

BISWAJIT SAHA

Austin, TX • 512-287-0579 • saha.biswajit.1992@gmail.com

LinkedIn: [linkedin.com/in/biswajit-saha-ut/](https://www.linkedin.com/in/biswajit-saha-ut/)

SENIOR FULL STACK SOFTWARE ENGINEER (AI & PLATFORM)

Full Stack Software Engineer with 7+ years of experience (approaching 8 years) building scalable, high-performance enterprise platforms across frontend, backend, and cloud infrastructure. Proven expertise delivering production AI-enabled systems, modern web platforms, and cloud-native microservices used by thousands of users. Strong focus on performance, scalability, reliability, and developer productivity. Experienced in React, Node.js, Python FastAPI, Kubernetes, cloud-native architecture, and modern AI systems including RAG pipelines, multi-agent orchestration, and LLM-powered applications.

TECHNICAL SKILLS

Frontend: React.js, Next.js, Angular, TypeScript, JavaScript, Tailwind CSS, Material UI, Bootstrap, HTML5, CSS3, D3.js

Backend & APIs: Node.js, Express.js, Python FastAPI, REST APIs, GraphQL, API design, authentication, authorization, rate limiting, streaming APIs

AI & LLM Systems: Retrieval Augmented Generation (RAG), multi-agent orchestration (CrewAI), LLM integration, prompt engineering, embeddings pipelines, vector search, document ingestion pipelines, semantic chunking, conversational AI systems, agent lifecycle orchestration, AI evaluation & validation workflows

Models & AI Platforms: GPT family models, Claude models, embedding models, vector database integrations, AI workflow orchestration

Databases & Storage: PostgreSQL, MongoDB, Neo4j, Redis, SQL, vector databases

Cloud & DevOps: Docker, Kubernetes, GitLab CI/CD, AWS, load balancing, autoscaling, container orchestration

Observability & Reliability: Prometheus, Grafana, ELK, Loki, logging, metrics, performance monitoring

Tools & Platforms: Swagger/OpenAPI, RabbitMQ, Istio, JIRA, Confluence, Git, Agile/Scrum

PROFESSIONAL EXPERIENCE

Dell Technologies — Austin, TX

Senior Software Engineer | Jul 2021 – Present

Digital Human & AI Agent Management Platform (Current)

Architecting and developing a full-stack platform that automates creation, deployment, and lifecycle management of AI-powered digital agents, transforming a previously manual, multi-team setup process into a self-service experience that deploys agents in under one minute.

Key Contributions - Designed and built a guided self-service platform enabling users to configure, deploy, and manage AI agents with validation and opinionated defaults to prevent misconfiguration. - Developed scalable React dashboards for agent management, conversation history, analytics, document management, and role-based access control. - Implemented backend services acting as the control plane for agent lifecycle management, authentication, authorization, rate limiting, and request orchestration. - Built asynchronous knowledge ingestion pipelines where uploaded documents are parsed, semantically chunked, embedded, and stored for retrieval. - Implemented Retrieval-Augmented Generation (RAG) pipelines ensuring responses are grounded in permission-filtered contextual data. - Integrated multi-agent orchestration workflows coordinating retrieval, reasoning, validation, and response formatting for complex requests. - Enforced access control not only at API level but also during retrieval, ensuring users access only authorized knowledge sources. - Deployed platform components on Kubernetes with independent horizontal scaling, load balancing, and rolling deployments. - Implemented observability using Prometheus and Grafana to monitor latency, error rates, usage, and system health. - Designed PostgreSQL schemas supporting configuration, logs, analytics, and access control management. - Established CI/CD pipelines enabling automated testing and deployment for faster releases.

Impact & Achievements - Reduced agent provisioning time from ~1 week to under 1 minute. - Improved page load performance by 40% via frontend optimizations and API tuning. - Achieved 99.9% platform uptime post-production deployment. - Supported 10,000+ concurrent interactions per month without degradation. - Increased internal platform adoption by 70% within 3 months. - Reduced complex query response times by 60%.

Graph Analytics for Repair Data (GARD)

- Developed graph-based analytics solutions using Neo4j and GraphQL for efficient repair data exploration.
- Built React and Next.js dashboards delivering actionable insights to business teams.

- Automated deployments using Docker and Kubernetes, improving scalability and resilience.
-

Dell Technologies — Austin, TX

Software Engineering Intern | Jan 2021 – May 2021 - Built applications using the GRAND stack (GraphQL, React, Apollo, Neo4j) for dynamic graph data exploration. - Optimized GraphQL queries improving performance and user responsiveness. - Collaborated within Agile teams using GitLab and JIRA workflows.

Lowe's — Charlotte, NC

Software Engineer | Jul 2017 – Dec 2018 - Modernized enterprise e-commerce systems by migrating legacy monolithic services to microservices architecture. - Led migrations to Spring Boot and Node.js services containerized via Docker and deployed on Kubernetes. - Reduced deployment times by 40% through CI/CD automation improvements. - Improved application response time by 30% using Redis caching strategies. - Migrated customer-facing experiences to React + Node architecture improving engagement and performance.

Wipro Technologies — Bangalore, India

Software Engineer | Nov 2015 – Jun 2017 - Developed enterprise MERN stack applications supporting large concurrent user bases. - Built microservices-based quiz platform supporting 10K+ concurrent users. - Delivered HR automation portals reducing manual workflows by 50%. - Containerized services reducing environment setup time by 60%.

EDUCATION

Master of Science (M.S.), Computer Science

The University of Texas at Austin | 2019 – 2021

Bachelor of Technology, Computer Science & Engineering

Biju Patnaik University of Technology | 2011 – 2015

CERTIFICATIONS

- AWS Certified Developer – Associate (2023)
- MERN Stack Certification – Wipro Technologies
- React – The Complete Guide (Next.js & Redux) – Udemy