

Patrick Selamy

New York, NY | (210) 704-7370 | pselamy@gmail.com | selamy.dev | linkedin.com/in/patrickselamy | github.com/pselamy

Forward Deployment Engineer / Agentic AI Control Planes, Integrations, and Observability

Summary

Software engineer focused on production agentic AI systems, forward-deployed integration work, secure tool runtimes, and observable distributed platforms. Former Meta and Google engineer who built early Gemini-based software-engineering agents, SWE-bench-style eval harnesses, vector-search benchmarks, and customer-facing platform systems. Runs a live autonomous-agent fleet with MCP-style tools, skills, observability, CI gates, human-review thresholds, and credentialed workflows.

Selected Agentic AI Deployment Portfolio

Agent Control Planes, Integrations & Governance

- Built a production multi-agent runtime where MCP servers expose callable capabilities and skills encode reusable procedures, with broker-held credentials, security-audit gates, and human-review thresholds.
- Designed recoverable orchestration loops, tool-use contracts, and dispatch/fan-out patterns for six always-on agents executing real workflows across documents, CI repair, repo changes, and long-running task decomposition.
- Built CI-backed verification and evaluation loops for autonomous engineering work, including SWE-bench-style harnesses, planning-vs-execution metrics, multi-substrate testing, and runner reliability checks.

Observability, Security & Enterprise Runtime Fit

- Instrumented agent fleets with OpenTelemetry traces, metrics, and logs across tool calls, token use, latency, trajectories, and recovery behavior, with live dashboards for operating long-running workflows.
- Modeled each agent as a reusable infrastructure module for identity, runtime, skills, schedules, and secrets, using GitOps and IaC patterns instead of one-off deployment work.
- Built trust-sensitive systems across internal security, enterprise identity, access management, and regulated healthcare/cloud environments.

Experience

Meta, Software Engineer

New York, NY | 2025 to 2026

- Expanded Facebook Marketplace third-party partner APIs and shipped-listing workflows for partner and seller-facing commerce.
- Shipped an ML-assisted listing flow that removed manual item-weight entry for shipped listings; predictions were evaluated against user-entered weights and carrier actuals, met the launch threshold near 80% accuracy, and lifted new-listing creation about 8%.

Google, Software Engineer

Pittsburgh / Austin / Remote / Atlanta | 2018 to 2025

- Worked across Google Shopping (Engineering Productivity), Google Play, Google Cloud Office of the CTO, and Internal Identity and Access Management.
- In Google Cloud Office of the CTO, built early agentic-coding systems: a Gemini-based software-engineering agent, SWE-bench eval harnesses, and planning-vs-execution metrics that guided funding and staffing; benchmarked vector-search frameworks at scale (ScaNN, FAISS, pgvector) to position GCP's AI competitively.
- In Google Play, helped build [Play Games Leagues](#) from the ground up on a five-engineer team and designed a configurable fuzzer for third-party Play APIs.
- Built Insider Risk systems for internal security, enterprise identity, and trust-sensitive access.

Earlier, Backend / Platform / Healthcare Engineering

Citi, Lumeris, ClearDATA, and others

- Built backend, platform, cloud, and healthcare systems in Python and Java across APIs, data workflows, and regulated environments.

Selected Systems & Public Work

- [Public agent-skills](#): reusable agent procedures, authoring patterns, and an MCP-vs-skills framework.
- [Operating Agent Fleets](#): an evidence-backed field guide for durable agent systems and bounded authority.
- [polymarket-insider-tracker](#): a Python tool that flags suspicious wallet activity on Polymarket prediction markets.
- AI instructor and speaker: led an [Ayiti AI](#) bootcamp session on MCP, AI-assisted coding, and deployment, and [addressed lawmakers on AI policy](#) at the NHAEON National Leadership Summit.

Technical Skills

Languages:	Python, Java, Kotlin, JavaScript/TypeScript, Hack, PHP, Bash, SQL
AI / Agent Systems:	AI agents, skills/playbooks, MCP-style tools, tool calling, evals, agent harnesses, autonomous coding loops, context engineering, RAG, model routing, secure tool execution, state management
Product & Systems:	API design, backend platforms, distributed systems, workflow orchestration, product metrics, experimentation, launch analysis, developer infrastructure
Security & Enterprise:	internal IAM, auth boundaries, insider-risk systems, credential brokering, trust-sensitive workflows, privacy/security review gates
Infrastructure:	Kubernetes, OpenTofu/Terraform, Argo CD, Docker, Linux, GitHub Actions, Actions Runner Controller, CI/CD, observability, multi-cloud
Data & Cloud:	PostgreSQL, Redis, Kafka, InfluxDB, Google Cloud, AWS