CYBERJA

CYBERSECURITY & GRC MAGAZINE

ISSUE 08 • NOVEMBER 2, 2026

THE TRUSTED ACCESS ISSUE

When developer tools, cloud secrets and AI services become the attack path

INSIDE THIS EDITION	PRACTITIONER VALUE
Malicious AI coding extension	Developer endpoint and supply-chain assurance
Cloud-hosted secrets exposure	Non-human identity governance
GitLab exploitation concerns	Risk-based patching and CI/CD security
Agentic AI accountability	Privacy, authorization and oversight
CyberJA workshop + 100-mark rubric	Ready-to-teach classroom exercise

CYBERJA ACADEMY & ADVISORY | richeaperry.com

Editorial note: This issue synthesizes publicly reported developments from October 27–November 2, 2026. Incident investigations may evolve. Verify official advisories before operational decisions.

From the Editor

The most dangerous identity in an organization may not belong to a person. It may be an application token, an API key, a trusted software extension, a CI/CD runner, or an autonomous agent with poorly bounded authority. This edition translates current developments into assessment questions, evidence requests and actionable security controls.

The central question: Can the organization identify, constrain, monitor and revoke every identity and integration that can exercise meaningful business authority?

Executive Intelligence Snapshot

Development	Why it matters	Priority
Cursor extension compromise	Trusted developer tooling may introduce persistence and expose credentials	Critical
Northflank security incident	Stored secrets and AI API credentials require urgent governance	Critical
GitLab CVE-2026-97229	Self-managed development platforms require remediation and compromise review	Critical
Chainalysis disclosure	AI-assisted intrusions increase incident-response pressure	High
UK ICO agentic AI consultation	Autonomous processing needs accountability and privacy controls	Strategic

Immediate leadership decisions

- Identify privileged developer environments and third-party software integrations.
- Verify ownership and rotation readiness for AI API keys and other workload secrets.
- Prioritize exposed CI/CD and source-control platforms in vulnerability management.
- Require independent authorization for autonomous actions involving sensitive data or security controls.

01 | The AI Coding Extension Supply-Chain Incident

Cursor reported malicious Claude Code extension versions 2.1.312 and 2.1.313 distributed through its extension marketplace on October 29. The affected packages attempted to establish SSH-based persistence. The malicious packages were not official Anthropic releases. Cursor estimated up to 5,000 potentially exposed users; this is not a confirmed compromise count.

Business and GRC implications

Developer workstations may hold repository credentials, cloud access, deployment tokens and production privileges. A compromised extension can therefore become an entry point into high-value systems. Inventory extensions, enforce approved sources, separate developer and production privileges, monitor endpoint persistence and protect credentials.

Evidence to request

- Installed extension versions and exposure window.
- Endpoint security telemetry and unauthorized SSH key review.
- Privileged credentials accessible from affected machines.
- Approved extension inventory and software-supplier assurance.

Cursor advisory: https://www.cursor.com/security

02 | Cloud Secrets and AI API Credential Exposure

Northflank warned customers on October 30 to assume environment variables and stored secrets may have been compromised after reports of unauthorized AI API-key usage. Customers were advised to revoke and replace potentially affected credentials. This precaution does not mean every customer experienced confirmed theft.

Why this changes IAM governance

API keys and workload secrets represent non-human authority. Their misuse may expose data, allow unauthorized actions, or incur significant AI service charges. Governance must establish identity, owner, business purpose, permission scope, storage, expiration, monitoring, rotation and revocation.

Control	Required evidence
Inventory	Application-to-secret mapping and named owners
Least privilege	Approved API scopes and business justification
Secret storage	Managed vault and access policies
Monitoring	API request and billing anomalies
Response	Tested revocation, rotation and recovery procedure

Northflank incident updates: https://status.northflank.com/

03 | GitLab: Vulnerability Management for Critical Delivery Systems

CVE-2026-97229 was added to the CISA Known Exploited Vulnerabilities catalog on October 29. Organizations running affected self-managed GitLab deployments should validate the latest applicable security guidance, patch or upgrade, and investigate suspicious activity. A patch alone does not establish that a previously exploited system is clean.

The CI/CD blast radius

A GitLab server may be connected to source code, build runners, secrets, container registries and production deployment workflows. Assess its business criticality by the downstream authority it possesses —not just by its server classification.

- Identify exposed self-managed instances and software versions.
- Verify fixes and supporting dependencies.
- Review audit logs, runner tokens, secrets and unusual repository changes.
- Document incident escalation and recovery requirements.

CISA KEV catalog: https://www.cisa.gov/known-exploited-vulnerabilities-catalog

GitLab security releases: https://about.gitlab.com/releases/categories/releases/

04 | AI-Assisted Intrusions and Cloud Financial Exposure

Chainalysis disclosed unauthorized access to some internal systems and customer-related information, describing offensive agentic techniques. The company reported no evidence of access to certain more sensitive categories, including clear-text passwords and MFA secrets, as of its statement. Distinguish confirmed exposure from suspected exposure in incident communications.

Separately, Wiz researchers reported unauthorized AI API usage associated with compromised cloud credentials. Such cases show that AI security incidents can produce both confidentiality risk and direct resource-consumption costs. Estimated list-price usage should not be presented as a verified customer bill.

Practical control recommendations

- Correlate cloud, API and identity telemetry.
- Establish AI service budgets, usage thresholds and anomaly alerts.
- Limit key permissions and apply source or network restrictions where supported.
- Prepare incident-response playbooks for credential theft and unauthorized AI usage.

Wiz research: https://www.wiz.io/blog

05 | Agentic AI Governance and Privacy

The UK Information Commissioner's Office is gathering evidence on agentic AI and data protection, with a November 20 consultation deadline. The consultation is not a new regulation. It highlights governance issues surrounding purpose limitation, transparency, security, lawful processing and accountability when software agents take autonomous actions.

Agent authority review

Question	Expected control
Who owns the agent?	Named accountable business owner
What may it access?	Scoped identity and least-privilege permissions
What may it do?	Approved actions and prohibited operations
Who validates high-risk changes?	Independent approval and policy enforcement
Can it be stopped?	Revocation and containment procedure
Can activity be reconstructed?	Protected, attributable audit evidence

ICO consultation: https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/

CyberJA Practical Workshop | 90 Minutes

The AI Developer Platform That Exposed Business Credentials

Scenario: SecureTech Solutions Ltd., a 600-person FinTech company, uses Microsoft 365, Azure, AWS, GitLab and AI-assisted development. Its cloud hosting provider reports a security incident. The SOC finds unusual AI API usage, a potentially compromised coding extension on a developer workstation, and no reliable inventory of non-human identities. The CISO requests a prioritized risk assessment.

Student assignments

1. Identify critical assets and business processes.

2. Inventory service principals, application secrets and AI API keys.

3. Assess confidentiality, integrity, availability and financial exposure.

4. Analyze the development-tool supply-chain risk.

5. Determine GitLab remediation and investigation priorities.

6. Recommend Zero Trust, IAM and monitoring controls.

7. Develop containment, rotation, recovery and communication actions.

8. Prepare a one-page executive risk report.

100-Mark Assessment Rubric

Criterion	Marks
Asset identification and business impact	10
IAM and credential governance	20
Cloud and AI risk assessment	20
Software supply-chain analysis	15
Vulnerability management	10
Zero Trust recommendations	10
Incident response and executive reporting	15
TOTAL	100

Expected evidence

Asset register; non-human identity inventory; cloud risk register; privileged access review; vulnerability remediation plan; incident response recommendations; executive summary.

30-Minute Action: AI API Credential Risk Assessment

Minutes 0–10: Identify one AI-connected application, its API keys, storage locations, owners and permissions.

Minutes 10–20: Review API usage, unexpected spending, credential exposure and rotation capability.

Minutes 20–30: Document gaps, rank risks, assign owners and set remediation deadlines.

Deliverable: A one-page assessment showing business purpose, credentials, access scope, observed usage, potential financial impact, controls and accountable remediation owner.

10 Practical On-the-Job CyberJA Learning Paths

Practical Cybersecurity Assessments: https://www.udemy.com/course/practical-assessments-for-cybersecurity-grc-professionals/?referralCode=8480B6A8357F2C94EB44

Cybersecurity & GRC Best Practices: https://www.udemy.com/course/best-practices-for-cybersecurity-and-grc-professionals/?referralCode=E004F3F78FD1646AD491

Practical Zero Trust Implementation: https://www.udemy.com/course/practical-zero-trust-implementation/?referralCode=CF353BFD22C572C9E64D

Holistic Cybersecurity Professional: https://www.udemy.com/course/the-ultimate-cybersecurity-professional/?referralCode=F12CD4F7B12FAB80FE27

Effective IAM Governance: https://www.udemy.com/course/effective-iam-governance-a-practical-approach/?referralCode=3102DF119BD566345928

Cloud Adoption Risk Assessment: https://www.udemy.com/course/cloud-adaption-risks-assessment/?referralCode=21DA464EAE30864C2B15

Management of AI Risk: https://www.udemy.com/course/management-of-ai-risk/?referralCode=60BC1FB19D94264222AA

Practical Requirements of a GRC Professional: https://www.udemy.com/course/mastering-practical-requirements-of-a-grc-professional/?referralCode=774E187098C83A8D28D3

Practical Compliance Management: https://www.udemy.com/course/practical-compliance-management-for-grc-professionals/?referralCode=6BD8E3854B84624F9AB0

IT Auditing Fundamentals: https://www.udemy.com/course/it-auditing-fundamentals-for-the-grc-analyst/?referralCode=A29F743269BE75DC5266

Continue Your CyberJA Journey

CyberJA Academy & Advisory: https://www.richeaperry.com/

CyberJA resource hub: https://linktr.ee/richeaperry/shop

Closing message: Governance maturity depends on knowing which human and non-human identities hold authority, continuously verifying their actions, and being able to revoke that authority when trust is compromised.