

Ayush Kumar

Agentic AI Engineer | AI Security & Enablement in Cybersecurity

+91-7091956935 | ayushkaps9462@gmail.com | linkedin.com/in/colonayush | github.com/COLONAYUSH | ayushkai.vercel.app | Bangalore, India | Open to Relocation

TECHNICAL SKILLS

Languages: Python, TypeScript, Go, Rust, SQL, Bash, HCL

Agentic AI & LLM Engineering: Claude Agent SDK, OpenAI Agents SDK, LangGraph, LangChain, CrewAI, AutoGen, MCP (Model Context Protocol), A2A Protocol, Multi-Agent Orchestration, ReAct, Tool Use / Function Calling, Structured Outputs (Pydantic), Human-in-the-Loop (HITL), Prompt Engineering, Claude API, OpenAI API, AWS Bedrock, LiteLLM

LLM Security, Red-Teaming & Governance: OWASP Top 10 for LLMs, OWASP Agentic AI Top 10, MITRE ATLAS, NIST AI RMF, ISO/IEC 42001, EU AI Act, Prompt Injection & Jailbreak Defense, NeMo Guardrails, Llama Guard, PyRIT, Garak, Promptfoo, NHI / Agentic Identity Governance

LLM Evaluation & Observability: Eval Harness Design, LLM-as-a-Judge, DeepEval, Ragas, LangSmith, Langfuse, OpenTelemetry, Prometheus, Grafana

AI / ML, RAG & Data: PyTorch, HuggingFace Transformers, Reinforcement Learning (PPO / DQN), Stable Baselines3, RAG / Graph RAG, Embeddings, FAISS, Pinecone, pgvector, Neptune Analytics, PostgreSQL, DynamoDB, Redis, Elasticsearch

Cloud, Infra & DevOps: AWS (Bedrock, SageMaker, Lambda, Step Functions, ECS / EKS / Fargate, Glue, Redshift, Aurora, Security Hub, CDK), Azure (Entra ID, Workload Identity, Function Apps, Logic Apps, PowerApps), Terraform, Ansible, Docker, Kubernetes, GitHub Actions CI/CD (OIDC), Temporal

Security Operations & Platforms: Splunk SIEM, Cortex XSOAR (SOAR), CrowdStrike EDR, Cloudflare WAF, Zscaler ZIA / ZPA, Wiz CNAPP, Qualys, Island, ServiceNow, MITRE ATT&CK, Detection Engineering, Threat Modeling, Secure SDLC, Zero Trust, OAuth 2.1, RBAC

EXPERIENCE

Lead Agentic AI Engineer — Cybersecurity

Mar 2025 – Present

Eli Lilly and Company

Bangalore, India

- Architected Lilly’s enterprise **Reference & Security-Mandated MCP Architecture** from the ground up — **15+ Enterprise Custom MCP Servers** exposing **240+ Agent-callable tools** on AWS ECS Fargate, hardened with 7-layer defense-in-depth (Cloudflare Edge → AWS WAF → ALB routing → network isolation → OAuth 2.1 → 3-tier RBAC → tool-level safety gates) and per-session credential isolation; adopted as the **Lilly-wide standard** for all AI-to-infrastructure integration.
- Led the **AI Ops initiative for Cyber Operations** — a **five-tier multi-agent AI platform** (Claude Agent SDK + LangGraph + Cortex) spanning orchestration, Triage, Supervisory, Process, and Tool-based agents across **10+ security services** (ServiceNow, Splunk, Cloudflare, CrowdStrike, Island); autonomously resolves **800+ tickets/month** across 25 platform services, cutting mean resolution time **88%** (45 min → 4 min), with Graph RAG retrieval reducing per-ticket token cost **85%** and Temporal saga-pattern rollback for safe infrastructure mutations.
- Designed Lilly’s **Agentic Identity (NHI) governance framework** — a three-category actor model (Human / Workload / Agent), per-agent Entra Workload Identities, 15-minute JIT credentials, and a **“Lethal Trifecta Separation”** control preventing any agent from simultaneously holding untrusted input, credentials, and mutation capability.
- Architected & shipped an **AI-native QC analytics platform** over ServiceNow, Jira, and Cortex XSOAR (**10K+ endpoints**) — exposed to agents and UI as a **42-tool Agentic MCP server** — featuring NLP ticket categorization, 7-dimension per-analyst quality scoring, priority-mismatch detection, predictive SLA/trend forecasting, and cross-team routing/handoff “lost-effort” analytics; cutting manual QC triage **~95%** and restoring accurate compliance reporting for VP/Director dashboards.
- Created Lilly’s **Agent Forensics & runtime governance model** to cybersecurity standard — full-trace agent decision logging, behavioral observability, and post-incident replay enabling forensic review and continuous policy enforcement across every production agent.
- Drove the org-wide **Claude Enterprise rollout** — led deployment, enterprise security baselines, and capability enablement that delivered governed generative-AI access across the entire organization.

Senior AI & Automation Engineer — Security Platforms

May 2024 – Feb 2025

Eli Lilly and Company

Bangalore, India

- Led the org-wide **Platform Infrastructure-as-Code initiative** — Terraform + Ansible automation for Cloudflare, Zscaler, and Splunk with GitHub Actions OIDC pipelines and a state-partitioning strategy that **eliminated deployment conflicts across 4 teams**; presented to leadership and adopted org-wide.
- Architected and built the **Splunk onboarding automation revamp** — an end-to-end process redesign on Azure (Entra ID, Function Apps, Logic Apps) and PowerApps that replaced manual runbooks with a self-service, approval-gated workflow, standardizing and accelerating data onboarding for the SOC.
- Engineered **25+ complex Splunk add-ons** and their onboarding across HEC, 3-legged-auth APIs, syslog, and universal forwarders, **setting the enterprise framework and quality bar** for all subsequent add-on development; automated Jira-driven workflow management, standardized playbooks for every ingestion pattern, and federated the AWS security data lake, establishing the operational foundation for all SOC analytics.

- Established the team’s **Security Platform DevX** practice — reusable Terraform modules, policy-as-code guardrails, and standardized templates and instrumentation that reduced operational toil and improved SLO adherence.

Engineer — Security Platform

Jul 2023 – Apr 2024

Eli Lilly and Company

Bangalore, India

- Built Lilly’s **enterprise Information Security data lake** end-to-end — unifying **50M+ data points** from Workday (SAD), CMDB, VRM, Qualys, and Acunetix via AWS Glue ETL pipelines, Lambda, Aurora, and Redshift to create a single source of truth for security analytics.
- Revamped the **enterprise Splunk logging pipeline** end-to-end, delivering **31 high-priority data onboardings** with custom ingestion models for reliable, standardized data flow into the SIEM.
- Shipped **real-time Cloudflare and Zscaler security-posture dashboards** delivering live health and posture visibility that accelerated Ops triage and gave leadership at-a-glance situational awareness.

Founding Engineer — SuperAGI

Apr 2023 – Jul 2023

SuperAGI

Bangalore, India

- Pioneered development of **SuperAGI**, an autonomous AI agent framework that garnered **21,000+ GitHub stars** — building the backend agent orchestration, tool-integration, and memory-management modules adopted by 15,000+ developers globally.
- Architected and developed **SuperCoder**, an agentic coding tool used by 5,000+ developers globally.
- Developed the AI-powered marketing-platform backend at Contlo.ai serving **20,000+ users**, including a ChatGPT plugin enabling end-to-end campaign management from conversational AI.

AI Research & Development Intern

May 2021 – Nov 2021

Samsung Research

Bangalore, India

- Designed an ML algorithm for optimal **5G Mobile Edge Server placement and task scheduling**, minimizing server count from 30+ to 11 — awarded a **Certificate of Excellence** (best project) and published a research paper on the algorithm.

PROJECTS

Multi-Agent Autonomous Threat Intelligence System | [Website](#) | [GitHub](#)

2025 – Present

- Built an open-source autonomous SOC platform (**55K+ lines Python, 45 modules, 2100+ tests, 130+ API endpoints**) with broader capabilities than commercial tools costing \$100K–\$500K/year.
- Engineered a **12-engine parallel detection pipeline** with weighted confidence fusion (ViT, BERT, transformer log-anomaly, DGA, UEBA, Sigma rules), a ReAct autonomous investigation agent with 15 tools and multi-LLM consensus, an RL-powered response engine (PPO/DQN), and Detection-as-Code with closed-loop analyst feedback.
- Implemented a **Security Knowledge Graph** with PageRank-based risk propagation, privacy-preserving federated threat intelligence (differential privacy + HMAC), and an AI self-defense layer (prompt-injection guard, model-integrity verification).

Agent Governance & Runtime Security Platform | *Python, Rust, MCP, OTel, gRPC*

2026 – Present

- Architecting an open-source **production agent security observability platform** — the “Wiz for AI Agents” — closing the gap between pre-deployment red-teaming (Promptfoo, PyRIT) and runtime agent governance in production.
- Core systems: **agent behavioral fingerprinting** with decision-pattern anomaly detection, **NHI credential lifecycle monitoring** with privilege-drift alerts, a **multi-agent trust protocol** with cryptographic attestation and delegation-chain verification, an **agent decision flight recorder** for full-trace forensics and replay, and **OWASP Agentic AI Top 10** continuous compliance scoring.

MCP Security & Observability Gateway | *Python, MCP, OpenTelemetry, OPA*

2026 – Present

- Building an open-source **security & observability gateway for the Model Context Protocol** — a drop-in proxy that inspects every tool call, enforces RBAC and tool-poisoning / prompt-injection guards, and emits full OpenTelemetry traces for end-to-end visibility into agent-to-MCP-server traffic.

EDUCATION

BMS College of Engineering

Bangalore, India

Bachelor of Engineering — Computer Science (CGPA: 9.2/10)

Aug 2019 – Jul 2023

CERTIFICATIONS & TRAINING

Certified Offensive AI Expert (HTB COAE) — Hack The Box | **Cloud Red Team Tactics: Attacking & Defending Azure (Advanced)** — Altered Security | **Splunk Certified Admin** — Splunk

PUBLICATIONS & AWARDS

“**MANTIS: A Multi-Agent Autonomous Threat Intelligence System for Real-Time Enterprise Cybersecurity**” — BSides Bangalore 2025 (Speaker) | **NullCon AI Paper of the Year** (DOI: 10.5281/zenodo.19161532)

“**Tamper-Evident ≠ Trustworthy: Forensically Sound Attribution of Autonomous-Agent Actions**” — Agentic-AI Forensics research (DOI: 10.5281/zenodo.20698154)

“**CyberWay: Secure Enterprise Standard for Development & Deployment of AI Connectors**” — Eli Lilly Internal Publication

Rising Star Award — Eli Lilly, among 10K+ employees | **LiFT Scholar** — The Linux Foundation (50 individuals globally) | **HPAIR 2020 India Representative (Harvard)** | **Open Source Summit Scholar** — The Linux Foundation