

AMIR HOSSEINPOUR

Offensive Security Engineer | Cloud, Container and AI Security

amir-resume.pages.dev

PROFESSIONAL SUMMARY

Offensive security engineer with six years in security, from leading a company's move to the cloud and building its security function from scratch to red teaming cloud, container and AI systems. At White Tuque I run penetration tests, red team and adversary emulation engagements across SaaS web applications, APIs and cloud infrastructure for 20+ enterprise clients, and design cloud detection and deception systems in AWS and GCP. I built the red team and penetration-testing pipeline for a major Canadian insurer across their on-premises and cloud environments, and ran six months of red team engagements for a Canadian independent power producer that ended in a resilience-checking playbook mapped to their compliance requirements. Deep in authentication and authorization (OAuth2.0, OIDC, Zero Trust), with 120+ validated vulnerabilities on Fortune 500 programs, including a high-severity race condition at PayPal. I ran the AI Fight Club at Palo Alto Networks, a hands-on workshop for 200+ practitioners on attacking and defending LLM and agent systems, and build offensive tooling in Python and Go. Six CVEs filed from embedded device research; 15+ conference talks; co-organizer of DEFCON Toronto (DC416).

CORE SKILLS

- **Offensive security:** penetration testing, red team and adversary emulation, manual exploitation beyond automated scanners, exploit development, proof-of-concept exploits, risk-rated findings, retesting fixes
- **Cloud and containers:** AWS and GCP security, cloud red teaming, Kubernetes attack paths, container escape, Docker API and registry abuse, IAM and least privilege
- **Authentication and authorization:** OAuth2.0, OIDC, Zero Trust, authorization bypass, IDOR, privilege escalation, multi-tenant isolation
- **AI/ML and LLM security:** prompt injection, insecure tool use and excessive permissions in LLM and agent systems; LLM and agentic tooling that automates reconnaissance, vulnerability discovery and pentesting workflows
- **Detection and deception:** cloud detection engineering, honeytokens, decoy IAM credentials and canary resources in AWS and GCP; incident response
- **Security programs and automation:** building security programs from scratch, risk-based prioritization, threat modeling, security design reviews, vulnerability management, automated security tests in CI/CD

PROFESSIONAL EXPERIENCE

White Tuque – Offensive Security Specialist

Toronto, ON | October 2024 – Present

- Built the red team and penetration-testing pipeline for a major Canadian insurer across their on-premises and cloud environments, with recurring adversary simulations and tabletop exercises in place of a single annual test.
- Ran six months of red team engagements for a Canadian independent power producer, then combined the results into a months-long series of working sessions with their IT team. The outcome was a resilience-checking playbook mapped to the compliance requirements they must meet.
- Design cloud detection and deception systems in AWS and GCP, from honeytokens, decoy IAM credentials and canary resources to detections built on CloudTrail and GCP Cloud Audit Logs, so attacker contact raises a high-confidence alert.
- Red team cloud and containerized environments, from IAM misconfigurations and overprivileged service accounts to Kubernetes attack paths, container escape and Docker API and registry abuse.
- Penetration test SaaS web applications and APIs for 20+ enterprise clients with manual exploitation beyond automated scanners, proof-of-concept exploits, risk-rated findings with actionable remediation guidance, and retests that confirm durable fixes; operate Cobalt Strike and Havoc C2 on red team engagements.
- Threat model designs and run security reviews with client engineering teams, review AI and agentic workflows for prompt injection, insecure tool use and excessive permissions, and mentor interns.

ASEC – Penetration Tester

Toronto, ON | May 2024 – October 2024

- Penetration tested SaaS web applications, REST and GraphQL APIs and AWS-backed infrastructure for fintech clients in Canada, the US, Australia and Europe, specializing in API, GraphQL and OAuth2.0/OIDC security under Nick Aleks, author of *Black Hat GraphQL* and former Head of Security at Robinhood.
- Wrote a custom Nuclei template for each finding so clients could run automated security regression tests in their CI/CD pipelines.

Ofogh Kourosh

Tehran

Lead Security Architect | 2023 – February 2024

- Led the company's move to the cloud, my main initiative as security architect, using the migration to improve manageability, network segmentation and security as the chain grew from 200 to 1,000 stores.
- Built the company's first detection and response function, enterprise IAM with least-privilege access, and incident playbooks from scratch.
- Owned the security roadmap and the company's security architecture, built on defense in depth, and set the security standards new systems shipped against.

Security Engineer | September 2021 – 2023

- Integrated security testing into CI/CD pipelines, built security automation and internal tooling in Python and Bash, and implemented network segmentation, hardening baselines, and secrets and access management across a multi-site environment.

Security Analyst | June 2020 – September 2021

- First-line incident responder: ran security monitoring, triaged and responded to security events, and built detections and incident response playbooks.

HackerOne – Security Researcher

Remote | February 2022 – Present

- 120+ validated vulnerabilities across Fortune 500 programs, including severe findings on PayPal, Airbnb, Sony, Booking.com and AT&T, among them a high-severity race condition in PayPal's checkout; built reconnaissance and proof-of-concept tooling in Python and Go.

DEFCON Toronto (DC416) – Co-organizer

Toronto, ON | 2025 – Present

- Find researchers and presenters and program the monthly talks and workshops for Toronto's security community.
- Ran the AI Fight Club at Palo Alto Networks, a hands-on workshop for 200+ practitioners on attacking and defending LLM and agent systems.

ORIGINAL SECURITY RESEARCH, CVES AND TALKS

- **Ecovacs Deebot fleet compromise:** soldered onto the UART, dumped the firmware, recovered cloud credentials and mapped the MQTT topic tree. The broker trusted whatever device identity a client claimed, so a six-bug chain could give root on any of the ~10 million robots in the fleet from any account. Six CVEs filed with MITRE; coordinated disclosure under Safe Harbor; presented at BSides Toronto 2026.
- **15+ conference talks**, including DEFCON Vancouver (Microsoft-hosted) on API attack chains and OAuth2.0 exploitation, SecTor 2025, BSides Toronto 2026, TASK Toronto 2025 and 2026, and InfoSec Hamilton.

OPEN-SOURCE TOOLS

- **Eidolon:** open-source (MIT) offensive security workspace that runs pentesting workflows with an AI agent in the loop, with per-engagement isolation, scope tokens, a hash-chained audit log and command gating.
- **API-Authentication-Checker:** Burp Suite extension that replays proxy history with Cookie and Authorization removed and flags endpoints reachable without authentication. Also **graphql_sdl_generator**, which rebuilds a GraphQL schema from .gql files for attack-surface mapping.

TECHNICAL SKILLS

- **Cloud and containers:** AWS, GCP, Kubernetes, Docker
- **Offensive tools:** Burp Suite Pro, Cobalt Strike, Havoc, Nuclei, ffuf, nmap, sqlmap, Metasploit, BloodHound, mitmproxy, Frida
- **Programming:** Python, Go (Golang), Bash, JavaScript, TypeScript
- **Frameworks:** OWASP ASVS, WSTG and Top 10; MITRE ATT&CK
- **Languages:** English, Farsi, Mandarin

EDUCATION

Bachelor of Computer Engineering, University of Sistan and Baluchestan