

Senior Software Engineer, AI

We are seeking a hands-on Senior Software Engineer to design, build, and ship AI-powered products across our portfolio — including eKYC identity verification, fraud detection, and AI agent chatbots. This is a senior individual-contributor role with strong technical leadership: you will own architecture decisions, set the engineering bar, and mentor teammates, while spending most of your time building.

You will work in a small, talented squad creating systems that combine AWS serverless technology, large language models, and machine learning within a regulated FinTech environment.

Responsibilities

- Design and implement end-to-end AI systems — inference pipelines, agent workflows, and tool-calling architectures — from design through production readiness.
- Build LLM-powered features on AWS Bedrock: context orchestration, system prompts, memory, retrieval (RAG), and structured inputs/outputs.
- Develop and optimize ML components for identity verification and fraud detection (OCR, face matching, anomaly and spoof detection).
- Build, scale, and secure microservices on AWS serverless (Lambda, API Gateway, DynamoDB, S3, Bedrock).
- Own evaluation frameworks, guardrails, and AI observability — logging, tracing, quality monitoring, and latency/cost-aware fallback strategies across models and providers.
- Oversee model deployment and performance monitoring for ML and LLM components in production.
- Contribute to the technical roadmap; drive architecture decisions aligned with product and compliance goals.
- Evaluate and integrate third-party services — verification vendors, OCR engines, biometric SDKs, and LLM providers — as needed.
- Collaborate with Product Owner, BAs, UX/UI designers, and QA to deliver seamless product experiences.
- Mentor engineers through code review, pairing, and enforcing development best practices; drive continuous improvement in CI/CD, observability, and cost efficiency.
- Ensure compliance with PDPA, PCI DSS, and other relevant standards, including responsible AI practices.

Requirements

- 3+ years of backend or full-stack engineering experience, including 2+ years hands-on work with AI/ML or LLM systems in production.
- Strong Python skills, with experience in NoSQL databases, REST APIs, and AWS serverless architecture.
- Demonstrated experience building LLM-based solutions: prompt engineering, agent-style architectures, and integration with providers such as AWS Bedrock, OpenAI, or Anthropic.
- Experience evaluating and monitoring AI systems — automated testing, quality metrics, and production observability.
- Proven track record designing scalable APIs and distributed systems.
- Understanding of applied ML integration (OCR, biometrics, document validation, or fraud/anomaly detection).
- Excellent problem-solving skills; able to translate technical decisions into business impact.

- Strong written and verbal communication skills, capable of explaining complex technical concepts to non-technical stakeholders.
- A team player who thrives in a collaborative, idea-driven, Agile environment, balancing hands-on coding with technical leadership.
- Actively follows AI advancements and industry best practices.

Nice to Have

- Experience building or maintaining Retrieval-Augmented Generation (RAG) systems, including vector databases and embedding pipelines.
- Familiarity with AI guardrails, model safety measures, or responsible AI best practices.
- Familiarity with KYC/eKYC, FinTech, or RegTech domains.
- Experience fine-tuning LLMs or building automated AI training pipelines.
- Node.js, Docker, or containerized deployment experience.