AJAY MAGAR

Lead Security Engineer · Application & API Security · AI Agent Security

Bengaluru, India · magarajay538@gmail.com · linkedin.com/in/ajay-magar
magarajay.com · github.com/Ajaymagar

SUMMARY

Security engineer with 5+ years in application and API security, currently securing DevRev's multi-tenant AI SaaS platform. I find API bugs such as injections, broken access control and SSRF, work with engineers to fix them, and build AI agents that take repetitive security work off people's plates: automatic security review of every API and infrastructure change (1,300+ endpoints covered), manager-approved access provisioning, SIEM alert triage, SOC 2 questionnaire automation from Drata, and external attack-surface monitoring.

EXPERIENCE

DevRev Oct 2023 – Present
Lead Security Engineer / MTS *Bengaluru, India*

Application & API security

- Reviewed 1,300+ public and internal API endpoints: authentication and access-control design (OAuth/OIDC, PATs, session flows, tenant isolation) and threat modeling of new features in tech specs.
- Found and triaged 30+ valid vulnerabilities, including multiple SSRF and access-control flaws (IDOR, cross-tenant access); raised fixes in the codebase and drove the rest to remediation with owning teams. Write-up: **magarajay.com/blog/webhook-ssrf**.
- Performed static and dynamic security reviews of the Android and iOS apps and SDKs (MASVS/MASTG).

Security automation & AI agent security

- Built an **AI API-security review agent** (Python, Playwright, Cursor/Claude Code) that syncs OpenAPI specs, live-tests endpoints with role-based test profiles stored in AWS Secrets Manager, and produces structured reports, cutting a typical API review from days to a few hours.
- Built an event-driven **security change review**: any new or modified API endpoint or request parameter, or infrastructure change automatically triggers an AI security review agent (on Amazon Bedrock AgentCore) with the full context of the change, so every change gets a security review without anyone requesting one.
- Built an **agentic access-request workflow**: an agent grants access to DevRev, Google Groups, BigQuery and internal tools only after verifying the manager's approval on the request issue via API. It runs entirely inside DevRev's infrastructure, and every grant is auditable.
- Threat-modeled LLM agents, MCP servers and tool-calling integrations (prompt injection, excessive agency, secrets exposure) and defined guardrails for internal agent use.

Compliance automation (SOC 2 / ISO 27001)

- Built a **grounded Compliance Agent**, live in Slack for the sales team, that fills security questionnaires and RFIs only from approved sources: 138 curated answers plus 26 policies and 206 SOC 2 / ISO 27001 controls auto-synced from Drata's API with a framework crosswalk. It cites every claim, escalates anything unsupported to a human, and passed 87 adversarial red-team questions with no critical failures.
- Support SOC 2 and customer-trust work: security questionnaires, vendor due-diligence and security policies.

Development loop, cloud & attack surface

- Integrated security into CI: Snyk Code and Snyk Open Source triage and policy management, secrets scanning, and GitHub Actions workflow review for *pull_request_target* and secret-exfiltration risks.

- Built **Cloud Security Hub**, an inventory of every public-facing AWS asset (public IPs, load balancers, CloudFront, API Gateway, Lambda URLs, Route 53). An offensive AI agent scans each asset daily for new bugs, the team is alerted whenever anyone exposes a new public resource in any region, and it has caught dangling-DNS / subdomain-takeover risks.
- Run ProjectDiscovery tools with custom templates across DEV/QA/PROD; triage Wiz cloud findings and CVEs; read-only AWS audits via IAM Identity Center.

Incident response, bug bounty & SOC

- Led the response to a **critical SQL-injection incident** reported through the bug bounty program: reproduced and automated the exploit, scoped the blast radius (error-based file read reaching a Kubernetes service-account token), drove the hotfix (removed the vulnerable parameter and disabled the analytics database's file-access functions), verified it with before/after tests, and wrote the postmortem. Write-up: **magarajay.com/blog/sqli-incident**.
- Manage DevRev's **Synack bug bounty program**: triage researcher reports, validate and de-duplicate findings, set severity, and route fixes to owning teams. Also triage third-party pentest findings (Trace).
- Built an **AI SOC triage agent** for Coralogix SIEM findings: it investigates each alert by correlating Okta, AWS CloudTrail (via a read-only security-audit role), GitHub and Google Workspace logs, then records a verdict with supporting evidence for analyst review.

Security training & CTFs

- Run security training and hands-on CTF events for 200+ engineers across multiple offices.

Paytail Oct 2021 – Sep 2023

Security Engineer *Remote*

- Performed web, mobile (Android) and API penetration testing of Paytail's consumer lending and merchant platforms, and worked with developers to fix findings.
- Reviewed source code and cloud configuration, and set up vulnerability scanning in the CI/CD pipeline.
- Supported security compliance and audit requirements for a regulated fintech (RBI guidelines, ISO 27001).

SKILLS

Security: Application & API security testing, threat modeling, OAuth 2.0/OIDC, IDOR/BOLA, SSRF, OWASP API Top 10, AI agent & MCP security, mobile (MASVS/MASTG), incident response, bug bounty program management, SOC triage, access governance, SOC 2 / ISO 27001 compliance automation

Languages: Python, Go, TypeScript, Bash

Tools: Burp Suite, Playwright, Snyk, Wiz, Oligo, Drata, ProjectDiscovery tools, NoseyParker, TruffleHog, MobSF, Frida, Coralogix, Okta, Amazon Bedrock AgentCore, Cursor, Claude Code

Cloud & CI/CD: AWS (IAM, Route 53, CloudFront, EKS, Secrets Manager), Kubernetes, Cloudflare, GitHub Actions, CircleCI

PROJECTS

- **evidence-collector** (open source, github.com/Ajaymagar/evidence-collector): read-only SOC 2 / ISO 27001 evidence collection from AWS, GitHub, Okta and Drata. Maps each check to the SOC 2 criterion it proves and writes a timestamped, SHA-256-hashed evidence pack for auditors instead of console screenshots. Write-up: magarajay.com/blog/soc2-evidence.
- **ClawdHunter**: an agent-driven bug bounty framework with a codified, attacker-first methodology for business-logic, authorization and tenancy bugs; runs across multiple AI coding agents.
- **Security blog**: write-ups on application security and compliance automation at magarajay.com.

EDUCATION

University of Mumbai, B.Tech, Electronics 2021