

Rob Cerda

San Francisco Bay Area, CA | me@robcerda.com
[linkedin.com/in/robcerda](https://www.linkedin.com/in/robcerda) | github.com/robcerda | robcerda.com

SUMMARY

Security engineer with 13 years at Netflix, from helpdesk to the Attack Emulation Team, operating at staff scope on a senior title. My operating model is risk-based prioritization: I aim detection, emulation, and controls at what actually threatens the business, and I set that direction for the teams around me. Currently doing threat-intel-driven red teaming, including an internally built threat-intelligence agent fusing multiple threat intel sources and vendors through custom MCP server integrations. I also build with LLMs: open-source security tooling for AI agents, including hardware-key-gated secrets management for agents running on untrusted hosts.

CORE COMPETENCIES

Risk-based prioritization: threat-likelihood × asset-criticality scoping for detections, emulations, and control investment; setting security direction across teams

Adversary emulation: MITRE ATT&CK-mapped emulation, full kill chain operations, threat-based evaluations, purple team collaboration, attack-path analysis, remediation tracking with engineering owners

Detection engineering: SIEM content as code (Terraform), detection rule lifecycle management, high-signal detections (coverage, precision), EDR telemetry, alert tuning

Data engineering for security: security telemetry pipelines, data modeling and enrichment for detection and intel workflows

Threat intelligence: multi-vendor intel fusion, intel-driven operations, agent-based intel fusion

LLM, agent safety & adversarial ML: production LLM-powered agents, secrets management for agents on untrusted hosts, adversarial ML

Mentoring: internal security engineers and interns across multiple roles

Cloud & automation: AWS, Terraform, Google Cloud, Google Workspace, infrastructure as code, Python

EXPERIENCE

Netflix — 2013 to Present

Security Engineer, Attack Emulation Team — Threat Intel / Red Team (current)

- Define emulation strategy by risk: full kill chain operations against crown-jewel assets, chaining multiple attack vectors the way real adversaries do, mapped to MITRE ATT&CK. Drive findings through to remediation with engineering owners.
- Designed, built, and operate an internal threat-intelligence agent that fuses multiple threat intel sources and vendors through custom MCP server integrations — moving intel from static reports into queryable, workflow-integrated tooling for analysts and emulation planning.
- Set the direction for what gets tested and why across security and engineering stakeholders, operating at staff scope: defining emulation priorities and intel requirements, not just executing tests.
- Automate portions of assessments — scoping, execution support, reporting — plus intel collection, enabling continuous assessment at scale. Custom Python tooling and MCP servers.

Detection Engineering

- Owned the detection rule lifecycle as code (Terraform-managed): development, measurement, tuning, and safe rollout across cloud, endpoint, and identity telemetry.

- Ran the detection backlog as a risk-ranked queue — threat likelihood against asset criticality — directing analyst attention and engineering effort at high-fidelity coverage first and deliberately deprioritizing low-risk noise.
- Built high-signal detections: tuned EDR telemetry and alert logic for precision and coverage on the ATT&CK techniques most relevant to the environment, suppressing low-risk noise.
- Closed detection gaps surfaced by real incidents and red team findings, feeding lessons back into content and sharing patterns with adjacent teams.

Corporate and Studio Security

- Secured the corporate fleet and Google Workspace estate company-wide: identity lifecycle, SSO/SAML, device posture, and access controls, prioritized by attack-surface risk rather than checklist coverage.
- Protected pre-release studio content with layered access controls and monitoring on high-value media assets.
- Ran security investigations end to end, from triage through containment and remediation.
- Evaluated, configured, and hardened third-party SaaS applications across the company.

Earlier roles: Information Security, Systems Engineering / Administration, Helpdesk — 2013 onward

- Progressed from helpdesk through systems administration (Google Workspace, identity lifecycle, SSO/SCIM, Python and Terraform automation) and general infosec (security reviews, risk assessments, vulnerability triage by exploitability and exposure, investigations, SaaS application security) to security engineering — six roles over 13 years at one company.

OPEN SOURCE

- **tapseal** — Hardware-key-gated secrets for AI agents on untrusted hosts: YubiKey-sealed, tap-to-deliver via phone, time-boxed plaintext. AI agent security, not just tooling. github.com/robcerda/tapseal
- **monarch-mcp-server** — MCP server for the Monarch Money API (414 GitHub stars).
- **MCP server portfolio** — Google services (custom OAuth scopes), Tesla Owner API, Cloudflare Zero Trust, Oura, Sportradar MLB, Rentcast.