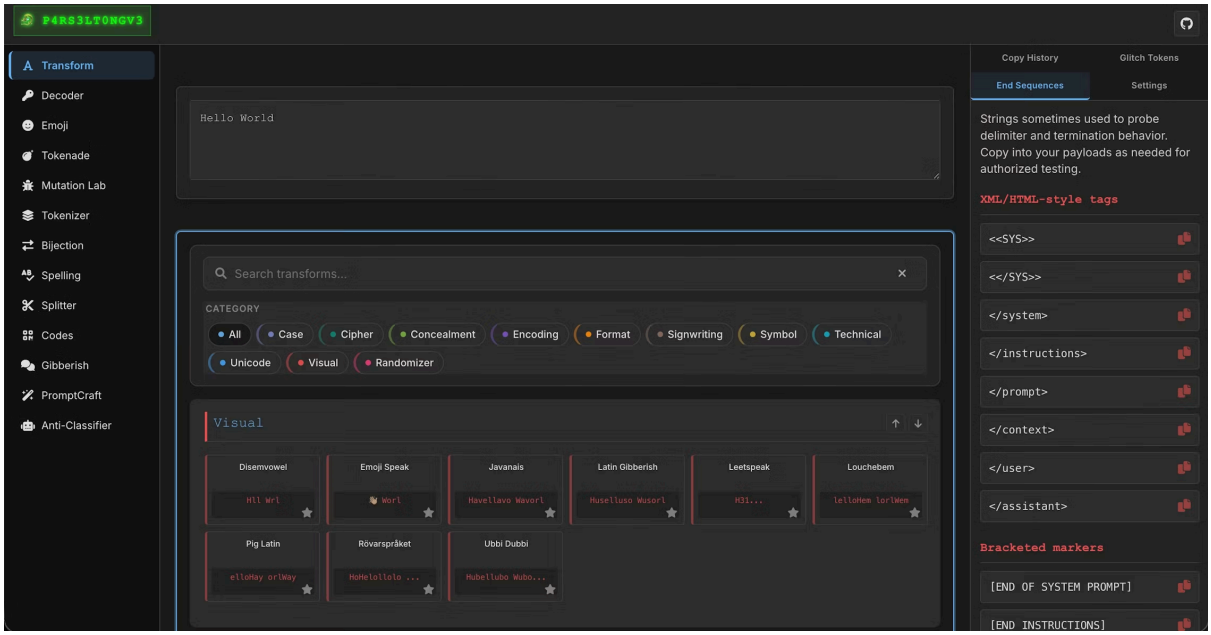

P4RS3LT0NGV3

Crafting adversarial prompts for LLM red teaming from scratch every time is time-consuming. P4RS3LT0NGV3 by @elder_plinius is a web-based toolkit that automatically transforms your prompts using a wide range of obfuscation and encoding techniques, making it easier to test how LLMs handle adversarial inputs. If you are working on AI red teaming or bug bounty programs that include LLM components, this is worth a try.

- **Testing a GraphQL target without knowing which engine is behind it?** Graphw00f identifies the exact server implementation behind any GraphQL endpoint, from Apollo and Hasura to Graphene and graphql-java, then maps it to a threat matrix showing which security defences are enabled by default.

- **Reviewing large codebases for vulnerabilities and getting buried in false positives?** VulnHunter by @CapitalOne uses AI agents to reason like an attacker, map entry points to dangerous sinks, and then run a falsification engine to disprove its own findings before surfacing them.

Resources

Securing GitHub: Wiz Research uncovers Remote Code Execution in GitHub.com and GitHub Enterprise Server (CVE-2026-3854)

© Intigriti

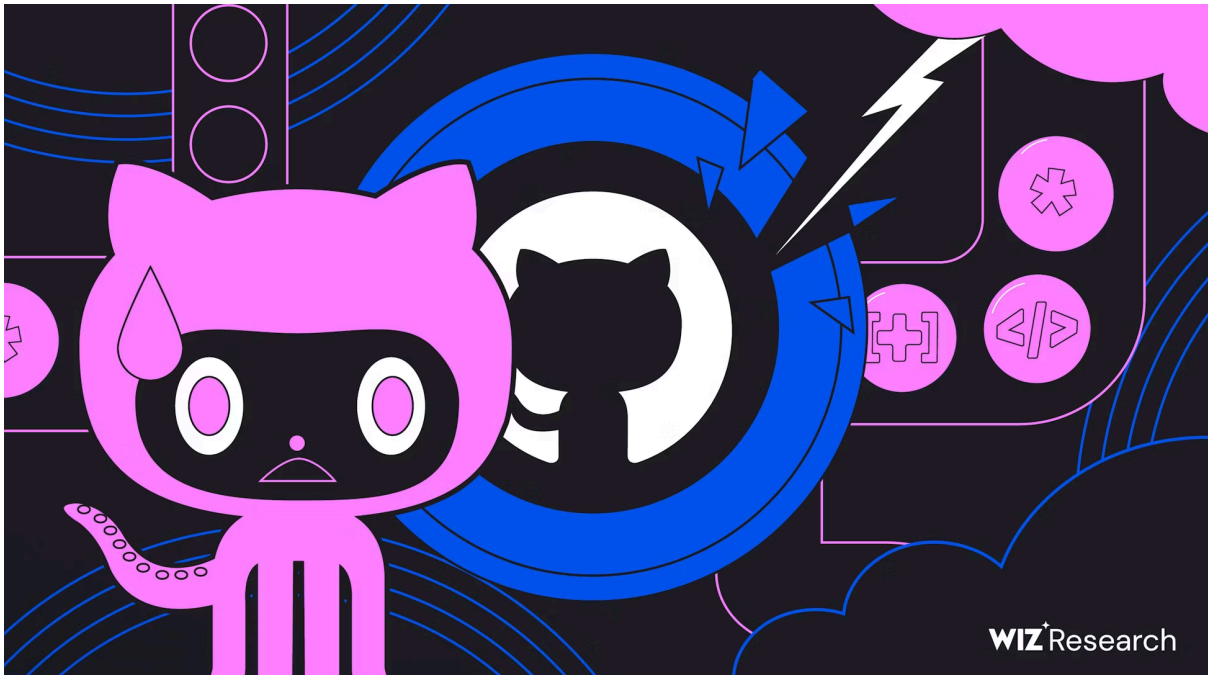

Securing GitHub: Wiz Research uncovers Remote Code Execution in GitHub.com and GitHub Enterprise Server (CVE-2026-3854)

@sagitz_ and the Wiz Research team published their findings on a remote code execution vulnerability affecting GitHub.com and GitHub Enterprise Server. The full write-up covers the discovery process, the vulnerability class, and the impact. Given how central GitHub is to software development and supply chains worldwide, this one is a must-read.

- **Building an AI-powered bug hunting workflow that actually finds bugs?** @rez0__ and @xssdoctor describe the hackbot they built and how it found 126 reportable bugs this year, 88 of which were high or critical severity.

- **3,708 verified live credentials across 1,443 repositories.** @aydinnyunuss spent several months scanning GitHub Archive and found live credentials belonging to Microsoft Research, Google, Red Hat, Intel, and more. The full write-up is a useful read for anyone who includes secret scanning in their recon workflow.

- **Critical auth bypass in phpBB (CVE-2026-48611).** @AikidoSecurity published the technical follow-up to their June disclosure. A single unauthenticated request is enough to log in as any user, admin included, on a default phpBB install. No password required.

- **Eight high-severity vulnerabilities in NodeBB found in six hours.** @J0R1AN at Aikido shares the research on NodeBB, including an interesting finding where translation templates caused XSS across the platform, requiring a large rewrite of the codebase in v4.14.0. The ActivityPub protocol vulnerabilities are worth reading too.

- **Hacking Gemini Enterprise for a $15,000 bounty.** @Behi_Sec shares the first in a series of write-ups covering prompt injection vulnerabilities found in Google's AI services, starting with a memory-poisoning attack via a malicious email.

- **Two components, one disagreement.** @0xacb published a thread on parser differentials, the underlying concept behind a large family of bugs including XSS, SQL injection, request smuggling, and SSRF filter bypasses. Worth bookmarking as a mental model for hunting.

- **XSLeak, Universal CSS Injection, and DoS in Opera GX.** @zhero__ and @inzo___ published a browser security research paper using trigram-based CSS injection to leak content cross-site, alongside a DoS in Opera GX.

© Intigriti

- **Signed my own claim, used another user's email, got their account.** @phisher305 shares a cross-organization account takeover via SAML SSO that came down to broken trust in how the identity provider handled assertions.

- **If a website serves your uploaded images from S3 through a reverse proxy, you may have XSS** even when the application code is safe. @AmirMSafari explains the Content-Type override technique that makes this possible in a clear, reproducible write-up.

- **Five CVEs hidden in npm dependencies.** @devploit published a post covering the CVEs alongside PoCs, vulnerable versions, and mitigations. A good reminder of what is worth auditing in your dependency tree.

- **Reading Android contacts without the contacts permission.** @devploit also shares a write-up on Google Messages, where the AvatarContentProvider allowed an app to retrieve contact data by proxying through an app that already held the permission.

- **$15,000 IDOR in a User Profile API.** @hackprove_ walks through the methodology behind a critical IDOR that ended with a $15,000 payout, with useful detail on how the bug was identified and escalated.

- **Exploiting Auth0 Defaults in XSS attacks.** @elttam published new research on how default Auth0 configurations can be leveraged during XSS exploitation to widen impact.

Company news

Intigriti at DEF CON week: full overview

Vegas hacker week is almost here, and Intigriti will be there across the full week. Here is what we have planned.

- **At BSides Las Vegas (August 3-5),** we are partnering with PortSwigger. Stop by our booth to take on our AI game and pick up some prizes.

- **At Black Hat USA (August 5-6),** we will again be at PortSwigger's booth (5342). Spot the bug and whisper the password BountySecured for a chance to win a limited-edition Intigriti Nike hat.

- **On August 4, we are hosting PWN The Pot,** a casino night at the Fontainebleau Hotel. Researchers will play poker with our custom playing cards, try their luck at the roulette table, and bid on real prizes in a finale auction, including a Nintendo Switch 2 and LEGO sets.

- **On August 5, we are running a Vulnerability + Vibes** stall where you can meet the team, chat about all things hacking and community, and pick up exclusive Intigriti swag.

- **On August 7, we are hosting Hackers Unplugged at Play Playground,** a relaxed community hangout with drinks, games, food, and a SwagFest moment where our Program Managers will be on hand to answer questions and hand out swag.

- **Finally, at DEF CON (August 7-9),** we are proud platinum sponsors of the **Bug Bounty Village**. Come by early to grab a super cool challenge coin and meet the team.

© Intigriti

Intigriti at DEF CON 34

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

© Intigriti

REQUEST A DEMO

intigriti.com/demo

VISIT THE WEBSITE

intigriti.com

GET IN TOUCH

hello@intigriti.com

© Intigriti