Michael Carlson

Staff Security Engineer | Application Security

Security engineer and force multiplier who builds AI-powered platforms, tooling, and automation that enable small security teams to operate at the scale of much larger organizations. Combines deep technical expertise across the full application security stack with a staff-level focus on team enablement — designing systems and workflows that raise the capabilities of every engineer on the team and reduce operational burden across the organization.

CORE COMPETENCIES

Application Security • Security Architecture • Penetration Testing • Threat Modeling • AI/LLM Security • Vulnerability Management • Supply Chain Security • CI/CD Security • Platform Engineering • AI/LLM Integration • Security Training & Enablement • Python

WORK HISTORY

Staff Security Engineer — Airtable | Aug 2022 – Present | San Francisco, CA

- Served as product security partner across multiple high-impact initiatives, translating security risk into practical requirements, mitigations, and delivery plans with engineering and product stakeholders
- Drove audit readiness by coordinating cross-team remediation and ensuring vulnerability action items remained within SLA for the annual audit cycle
- Authored security risk analyses and RFCs for emerging threat areas (including AI-assisted development risks), aligning stakeholders on prioritized controls and implementation roadmaps
- Built and evolved internal security automation platforms to scale AppSec coverage under constrained headcount, improving prioritization quality and reducing manual review burden
- Rolled out software supply-chain malware and dependency-risk scanning across the enterprise, including coverage for AI-generated code paths and third-party ecosystem risks
- Replaced legacy onboarding with a modern interactive AppSec training platform (88 modules across 8 security domain tracks), improving consistency and engineer feedback

Staff Security Engineer — Robinhood | Jul 2021 – Jul 2022 | Menlo Park, CA

- Managed the security services engineering team of 3, owning the authentication sidecar service and data tokenization service that protected user credentials and PII across the platform
- Transitioned to IC6 as technical lead for crypto custody security, designing the security architecture for cold, warm, and hot wallet solutions safeguarding billions in customer cryptocurrency assets
- Led security design of a custom in-house cold wallet solution for secure offline asset storage, including key ceremony procedures and multi-signature authorization workflows

Security Engineering Manager — BitMEX | Apr 2019 – Jul 2021 | San Francisco Bay Area

- Managed a team of 5 security engineers building custom security services and infrastructure for a cryptocurrency derivatives exchange
- Led development of a custom multiparty additive homomorphic signature cold wallet solution for secure cryptocurrency custody

- Built and deployed a custom osquery fleet for endpoint visibility and threat detection
- Designed and implemented an mTLS certificate authority integrated with the company's SSO solution
- Drove Kubernetes security hardening across the platform

Senior Application Security Engineer — Workday | Jul 2011 – Apr 2019 | Pleasanton, CA

- Progressed from application security engineer to senior IC over 8 years, spanning three distinct functions: AppSec, penetration testing, and security architecture
- Built the internal penetration testing function from scratch, grew and led the team to 4 engineers conducting continuous assessments of Workday's enterprise SaaS platform
- Collaborated on design of enterprise authentication services handling millions of user sessions, including mobile PIN authentication and key wrapping iteration
- Designed the security architecture blueprint for production deployment that was adopted as the org-wide standard, establishing security patterns for new services across the company

Information Technology Consultant — CSU Chico Distributed Learning Technologies | Sep 2009 – Jun 2011

- Built integration tooling in Ruby on Rails connecting the Blackboard Learn LMS with the PeopleSoft HR system

SECURITY ARCHITECTURE & ASSESSMENT

Custom Domains Security Architecture (Airtable)
Served as security DRI throughout the project lifecycle. Designed session isolation architecture, secure token exchange patterns (authorization code exchange over direct JWT), replay attack prevention with one-time codes and semaphore concurrency control, JWT security hardening with audience validation, and secure OIDC implementation with PKCE.

AI Product Penetration Testing (Airtable)
Led security risk analysis for the company's flagship AI product. Conducted penetration testing that uncovered critical data integrity violations, attacker persistence vectors, and data exfiltration risks. Authored architectural recommendations including a multi-agent model with constrained Worker LLMs and human-in-the-loop verification for dangerous operations.

LLM Security Rules Framework for AI-Assisted Development (Airtable)
Authored 34 always-active IDE security rule files governing every line of AI-generated code in the monorepo. Rules enforce admin action authorization across 10+ internal roles, encrypted database layer security with mandatory parameterized query patterns, authentication and session management, credential handling, deserialization security, input validation, and library approval policy. These same rules feed into automated PR evaluation, creating a unified security-by-default posture across both AI code generation and code review.

Application Security Risk Library (Airtable)
Built a centralized risk cataloging system mapping security risks to mitigating programs, tracking maturity scores, KTLO effort, and improvement opportunities — enabling data-driven conversations with leadership about security investment prioritization.

KEY PROJECTS & IMPACT

GridGuard — Vulnerability Management Platform

Designed and built a comprehensive vulnerability management platform that aggregates, triages, and tracks security findings across the entire organization.

- **Scale:** Manages 40,000+ vulnerabilities across 8 scanner types including SAST, SCA, secret scanning, container scanning, bug bounty, and penetration testing tools
- **AI-Powered Triage:** Integrated Claude AI for automated vulnerability assessment, exploitability analysis, and intelligent prioritization — ensuring on-call teams focus on the most impactful issues while routine findings flow into quarterly patching
- **Automated Team Assignment:** Built 145 assignment rules with 6 match types and AND/OR logic, automatically routing vulnerabilities to 19 engineering teams
- **Proactive Monitoring:** Deployed Terraform-managed Datadog dashboards and monitors covering workflow failures, API errors, and vulnerability volume spikes
- **Reliability Engineering:** Implemented comprehensive backup/restore system, exponential backoff retry logic, automatic duplicate detection, and centralized metrics services

Impact: Transformed vulnerability management from a manual, reactive process into a fully automated pipeline across 19 teams. AI triage and intelligent prioritization ensure critical vulnerabilities get immediate engineering attention while the vast majority are deferred to quarterly patching cycles — still remediated, but without interrupting sprint work. Estimated 2,000+ engineering hours/year saved by eliminating ad-hoc investigation and triage across on-call rotations (19 teams × ~2 hrs/week of triage overhead eliminated).

GitGuard — AI-Powered Security PR Review System

Designed and built an AI-powered system that evaluates pull requests for security impact, improving signal-to-noise ratio and enabling the security team to focus on PRs that actually matter.

- **AI Security Analysis:** Leverages Anthropic Claude API for intelligent PR risk scoring with attack path analysis and OWASP risk mapping
- **Dynamic Rule Engine:** Security rules loaded from Airtable enable non-engineer team members to update evaluation criteria without code changes
- **Smart Model Switching:** Implements cost-optimized model selection based on PR complexity
- **Opinionated Architecture:** Established the team's reference MVC pattern with dependency injection, schema-driven development, and sandbox/production isolation

Impact: AI handles initial triage of every PR, surfacing only those with genuine security implications. Any team member can review flagged PRs with full context, lowering the expertise barrier and saving the equivalent of a full-time analyst's workload (~$200K/yr in capacity). Freed the team to take on proactive security architecture work instead of reactive PR review.

Appsec-Toolbox — Security Engineering Automation Suite

Built a collection of CLI tools and AI-powered skills that automate day-to-day security engineering workflows.

- **Claude Code Skills:** Custom AI skills for GitHub Actions audit, PR security review, OpenSearch security investigation, Airtable task management, Dependabot triage, and Copilot review categorization
- **Incident Response Automation:** Automated secret verification, malware package scanning, security header scanning, and SBOM generation

- **Team Enablement:** Tools designed so any security engineer can perform senior-level investigations using AI-augmented tooling

Impact: Democratized security operations across the team. KTLO tasks that previously required hours of manual work are completed in minutes.

OSS License Attribution System

Built an automated compliance system addressing legal requirements to publish open source software licenses across all client-facing applications.

- **Industry-Standard SBOM:** Migrated from fragile custom parsing to CycloneDX SBOM generation via cdxgen
- **6 Platform Coverage:** Automated weekly workflows for Web, Desktop, Android, iOS, and SDK platforms with platform-specific parsing
- **Zero-Touch Operation:** Runs autonomously on staggered weekly schedules with comprehensive error handling and centralized dependency exclusion rules

Impact: Eliminated a recurring multi-day manual compliance effort. Legal team has continuous, up-to-date license attribution without any security team involvement.

Launchpad — Security-Hardened Application Boilerplate

Built a production-grade, security-first TypeScript boilerplate implementing 15+ hardened security patterns as secure-by-default infrastructure for backend APIs, fullstack applications, and AI/agent systems.

- **Security-by-Default Architecture:** 6 custom ESLint rules enforcing auth on every route, body validation on mutations, strict Zod schemas, and database access layer boundaries — developers cannot accidentally create unprotected endpoints
- **Production Cryptography:** Dual-mode libsodium encryption at rest (passphrase-based with Argon2id KDF and raw-key mode), timing-safe HMAC with zero-downtime key rotation, and double-submit CSRF tokens with signed cookies
- **Resilience Patterns:** Circuit breakers with Redis-coordinated half-open probes, graceful shutdown with SSE stream draining, Redis-backed session management with dual TTL (absolute + idle), and structured observability via Pino and Datadog StatsD
- **Supply Chain & CI/CD Security:** Multi-layer pipeline with Socket.dev scanning, license policy enforcement (GPL/AGPL/SSPL blocking), CodeQL static analysis, and GitHub Actions enforcement of GHAS features
- **Minimal Attack Surface:** Distroless container images (zero shell/package manager), multi-stage Docker builds, and Terraform IaC modules for AWS deployment

Impact: Provides a battle-tested foundation that eliminates weeks of security infrastructure setup for new projects, embedding security patterns and lessons learned from production systems into a reusable, clone-and-trim template.

Additional Engineering Contributions

Supply Chain Security in CI/CD: Inherited a pipeline with no supply chain scanning. Designed a multi-layered dependency security pipeline blocking PRs with known-vulnerable dependencies and enforcing a 7-day supply chain age policy to mitigate attacks like xz-utils-style compromises. Includes Datadog

metrics, bypass labels with audit trail, and developer-friendly remediation playbooks. Combined with Socket.dev pre-push scanning for defense-in-depth.

Application Security Training Platform: Developed a dual-mode interactive training system (88 slides across 8 security domains) purpose-built for the internal codebase and infrastructure. Content references real internal code paths and validated method signatures with regular accuracy audits. Replaced traditional slide decks, enabling self-paced onboarding that scales across the entire engineering org.

NewsGuard — Threat Intelligence Platform: Automated threat intelligence aggregation with AI-powered impact analysis, time-series rollups, and weekly Slack digests to the security team.

CodeQL Optimization: Resolved persistent timeout and resource exhaustion issues by implementing directory-based chunking, scan segmentation, and matrix parallelization — restoring reliable SAST coverage.

THE FORCE MULTIPLIER EFFECT

Challenge	Solution	Scale Achieved
Manual vulnerability triage	AI-powered categorization and exploitability analysis	40,000+ vulns prioritized across 19 teams
PR security review bottleneck	GitGuard AI evaluation with dynamic rules	Every PR evaluated, team reviews only flagged items
Manual team assignment	145 rules routing to 19 teams	40,000+ vulns auto-assigned
Compliance license tracking	Automated SBOM generation across 6 platforms	Zero manual effort, weekly cadence
Threat intelligence gathering	AI-filtered RSS aggregation with Slack delivery	Continuous monitoring, no analyst time
Security training delivery	Codebase-specific interactive platform	Scales to entire engineering org, no presenter needed
KTLO operational burden	10+ CLI tools and AI-powered Claude Code skills	Any team member productive in minutes
Silent automation failures	Terraform-managed Datadog metrics and monitors	Issues detected before stakeholder reports
Onboarding & knowledge transfer	Training platform + opinionated boilerplate	New engineers productive in days, not weeks

EDUCATION

California State University, Chico — BS, Management Information Systems

CERTIFICATIONS

GIAC Strategic Planning, Policy, and Leadership (GSTRT)

GIAC Security Leadership Certification (GSLC)

GIAC Web Application Penetration Tester (GWAPT)

GIAC Certified Incident Handler (GCIH)

TECHNICAL STACK

Languages: Python (primary), TypeScript, JavaScript, HTML/CSS, Bash
AI/ML: Anthropic Claude API, Claude Code skills, Airtable AI fields, LLM prompt engineering
Backend: Node.js, Fastify, PostgreSQL, Redis, libsodium, Zod
Frontend: React, Vite, Tailwind CSS
Platforms: Airtable (as application backend), GitHub Actions, Datadog, Slack API
Security Tools: CodeQL, Socket.dev, Dependabot, custom ESLint security rules, SAST/SCA/DAST scanners, bug bounty platforms, container security
Infrastructure: Terraform, Docker (distroless containers), CycloneDX/cdxgen (SBOM), OpenSearch