ALEX • KASKASOLI

README

Ten years in offensive security and security engineering, the last three at Google DeepMind. I started with four years of red teaming and pentesting finance and critical-infrastructure clients, often exploiting cloud, CI/CD and DevOps tooling, and kept attacking systems after that: internal offensive testing at TrueLayer and Kaluza, and published novel Kubernetes and cloud attack techniques, including a CVE and GitOops, a BloodHound for CI/CD.

At DeepMind I founded and led the AI cybersecurity safety evaluations, running agents against vulnerable-by-design targets in sandboxed k8s environments; the framework I started went on to run most of DeepMind's safety evaluations. I also adversarially tested the SynthID watermarks, and threat modelled production ML deployments with the engineering teams that owned them.

As Head of DevOps and Security and as a Tech Lead Security Engineer, I owned infrastructure and security end to end across AWS and GCP, so I know how the systems I attack are built, and how to land fixes engineers adopt rather than route around.

RESEARCH

Core contributor: Evaluating Frontier Models for Dangerous Capabilities (2024); Gemini 1.0 (2023)
Co-author: A Framework for Evaluating Emerging Cyberattack Capabilities of AI (2025)
Co-author: Scalable Watermarking for Identifying Large Language Model Outputs (Nature, 2024)
Contributor: SynthID-Image (2025); Gemini 2.5 (2025)
Blog at https://alex.kaskaso.li: open-source offensive tools, an AWS attack nominated for PortSwigger's "best hacks of 2018", CI/CD security thinking cited on Google's Kubernetes podcast, and posts on sandboxing and hardening AI agents.

SKILLS

Offensive Security

Red teamed finance and critical-infrastructure environments; web app and infra pentests. In-house testing at TrueLayer and Kaluza. Burp, Cobalt Strike, BloodHound, Metasploit.

AI & Agent Security

Built DeepMind's cyber safety evals, with agents attacking vulnerable-by-design targets. Adversarially tested SynthID watermarks. Threat modelled production AI systems.

CI/CD & Supply Chain

Built GitOops (600+ stars) to map CI/CD attack paths across GitHub orgs; spoke on it at BSides London. Exploited CI/CD on red team engagements, and hardened pipelines with SAST, DAST and dependency scanning.

Cloud & Kubernetes

CVE for RCE/VM escape in minikube. Hardened multi-account AWS and GCP orgs, and ran production clusters. Built sandboxed k8s environments for hostile AI agents.

Detection & Response

Debriefed defenders on detection gaps after red team ops. Centralised and normalised org-wide AWS logs for SIEM ingestion. Handled production incidents.

Coding

Offensive and production tooling in Python and Go; worked with JavaScript, C. Now coding daily with AI harnesses (Claude Code, opencode).

EXPERIENCE

Google DeepMind — 3 years — *London, UK*

Security SWE in a cross-team function. Founded and owned the AI cybersecurity safety evaluations, building the framework and CTF targets and running the vendor engagements behind them. Adversarially tested image/text watermarks; threat modelled production ML deployments, driving fixes to resolution.

Kaluza / OVO Energy — 2 years — *London, UK*

Tech Lead Security Engineer for a 100+ account AWS and GCP organisation. Ran internal offensive testing, set security standards for 20+ engineering teams and led org-wide remediation programmes, fixing vulnerabilities with the teams that owned them.

TrueLayer — 1 year — *London, UK*

Security Engineer promoted to Head of DevOps and Security, leading a team of four. Ran internal offensive testing, owned external pentest engagements end to end, and turned ISO 27001 requirements into infrastructure-as-code policy.

MWR InfoSecurity — 4 years — *London, UK*

Security Consultant running red team engagements, plus web app and infrastructure pentests, for finance and critical-infrastructure clients. Moved laterally through CI/CD and DevOps tooling; debriefed defenders on detection gaps.

CERTIFICATES

OSCP, OSCE, CRT, CCT
Graduate-level MicroMasters credential in Statistics and Data Science (MITx)
BSc Ethical Hacking & Countermeasures (University of Abertay)

CONTACT

alex@kaskaso.li