Esteban Castorena

EstebanCastorenajr@gmail.com
512-680-6526

Senior Security Engineer & AI Agent Builder

Portfolio: Estebomb.com **Linkedin**: linkedin.com/in/Esteban-Castorena **Github**: github.com/EstebanCastorena

Career Profile

Security engineer with production AI agent systems experience. 5+ years in detection engineering and incident response, combined with hands-on experience designing multi-agent LLM pipelines, AI-powered security automation, and cloud infrastructure on AWS. Driven by the belief that adversarial thinking and detection engineering are essential to building safer AI.

PROJECTS & VENTURES

VividScripts.ai - Multi-Agent Story-to-Video Pipeline *(Built with Claude Code)*

- Architected a 16-step AI agent pipeline coordinating specialized agents through typed dataclass contracts and a blueprint cascade architecture that compresses story-level intelligence for downstream agents
- Built LLM orchestration layer with unified dispatcher, and 5-layer output sanitization for handling malformed LLM responses in production. Deployed to AWS with Terraform (48 resources), OIDC-based CI/CD with zero stored credentials

Aquasafe Security - AI Email Security Inbox *(Founded 2023)*

- Built an AI-driven phishing detection product integrating with Gmail and Microsoft - analyzes every inbound email in real-time, provides threat context, and ensures users review flagged messages safely
- Engineered cloud infrastructure on AWS and GCP for real-time threat detection; developed Python automation scripts that reduced mean time to response from 30 minutes to under 1 minute

Technical Writing - estebomb.com/blogs

"Building a Multi-Agent Pipeline" | "The Consensus Model - Multi-Agent Debate" | "What Detection Engineering Taught Me About AI Safety"

EXPERIENCE

Expel - Austin, Texas (*2024 – Present)*
Senior Detection & Response Engineer (promoted early 2026)

- Designed an AI orchestration system using Claude agents that automates end-to-end security integration development from data source analysis to detection logic and SOC enablement, reducing development time from months to minutes
- Built an AI-powered rule review framework automating vendor rule analysis across all SIEM platforms, reducing review time from 8+ hours to under 5 minutes while improving quality - adopted org-wide
- Developed a multi-agent Consensus Model for MITRE ATT&CK classification - agents with distinct analytical biases independently evaluate detection rules, debate, and converge on mappings across 2,000+ rules
- Built a GenAI context engine for automated alerts - enriches escalations with investigation context and plain-language explanations, enabling faster customer action and eliminating SOC rework cycles
- Led delivery of 4 major integrations across network, endpoint/XDR, and email
- Technical lead during critical incidents: restored identity signal coverage during vendor API deprecation, managed 3-week data replay coordinating with SOC to prevent analyst overload

Reliaquest - Tampa, Florida (*2021 – 2023)*
Senior Security Analyst - Incident Response

- Threat detection, investigation, and response across enterprise environments serving Fortune 500 customers

TECHNICAL SKILLS

Languages: Python, SQL (BigQuery, KQL, SPL) | **Cloud:** AWS (ECS, S3, CloudFront, GovCloud), GCP, Terraform, GitHub Actions, Docker | **AI/ML:** Claude API, Claude Code, multi-agent orchestration, prompt engineering, consensus model design | **Security:** Detection engineering, incident response, MITRE ATT&CK, threat modeling, SIEM (Splunk, Sentinel, Datadog), DLP, XDR

Certifications: eCPPTv2, eJPT, Splunk Core Certified Power User, CompTIA Security+

EDUCATION

BBA in Cyber Security, The University of Texas at San Antonio - San Antonio, TX May 2021