Shirisha Gujja

San Jose, CA | shirishagujja456@gmail.com | +1 (669) 360-9706 | linkedin.com/in/shirisha-gujja-85a487215

Education

San Jose State University Jan 2026 – Present
Master of Science in Software Engineering San Jose, CA

Vasavi College of Engineering Dec 2020 – May 2024
Bachelor of Engineering in Computer Science Hyderabad, India

Skills

- **Programming Languages:** Python, Java, TypeScript, JavaScript, SQL, C
- **Frameworks & Technologies:** FastAPI, Django, React.js, Next.js, Node.js, Express.js, SQLAlchemy, REST APIs, GraphQL
- **Databases, Cloud & DevOps:** PostgreSQL, MySQL, MongoDB, MS SQL Server, Redis, AWS (EC2, S3, IAM, SES), Docker
- **AI & Tools:** OpenAI GPT-4, Claude, GitHub Copilot, Pandas, NumPy, Scikit-learn, OpenCV

Experience

S&P Global Sep 2024 – Dec 2025
Software Development Engineer Hyderabad, India

- Designed and optimized MCP-based tool interactions for LLM-powered enterprise features, enabling models to retrieve structured data from GraphQL services and perform context-aware application workflows.
- Developed selective GraphQL retrieval mechanisms for LLM tool calls, reducing unnecessary data transfer and improving the efficiency and reliability of model-to-backend interactions.
- Developed frontend React UI components and reusable widgets to display AI-generated insights and backend data in a clear, user-friendly interface.
- Trained Copilot using historical support data to streamline workflows and eliminate repetitive manual tasks.
- Technologies: **React, GraphQL, MCP, LLMs, Copilot, REST APIs, Git**

S&P Global Jan 2024 – Jul 2024
Software Development Engineer Intern Hyderabad, India

- Led a proof-of-concept to optimize a critical refresh job that runs daily and was originally implemented as MS SQL Server stored procedures by re-architecting it on Amazon Redshift with PostgreSQL-based stored procedures.
- Engineered modular Python ETL scripts to orchestrate Redshift workflows, incorporating error handling, logging, retry logic, and automated data-validation checks.
- Reduced daily execution time of the refresh job from 9 hours to 1 hour 27 minutes, achieving a 79% performance improvement, and presented the solution for adoption as the standard refresh process.
- Technologies: **Python, Amazon Redshift, PostgreSQL, MS SQL Server, AWS lambda, ETL**

Projects

FitForge Website | GitHub

- Built and deployed a full-stack AI-powered fitness platform using Next.js, TypeScript, FastAPI, PostgreSQL, and Docker, enabling users to manage workouts, nutrition, body progress, and fitness goals.
- Architected six modular backend domains covering authentication, workouts, nutrition, progress tracking, profiles, and AI coaching, with dedicated database models, validation schemas, service layers, and REST API endpoints.
- Implemented secure authentication using JWT access tokens, refresh token rotation, HTTP only cookies, password hashing, session revocation, and automatic token refresh for uninterrupted user sessions.
- Integrated the OpenAI API to generate personalized workout and meal plans and provide context aware fitness coaching, with structured response validation and graceful handling of rate limits and service failures.

Meridian GitHub

- Built a full stack personal finance and investment platform using Next.js, TypeScript, FastAPI, PostgreSQL, and Redis, enabling users to manage accounts, transactions, budgets, savings goals, investments, net worth, and cash flow forecasts.
- Designed an event driven transaction processing pipeline using the transactional outbox pattern to categorize transactions, detect unusual spending, and deliver real time alerts through independent services.
- Integrated Plaid for secure bank account linking and transaction synchronization, encrypting access tokens at rest and enforcing user level authorization across sensitive financial data.
- Integrated the OpenAI API to improve merchant categorization and generate monthly spending insights from transaction data, with rule based fallbacks when the AI service is unavailable.