



Staff Security Engineer – Red Team (AI)

Joshua Ho

STRENGTHS

AI Security Testing

Evaluated LLM applications and agents, then connected model behavior with safety controls and useful defensive telemetry.

Adversary Emulation

Built repeatable Caldera scenarios from ATT&CK and ATLAS techniques; teams gained clearer tests for missed detections.

Stakeholder Leadership

Brought Detection Engineering, Threat Intelligence, and Risk Management into one exercise rhythm; decisions became easier.

Operator Tradecraft

Applied C2, payload delivery, privilege escalation, and lateral movement within approved guardrails and documented rollback plans.


Clear Reporting

Turned technical evidence into risk-ranked findings and practical playbook changes; executives knew what needed action first.

SKILLS

- Red Team Operations
- Purple Team Operations
- Adversary Emulation
- AI Security
- LLM Architecture
- Agentic AI
- MITRE ATT&CK
- MITRE ATLAS
- MITRE Caldera
- Cobalt Strike
- Sliver
- Metasploit
- Empire
- BloodHound
- Detection Validation

LANGUAGES



MY CAREER

V.E.

- Staff Security Engineer – Red Team (AI) at BluePeak Assurance (8 Months)
- Principal Offensive Security Engineer at Canyon Peak Technologies (1.9 Years)
- Senior Red Team Engineer at HarborStone Financial Services (1.9 Years)
- Offensive Security Engineer at Northstar Cloud Security (1.9 Years)
- Security Analyst – Offensive Security at Redwood Cyber Defense (1.9 Years)

SUMMARY

Staff-level offensive security engineer with more than eight years of experience turning realistic adversary behavior into measurable improvements across enterprise identity, endpoints, networks, cloud, SaaS, and AI environments. Led Red Team, Purple Team, and adversary emulation operations from objective setting through safe execution, deconfliction, evidence review, and executive reporting. Hands-on work covers LLM architecture, agent behavior, MITRE ATT&CK, MITRE ATLAS, MITRE Caldera, Cobalt Strike, Sliver, Metasploit, Empire, BloodHound, AWS, Azure, Kubernetes, PyTorch, TensorFlow, Hugging Face, and LangChain. Partners closely with Detection Engineering, Threat Intelligence, Incident Response, and Risk Management teams. Ready to help Redwood Assurance Group improve AI security, telemetry coverage, response readiness, and control validation through disciplined offensive operations.

EXPERIENCE

Staff Security Engineer – Red Team (AI)

BluePeak Assurance

January 2026 - Present

Austin, TX

Leads AI-focused Red Team, Purple Team, and adversary emulation operations across enterprise identity, endpoints, cloud workloads, SaaS applications, and internal AI agents. Owns engagement objectives, safety controls, deconfliction, operator tradecraft, detection validation, and stakeholder reporting.

- Led AI adversary operations across identity, endpoints, cloud, SaaS, and internal agents.
- Established LLM evaluation procedures covering context management, tool use, data exposure, and safety controls.
- Built MITRE Caldera emulations linking ATT&CK and ATLAS scenarios with repeatable telemetry tests.
- Directed Cobalt Strike, Sliver, and BloodHound exercises with authorization, deconfliction, and rollback controls.
- Partnered with Detection Engineering and Risk Management on detection tuning and response playbooks.
- Applied PyTorch, TensorFlow, Hugging Face, and LangChain prototypes to accelerate evidence analysis.

Principal Offensive Security Engineer

Canyon Peak Technologies

January 2024 - December 2025

Denver, CO

Owned enterprise adversary emulation across hybrid AWS and Azure environments, guiding offensive security engineers through cloud, container, Kubernetes, serverless, AI, and mobile security initiatives. Connected technical findings with control validation and incident response readiness.

- Owned AWS and Azure emulation plans spanning cloud identities, containers, Kubernetes, and serverless functions.
- Designed Purple Team validation plans with success criteria, telemetry requirements, safety boundaries, and evidence standards.
- Maintained Caldera content with Metasploit, Empire, Sliver, and BloodHound scenarios for control testing.
- Evaluated AI assistants with Hugging Face and LangChain, documenting context behavior and monitoring opportunities.

- Guided mobile reverse engineering through obfuscation and anti-emulator protections for authentication improvements.
- Mentored five engineers on tradecraft, evidence handling, engagement safety, and concise reporting.

Senior Red Team Engineer

HarborStone Financial Services

January 2022 - December 2023

Charlotte, NC

Ran authorized Red Team and Purple Team engagements across financial applications and corporate infrastructure. Translated ATT&CK techniques into repeatable emulations, validated telemetry, and presented risk-ranked findings that improved defensive action.

- Ran authorized Red Team engagements from scope confirmation through execution, validation, and stakeholder readout.
- Converted ATT&CK techniques into Caldera plans using Cobalt Strike, Metasploit, and BloodHound.
- Tested payload delivery, privilege escalation, lateral movement, persistence, and C2 under documented guardrails.
- Reviewed Windows, Linux, AWS, and network telemetry with Detection Engineering for missed events.
- Built Python utilities and TensorFlow experiments that accelerated evidence triage for machine-learning workflows.
- Presented risk-ranked findings with practical remediation paths for infrastructure and Risk Management teams.

Offensive Security Engineer

Northstar Cloud Security

January 2020 - December 2021

Raleigh, NC

Delivered penetration testing and adversary simulation across cloud, web, API, endpoint, container, and mobile environments. Progressed from supported execution into independent engagement ownership through stronger rules of engagement, evidence quality, and reporting.

- Delivered penetration tests across AWS, Azure, web applications, APIs, and corporate endpoints.
- Developed Metasploit, Empire, and BloodHound workflows for identity and lateral movement validation.
- Correlated operator activity with SIEM and endpoint telemetry during Purple Team exercises.
- Containerized testing tools and assessed Kubernetes, microservices, and serverless permission weaknesses.
- Conducted mobile assessments involving certificate pinning, obfuscation, and anti-emulator protections.
- Improved rules of engagement, evidence quality, and technical reporting for independent ownership.

Security Analyst – Offensive Security

Redwood Cyber Defense

January 2018 - December 2019

Richmond, VA

Performed supervised vulnerability validation and penetration testing across operating systems, network services, web applications, cloud configurations, and internal business systems. Built a foundation in offensive tradecraft, ATT&CK mapping, evidence handling, and remediation reporting.

- Performed authorized testing across Windows, Linux, network services, web applications, and business systems.
- Used Metasploit, Nmap, BloodHound, and scripting to verify exploitable paths with clear evidence.
- Built foundational C2, payload delivery, privilege escalation, persistence, and lateral movement skills.
- Mapped assessment findings to ATT&CK and organized evidence for repeatable detection testing.
- Reviewed AWS and Azure permissions, storage exposure, security groups, logging, and administrative paths.
- Produced remediation reports that helped owners prioritize patching, access changes, and monitoring.

PROJECTS

AI Adversary Evaluation Program

2026

Created a controlled program for testing LLM applications and agents across prompt handling, context management, tool invocation, data exposure, safety controls, and defensive telemetry.

Hybrid Cloud Detection Validation

2025

Mapped enterprise adversary emulations across AWS, Azure, Kubernetes, containers, and serverless resources, giving control owners repeatable evidence for detection and response improvements.

Mobile Reverse Engineering Lab

2023

Built an authorized lab for mobile applications with obfuscation and anti-emulator protections, connecting reverse engineering findings with authentication and fraud-monitoring controls.

LEADERSHIP & AWARDS

- 2025 — Recognized for building repeatable AI adversary emulations mapped to MITRE ATT&CK and MITRE ATLAS.
- 2024 — Selected to mentor five offensive security engineers on tradecraft, evidence quality, and safe operations.
- 2023 — Earned peer recognition for clear risk-ranked reporting that improved incident response readiness.

EDUCATION

Bachelor's Degree in Cybersecurity

Virginia Commonwealth University GPA: 3.8

Richmond, VA

Coursework: Network security, secure software development, digital forensics, ethical hacking, incident response

CERTIFICATIONS

- Offensive Security Certified Professional (OSCP) 2021
- Certified Red Team Operator (CRTO) 2023
- Certified Information Systems Security Professional (CISSP) 2024
- MITRE ATT&CK and ATLAS Practitioner Workshop 2025

TECHNICAL SKILLS

- **Red Team Platforms:** MITRE Caldera, Cobalt Strike, Sliver
- **Offensive Frameworks:** Metasploit, Empire, BloodHound
- **AI Frameworks:** PyTorch, TensorFlow, Hugging Face
- **AI Agent Tools:** LangChain, LLM applications, AI agents
- **Threat Frameworks:** MITRE ATT&CK, MITRE ATLAS, adversary emulation
- **Cloud Platforms:** AWS, Azure, GCP
- **Operating Systems:** Linux, macOS, Windows
- **Cloud Native:** Containers, Kubernetes, serverless functions
- **Operator Tradecraft:** C2, payload delivery, privilege escalation
- **Movement and Persistence:** Lateral movement, persistence, operational security
- **Detection Operations:** Detection engineering, telemetry analysis, SIEM
- **Response Operations:** Incident response, response playbooks, control validation
- **Programming and Scripting:** Python, scripting, programming
- **Mobile Security:** Mobile reverse engineering, obfuscation, anti-emulator testing
- **Infrastructure Security:** Networks, protocols, firewalls, databases

SKILLS

- Red Team Operations
- Purple Team Operations
- Adversary Emulation
- AI Security
- LLM Architecture
- Agentic AI
- MITRE ATT&CK
- MITRE ATLAS
- MITRE Caldera
- Cobalt Strike
- Sliver
- Metasploit
- Empire
- BloodHound
- Detection Validation

PROFESSIONAL AFFILIATIONS

- Member, Information Systems Security Association, Austin Chapter.
- Member, Red Team Alliance and offensive security practitioner communities.

LANGUAGES

- English (Native)
- Mandarin Chinese (Intermediate)

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST