ARAVIND SUNDARESAN

Senior Infrastructure & ML Systems Engineer

+91 97890 28238 | aravindsundaresan099@gmail.com | linkedin.com/in/aravind-sundaresan | github.com/Aravind0403 | Hyderabad, India

SUMMARY

Infrastructure and ML systems engineer with 7+ years building reliable, high-throughput distributed platforms at production scale. At Microsoft, eliminated 50+ retry storms across 17,000+ microservices via adaptive concurrency control and owned end-to-end BCDR achieving near-zero RPO. At Amazon, built a distributed orchestration system enabling 24/7 hardware-software integration across 50+ physical devices. Since late 2023, conducting independent AI infrastructure research — shipped three open-source systems: **Clairvoyant** (LLM inference scheduler, arXiv preprint, MLSys 2027 target), **ACO** (adaptive compute scheduler, P99 <10ms), and **ServiceScope** (LLM-powered dependency mapper). Python · Go · Java · PyTorch · ONNX · Kubernetes · AWS · Azure.

INDEPENDENT AI INFRASTRUCTURE RESEARCH (NOV 2023 – PRESENT)

Clairvoyant — LLM Inference Optimisation Proxy *Go · ONNX Runtime · XGBoost · Python · OpenAI-compatible*

- Designed and **open-sourced** a Go HTTP proxy eliminating Head-of-Line Blocking in LLM serving — ONNX-exported XGBoost classifies output token length in **0.029ms** (100× below generation latency), reorders dispatch via **Shortest-Job-First scheduling**, achieving **100% correct ordering** validated on Apple M1 and NVIDIA RTX 4090; formatted as an **OpenAI-compatible reference architecture**
- Evaluated across **7 public LLM datasets**; identified GPT brevity bias in curated corpora (Alpaca, CodeAlpaca) rendering them unusable as SJF training sources; natural conversation logs yield **76–96% ranking accuracy** (21–29 pp above 3-class baseline) — arXiv preprint in preparation, **MLSys 2027** submission target; Hugging Face artefact releasing May 2026

ACO — Adaptive Compute Scheduler *Python · FastAPI · PyTorch · LSTM · NumPy · Linux*

- Achieved **P99 <10ms** scheduling latency and **+28% resource utilisation** over first-fit baseline via Ant Colony Optimisation + **PyTorch LSTM** spike prediction — intent-aware routing enforces GPU/CPU/ARM64 affinity constraints; **95%+ SLA** under burst load, **202 tests, 0 failures** (Alibaba 2018 + Google Borg 2019 cluster traces)

ServiceScope — LLM-Powered Microservice Dependency Mapper *Python · FastAPI · Ollama · Neo4j · PostgreSQL*

- Maps HTTP dependencies across Python repos at **190 files/sec** using AST extraction + **Ollama gemma3:4b** local LLM inference; **0% LLM inference failure rate**; validated on django (2,886 files) and robusta (394 files) — answers blast-radius questions without external API calls

PROFESSIONAL EXPERIENCE

Microsoft R&D India *Aug 2021 – Nov 2023 · Hyderabad*

Software Engineer II — Platform & Infrastructure

Metadata Platform — 17,000+ Microservices

- Eliminated **50+ production retry storms** by root-causing concurrent write contention and implementing **Redis**-based traffic reshaping with adaptive concurrency control (95/5 optimistic/pessimistic split per live conflict rate) — drove adoption across the full distributed platform; served as **on-call DRI** achieving **<15 min MTTR**
- Designed and owned **BCDR end-to-end** — active-passive failover across two Azure regions via Traffic Manager, Service Bus geo-DR, idempotent message replay, and Cosmos DB automatic region promotion, achieving **near-zero RPO**
- Reduced operational toil by **∼70%** via distributed service ownership validation platform (Azure Service Bus, Cosmos DB, C# Azure Functions); automated multi-region provisioning with **Azure Bicep + ARM templates**

Developer Velocity — Windows Platform

- Built **C++ Test-in-Production telemetry framework** adopted across **7–10 teams** (Windows India chapter) — 90% instrumentation coverage as CI/CD gate; implemented 4-stage ARM64 pipeline cutting setup **60 → 20 min (3×)**

Amazon Development Centre *Sep 2017 – Aug 2021 · Chennai*

Software Development Engineer — Systems

- Built **Java distributed orchestration system** for 50+ physical Alexa devices (Echo, Fire TV, Fire tablets) — category-aware routing, deterministic state management, ADB-based automatic recovery; enabled **24/7 continuous validation** with zero test blocking from hardware failures
- Improved defect triage efficiency by **∼30%** with structured ADB diagnostic tool adopted by **7+ Amazon device teams**; automated Alexa Comms CI/CD pipeline eliminating all manual qualification steps

EDUCATION

B.Tech in Information Technology *2016*

Anna University

SKILLS

Languages	Python (primary) · Go · Java · C# (.NET) · C++ · Bash · SQL
ML & AI	PyTorch · ONNX Runtime · XGBoost · LSTM · Ollama (gemma3:4b) · local LLM inference · prompt engineering
Distributed	adaptive concurrency control · BCDR design · event-driven architecture · idempotent processing · retry storm elimination · resource scheduling
Cloud & Infra	AWS (EC2, SQS, S3) · Azure (Functions, Service Bus, Cosmos DB, Redis, Traffic Manager) · Docker · Kubernetes (k3s · Helm · ArgoCD) · Linux (Ubuntu · systemd)
Reliability	on-call incident response · <15 min MTTR · SLO/error budget · Prometheus · Grafana · Loki · Open-Telemetry · Application Insights
IaC & CI/CD	Azure Bicep · ARM templates · Terraform · Azure DevOps Pipelines · ArgoCD (GitOps) · Helm · CI gate design