

San Lee

Senior Product Associate

hi@sanlee.me · linkedin.com/in/leesan · sanlee.me

SUMMARY

Product professional at JPMorganChase, working on enterprise collaboration products (SharePoint Online and OneDrive). Behind that seat are seven years in infrastructure and platform engineering at the same firm, an MBA, and service as a U.S. Army veteran. I work across the technical and product seam, turning stakeholder needs into sequenced, testable delivery and trusting evidence over intuition when the two disagree. Outside the firm I build LLM applications (evals, RAG, agentic tool-use) in the open, run the way a technical program manager runs a program: AI agents are the build capacity; direction, contracts, and evidence gates are mine.

EXPERIENCE

JPMorganChase

Jun 2026 – Present

Senior Product Associate, Content Creation & Collaboration · Jersey City, NJ

- Own planning, build, and delivery of collaboration features across SharePoint Online and OneDrive for the entire firm — available to 300,000+ employees — moving them through a structured release lifecycle.
- Gather and document requirements across lines of business; report roadmap, delivery status, risks, and dependencies to stakeholders and senior leadership.
- Apply capacity management and prioritization to a high-demand, many-to-one user base — tracking execution in Jira and partnering closely with product engineering.

Seoul National University, Graduate School of Business

Aug 2024 – Aug 2025

Full-time MBA in Seoul, South Korea — career break for graduate study, then relocated back to the U.S. Details under Education.

JPMorganChase

Oct 2018 – Jul 2024

Senior Associate, Software Engineer III · Jersey City, NJ

- Built and ran microservices on a 4–5 engineer team working toward zero-touch provisioning (ZTP) of network devices — fully shared ownership, everyone with a hand in every service rather than an isolated slice.
- Practiced test-driven development and pair programming; designed and documented REST APIs (Swagger/OpenAPI), containerized services with Docker and Kubernetes, built event-driven flows with Kafka, maintained CI/CD pipelines, and remediated security scan findings to firm standards.
- Worked with SRE on observability — instrumenting endpoints and tracing API calls with OpenTelemetry to troubleshoot latency — and shared the team's operational load: documentation, on-call, and weekend failover and resiliency testing.
- Specced and ordered infrastructure, right-sizing it to “just enough” to keep costs in check.
- Pushed back on stakeholders with clear tradeoffs to keep scope and budget grounded.

JPMorganChase

Sep 2017 – Oct 2018

Analyst, Merchant Services Command Center Engineer · Jersey City, NJ

- Owned real-time monitoring and incident response for mission-critical merchant network systems supporting \$1T+ in annual transaction volume.
- Coordinated cross-functional triage across engineering, infrastructure, and operations — prioritizing by business impact and risk — and drove systemic fixes that reduced recurring incidents.

U.S. Army, NJ Army National Guard

Sep 2011 – Sep 2017

Corporal, Infantryman / Team Leader · Freehold, NJ

- Led a 3-person infantry team on a security-force (SECFOR) deployment to Qatar (Operation Enduring Freedom, 2014–2015) — entry-control-point and guard-tower operations; honorably discharged.

SELECTED PROJECTS

Built and documented in the open — every claim below links to its writeup and the public code behind it. sanlee.me

- **LLM eval harness (defense-news classifier)** — built the eval before trusting the model: an LLM-as-judge scored against a hand-labeled gold set caught two changes I'd have shipped on instinct (a “cleaner” prompt, an unproven retrieval add) that actually regressed accuracy. The product call: gate every change on evidence, not on how confident it feels. → sanlee.me/projects/defense-news-classifier
- **kb-agent** — RAG + tool-use agent orchestration, answering from a live knowledge base over versioned contracts instead of a static index, so answers can't silently drift out of sync with the source of truth; also ships as an MCP server any MCP host can mount. → github.com/sanlee-ys/kb-agent
- **Cross-repo LLM system** — evidence-gated model routing (cheaper model by default, escalate only where an eval proves the lift), wired across four services with frozen wire contracts and an idempotent async pipeline. Built the premium-tier upgrade, measured it, and shipped the negative result: it moved zero rows at ~1.97× the cost, so the shipped path stays single-model. → sanlee.me/projects/the-system

EDUCATION

Seoul National University, Graduate School of Business

Aug 2024 – Aug 2025

MBA, Management Track · Seoul, South Korea

Rutgers University, School of Communication and Information

Sep 2010 – May 2016

B.A., Information Technology & Informatics · New Brunswick, NJ

LEADERSHIP

Vector Regional Lead, Jersey City — ran the local chapter of JPMorganChase's global, multi-location technical knowledge-sharing program, organizing sessions for engineers in the office. (Sep 2023 – Jul 2024)

Int'l Relations Committee Leader, SNU MBA Student Council · VP, SNU MBA Sports Club (KRW 2M budget, 33 members)

SKILLS & CERTIFICATIONS

AI / LLM: LLM application development — structured output & tool-use, RAG, agentic tool-use, MCP servers, eval-driven development, LLM-as-judge, model routing; Anthropic Claude API; Claude Code (agentic coding); Microsoft Copilot; prompt engineering

Product & Program: cross-functional program management, roadmap planning, stakeholder communication, dependency & incident management, tradeoff evaluation, compliance & audit readiness

Technical: Python (FastAPI), Java, SQL, AWS, Kubernetes, Docker, Kafka, network device provisioning, Jenkins, GitHub Actions, CI/CD, OpenTelemetry · Tools: Jira, Jira Align, Confluence
Certifications: AWS Cloud Practitioner (2023) · Languages: Native English, Intermediate Korean

Personal site — views and projects are my own, not my employer's.