

# Ran (Eddie) Dang

Staff Software Engineer — AI Infrastructure & Developer Tools  
EDDING.SUREE@GMAIL.COM | [linkedin: ran-dang](#) | [github: eddieran](#)  
Location: Singapore | Open to: North America, Singapore (I-140 Approved)

## SUMMARY

---

- **Infrastructure at Scale:** 8+ years building production systems at hyper-growth companies (TikTok, OKX). Designed and shipped a company-wide SLO platform for 1000+ microservices, topology-aware anomaly detection, and automated RCA—the reliability backbone for billions in daily transaction volume.
- **AI Agent Infrastructure:** Introduced autonomous agent-driven development workflows adopted by 4–5 engineering teams. Independently built and shipped an AI agent skill manager (Go, v3.0, supporting 10 agent platforms) and an AI agent orchestration framework (Python, published on PyPI).
- **Restless Builder:** Shipped 5+ independent projects in the past year—all built with AI as a force multiplier. Currently exploring model distillation and fine-tuning to build cost-efficient local models for specialized tasks.

## INDEPENDENT PROJECTS & OPEN SOURCE

---

**SkillPM** | *Go, v3.0, 312 commits* | [github.com/eddieran/skillpm](https://github.com/eddieran/skillpm)

- Universal skill package manager for AI agents—atomic installs, rollback-safe sync, procedural memory, and security scanning across Claude, Codex, Gemini, Copilot, Cursor, and 5 other agent platforms.
- Full engineering lifecycle: CI/CD via GitHub Actions, Goreleaser for multi-platform binaries, Homebrew distribution, comprehensive test suite.

**Lindy Orchestrator** | *Python, PyPI* | [github.com/eddieran/lindy-orchestrator](https://github.com/eddieran/lindy-orchestrator)

- AI agent orchestration framework based on harness engineering: decomposes goals into dependency-ordered task DAGs, dispatches agents to isolated modules, validates output via pluggable QA gates.
- Multi-provider (Claude CLI, Codex CLI), live DAG dashboard, inter-agent mailbox, entropy scanner, issue tracker integration. No database—git-native coordination.

## PROFESSIONAL EXPERIENCE

---

**Staff Engineer / Tech Leader, OKX**

July 2024 – Present

*Infrastructure & Reliability Engineering*

*Singapore*

- **Autonomous Agent Operations:** Before OpenAI's Symphony was open-sourced, independently built [Lindy Orchestrator](#)—a stable AI agent orchestration framework grounded in harness engineering theory. After Symphony's release, studied its architecture deeply and integrated enterprise-grade capabilities: multi-tracker support, heterogeneous agent backends, internal CI/CD pipeline and swimlane deployment integration. Now adopted by 4–5 engineering teams with agents working near-autonomously around the clock, zero human intervention required, with quality controlled via automated QA gates.
- **Company-wide JVM Platform Migration:** Led the strategic migration of 800+ Java services from commercial ZingJDK (C4GC) to OpenJDK 21 (Generational ZGC). Built a custom JVM benchmark framework covering GC pause time, allocation rate, tail latency, and memory pressure across Intel/AMD architectures at 32G/96G heap sizes. Delivered 30+ ZGC tuning parameter guidelines and full toolchain compatibility validation (async-profiler, JFR, Virtual Threads observability). Result: GC pause reduced from tens-of-milliseconds (G1GC) to <1ms, saving annual commercial licensing costs.
- **0-to-1 SLO Platform:** Designed and shipped a comprehensive Service Level Objective system from scratch—data pipelines, metrics aggregation, alerting, and dashboards—now the reliability backbone for 1000+ microservices.
- **Automated Incident Response:** Architected topology-aware anomaly detection and a 5-minute automated RCA system, transforming incident response from reactive firefighting to proactive resolution across billions in daily trading volume.
- **High-Performance Routing Engine:** Optimized an industrial-grade routing engine integrating 300+ external protocol APIs and 5M+ data points; enhanced pathfinding algorithms for optimal route computation with sub-second latency.

**Staff Software Engineer, TikTok**

Jan. 2023 – July 2024

*Content Trust & Safety Platform*

*Singapore*

- **Configurable Platform (0-to-1):** Built an event-driven governance strategy platform enabling multi-country regulatory compliance with minimal per-market engineering overhead—directly enabled TikTok Shop’s global expansion.
- **Real-Time Risk Analytics:** Championed a parallel risk measurement system providing leadership with real-time visibility into content risk levels, informing key business decisions on market entry and policy enforcement.
- **Engineering Culture:** Established the team’s engineering culture from the ground up, defining development specifications and quality gates that sustained high-velocity feature delivery.

*Earlier Experience:* Senior SWE, **Meituan** (2021–2023) — led operational infrastructure digital transformation for Mobike | Senior SWE, **YuanFuDao** (2019–2021) — scaled infra to 20K concurrent users, RFC-driven technical strategy | SWE, **Meituan** (2016–2019) — local services platform engineering.

## EDUCATION

---

**Beijing Institute of Technology**

*Bachelor of Science, Internet Of Things, Computing Academy*

Beijing, China

2012 - 2016

## EXPERTISE

---

**Languages:** Python, Go, Java (JVM internals), TypeScript, Rust

**Infrastructure:** Kubernetes, Docker, AWS, GCP, Prometheus, Grafana, InfluxDB, OpenTelemetry

**Backend & Data:** FastAPI, Kafka, Flink, PostgreSQL, Redis, Elasticsearch

**AI & Agent:** Claude (MCP, Agent SDK, multi-agent orchestration), LLM tool use, prompt engineering

**Other:** Bilingual (English/Mandarin), Cross-functional collaboration, API design