HARSH JANNAWAR

hjannawar014@gmail.com (425) 780-8648 San Francisco, CA harshjannawar.com linkedin.com/in/harsh-jannawar

EDUCATION

University of Washington 09/2024 – 06/2026 | Seattle, WA
Master of Science in Cybersecurity Engineering, GPA: 3.93/4

Symbiosis Skills and Professional University 08/2020 – 05/2024 | Pune, India
Bachelor of Technology: CSIT (Cybersecurity), GPA 3.65/4

PROFESSIONAL EXPERIENCE

AI Cybersecurity Engineering Intern 06/2025 – 08/2025 | California
SecureAIs
- **Drove 40+ paying users by improving data protection accuracy from 65% to 95%:** Developed a PII detection and redaction pipeline using Presidio, regex-based validation, and packet analysis across 10k test samples.
- **Reduced security review cycles by 60% through CI/CD automation:** Integrated SAST, DAST, dependency scanning, and secret detection into engineering pipelines to provide continuous security feedback during development.
- **Identified 20+ critical vulnerabilities across AI platforms:** Conducted adversarial testing and secure code reviews to uncover issues including authentication bypass, prompt injection, and token leakage prior to production release.
- **Implemented end-to-end authentication flows across engineering services:** Built and tested secure login, session handling, token validation, role-based access control, and protected API routes to reduce unauthorized access risk across application workflows.

Security Analyst Intern 06/2023 – 05/2024 | India
SecureThings
- **Cut asset analysis time from 9+ hours to under 30 minutes**: Built an event-driven scanning platform with Dockerized asynchronous Python workflows running Nmap and Nuclei, deployed via AWS Lambda and ECS, normalizing findings against CVE, CWE, and EDB-ID before storing in MongoDB and S3.
- **Built a lightweight device monitoring tool for resource-constrained security telemetry**: developed data collection agents in Python and Bash for devices with limited RAM and storage, streamed logs directly to an EC2-hosted pipeline, and built a dashboard to analyze security events.
- **Triaged 15+ high-severity vulnerabilities with validated exploits:** Performed application and infrastructure security assessments using Nessus, Nmap and Burp Suite to identify and validate critical risks.
- **Remediated 25+ cloud and application security gaps:** Strengthened AWS and Azure environments by improving IAM permissions, container security controls, secrets handling, and exposed service configurations.

Security Intern 06/2022 – 07/2022 | India
Pune Metro Rail Project
- **Reduced attack surface across 15+ findings**: Reviewed IAM roles and security group configurations across 20+ cloud resources, identifying over-permissive access and exposed services and recommending least-privilege changes.
- **Improved audit readiness and incident response coverage by 30%**: Mapped 12+ controls to monitoring signals in CloudWatch and GuardDuty, supporting alignment with ISO/IEC 27001:2022 and the NIST Cybersecurity Framework.

CORE COMPETENCIES & TECHNICAL EXPERTISE

Technical Skills: Security development, Threat modeling, Application security , Penetration testing, Secure code review, Authentication and authorization, DevSecOps, vulnerability management, Python automation, detection engineering, AWS, Docker, container security, IAM

Programming: Python, TypeScript, Next.js, SQL, JavaScript, Node.js, Bash Scripting

Frameworks and Database: ISO 27001/42001, NIST, SOC 2, OWASP TOP 10, MongoDB, Postgres, SQL

Tools: Burp Suite, Nmap, Nessus, Nuclei, Metasploit, Splunk, Git

Certifications: CompTIA Security+, CompTIA Pentest+ (THM), DevSecOps, TCM Practical Ethical Hacking, Oracle Cloud Infrastructure Foundations

PROJECTS

Inline Proxy - Security Gateway for AI Agents (Python, asyncio) 08/2026
- **Blocks the full prompt-injection exfiltration chain, verified end to end**: Transparent MCP proxy enforcing per-tool policy like allow, block, redact, human-approval and more on every agent tool call; taint tracking stops a poisoned GitHub issue from publishing a repo's .env.
- **194 tools policed at ~0.1 ms p99, 0% false positives across 200+ real tool descriptions**: Default-deny packs for GitHub, Jira, and Slack, each extracted from upstream server source; OIDC identity, Ed25519-signed policy bundles, and SIEM audit streaming to Splunk/S3.

Aegis LLM - Automated LLM Security Testing Suite 03/2026
- **Automated security testing for LLM applications mapped to the OWASP LLM Top 10**: Runs thousands of injection and jailbreak prompts against a target, scores each with LLM-as-a-Judge, and diffs results against baseline models so findings surface as actionable regressions.
- **Generates novel attacks and profiles targets before testing**: Built ThreatForge, a genetic-algorithm prompt-evolution engine that breeds new attack variants, plus a recon module that fingerprints the target model, detects RAG usage, and enumerates connected tools.

ACHIEVEMENTS & TECHNICAL CONTRIBUTIONS

Published Research: AI Security Compliance and Testing Framework for LLM Systems
M.S. Thesis, University of Washington, 2026 | ProQuest Dissertations & Theses, 32738510

Published cybersecurity and technical blogs on Medium

Ranked Top 1% on TryHackMe
Hands-on experience across DevSecOps, penetration testing, red team, and blue team paths and competed in CTF challenges.

Won Capture The Flag Competition at UWB GreyHats