

# Cyberverse

## Inside this issue →

Your AI agent missed the security training.

School's open. Files are locked.

Inside the BeyondTrust zero-day.

## BLACK OCTOBER

Real threats. No tricks.

astra

October was supposed to be about cybersecurity awareness. Then Black October arrived and gave security teams considerably more awareness than they'd asked for.

Somewhere, a CISO is watching a horror movie to relax. After Black October, the ghosts on screen probably feel manageable and less scary.

At Astra, we took the testing theme personally and looked under the costume of our own security claims.

We put Astra Autonomous Pentest up against the same app XBOW benchmarked, using the same methodology and an independent judge. The results? Not all treats for either of us. But a benchmark that only shows your good side isn't a benchmark, it's a dating profile.

On a slightly less terrifying note, we put Astra on a billboard! A slightly bigger announcement than our usual LinkedIn post.

Between benchmarking, shipping new features, and roundtables in Chennai and Singapore, it's been quite a month. One thought ties it together:

**Security claims are best tested, not blindly trusted. Even when they're our own.**

Until next month, may the only skeletons you find be Halloween decorations, not vulnerabilities in your code.

**Happy Halloween!**

*Shelton Jacob*

Editor-in-Chief,  
Astra Security



# BLACK OCTOBER

## EVERYTHING YOU MISSED WHILE PATCHING EVERYTHING ELSE

The naming came late and crude. On October 22, a widely shared post called the previous 48 hours the worst two days in cybersecurity history and gave the moment a name: Black October. It stuck because it was accurate.

But the frame is slightly off. Black October sounds like one catastrophe. What happened was roughly twenty incidents in one week, almost all pointed at the same part of the stack, the dependencies, platforms, package registries, and dev tools engineering teams install once and forget. Almost none of the ones that mattered touched anyone’s own code. **Three patterns connected them.**

Black October made one thing clear:

**your code doesn’t have to change for your risk to change.**

A routine update, a trusted package or a stolen platform key can do that for you.

If you missed a few headlines while dealing with the others, understandable. Here’s the catch-up nobody wanted.

1

### Exploitation outran

**patching:** Amazon detected PHP2Shell attacks four days before a fix arrived.

2

### Trusted distribution channels became attack

**routes:** the malicious OVHcloud CLI needed just 94 minutes of availability.

3

**And stored secrets became exposed credentials:** Railway’s encryption could not protect every secret once attackers reached the means to decrypt them.



# BLACK OCTOBER

10  
OCT

 **shopify**  
Checkout hijacks redirect shoppers through lookalike payment pages.

12  
OCT

 **GitLab**  
Discloses critical RCE. Refrigeration systems reportedly manipulated at ~700 French supermarkets.

13  
OCT

  **Microsoft 365**  
Stolen ConnectWise credentials expose 162 customers' Microsoft 365 environments.

14  
OCT

 **zeabur**  
Breach affects 321,810 accounts and 3.2 million environment- variable records.

15  
OCT

 **RankMath**  **eMudhra**  
Weak eMudhra intermediate CA keys reported. Backdoored Rank Math PRO update reaches official distribution.

16  
OCT

 **RankMath**  **opencode**  
Rank Math backdoor window closes after 11 hours. OpenCode Console API keys exposed.

17  
OCT

 **amazon**   
 **eMudhra**  **UpdraftPlus**   
Amazon honeypots detect PHP2Shell exploitation before disclosure. Plex exploitation reportedly begins. Browsers block eMudhra certificates. Compromised UpdraftPlus/ TeamUpdraft key distributes malicious plugins.

19  
OCT

 **OVHcloud**  
Malicious CLI release available on GitHub for 94 minutes.

20  
OCT

  **Railway**   
PHP warns of two critical PHP2Shell RCE flaws. Railway attack targets AI provider keys. GMO Internet confirms Onamae Desktop Cloud malware.

21  
OCT

 **Railway**     
  **LiteLLM** **NEXT.js**  
PHP ships emergency patches. Node.js fixes 10 CVEs including high-severity Brotli RCE. Railway urges secret rotation after potential decrypted secrets in up to ~1.85% of workspaces. LiteLLM supply-chain attack. GitLab and Next.js issue critical fixes. SBI reports exposed records.

22  
OCT

**Linux**  **Railway**   
Railway invalidates all API tokens. Six Linux kernel zero-days accidentally disclosed. F5 NGINX advisory warns of potential RCE and denial of service.

23-24  
OCT

 **Railway**  **Bun**  **python**  **envoy**  
Brotli concerns extend to Python, Bun and Envoy ("HTTP2Shell" comparisons). Another PHP RCE, CVE-2026-115560, emerges. Railway shares further incident details.

OCTOBER 2026

# TOP CYBERSECURITY NEWS

## 01 RANSOMWARE · EDUCATION SECTOR Ransomware is up 48%, and education is the new healthcare.

Check Point put global attacks at **2,803 per organization per week** in September, up 48% year over year. Bitdefender logged **790 ransomware victims** the same month. On October 23, Osaka Metropolitan University confirmed an attack affecting ~500 servers, disrupting classes, email, and admin systems.



ZERO DAY · CVE-2026-102764

## 02 BeyondTrust zero-day exploited since August.

Mandiant linked exploitation of CVE-2026-102764, a critical flaw in BeyondTrust’s Remote Support appliances, to a China-aligned group it tracks as ARMORY CASTLE. BeyondTrust notified Okta on October 2. The patch exists. Not everyone has deployed it.

NODE.JS RCE · CVE-2026-111497

## 04 CISA confirms Node2Shell is being exploited.

On October 28, CISA added CVE-2026-111497, the Node.js Brotli RCE from Black October’s October 21 patch wave, to its active-exploitation list. If your Node.js rollout is still in progress, the window has officially closed.

ZERO-DAY · CVE-2026-104935

## 03 Ivanti Connect Secure zero-day pivoted to Active Directory.

Volatility tied CVE-2026-104935 to a group that moved from the Ivanti VPN into AD and walked out with domain credentials, local administrator passwords, and BitLocker recovery keys. A VPN compromise is almost never only a VPN compromise.

RANSOMWARE · EDUCATION SECTOR


## 05 Ransomware is up 48%, and education is the new healthcare.

Check Point put global attacks at 2,803 per organization per week in September, up 48% year over year. Bitdefender logged 790 ransomware victims the same month. On October 23, Osaka Metropolitan University confirmed an attack affecting ~500 servers, disrupting classes, email, and admin systems.

The pattern this month was not about discovery. It was about who the attacker talked to.

CYBERSECURITY AWARENESS MONTH

# YOUR AI AGENT HASN'T DONE THE SECURITY TRAINING



October 2026: you’ve taught employees to question suspicious instructions. What about the software authorised to act on them?



Think before you click.



Verify the sender.



Report anything suspicious.

Good advice.  
**But who does the thinking when you delegate the clicking?**

An AI agent that reads support tickets, searches internal records and sends replies needs more than useful instructions. It needs boundaries around what it can access and do.



**Your newest colleague won’t attend the lunch-and-learn. It may already have access to the CRM.**

## Have you tried turning awareness into security?

Here, we mean agents that can use tools to take actions and not simply generate answers. That distinction matters: a misleading response is one problem; a misleading response that triggers a customer-data export is another.

The challenge for October 2026 is to pair human awareness with enforceable limits on software acting on our behalf.

Start with a scenario.

1

# THE AGENT THAT LOGGED INTO YOUR BANK

2026 · BANKING · THIRD-PARTY AI AGENTS

According to Axios, U.S. banks are restricting third-party AI agents accessing customer accounts. Customers may willingly provide their credentials, even where banks prohibit that access.

The harder question comes after login: is the agent following the customer’s instructions, exceeding them, or acting without authorisation? Sources told Axios that banks struggle to distinguish these situations and described activity that appeared independent of user instructions.

Those accounts warrant care: Morgan Stanley told Axios it had not observed “rogue AI agents accessing our platforms.”

 THE 2026 QUESTION

Does permission to access an account also mean permission to act on it?

 THE CONTROL

- Use approved, scoped access rather than shared passwords.
- Keep information retrieval read-only.
- Require explicit approval for transfers, purchases and account changes, with a record of what the customer authorised and a way to revoke access.



“Check my balance” should not become “make yourself at home.”

2

# THE CREDENTIAL THAT HAD TOO MUCH FREEDOM

2024 · CHANGE HEALTHCARE

UnitedHealth’s CEO testified that attackers used compromised credentials to access a Citrix portal without MFA, then moved laterally and stole data.

This was not an AI incident. It is a reminder to examine the access behind a successful login.

 THE 2026 QUESTION

If an agent is manipulated, what can its credentials unlock?

 THE CONTROL

- Use a separate, narrowly scoped identity for each workflow. Prefer short-lived credentials where supported. Enforce authorisation in the connected service and test revocation.
- MFA remains essential for human access. It does not determine what an agent may do with an already-issued token.



“It uses my account” is a permissions decision, not a setup detail.

3

# THE FIX THAT NEEDED A SECOND LOOK

2017 · EQUIFAX

According to the FTC, Equifax circulated a patch alert to more than 400 employees. An improperly configured scan missed the vulnerable software, and the company failed to verify that patching was complete.

The lesson for automated remediation is straightforward: an instruction, an attempted fix and a verified result are three different things.

## THE 2026 QUESTION

The agent says the vulnerability is fixed. What evidence closes the ticket?

## THE CONTROL

- Review the change, run an independent security check and confirm the fix reached the affected environment. Keep a named owner accountable.



“Done” is a status message. Evidence is better.

# BEFORE YOU GIVE YOUR AGENT THE KEYS



1

## Engineer the harness

The model needs a well-built harness: the software that manages its context, tools and execution. Isolate risky operations, set execution limits and verify results before marking work complete. A confident “done” is not a test result.

2

## Put permissions outside the prompt

“Never disclose sensitive data” is useful guidance. The connected tools must independently enforce access, destinations and approval requirements.



3

## Make approval mean something

Show the reviewer the exact action, recipient and affected data. “Allow agent to continue?” is a button, not enough context for a decision.

4

## Rehearse the bad day


Test a malicious ticket. Try an out-of-scope request. Revoke access mid-workflow. Check whether the logs explain what happened. Build recovery where possible. Prevent irreversible actions before execution: you cannot roll back a disclosed secret.

# Keep the posters. Give them backup.

Awareness helps employees question suspicious instructions. Agent security also requires controls that hold when the software fails to question them.

This October, pair the training campaign with an access review: every agent, every connected tool, every identity and every consequential action.

**Give your agent a job description.**  
**It doesn't need the master key.**



# MEANWHILE

# AT ASTRA



# ASTRA VS. XBOW



There's no shortage of security vendors claiming they find more vulnerabilities than everyone else. So we figured we'd put our own claims to the test.

## The test

We ran Astra Autonomous Pentest against **Photoview v2.4.0**, the same application used in XBOW's benchmark, following Doyensec's published methodology.

## The result

27

verified vulnerabilities found by Astra, versus 7 by XBOW

11

vulnerability classes covered, versus 5 by XBOW

3.9x

more verified findings


**Astra's precision 57.4%. XBOW's was 100%.**

We found more real vulnerabilities, but we also generated more false alarms. That's a trade-off we're actively working to improve.

But the bigger question is **what matters more to your team: broader vulnerability discovery or less noise during triage?**

[See the full benchmark →](#)

# HERE'S WHAT WE SHIPPED IN Q3.



## Astra MCP: Security testing, meet your AI assistant.

If you practically live in Cursor, Claude, VS Code, or ChatGPT, this one's for you. You can now connect your AI assistant to Astra to run scans, triage vulnerabilities, and generate reports using plain language. Same permissions, OAuth sign-in, and fewer trips to the dashboard for routine work

▶ Run scans

≡ Triage vulnerabilities

📄 Generate reports

## SOC 2 Coverage Report: Because auditors love evidence.

And understandably so. Our new SOC 2 Coverage Report maps pentest scenarios to the relevant Trust Services Criteria, showing what's Validated and what's At Risk. You get the findings, impact, proof of concept, and remediation details together in one PDF. Now that's less time assembling evidence and more time working on what actually needs fixing.



**Validated**

Controls and evidence mapped



**At Risk**

Gaps clearly identified

## Postman Credential Verification: No more expired-token surprises.

You know what's frustrating? Starting a pentest only to discover that a token in the Postman collection has quietly expired. We've made that a

a little less painful for you though. Astra now verifies that your credentials work against each endpoint before testing begins. If something's broken, you'll know before it holds up the engagement.



**Expired token**  
Before



**Verifies credentials**  
Now

**Cloud Scanning: Less noise, more signal.**

Not every cloud resource deserves the same attention. You can now connect your cloud account once and use tags to tell Astra exactly what to scan. Want to include production and exclude test environments? Done. New resources inherit your tagging rules automatically. Because your pentest report shouldn't be longer just because you have more test infrastructure.



**Production**  
Include



**Test environment**  
Exclude



**New resources**  
Inherit tagging rules

**And that's not all!**

We've also shipped a cleaner dashboard, improved integrations, activity logs, and plenty of smaller improvements that make everyday security work easier.

ONE MORE THING WE'RE WORKING ON:

**White-box Autonomous Pentesting**

is on our roadmap for later this quarter. More on that when we're ready to show you what's under the hood.

See everything we shipped in Q3 →

# WE RENTED A **BILLBOARD** IN BENGALURU.

Go hack yourself with Astra →

Subtle, we know. We thought Bengaluru needed a little security reminder.

Was a post enough? Probably. But we took it off the feed and onto the streets anyway. That was a full-blown “Look, Ma, I made it!” moment for me.

**Shikhil Sharma**

Co-founder & CEO, Astra Security



# Events


We've spent a fair bit of time talking to security leaders lately, both in person and virtually. And if there's one thing I enjoy more than building security products, it's hearing from the people who actually have to use them.

Here's what's been happening, and what's coming up.



# The Human in the Loop





See you in **the next orbit**