Sarath Kumar Dunga

Full Stack Developer

sarathk2606@gmail.com | (602) 703-1421 | LinkedIn | Portfolio | GitHub

Summary

Full Stack Developer with 3+ years of experience building scalable, cloud-based web applications and microservices using Python, Node.js, React, and .NET Core. Skilled in frontend and backend development, RESTful/GraphQL APIs, AI/ML integration, database design, and secure cloud deployments. Experienced with DevOps, CI/CD pipelines, and performance optimization, delivering reliable, user-focused software solutions.

Technical Skills

- **Programming & Frameworks**: JavaScript (ES6+), TypeScript, Python (Flask, Django, FastAPI), Node.js, Express.js, .NET Core, ASP.NET Core, C#, HTML5, CSS3, Bash, SQL (T-SQL, PL/pgSQL).
- **Frontend Development & UI/UX**: React.js, Redux, Next.js, React Router, Bootstrap, Tailwind CSS, Material UI, Chart.js; responsive and interactive dashboards, real-time data visualization.
- **Backend Development & APIs**: REST, GraphQL, WebSockets, CQRS, Microservices architecture, JWT, OAuth2, OpenID Connect, Secure API Gateways, Event-Driven Architecture (Kafka, Azure Event Hub), API versioning, unit & integration testing.
- **AI, ML & Analytics**: TensorFlow, PyTorch, Keras, scikit-learn, NLP & text analytics (OpenAI GPT, BERT, Hugging Face, spaCy, NLTK, AWS Comprehend & Titan), data preprocessing, feature engineering, predictive analytics, recommendation systems, model deployment & API integration.
- **Databases & Cloud**: PostgreSQL, MySQL, MongoDB, Firebase, MS SQL Server, Cassandra, Redis, Snowflake; AWS (EC2, S3, Lambda, Kinesis, DynamoDB, EKS, CloudWatch, CodePipeline), Azure (AKS, Azure SQL, Storage, Event Hub), GCP (GCS, Firestore, Kubernetes Engine, Firebase Hosting).
- **DevOps, Testing & Monitoring**: Docker, Docker Compose, Kubernetes, Helm, Terraform, Jenkins, GitHub Actions, GitLab CI/CD, MSBuild, NuGet, SonarQube; xUnit, NUnit, Moq, Jest, Mocha, Cypress, Postman, Swagger; Prometheus, Grafana; Security & compliance: RBAC, API encryption, GDPR, ISO 27001, SOC 2, PCI, SSL/TLS, CORS, API Rate Limiting.

Professional Experience

Full Stack Developer, *eBay*   02/2025 – Present | Remote, USA

- Engineered an AI-powered product recommendation engine using Python microservices, React frontend, and ML models, delivering personalized shopping suggestions, boosting user engagement by 12%, and enhancing overall experience across eBay's global marketplace.
- Orchestrated scalable event-driven backend services in Python with Flask and FastAPI, crafting REST APIs for order processing, payments, and inventory workflows, reducing latency by 15% and increasing reliability for millions of daily transactions on eBay platforms.
- Architected interactive React dashboards for sellers, integrating analytics and customer behavior insights, improving seller engagement by 10%, streamlining order management, and enabling real-time visibility into listings, buyer trends, and marketplace performance metrics.
- Leveraged AWS S3 and Lambda for real-time data processing, handling product feeds and user behavior streams, enabling automated notifications, monitoring anomalies, and improving processing efficiency by 18%, while maintaining secure, scalable cloud infrastructure.
- Collaborated with DevOps to orchestrate CI/CD pipelines using GitHub Actions, Docker, and Kubernetes for Python microservices, implementing automated testing with PyTest, ensuring smoother deployments, high availability, and full observability across eBay systems.

Full Stack Developer, *Reliance Jio*   09/2021 - 12/2023 | Kolkata, India

- Developed a scalable telecom web portal using React.js frontend and Node.js backend, integrating real-time subscription management and usage analytics, streamlining workflows and improving operational efficiency by approximately 15% across Kolkata user base.
- Built interactive dashboards with React and Chart.js, providing live network monitoring, fault detection alerts, and traffic analytics, enhancing system visibility and decision-making efficiency, reducing incident response time by around 12%.
- Engineered RESTful APIs and microservices architecture using Node.js with Express and MongoDB for billing, authentication, and subscription management, improving maintainability, modularity, and reducing service errors by roughly 13%.
- Optimized storage and data workflows using GCP Firestore and Cloud Storage, implementing encryption and access controls for sensitive user data, improving storage efficiency and reducing data security incidents by nearly 14%.
- Implemented comprehensive unit and integration testing with Jest and Mocha, alongside CI/CD pipelines using GitHub Actions, improving code reliability and deployment efficiency, reducing post-release bugs by approximately 16%.
- Collaborated with DevOps teams to containerize microservices using Docker and deploy on GCP Kubernetes Engine, ensuring high availability, scalability, and resilience, reducing service downtime incidents by nearly 16%.

Education

Arizona State University, Tempe, AZ, USA
Master of Science in Software Engineering   01/2024 – 12/2025
Indian Institute of Engineering Science and Technology (IIEST), Shibpur, India
Bachelor of Technology in Electronics and Telecommunication Engineering   07/2019 – 05/2023

Projects

GistiFi AI Summarizer | Chrome Extension | GitHub Link | Video Demo

- Developed a dual-mode Manifest V3 Chrome extension using Gemini 2.0, enabling multi-format content summarization and LeetCode-specific productivity features, boosting user efficiency 3X, achieving 8X growth, and preparing Version 2.0 with enhanced user upgrades.

HackASU: DevAI Manager | AI-Powered Developer Agent | DevPost | GitHub Link | Video Demo

- Designed a Custom MCP server integrating GitHub and Jira, leveraging Anthropic Claude for commit narrative synthesis, implementing anomaly detection to find workflow blockers, and optimizing real-time developer productivity through AI-powered insights and performance benchmarking.

Elastic Face Recognition Application | GitHub Link

- Built a distributed AWS Python application with auto-scaling EC2 instances for high-concurrency face recognition, achieving 96.12-second inference for 100 requests and a custom scale-in algorithm reducing latency to 0.24 seconds during traffic shifts.