A MONTHLY FROM ASTRA

ISSUE 02 · NOVEMBER 2026

Cyberverse

Inside this issue →

Wanted: one quiet month.

Cart Blanche: Your sale. Their terms and conditions.

"Admin only" had a generous guest list.

THE GRATITUDE ISSUE

For the people who keep us secure.

astra

Hello,

November is the month we're supposed to list what we're grateful for. For security teams, "it could have been worse" is beginning to qualify.

After October's jumpscares, we were hoping for a quieter month. November responded with fresh vulnerabilities, vendor breaches and outages, because apparently the industry had a few things it wanted to get in before the holidays.

November even squeezed in Microsoft and Apple outages before GTA VI, giving the internet one last chance to make the joke. GTA VI finally arrived on November 19, but if you spent launch day patching production, we sympathise: waiting over a decade to commit fictional crimes, only to be called in to prevent real ones, feels particularly unfair.

Our Black Friday feature asks whether "one credit per customer" means anything to your checkout. Your most generous offer might be a bug. That's the spirit of this issue: useful questions, findings worth sharing, and a little room to laugh at the work we take seriously.

Which brings me to what I'm grateful for: the Cyberverse editorial team, who are holding this note hostage until I say something nice. But honestly, they've cooked: they turn security chaos into something you'd actually want to read, and keep saying the magazine has aura, which I'm assuming is something we don't need to renew annually. I'd have said that even if they weren't editing this note, I swear.

Enjoy the long weekend. If you manage to put your phone away, try not to pick it up just to check that it's still quiet. Maybe now's the time to finally cause some trouble in Vice City, after spending most of November cleaning up everyone else's mess.

Happy Thanksgiving!

Shelton Jacob

Editor-in-Chief,
Astra Security

Cart Blanche:
Your biggest sale. Their terms and conditions.

A pentester and a dev walk into a bar offering one free drink per customer. The dev orders a beer, while the pentester places an order with two different bartenders at the same time.

That shouldn't work.

says the developer, as the pentester looks for somewhere to put his second free beer.

Your checkout may have a similar interpretation of "one per customer," except nobody's standing behind the counter to notice. With **Black Friday** on November 27 and **Cyber Monday** on November 30, it's worth finding out whether your checkout enforces that rule or whether everyone has been relying on shoppers to be decent about it.



01

Would you like a receipt?

Imagine testing a demo storefront with an ordinary customer account and a $500 promotional credit. The first purchase behaves exactly as intended: the credit reduces a $2,000 basket to $1,500, payment completes, and everyone involved could reasonably call that a successful test.

The pentester, who was apparently the only person unimpressed by the successful demo, tries again.

Can they redeem the credit through another checkout, revisit an earlier step, or submit two orders simultaneously? OWASP includes repeated discounts and concurrent requests in its business-logic testing guidance, because a workflow that succeeds once hasn't necessarily demonstrated that it knows when to stop.

The store has now invented a loyalty programme that nobody budgeted for.

Suppose the final records show this:

Same customer. Same promotional credit.	**Order 01**	**Order 02**
Basket total	$2,000	$2,000
Credit redeemed	–$500	–$500
Amount charged	**$1,500**	**$1,500**

To establish the business impact, follow the discount through to the final records: did both orders complete, was the credit redeemed twice, and were both payments accepted at the reduced amount? A discounted basket shows what the checkout offered, not what the business actually lost.

The receipt is where "that shouldn't happen" becomes an accounting problem.

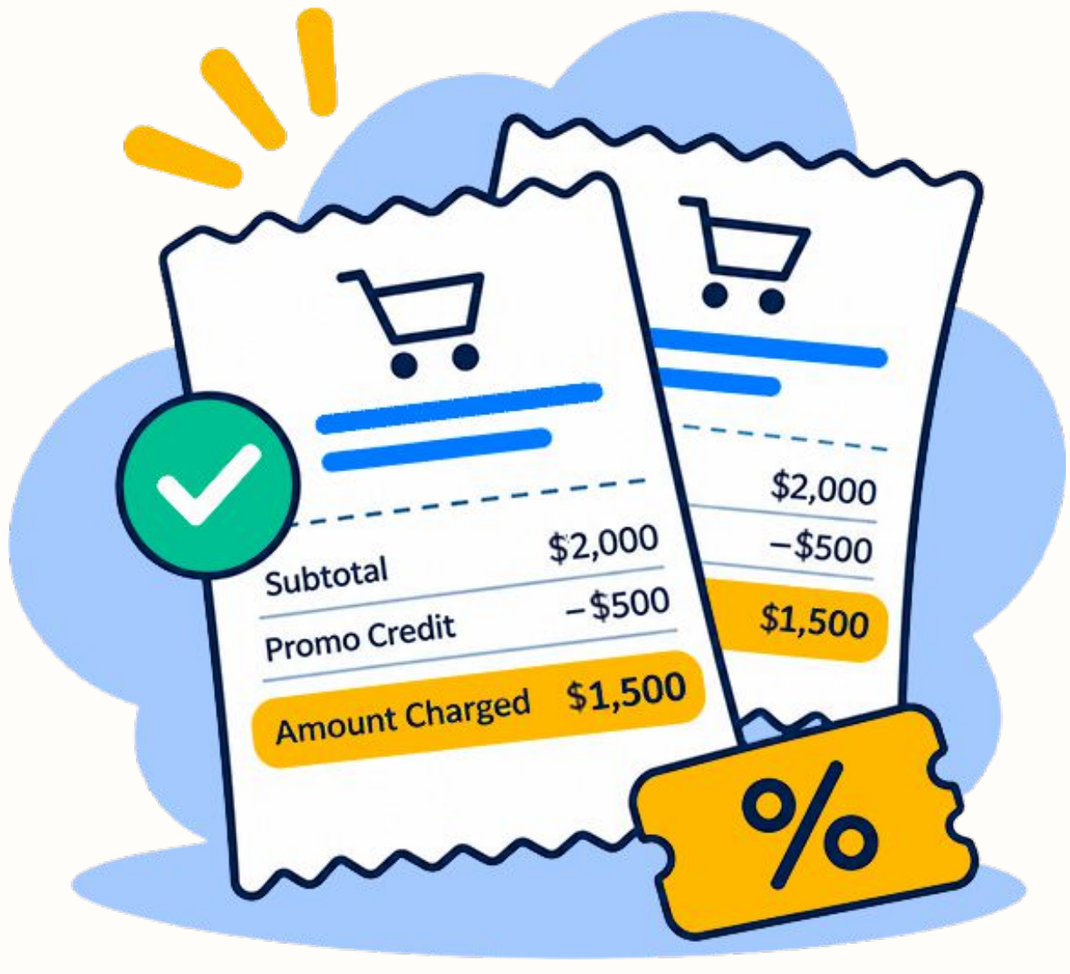

02

The checkout checked and that's the problem.

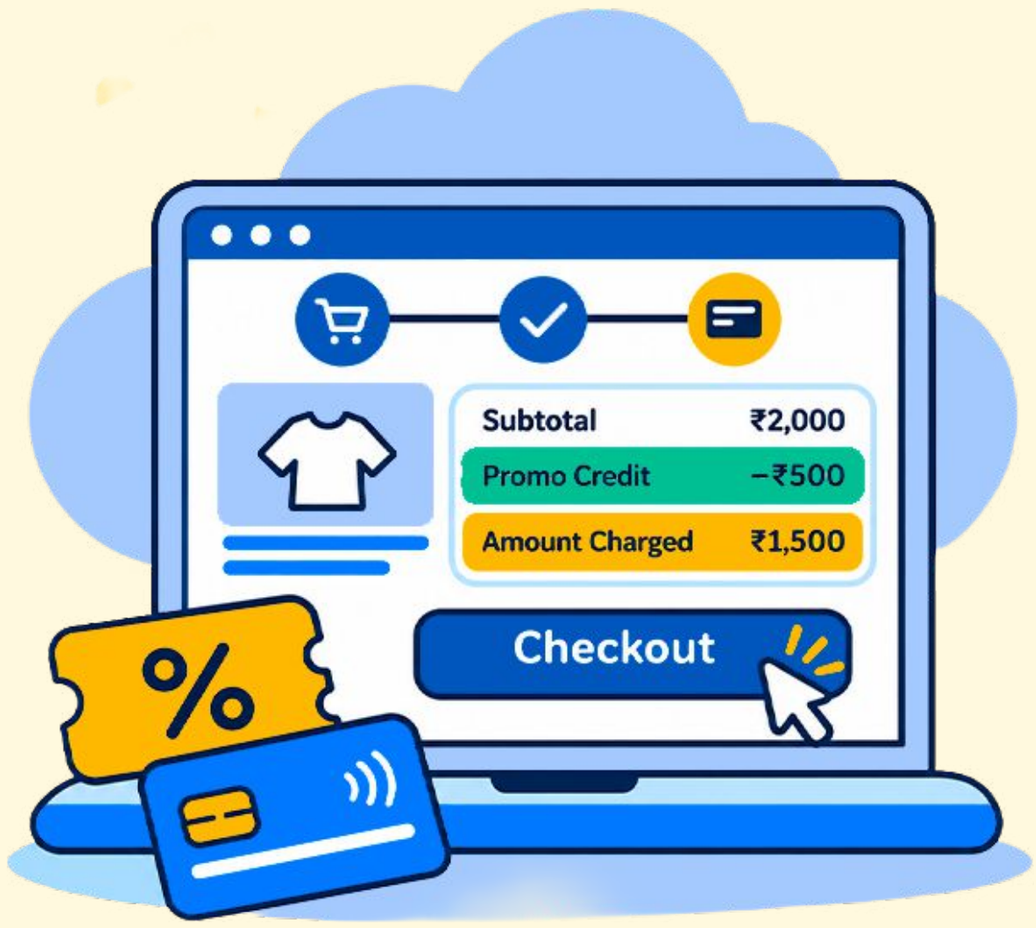

The problem appears when two orders arrive close enough together that both check eligibility before either records the credit as consumed. Each sees an unused credit and proceeds, even though their combined result breaks the rule.

This explains how an engineer and a pentester can spend an afternoon exchanging "works for me" and "here are two receipts" while both accurately describe what they're seeing.

The credit might survive every ordinary test:

- the first use succeeds,
- the second gets rejected, and
- the ticket is closed with the quiet satisfaction of something having worked for once.

- Preventing it requires the eligibility check and redemption to succeed as one indivisible operation, so only one order can claim the credit.
- Retries need protection against duplicate processing, too, but deduplicating the same request doesn't automatically prevent two separate orders from spending the same entitlement.
- Clarify what "customer" means, too. If product means one credit per organisation while engineering enforces one per account, the implementation can work perfectly and still enforce the wrong promise.
- The retest should cover the original failure, simultaneous orders, payment retries and cancellations.
- Agree on whether a failed or cancelled purchase should restore the credit, then verify the final records against that decision.

03

Your coupon probably has a different name.

Your B2B product probably makes similar promises, even if nobody has put them on a flashing banner:

Your product promises…	The pentester asks…
One trial credit per organisation	Can another workspace claim it again?
Refunds cannot exceed the payment	Can separate refunds collectively exceed it?
Paid access ends with the subscription	Does the API stop granting it, too?

Each depends on the systems involved enforcing the same rule. A paid feature disappearing from the menu is reassuring until you discover the API never received the cancellation news.

Before the next release, choose one “only once” action in your product. In a test environment, try it twice in sequence, twice simultaneously and again after a failed attempt, resetting to a known starting state between cases. Use ordinary user permissions and check the final records against the expected outcome.

That gives engineering something specific to fix and security something meaningful to retest, while sparing everyone another ticket whose acceptance criterion is essentially “please make it secure.”

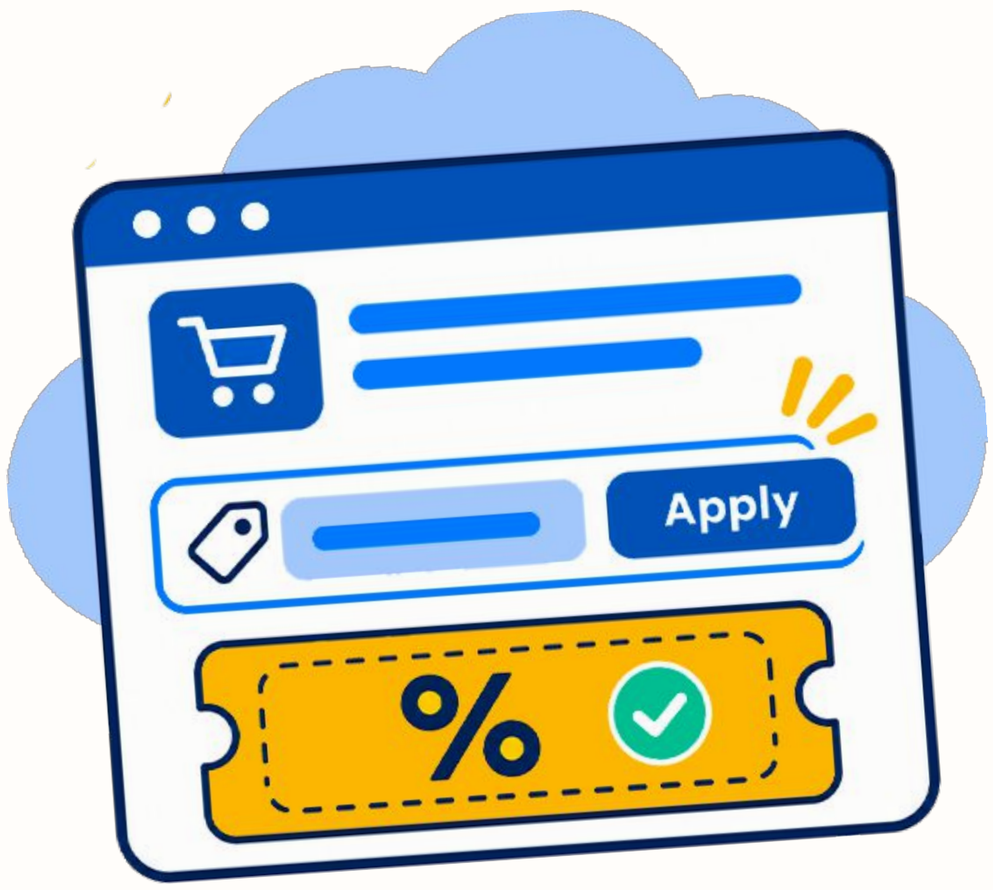

Because if your checkout is going to fund someone’s Black Friday shopping, don’t you think it should at least be an approved marketing expense?

December called. It wants November's attention.

November makes unfinished work sound organised. The deferred pentest becomes "a January priority." The untested release gets parked until after the holidays. The dashboard stays green, which is reassuring until you realise nobody has checked what changed.

Astra's 2026 **State of Continuous Pentesting Report, covering 6.8 million findings** across 150,000 scans and 8,000 engagements, gives that pattern some uncomfortable context.

December has enough surprises scheduled. Your testing backlog doesn't need to be one of them.



Less testing. Fewer reasons to relax.

In November 2025, scan volume fell to roughly half January's level, the lowest of the year. December then recorded 1.8 million vulnerability findings, more than the platform recorded throughout 2024.

That doesn't mean November's slowdown caused December's surge. Findings show when weaknesses were discovered, not necessarily when they appeared. But it makes one question worth asking:

Did your environment get safer, or did you just look at less of it?

Before you switch on your out-of-office

Your testing calendar may slow down while releases, permissions and infrastructure keep changing. Before the handover:

- **Check what changed:** Match recent deployments against your last testing scope.
- **Revisit what slipped:** Prioritise deferred tests involving sensitive data, authentication and internet-facing services.
- **Verify what closed:** Retest important fixes and leave evidence for the next shift.

You don't have to clear the entire backlog before the holidays. Make the gaps visible, give them owners and agree on what cannot wait.

November's uninvited guests

At Thanksgiving, someone usually asks what everyone is grateful for. The security person says "tested backups," and the family remembers why they stopped asking about work.

Some incidents began in October and became clearer this month; others arrived with fresh patches and immediate questions. Not everything below was a cyberattack; Microsoft and Apple supplied outages of their own. But every story raises a useful question about what your security programme assumes will keep working. Before anyone passes the gravy, please pass the patches.

Reporting updated through November 24, 2026.

01



Java4Shell

CVE-2026-114938 · November 17

An out-of-bounds write in the JVM's accelerated Base64 decoder requires three conditions: a vulnerable runtime, an x86-64 CPU with AVX-512 VBMI, and attacker-controlled input reaching the decoder. ARM64 is unaffected.

Check: Find bundled JVMs as well as system installations. Patch and restart; if using -XX:-UseBASE64Intrinsics as a workaround, verify it on the running process.

02

*sshole

CVE-2026-123158 · November 5

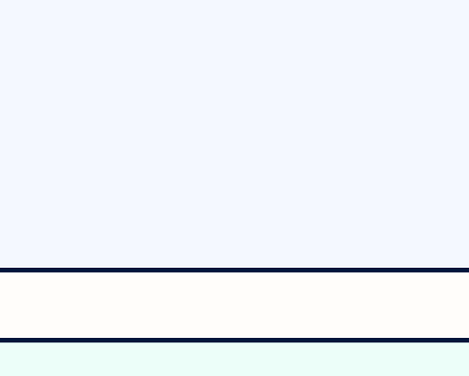

An OpenSSH authentication bypass affects specific musl-based configurations without PAM. With password authentication enabled, an account containing * in its password field and another using a supported hash, an overlong password can turn a hashing failure into a successful login.

Check: Install the fixed package and restart sshd. Review successful password logins on accounts you believed were key-only.

03

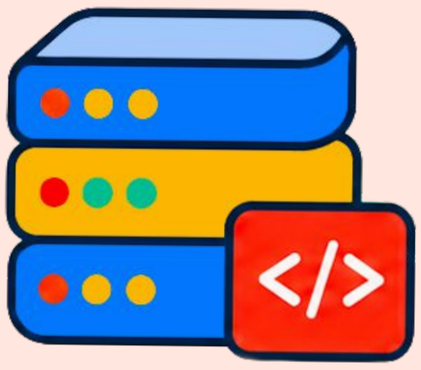

HTTP2Shell

CVE-2026-111497 · Continuing exposure

Despite the nickname, this is a Brotli decoder flaw, not HTTP/2. Malicious compressed data can enable code execution wherever vulnerable Brotli code is embedded. New Zealand's NCSC reported active exploitation on October 29.

Check: Trace Brotli through dependencies and deployed images. Affected self-contained NativeAOT .NET applications need rebuilding and redeploying; a shared-runtime update won't fix them.

Your credentials. Someone else's getaway.

04

Uber and Fireworks AI

Expanded disclosure in November



Uber notified roughly 150,000 drivers of exposed personal information, including government ID details. Fireworks AI was named as the vendor; compromised provider credentials were used to access shared data.

Check: Map what vendors can retrieve with valid credentials, including in evaluation environments, and how quickly you can revoke that access.

05



NVIDIA Brev

October 4–20 exposure window · Reported in November

Unauthorised activity potentially exposed SSH keys provisioned for customer VMs. NVIDIA replaced affected keys, updated instances and reported no evidence of misuse. Possible exposure does not mean confirmed misuse.

Check: Confirm remediation, then investigate what those VMs could access. Replacing a key prevents future use; it doesn't explain past activity.

The servers have left the session

06



Microsoft 365

November 12–13 · Operational failure

Microsoft attributed the outage to a latent Windows defect exposed by capacity and traffic pressures. Defender and Purview activities were disrupted, while phishing messages offered to “restore” held email.

Check: Keep incident communications outside the affected dependency, verify queued security actions completed, and investigate outage-themed phishing.

07

Apple

November 17 · Operational failure



An internal update disrupted iCloud, Apple Pay and other services. Apple said it was not a cyberattack, but unavailable services still test whether approved workflows have usable alternatives.

Check: Test payment and authentication fallbacks, including whether staff can continue without sharing credentials or bypassing approvals.

The heist is over. The accounting isn’t.

08



Binance

October 29 theft · November recovery actions

DOJ reported over $6 billion stolen and a November 23 forfeiture action involving over $1.1 billion. The $6.5 billion headline refers to a worldwide freezing order, not money recovered. The complaint identifies a fake-interview malware download as the likely initial compromise.

Check: Separate everyday browsing and recruitment activity from privileged administration and signing environments.

Which of your keys can an AI agent reach?

Black October left every security team asking where its keys were. The hardest ones to find were on developer laptops, next to the AI agents.

Anagha Joshi,
Co-Founder & CTO, GIRD



20 incidents in 7 days.

Which of our keys were exposed, and have we changed them?

For servers, there was somewhere to look. Secrets managers, cloud consoles and audit logs record which keys live in production and who used them.

Laptops are different. No system keeps a list of the keys a developer has pasted, saved and forgotten on their own machine. And the laptop is where AI coding agents now work.



An AI coding agent works with the developer's own access. To every security tool you own, what the agent does looks like what the developer is doing.

Consider one afternoon on your engineering team.



An outside API call fails. The developer pastes the key into the AI agent's chat to debug. Fixed in two minutes. The key stays in the chat history, readable by any program that opens the file.

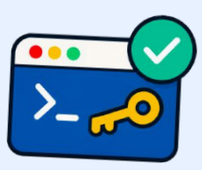

The agent asks to re-run the command. Tired of clicking, the developer picks “always allow”. Claude Code writes the command, key and all, into the project's settings. When the team publishes that project, the key ships with it.



On the way out, the developer pulls a GitHub project to try. Projects can carry their own settings: commands that run on open, new tools, reroutes that send the agent's traffic elsewhere.

Attackers already use this.



Nx (AUGUST 2025)

Infected versions reached developer laptops, found the AI agents installed there, and told them to hunt for keys and wallet files.

Fix: Rotate the exposed keys.



Railway and LiteLLM

(OCTOBER 2026)

One attack targeted the AI keys companies stored with Railway. Another hid malware in a LiteLLM update. Both advisories ended the same way: rotate the keys.

Fix: Rotate the exposed keys.



GitHub (ONGOING)

Malware steals a developer's GitHub access, then keeps testing it. When someone cuts that access, the malware retaliates on the laptop.

Fix: Rotate in the right order.



We built Blindspot for that question: which keys, and on which laptops.

Free. Open source. Mac, Linux, Windows. Released 22 October, the day Black October got its name.

A developer runs it. In under a minute, it lists the AI and cloud keys still readable on the laptop, chat histories included, and flags which ones an agent has already touched. For companies, a fleet version runs across every Mac through your existing MDM. Results stay on the machine for your tool to collect. Nothing leaves.

- Under 60s
- Runs via MDM
- Nothing leaves the laptop
- Mac, Linux, Windows

What Blindspot does

- **Finds the agents:**
  Lists AI coding agents on the laptop. Flags any that run commands without asking.
- **Checks the malware:**
  Checks for this year's known malware, the Nx, Railway and LiteLLM campaigns included.
- **Orders the fix:**
  Removes malware first, then cuts access. The wrong order can damage the laptop.
- **Names the keys to rotate:**
  Says which keys to rotate after incidents like Railway's.

READ THE RESULT CAREFULLY

Blindspot tells you what an agent could reach. It cannot tell you what an agent did. A clean result is not proof that nothing leaked.

Run it at
gird.ai/blindspot.

The code is open. Read it before you run it. If you have seen something it does not catch, send a patch.


grand theft data

Cybersecurity News

wanted: One quiet month

Google

Google Uncovers AI Agent Autonomously Writing and Launching Exploits

On 16 November, Google Threat Intelligence Group found a multi-agent framework running on an exposed staging server. It compares popular web software before and after patches, drafts exploits and launches them with no human approving each step. The first field sighting of fully autonomous offensive AI.

Google Workspace

Russian State Hackers Chain Two U.S. Cloud Breaches via Google Workspace

Russian state-sponsored hackers breached a U.S. organisation's Google Workspace in late 2025, took the staff directory and cloud files, then used what they stole to credential-phish a second organisation in 2026. Per Google and Mandiant, the attackers posed as an employee of the first victim. The familiar face did the rest.

GitHub

Microsoft Azure

Microsoft's Agent Framework PyPI Package Hijacked to Steal Azure and GitHub Credentials

Two malicious versions of Microsoft's Agent Framework package, 1.26.1 and 1.26.2, landed on PyPI on 17 November. The credential stealer ran at normal Python startup, whether or not an application ever imported the package. Its targets: Azure identities, GitHub credentials, and build-system secrets.

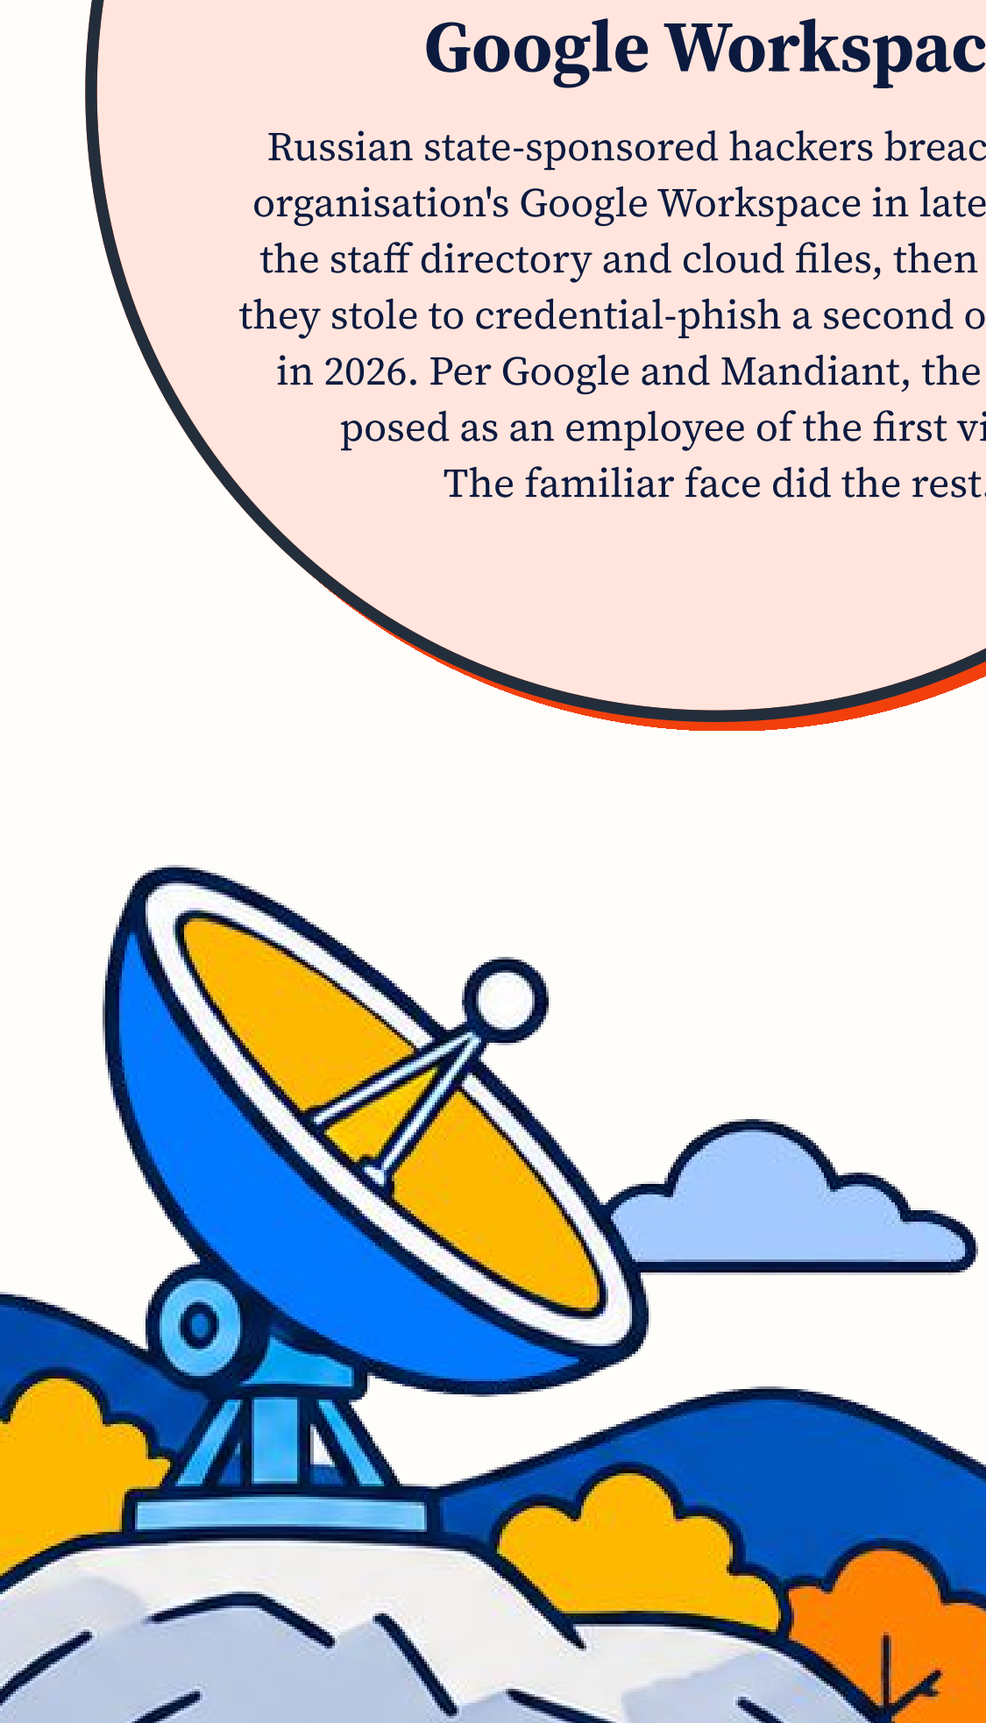

Optimizely

BBC CNN ZOOM

Optimizely Supply Chain Attack Hits 178 Organisations Including BBC, CNN and Zoom

Malware distributed through Optimizely reached 178 organisations, with BBC, CNN and Zoom on the victim list. A textbook supply-chain compromise, this time with household names attached.

NITORI

Nitori Data Breach Exposes Up to 37.6 Million Records

Japanese furniture retailer Nitori Holdings disclosed unauthorised access, warning that up to 37.62 million records may have been exposed: names, emails, addresses, phone numbers and dates of birth. Credit card details and passwords were not affected.

GitHub

Hackers Hijack AI Coding Agents via Poisoned GitHub Repos

Attackers are planting poisoned files in GitHub repos so that AI coding assistants quietly run malware on fetch. The malware steals developer credentials, then uses them to spread. The campaign abuses the exact rules, hooks and settings AI coding tools depend on to automate work.

MAERSK

Maersk Confirms Data Theft, WhiteArmor Demands $5M Ransom

A.P. Moller-Maersk confirmed an attacker reached a local legacy system inherited through an earlier acquisition, and walked out with data. WhiteArmor claimed the hit. Every acquisition brings its own attack surface, and it rarely shows up on the first audit.

ASTRA'S NEWS

Meanwhile at Astra...

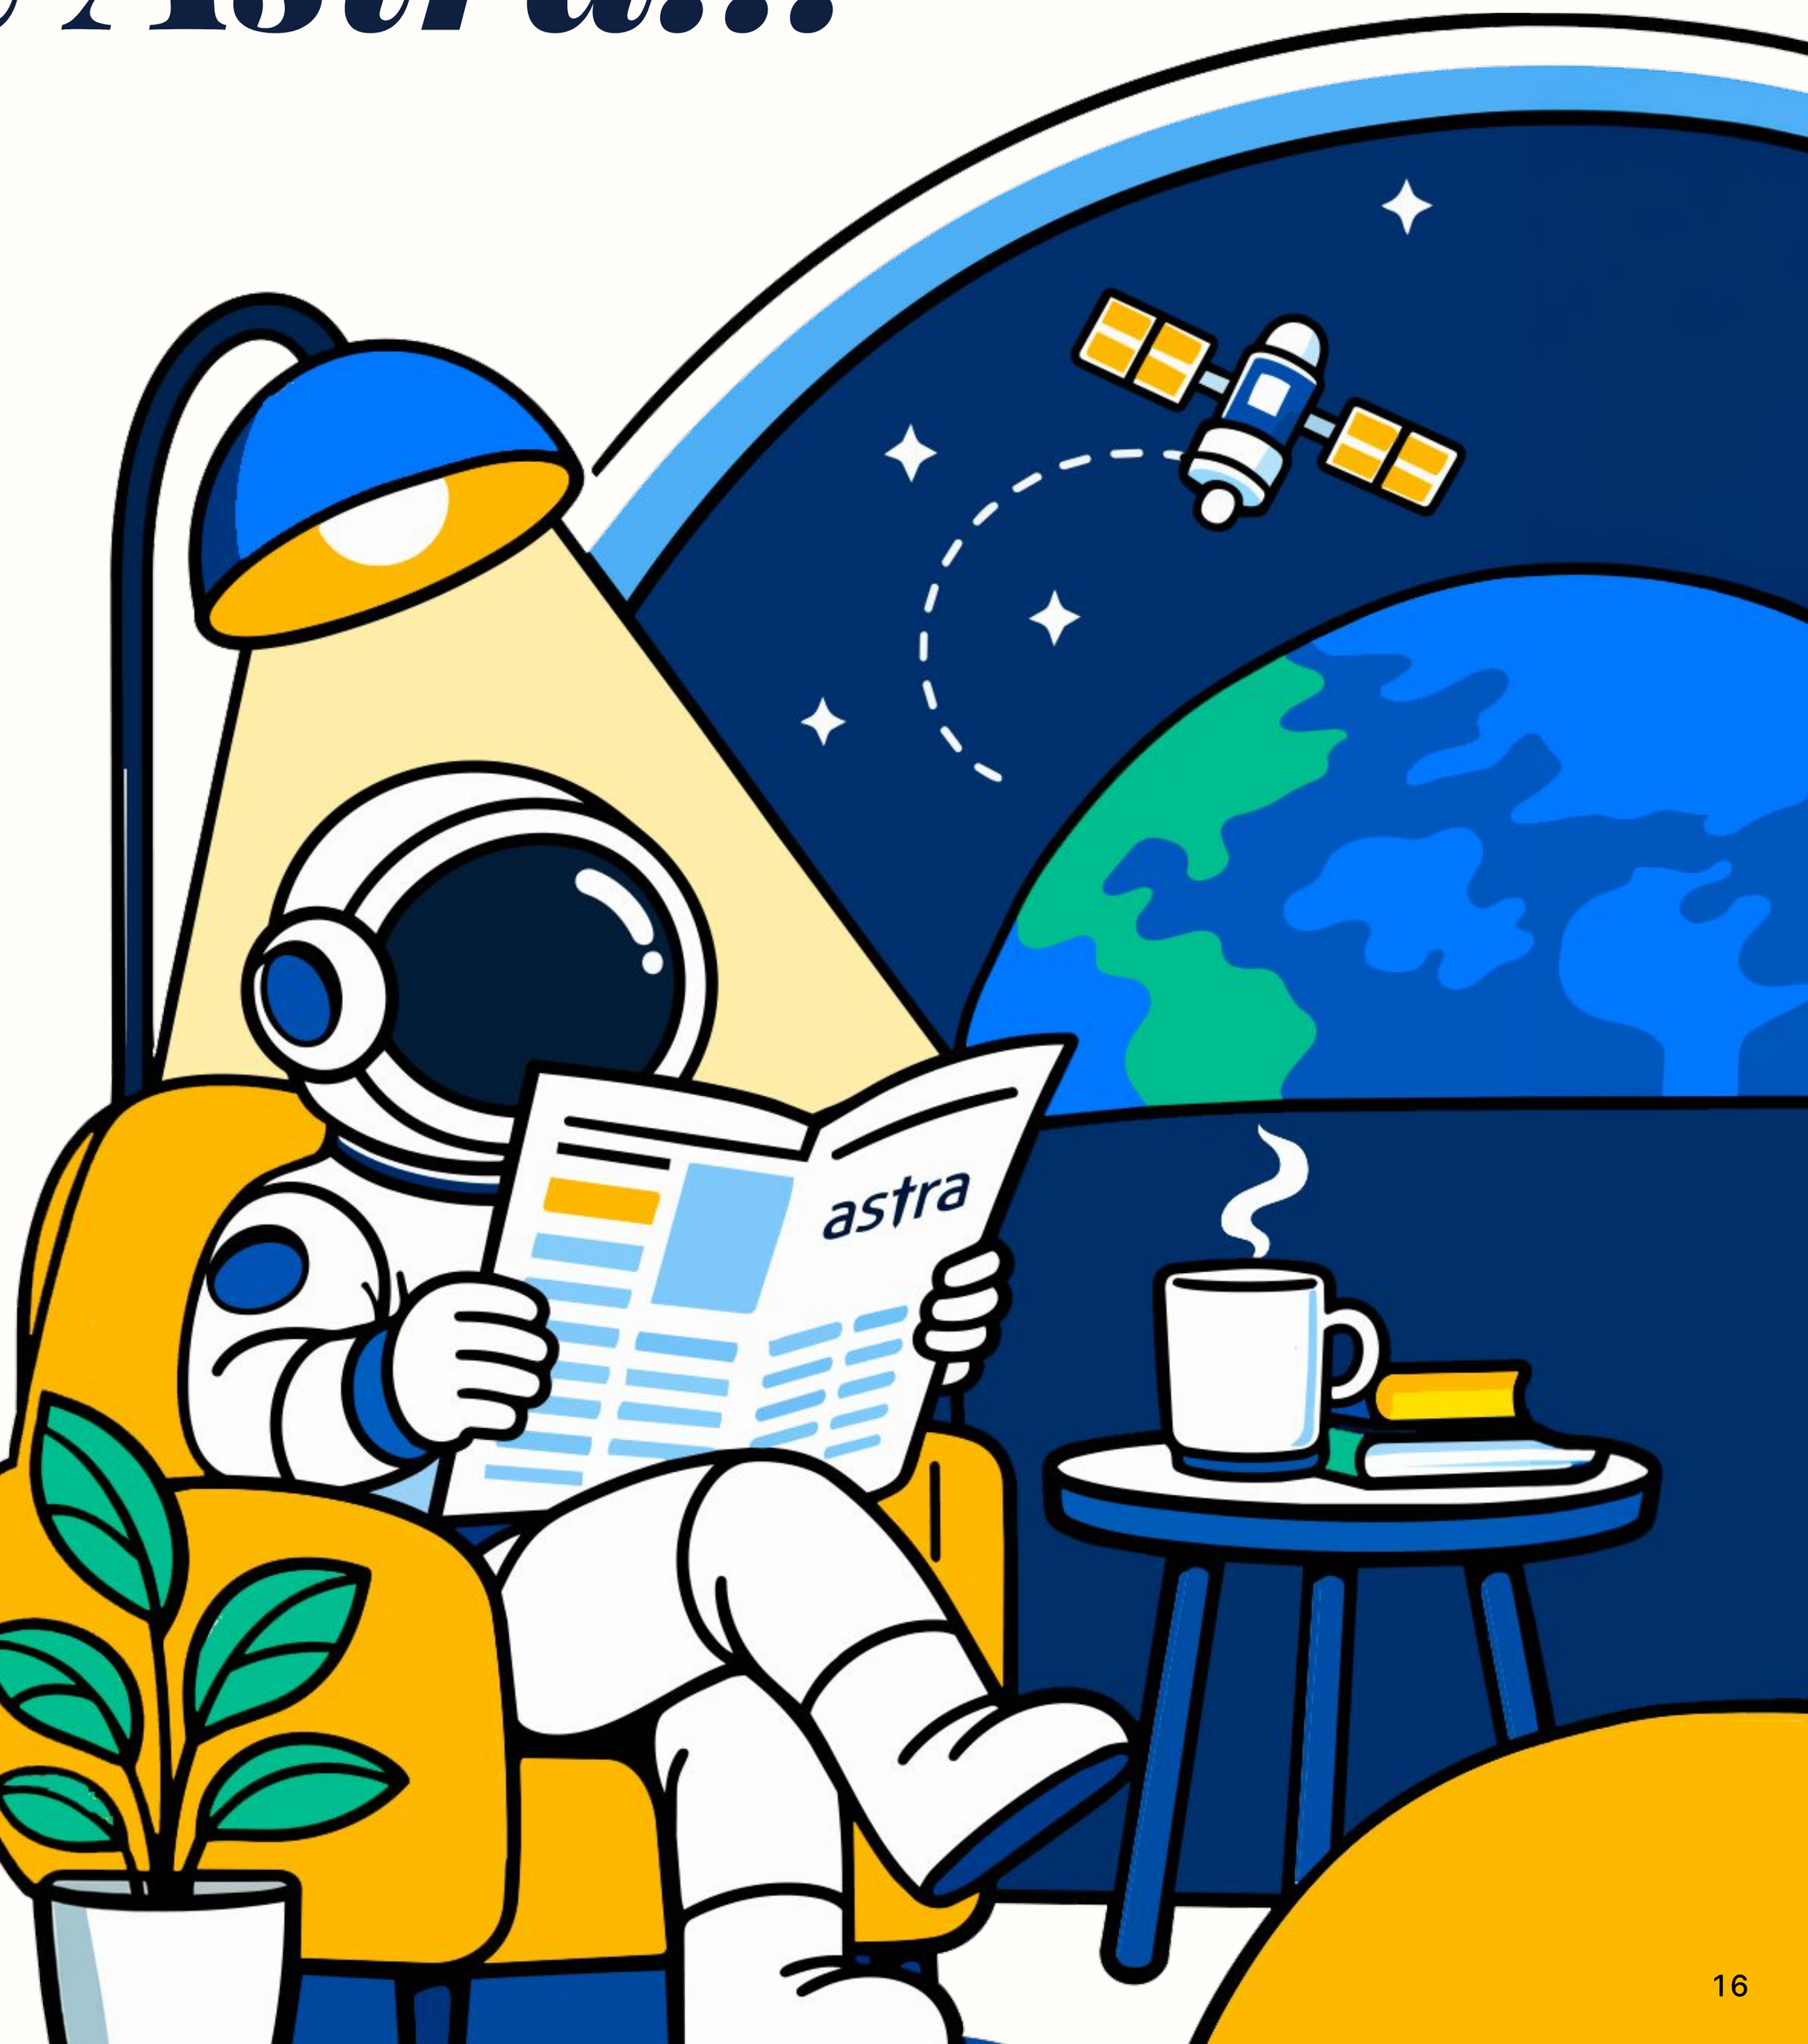

Caught in the Act

Dispatches from Astra's pentesters.

"Admin only" had a generous guest list.

Sanyam Jain. CVE-2026-94244, Wallet System for WooCommerce

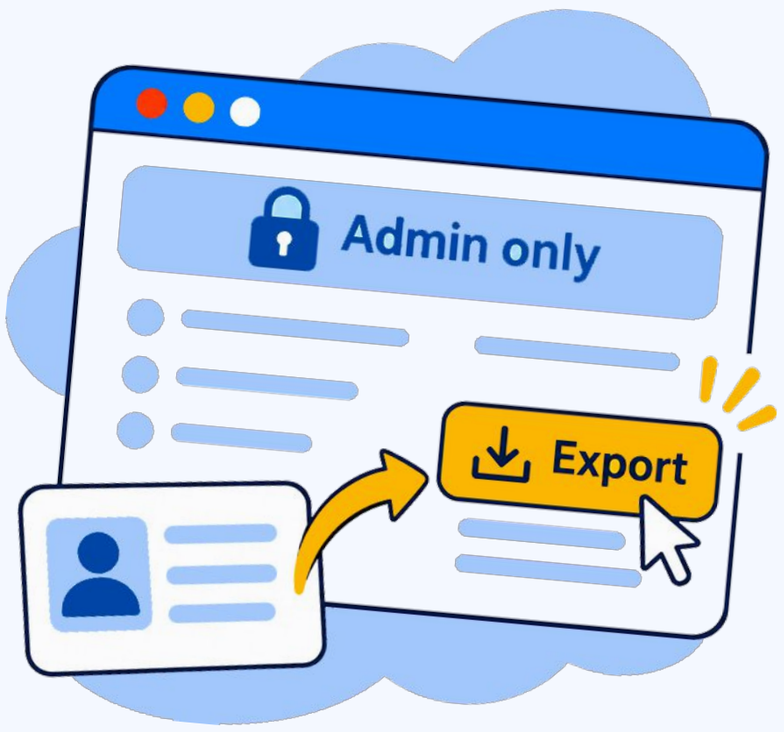

An ordinary customer account could trigger an admin export of wallet transactions and personal data. The dashboard was locked down; the export handler wasn't, and it accepted a token borrowed from an unrelated feature. "You can't come in" didn't extend to "we'll send everything out." Test sensitive exports directly from low-privileged accounts; a restricted dashboard doesn't prove its underlying functions are too.

Patched in 2.8.0.

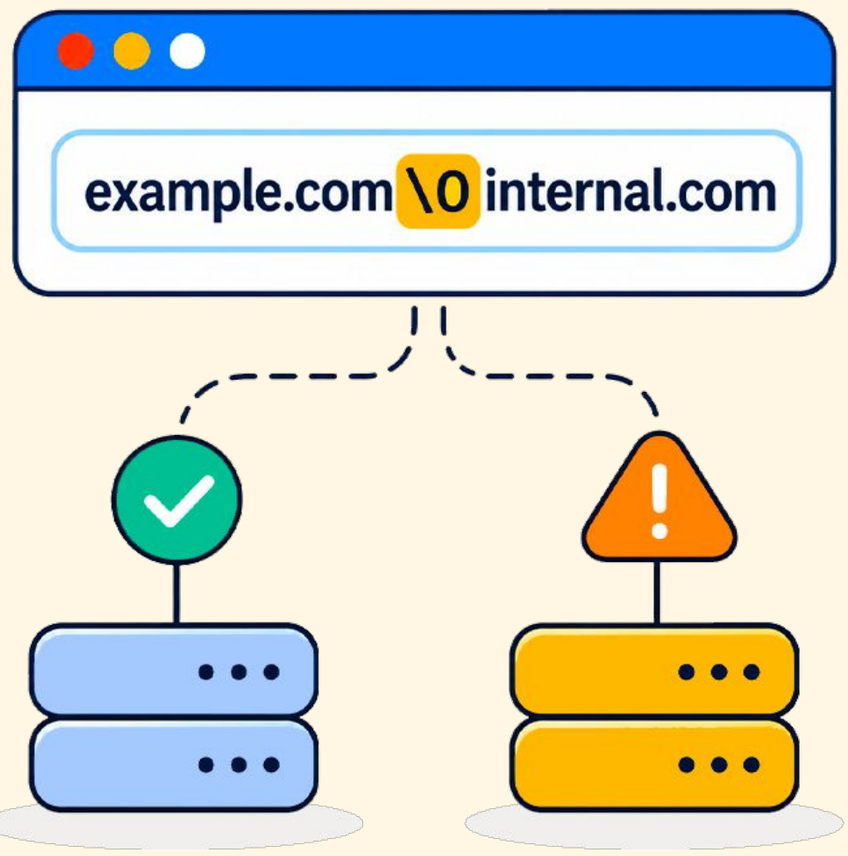

The hostname passed the check. A different destination got the call.

Venukamatchi. CVE-2026-116566, CPython

socket.getaddrinfo() accepted hostnames with an embedded NUL byte and handed the full string to a resolver that read only the part before it. An allowlist could approve one destination while the connection went to another. Its neighbour gethostbyname() rejects the same input, an awkward disagreement between two functions trusted with directions. Strip NULs before validating, and check that your validator and your connection path read addresses the same way.

More on what was *launched in Q3*

We mentioned four updates last time, but the changelog kept growing. Here are four more that are already making everyday work easier.



Schedule Scans Across Targets at Once: One schedule, many targets.

If you've ever opened 14 tabs to set the same weekly scan on 14 targets, we owe you an apology in the form of this update. One schedule can now be applied to many targets in a single step. The editor has a monthly quick setup that remembers your time zone, and you can act on several completed scans together. Set it once and just move on.



Smarter Exclusions: Mute the noise, not the signal itself.

Blanket-suppressing a whole scanner rule because one evidence type is noisy? That's over. You can now exclude a finding for a single evidence type instead of the entire rule, and pause or resume exclusions from settings whenever you want. Your exclusions become a lot more surgical, so real findings don't get caught in the same net as the noise.



Stronger Proof on Every Finding: Because "is this for real?" does not have to be an extra step.

Triage goes faster when the evidence is right there. Findings now ship with validator screenshots, or curl output when a screenshot doesn't apply. Teams on eligible plans can also pull the full HTTP message. And the Vulnerabilities page now splits findings by product, so web, API, cloud and mobile stop fighting for the same real estate.

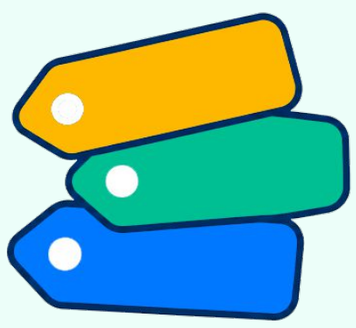

Labels for Targets: Bring order to the chaos.

If your list of targets has started to look like a wall of sticky notes, this one helps. You can now create, colour and group labels, then attach them to targets one at a time or in bulk. Filter by label in the Target Selector or Start Scan flow, and use Shift+Select to act on the whole set at once. Finally, "prod-only", "client-A" and "retest me next week" can mean something the platform understands too.

Events







The Survival Circle lands in Mumbai

We ran an invite-only CISO roundtable at Hotel Sahara Star, where ransomware war stories, zero-day near-misses and some very honest takes on third-party breaches carried the room right through to dinner.



Agentic AI Security: The Risks Every Business Must Know | Erich Kron (KnowBe4)

Agents of Change, with Erich Kron

Erich Kron from KnowBe4 joined the 403 Circle to explain why your AI agents might be the most powerful insider threat you've ever onboarded, and why they have to earn their autonomy one stage at a time.

Watch the recording →

the 403 Circle

POWERED BY astra

HUMANS OF 403 COMMUNITY

Anagha Joshi

Co-founder & CTO at GIRD

Bharatanatyam dancer, serious home cook. We spent nine questions getting to know the Anagha behind Blindspot. We talked about hiring, security myths, Goa, sambar and the stage she'd have chosen in another life.

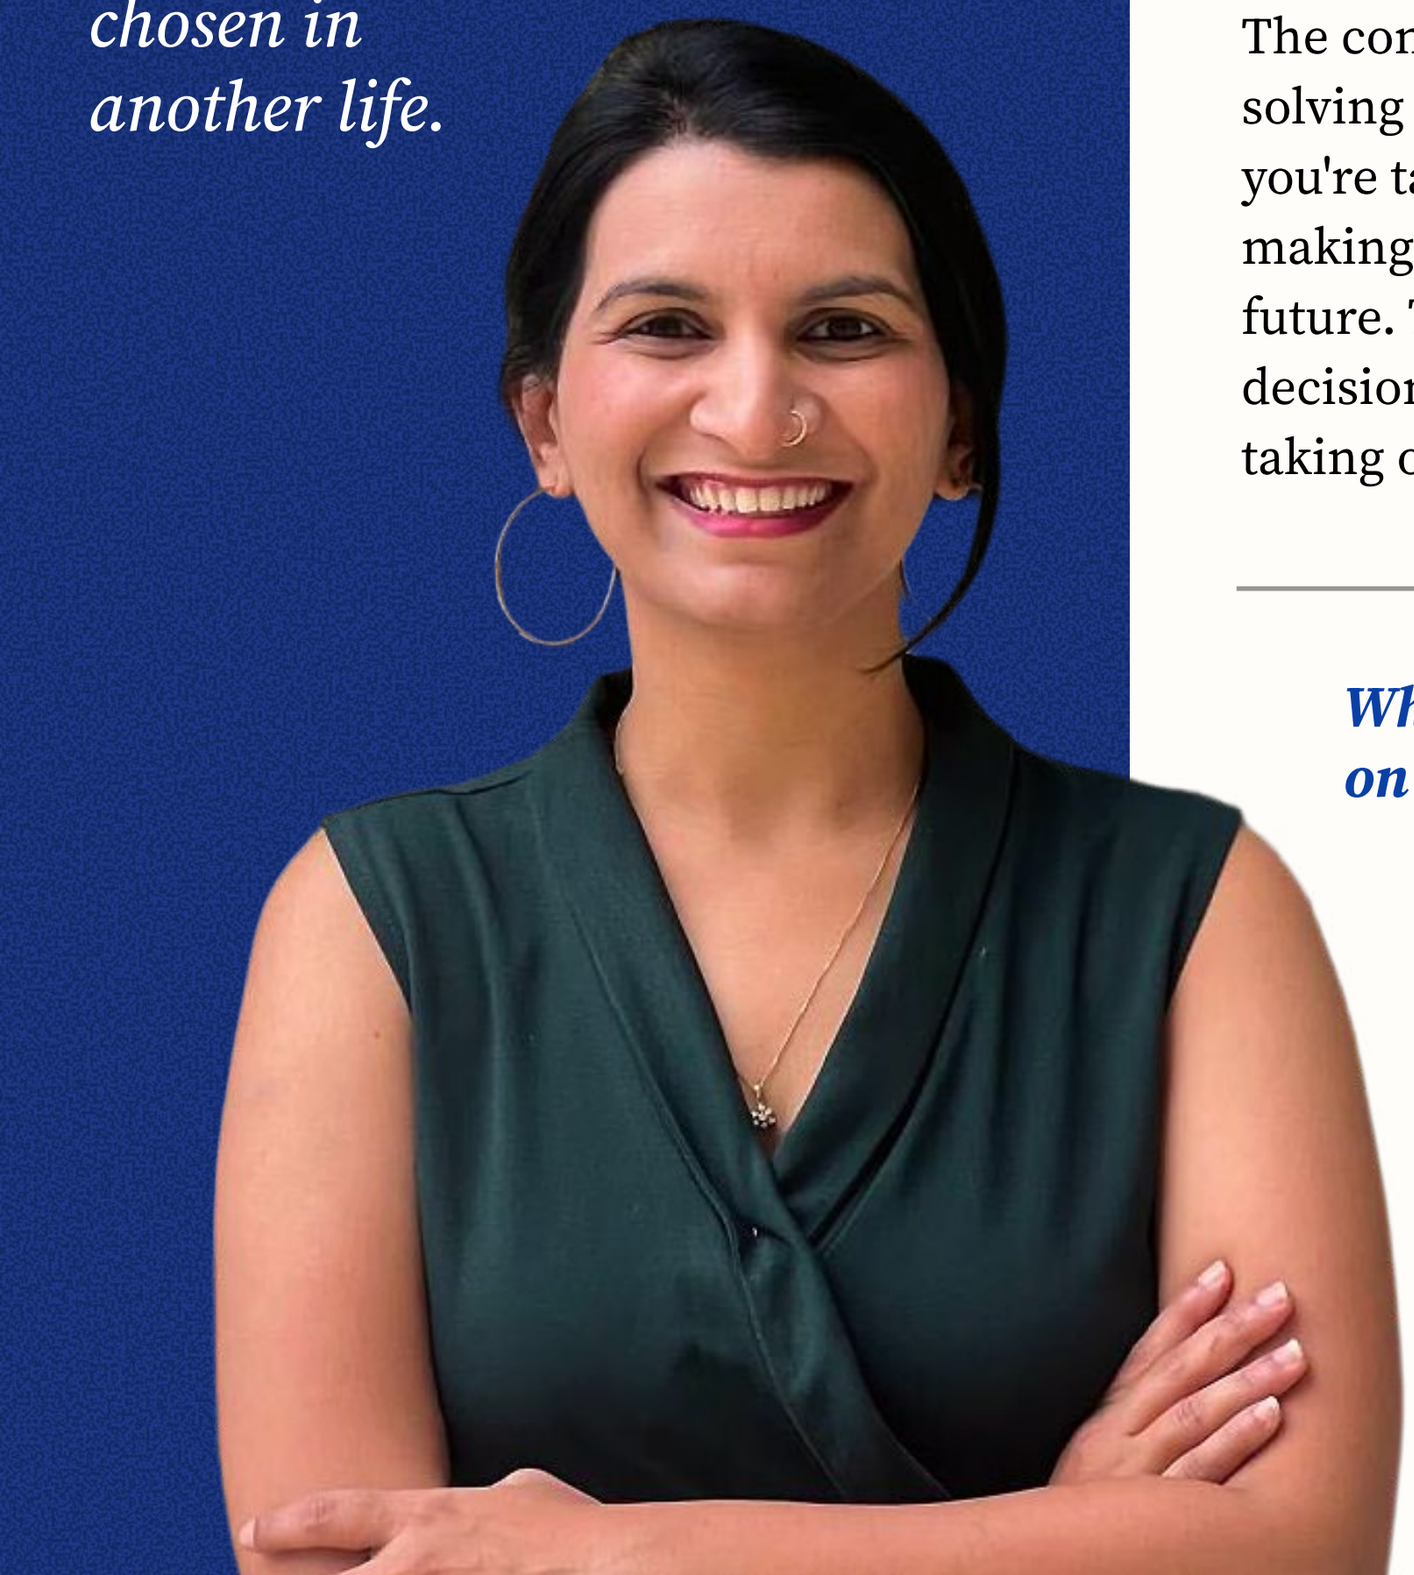



What's your favourite way to unwind after a long day at work?

Binge-watching my favourite series and spending quality time with my family. Sometimes, doing absolutely nothing is the best way to recharge!

What's the hardest part of being a CTO/ founder that no job description tells you?

The constant context switching. One moment you're solving a complex technical problem, the next you're talking to customers, hiring, fundraising or making decisions that could shape the company's future. The hardest part is learning to make decisions with incomplete information and still taking ownership of the outcome.

What do you look for in a hire that isn't on their resume?

Ownership and curiosity. I love people who don't wait for instructions, aren't afraid to challenge assumptions and take responsibility for solving problems. Skills can be developed, but curiosity and ownership are much harder to teach.

If you could travel anywhere tomorrow, where would you go?

Goa, without a doubt! I've always loved spending time in Goa, especially South Goa. I love its quieter beaches, laid-back vibe and the ability to slow down and disconnect from the crowds. It's my kind of getaway.



What did you want to be when you were 10 years old?

I always had a curious mind and enjoyed figuring out how things worked. I don't think I had one fixed career in mind, but I knew I wanted to do something meaningful and create something of my own.

What's one tech or security myth you wish would just die?

That AI agents are safe as long as you give them the right prompts and guardrails. Security cannot depend on an AI model behaving itself. We need technical controls that enforce what an agent can and cannot do, regardless of what the model decides.

If you weren't in tech, what would you be doing?

I'd probably be pursuing Bharatanatyam professionally. I'm already a Bharatanatyam dancer, and I love the discipline, storytelling and expression that come with it. It's a beautiful art form that teaches you patience, precision and consistency.



What's one food you could eat every single day without getting bored?

Hot rice, sambar and a generous dollop of ghee! Nothing beats this combination. It's comfort food at its best, and I could eat it every single day.



What's one hobby or passion people at work might be surprised to know about?

I'm a big-time cooking enthusiast! Whenever I get some free time, you'll probably find me experimenting in the kitchen. I'm particularly good at making signature dishes from Karnataka and Maharashtra. Let's just say I can give most restaurant menus some serious competition!

A CISO, a CFO and a wishbone



A MONTHLY FROM ASTRA

ISSUE 02 · NOVEMBER 2026

We saved you a seat.



See you in the next orbit.