VAIBHAV AGRAWAL

Email: agrawalvaibhav2004@gmail.com
Location: Bengaluru, India
LinkedIn — **Github**

SUMMARY

Final-year Computer Science (Cyber Security) student with hands-on experience building defensive security tooling end to end — a self-hostable SIEM with multi-source log ingestion and Sigma-based detection, an OWASP-mapped security scanner for LLM applications and tool-using agents, and an automated threat-intelligence pipeline that enriches CVEs with KEV and EPSS signals. Works across detection engineering and application security, with a particular interest in validating that detections behave as intended before they are relied on. Seeking security engineering, detection engineering, or SOC analyst roles.

TECHNICAL SKILLS

- **Security & Detection:** MITRE ATT&CK, Sigma rules, OWASP Top 10 (Web & LLM), CWE, CVSS v3, EPSS, CISA KEV, threat hunting, incident response, log analysis
- **Offensive & Analysis:** Burp Suite, Metasploit, Nmap, SQLmap, Wireshark, ffuf, subfinder, adversarial LLM red-teaming
- **Languages:** Python, TypeScript, JavaScript, Java, Bash/Shell, SQL, YAML
- **Platforms & Data:** FastAPI, React.js, Docker & Docker Compose, GitHub Actions, OpenSearch/Elasticsearch, PostgreSQL, MongoDB, Linux
- **Protocols & Systems:** TCP/IP, HTTP/HTTPS, syslog (RFC 3164/5424), TLS, JWT auth, RBAC, grep/regex

TECHNICAL PROJECTS

RedCell | *Security Scanner for LLM Applications & MCP Agents* GitHub Link
Technologies: Python, Model Context Protocol, OWASP LLM Top 10, GitHub Actions, pytest

- Built a pip-installable security scanner running 20+ adversarial probes mapped to the OWASP Top 10 for LLM Applications (2025), with a zero-setup offline demo target and a `--fail-on` severity gate for CI pipeline integration.
- Implemented excessive-agency testing against live MCP servers over stdio: enumerates exposed tools, flags destructive ones from MCP annotations and signatures, and (opt-in) confirms exploitability by attempting the ungated call — separating advisory exposure from confirmed finding.
- Designed indirect-injection detection that distinguishes a model *obeying* an injected instruction from merely quoting it, using out-of-band success signals that cannot appear in a faithful summary; validated at 21/22 detection on vulnerable mocks with **zero false positives** on hardened controls.

Argus | *Self-Hostable SOC & SIEM Platform* GitHub Link
Technologies: FastAPI, React.js, OpenSearch, PostgreSQL, Docker, Sigma, JWT

- Architected a self-hostable SOC/SIEM with a config-swappable storage backend (OpenSearch + PostgreSQL for production, SQLite for zero-infra demo), fully containerized via Docker Compose for one-command deployment.
- Built a real-time, multi-source ingestion pipeline (HTTP/bulk NDJSON, RFC 3164/5424 syslog, file tailing) that streams events into a detection engine evaluating Sigma-style YAML rules compatible with public SigmaHQ rulesets.
- Implemented correlation-based alerting, SPL-like threat hunting, and a React incident-response workflow (JWT auth, role-based access control), distilling high-volume detections into an actionable analyst view mapped to MITRE ATT&CK.

OSINT CyberBot | *Automated Threat Intelligence Pipeline* GitHub Link
Technologies: Python, GitHub Actions, LLaMA 3.1, FastAPI, NVD/KEV/EPSS APIs

- Engineered a serverless OSINT pipeline on GitHub Actions ingesting 13 authoritative RSS feeds, using Groq LLaMA 3.1 with Pydantic-validated JSON extraction, retry/backoff, and model fallback to parse unstructured reports into structured, queryable intelligence.
- Enriched CVEs with NVD CVSS, CISA KEV, and FIRST EPSS signals into a weighted 0–100 risk-priority score, with three-layer deduplication (URL, entity-overlap, TF-IDF cosine) across vulnerability and threat-actor lifecycles.
- Shipped a read-only FastAPI service hardened with rate limiting and guardrails against injection, SQLi, and path traversal, plus a dashboard with CVE dossiers and a force-directed threat graph.

Sentinel | *Web Vulnerability Assessment Scanner* GitHub Link
Technologies: Python, PostgreSQL, OWASP Top 10, CWE

- Built an automated scanner detecting OWASP Top 10 classes including injection and security misconfiguration, mapping each finding to its CWE identifier so results point at root cause rather than symptom; backed by a regression suite giving every detector a known-positive and known-negative case.

CERTIFICATIONS

Google Cybersecurity Professional (Coursera/Google) | **Google Cloud Cybersecurity** (Google Cloud) | **Threat Intelligence & Hunting** (IBM SkillsBuild) | **Security Operations Center in Practice** (IBM SkillsBuild)

EDUCATION

Ramaiah Institute of Technology 2023 – Expected 2027
Bachelors in Engineering in Computer Science (Cyber Security Specialization) Bengaluru, India