

Vidya Sagar Reddy Desu

Bentonville, AR | (607) 297-8543 | d.vidya381@gmail.com | linkedin.com/in/vidyasagarreddydesu | github.com/vidya381

SUMMARY

I build backend systems that scale. I specialize in Java and Spring Boot, designing microservices architectures that handle complex workflows and high-volume transactions. I've built real-time data pipelines with Kafka for processing financial data, optimized SQL queries across relational databases, and worked with NoSQL systems like Cassandra and MongoDB. I'm solid with Docker and Kubernetes for containerization and cloud deployment on AWS and Azure. Problem-solving is what I enjoy most. I like debugging performance bottlenecks, architecting solutions around constraints, and figuring out why systems break under load. Understanding the "why" behind failures helps me build more resilient systems. I'm currently learning Go and Next.js to expand into full-stack development and round out my skillset.

EDUCATION

Binghamton University, State University of New York

Master of Science in Computer Science

Jan 2023 - Dec 2024

Sreenidhi Institute of Science and Technology, Hyderabad, India

Bachelor of Technology, Computer Science and Engineering

Aug 2013 - May 2017

TECHNICAL SKILLS

Programming Languages: Java, Python, JavaScript, TypeScript, C, C#, Go (learning)

Frameworks & Libraries: Spring Boot, Spring, Hibernate, JUnit, Mockito, NextJs (learning)

Web Technologies: HTML5, CSS3, REST APIs, GraphQL, OAuth

Databases & Data: MySQL, Cassandra, Cosmos DB, Oracle SQL, MongoDB, PostgreSQL

Cloud & DevOps: AWS, Azure, Docker, Kubernetes, Jenkins, TeamCity, CI/CD, Azure Blob Storage

Architectures & Patterns: Microservices Architecture, MVC, Design Patterns, Event-driven Architecture

Development Tools: Git, JIRA, IntelliJ IDEA, Visual Studio, Maven, SonarQube, Grafana, JMeter, Confluence

Messaging & Streaming: Kafka

PROFESSIONAL EXPERIENCE

Binghamton University

Graduate Researcher | Feb 2025 – Nov 2025

Designed and deployed a Bulls and Cows code-breaking game using Spring Boot (Java) backend with PostgreSQL and Hibernate, paired with vanilla HTML/CSS/JavaScript frontend.

- Implemented JWT authentication with BCrypt encryption, designed REST APIs for concurrent game state management
- Built achievement system with 19 unlockable badges and leaderboard with real-time rankings
- Created responsive user profiles tracking statistics (win rates, streaks, best scores) with light/dark mode toggle
- Deployed to production on Render using containerization and managed service deployment

Graduate Teaching Assistant | Aug 2023 - Dec 2024

Taught advanced database systems to 35 graduate students. Ran weekly office hours where I helped students understand SQL optimization, NoSQL databases, and AWS infrastructure through real examples and debugging sessions.

- Explained complex concepts in ways that stuck with students. Got consistent feedback that the explanations made sense
- All 35 students passed the course
- Designed hands-on AWS projects that showed students how database design actually works in production

Prime Healthcare Management Inc

IT Analyst Intern | Sep 2024 - Nov 2024

Built backend services for a Research Application Portal on Azure. Focused on designing APIs and implementing authentication for researchers to submit and manage nursing research proposals.

- Designed and implemented GraphQL APIs using Hot Chocolate (.NET C#) to handle research proposal submissions and approvals
- Built role-based access control so researchers could only access their own proposals and documents
- Integrated Azure Blob Storage to securely store research documents and approval files for nursing students
- Normalized MySQL database schema to improve query performance and data integrity
- Wrote unit tests with JUnit and Mockito to ensure backend reliability

IBM India Private Limited

Sales Workstation – Barclays Bank | Apr 2018 - Dec 2022

Worked on modernizing Barclays Bank's global trading platform. The old system was clunky and slow, so we rebuilt the backend using Java Spring Boot to handle complex workflows more efficiently.

- Rebuilt REST APIs to support modern trading workflows instead of legacy patterns
- Integrated Oracle and Cassandra with Kafka to move trading data in real-time. This cut latency by 60%
- Handled complex data transformations and state management across multiple trading services
- Built unit tests with JUnit and Mockito so we caught problems before they hit production

Software Engineer - Billing Platform | Aug 2019 – Aug 2020

Built backend APIs for AT&T's billing system that processed millions of transactions. This was a critical system handling billing for thousands of customers, so reliability and performance were everything. We had to make sure billing ran smoothly every single day.

- Designed REST APIs that stayed up 99.9% of the time even under heavy load
- Connected services with Kafka so billing updates flowed instantly across the system instead of sitting in queues
- Optimized SQL queries and monitored performance in production to keep the system fast and stable
- Coordinated with distributed teams across US and India to ship features on schedule

Software Engineer - Service Delivery Transformation | Apr 2018 – Aug 2019

Built an enterprise validation system for AT&T to automate data validation workflows. The old process was manual and slow, so we built microservices using Spring Boot and Kafka to handle validation at scale. This system processed tons of data every day without breaking.

- Designed Kafka pipelines to move validation events between services in real-time
- Implemented Drools rule engine to automate validation logic and cut processing time by 15%
- Optimized MySQL queries and schema design so the system could validate large amounts of data without slowing down
- Wrote unit tests with JUnit and Mockito to catch bugs earlier in development. Improved defect detection by 35%

ACADEMIC PROJECTS

Stock Price Forecasting

Built a machine learning pipeline to predict stock prices using historical data. Implemented ARIMA and LSTM neural networks in Python with TensorFlow and Keras to forecast financial time series.

- Collected and preprocessed historical stock data using Python ETL scripts
- Trained ARIMA and LSTM models to predict price movements
- Evaluated model accuracy and created visualizations to compare predictions against actual prices

AWS – Databricks

Designed a data warehouse on AWS to store and process large datasets. Built an ETL pipeline to transform raw data into usable formats for analysis.

- Created CloudFormation stacks to set up S3 storage and Redshift data warehouse infrastructure
- Built ETL pipeline using PySpark and Kafka for real-time data processing and loading
- Optimized data storage and reduced query response times by 50%

PERSONAL PROJECTS

MySpendo (Personal Expense Tracker) | Live: myspendo.vercel.app

Built a full-stack personal finance application using Go backend and Next.js frontend. Designed to help users track income and expenses, manage budgets, and visualize financial data through comprehensive analytics.

- Implemented JWT authentication with bcrypt password hashing and built 25+ REST API endpoints for transactions, categories, budgets, recurring rules, and analytics
- Designed normalized PostgreSQL database with 5 tables, proper foreign keys, and indexed columns for query optimization
- Created recurring transaction automation with hourly background jobs that handle edge cases like month-end dates and leap years. Used PostgreSQL advisory locks to prevent duplicate processing
- Built responsive Next.js frontend with interactive charts using Recharts, transaction filtering with pagination, and real-time budget alerts
- Deployed backend on Oracle Cloud with systemd service, frontend on Vercel, and database on Neon. Implemented security measures including rate limiting, input sanitization, and prepared statements to prevent SQL injection

Weather Insight (Weather Analytics Platform) | Live: weather-insight-ml.vercel.app

Built a weather tracking app with machine learning that detects temperature anomalies and predicts trends across multiple cities. Used FastAPI backend with React frontend to handle real-time weather data and run statistical analysis.

- Applied Z-Score analysis to flag unusual temperature spikes, Linear Regression for 7-day predictions, and K-Means clustering to group similar weather patterns across cities
- Set up hourly background jobs that automatically collect weather data for user-selected cities and store 180 days of historical readings
- Built caching layer that reduced redundant API calls by 66% using 10-minute TTL and request deduplication
- Deployed on Oracle Cloud with PM2 managing the FastAPI workers, Vercel hosting the React frontend, and Neon handling PostgreSQL database

CERTIFICATIONS AND ACHIEVEMENTS

- Microsoft Azure Fundamentals (AZ-900), April 2021
- **Open-Source Contributor:** Merged pull requests to `kubernetes/kubernetes`, `kubernetes/kubeadm`, `microsoft/pylance-release`, `oracle/opengrok`, and `springdoc/springdoc-openapi` repositories
- Udemy: Master Microservices with Java, Spring, Docker, Kubernetes
- Udemy: A Fullstack Saga – Spring Boot, Angular, Kubernetes
- IBM Best Graduate Hire Award 2019 – recognized as a top-performing new graduate