Tyler Almeida

AI & OFFENSIVE SECURITY ENGINEERING

Remote, USA · TylerAlmeida.com · Tyler@TylerAlmeida.com

I build and refine the systems that surface vulnerabilities continuously, and build the governance structures that let their findings land. Bug discovery is now cheap; prioritization and risk-targeted remediation is the new gold rush.

250K+ code changes reviewed · 2k+ exploitable vulnerabilities surfaced · 150+ critical/high resolved before production · autonomous testing across 900+ repositories

SELECTED SYSTEMS BUILT

Autonomous Offensive Security Platform

2025–2026

Co-creator and core engineer · 3-person team · adopted as the reference model for AI across the security organization

- Built a **fully autonomous loop** at all layers of development that continuously discovers, validates, and files security defects to the issue tracker with no human in the loop. Operationalizes targeted scanning across network testing, dynamic application testing, continuous code review, new product launch testing, mobile application security testing, and red-teaming profiles.
- Architected core platform behavior from first principles: assessment contextualization, agent-to-agent hand-off contracts, and the scale-hardening and reliability work required for unattended continuous operation.
- Designed **targetable, code-first** testing chain: a threat-modeling and recon agent reads the codebase, identifies attack surface and priority focus, then hands scoped targets to a tester agent operating under its own exploitation prompt. Language and stack agnostic by design.
- Led the **token-caching redesign** at the internal gateway, self-initiated after identifying an approaching cost ceiling: scan cost down ~55% (code review) and up to ~95% (platform), the unit economics that made always-on testing viable rather than a periodic campaign. Recognized by the CTO and CISO as a flagship AI capability in a company-wide engineering showcase.

Continuous Code-Review Scanning

2026

Continuous security review of every code change for $0.20 per PR using state of the art reasoning models

- Scaled to a reliable **~17K code changes/week**; 250K+ reviewed to date and **2k+ exploitable vulnerabilities surfaced**, ~20% remediation rate focusing on top severities. Prototyping an autonomous remediation proposal (PR) pipeline to drive this rate up.
- Redesigned the inter-agent data schema so findings carry reconstructable source references and decision-ready context: summary, description, and each severity factor with its own written justification. Prior state was effectively plain-text notification; this is what made findings actionable by owning teams.
- Migrated from a low-code prototype onto the core platform for atomic post-merge scanning at **~50% lower cost** with materially more capability; scales from documentation edits to sensitive business logic without quality loss.

Frontier-Model Scanning Surge: Execution, Triage & Remediation Pipeline

2026

Owner of the pipeline that turned a two-week surge across ~900 repositories into landed remediation · one of three engineers at the firm operating a government-restricted frontier model

- **Significant valid critical/high vulnerabilities identified and resolved with no business downtime**; all critical priority issues remediated within 3 weeks.
- Built the ingestion pipeline end-to-end: parsed **~30k unstructured model findings with no predefined schema** (raw text), mapped each to its source file and a durable code reference, deduplicated against all prior findings, scored against the internal severity model, and generated **~5K triaged security tickets**.
- Engineered the batching harness (Python + Make) to process **15 sandboxed repositories concurrently**, clearing all in-scope repositories in **2 weeks** (throughput up 8x), and built the program's reporting dashboard webapp from scratch on a rapid turnaround.
- Performed reachability analysis and progressive filtering to isolate genuine criticals. The triage layer is what separated true criticals from tens of thousands of low-signal findings.

EXPERIENCE

Coinbase

Remote · 2021–Present

Senior Offensive Security Engineer (2024–Present) · *previously Security Risk & Policy*

PROGRAM OWNERSHIP, STRATEGY & GOVERNANCE

- Defined the testing model itself: which surfaces are scanned on every code change, which get deep targeted testing, at what cadence, and built the red team identities and controls, approved by legal, HR-systems, IT, and compliance, that made recurring testing against production possible without disruption.
- **Accelerated other teams' 0→1 in AI security**: contributed to four distinct AI security tools now owned by privacy, blockchain security, secure design, and remediation functions.
- Owned the external pentest and compliance testing program: vendor lead time **6 weeks → 8 days**, ~20% lower cost per engagement YoY, audit-mandated pentests negotiated **44 → 4**, and externalized 15 regulatory assessments at ~2x historical volume, freeing the internal team for high-risk targets. Mobilized ~10 external engagements within 48 hours during a major security event.
- **Global Security Policy & Governance:** Drove the coinbase-wide shift to customer-user-journey-based regulatory testing (DORA/MiCA), unblocking EU licensing workstreams. Overhauled the Policy & Standard review process from a 2-month review to a streamlined, 2-week update, and drove over 100 Policy gap remediations within a 3-month rotation.
- **Security Risk:** Developed the internal security risk register, implementing a lightweight homegrown FAIR risk quantification methodology into the existing ticketing system, allowing Security and organization leadership to make risk-informed business decisions.

OFFENSIVE OPERATIONS & RED TEAM

- Designed and executed an **end-to-end data-exfiltration red team**: 86 scenarios designed, 36 executed and validated. Results informed the enterprise DLP program charter and generated new critical detection coverage.
- Designed and led endpoint-focused red teams; findings drove same-week fleet-wide hardening: access restrictions, network gating of sensitive applications, and recall of several hundred devices for reconfiguration.
- (Precursor to Autonomous Pipelines) Built **AI-orchestrated pentesting agents**: Claude Code plugin shared widely with the Offensive Security team which 3x pentest capacity, 75%+ cost reduction per engagement, and streamlined/standardized reporting capabilities. Multiple critical production vulnerabilities identified over the 30-day pilot.

Protiviti — Manager, Security Risk & Privacy Consulting

San Francisco · 2016–2021

- Stood up security and privacy programs from zero for Fortune 500 and high-growth clients across NIST CSF, ISO 27001, SOC 2, GDPR, and NYDFS — converting one-off assessment findings into sustained, scalable programs.

Oracle — Enterprise Solution Engineer, Database Security & Integration

Redwood Shores · 2013–2016

- Secure data architecture for Fortune 500 clients; $7M+ in technical sales.

SKILLS

Deep · AI/agentic security systems · autonomous scanning pipelines · LLM security & adversarial testing · finding triage and signal-to-noise engineering · security automation architecture

Working · application & network penetration testing · red team operations · cloud security · Python · TypeScript · threat modeling · vulnerability management

Domain · financial services · regulated environments (DORA, MiCA, PCI, SOC 2) · Web3/crypto

CERTIFICATIONS & EDUCATION

CISSP (ISC2) · **GEVA**, Enterprise Vulnerability Assessor (GIAC) · **CISA** Certified Information Systems Auditor (ISACA) · Certified Blockchain Security Professional

BS, Management Information Systems — California Polytechnic State University, San Luis Obispo