AM

AJAY MAGAR

Senior Security Engineer · Application & API Security · AI Agent Security · Security Automation

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, security questionnaires, and external attack-surface monitoring.

Work History

DevRev – Security Engineer
Bengaluru, Karnataka
10/2023 – Present

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.
- 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 a **grounded Compliance Agent** that answers security questionnaires and vendor due-diligence requests only from approved policies and evidence, cites every claim, and routes unsupported items to a human. Includes an evaluation harness and an answer-library sync.
- 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.

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

Skills

- Application & API Security Testing
- Threat Modeling & Design Review
- OAuth 2.0 / OIDC, AuthZ Models
- IDOR / BOLA, SSRF, OWASP API Top 10
- AI Agent & MCP Security
- Mobile Security (MASVS / MASTG)
- Security Automation & AI Agents
- Attack Surface Management
- SOC Triage & Log Investigation
- CI/CD & Supply-Chain Security
- Bug Bounty Program Management
- Access Governance & Least Privilege
- SOC 2 & Security Questionnaires
- Python, Go, TypeScript, Bash

Tools

Burp Suite, Playwright, Snyk, Wiz, Oligo, ProjectDiscovery tools, NoseyParker, TruffleHog, MobSF, Frida, AWS (IAM, Route 53, CloudFront, EKS, Secrets Manager), Kubernetes, Cloudflare, GitHub Actions, CircleCI, Coralogix, Okta, Grafana/Loki, Amazon Bedrock AgentCore, Cursor, Claude Code.

Education

2021

University of Mumbai
Bachelor of Technology (B.Tech): Electronics

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.

Bug bounty, SOC & GRC

- 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.
- Support SOC 2 and customer-trust work: security questionnaires, evidence collection and security policies.

Paytail – Security Engineer
Remote
10/2021 – 09/2023

- 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).

Projects

- **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 and API security at **magarajay.com**.