

NAVANEETH RAO

(469) 450-9826 | navaneethbv@gmail.com | linkedin.com/in/navaneethrao

PROFESSIONAL SUMMARY

Staff Software Engineer with 13+ years building distributed cloud platforms, ML serving systems, high-throughput backend systems, and developer infrastructure at scale. Consistent record of business impact: \$3M+ in cost savings, 30% efficiency gains, and 99.99% uptime across mission-critical systems. Experienced in system design leadership, architecture governance, and cross-team technical mentorship. Deep expertise in Java, Kubernetes, AWS, and GCP across ML inference, financial systems, and platform engineering. Combines technical architecture depth with business acumen (UIUC MBA, 3.95 GPA) to align engineering strategy with enterprise goals.

PROFESSIONAL EXPERIENCE

Staff Software Engineer | Walmart Inc. | Apr 2024 – Present · Seattle, WA · Remote

- Migrated Walmart Search internal cloud-native platform clusters from single-tenant to secure multi-tenant Kubernetes-based model-serving architecture, improving platform availability to 99.9% and reducing inference latency by 25%.
- Defined platform architecture standards and deployment best practices for ML model serving, reducing deployment cycle time by 40% and eliminating post-release rollback incidents across the Perceive platform.
- Designed a multi-tier inference routing system differentiating low-latency cache-serving from deep model execution, optimizing query routing and failover across multi-tenant, multi-locale environments supporting hundreds of millions of daily Walmart Search queries.
- Engineered a query normalization and tokenization layer that improved downstream model prediction accuracy for Walmart Search query understanding, with standards later adopted across multiple Search engineering teams.
- Established architecture governance and distributed systems training programs, scaling engineering velocity across multiple Search platform teams and reducing onboarding ramp time for new engineers joining the multi-locale org.

Technologies: *Java, Python, Kubernetes, Spring Boot, GCP, CI/CD*

Staff Software Engineer | PayPal | Feb 2023 – Apr 2024 · Remote

- Engineered a Financial Risk Reporting platform on GCP processing 10M+ daily transactions, enabling scalable fraud detection and chargeback analytics for PayPal merchants.
- Orchestrated a Databricks-based analytics platform for terabyte-scale real-time processing, improving business operations efficiency by **30%** across merchant analytics and decision-making workflows.
- Migrated core data workflows from RDBMS to AWS DynamoDB, cutting operational costs by **30%** while improving horizontal scalability and availability.
- Spearheaded data-layer modernization strategy for core legacy systems, establishing ORM-based frameworks that reduced feature-to-production lifecycle by 50% and unblocked multiple engineering teams from database bottlenecks.
- Designed scalable event-driven data pipelines and distributed processing workflows supporting high-volume financial risk analytics across PayPal merchant systems.

Technologies: *Java, GCP, AWS (DynamoDB, EKS, Lambda), Databricks, Spring Boot, Datadog*

Senior Software Engineer | Ford Motor Company | Oct 2019 – Feb 2023 · Detroit Metro Area

- Designed and built the Azure Right to Provision service for the Ford Cloud Portal, automating subscription creation and budget management with Java and Node.js — enabling self-service cloud provisioning for 5,000+ engineers and eliminating manual overhead.
- Developed a Role-Based Access Control system with automated Cloud Billing in TypeScript and Kubernetes, securing budget governance for 5,000+ employees and eliminating unauthorized spend exposure.
- Led the end-to-end delivery of Cloud Portal 2.0, Ford's central internal Developer Experience (DX) platform; coordinated cross-functional teams to integrate multi-tenant analytics and self-service governance tooling for 5,000+ engineers.

- Conducted architecture reviews and established platform engineering standards adopted across five internal teams, improving deployment consistency and reducing cross-team integration friction.
- Built an observability system in Java, Kafka, and Docker capturing platform metrics, reducing mean time to detection (MTTD) for incidents and enabling data-driven capacity planning.

Technologies: *Java, TypeScript, Azure, GCP, Kubernetes, Kafka*

Software Engineer / Team Lead | IBM | Aug 2016 – Sep 2019 · Austin, TX

- Led engineering of monitoring and alerting infrastructure across 50+ data centers and 10,000 servers using Icinga, Grafana, and PagerDuty — achieving **99.99% uptime** for IBM Cloud and protecting enterprise client SLAs.
- Built a 500 TB/day logging pipeline through Logstash, Kafka, Elasticsearch, and Kibana, enabling real-time operational insights that directly informed \$3M+ in infrastructure cost reductions.
- Led a cross-functional team of 10 engineers, driving sprint planning, delivery coordination, and operational execution across cloud infrastructure initiatives.
- Served as SME for Level 1 escalations; led root cause analysis across Jira, GitHub, and ServiceNow, consistently meeting a 15-minute SLA for customer escalations.

Technologies: *Java, Kafka, Kubernetes, Elasticsearch, Icinga, PagerDuty*

Service Engineer Intern | Microsoft | May 2015 – Aug 2015

- Overhauled the in-house Azure billing system in Java, delivering a 40% performance improvement and increased pipeline reliability.
- Implemented Cassandra data modeling and query optimization for a real-time analytics platform, achieving a 50% improvement in data retrieval speed.

Technologies: *Java, Azure, Cassandra, Spring Boot, REST APIs, CI/CD*

Senior Systems Engineer | Infosys | Aug 2012 – Aug 2014

- Delivered web-security projects for ANZ Bank using Java and Spring Framework, and developed an Android OTP authentication app that reduced unauthorized access by 50%.
- Trained 5 interns in the Finacle e-Banking Framework and RESTful API design; all interns shipped features to production.

Technologies: *Java, Spring Framework, JavaScript, SQL, Jenkins, Finacle*

EDUCATION

Master of Business Administration (MBA)

Gies College of Business, University of Illinois Urbana-Champaign · Oct 2022 – May 2024 · GPA: 3.95 / 4.0

Master of Science, Computer Science — Specialization: Machine Learning & Big Data

The University of Texas at Dallas · Aug 2014 – Jun 2016 · GPA: 3.9 / 4.0

Bachelor of Engineering, Information Science

Visvesvaraya Technological University · Sep 2008 – May 2012

TECHNICAL SKILLS

Languages: Java, Python, TypeScript, JavaScript

Frameworks: Spring Boot, Node.js

Cloud & Infra: AWS (DynamoDB, EKS, Lambda, ECS, RDS, EC2), GCP (BigQuery, Cloud Run, GKE), Azure (ARM, CosmosDB)

Platforms & Tooling: Kubernetes, Docker, Kafka, Databricks, OpenShift, Jenkins, Tekton, CI/CD

Observability: Grafana, Prometheus, OpenTelemetry, Datadog, Elasticsearch/Kibana, Splunk

Databases: MongoDB, PostgreSQL

Leadership: Distributed Systems Design, Platform Architecture, Technical Mentorship, Architecture Reviews, Cross-functional Engineering Leadership

AI/ML Infrastructure: ML Serving Systems, Model Inference Routing, Tokenization Pipelines, Feature Engineering, Large-Scale Data Analytics