

ABHISHEK NANDAKUMAR

AI Software Engineer

Email: abhismail998@gmail.com

Location: Tempe, AZ, USA (open to immediate relocation)

Nationality: Indian

Mobile: +1 (602) 820-2950 (USA)

LinkedIn: <https://www.linkedin.com/in/abhin1998/>

Portfolio: <https://www.aboutabhi.com/>

GitHub: <https://github.com/Abhishekn1947>



PROFESSIONAL SUMMARY

Versatile AI Software Engineer with over four years of experience building scalable, **AI-powered web applications** while leading technical solution design for multiple clients. Expert in **Angular.js, React.js, TypeScript, Python, and AWS cloud services (RDS PostgreSQL, Lambda, API Gateway, S3)**, with strong proficiency in **integrating AI/LLM APIs**, developing responsive high-performance UIs, and **engineering dynamic, auto-scalable backends using GenAI microservices**. **Authored and responded to RFPs by crafting detailed technical architectures**, PostgreSQL-optimized database solutions, resilient recovery strategies (aligned with RPO/RTO targets), and **client-tailored proposals-including accurate pricing models based on AWS usage forecasts and cost optimization-to secure and deliver projects across diverse industries**. Skilled in Agile methodologies, cross-functional collaboration, and stakeholder engagement to drive product excellence, operational efficiency, and business impact. Open to immediate relocation.

WORK EXPERIENCE

- AI Software Engineer (GenAI, Agentic Systems, Python, Angular, AWS Bedrock)** **September 2025 – Present**
Baltimore, Maryland, USA
 - Spearheaded **agentic AI modernization using Claude LLM on AWS Bedrock** as the core reasoning engine, deploying multi-agent systems for zero-hallucination legacy code transformation and full business logic preservation across 81 COBOL programs.
 - Designed and orchestrated specialized Claude AI agents (via in-house builder) for RAG-powered tasks:** ingested COBOL sources/copybooks/BMS maps, chunked into logical units (e.g., PERFORM blocks), embedded for semantic search, stored in knowledge bases, retrieved relevant context, and augmented prompts for accurate analysis, extraction, and migration.
 - Engineered RAG pipelines with ingestion (legacy artifacts), chunking/splitting, embedding, vector storage/retrieval, and generation steps**, combined retrieved chunks with task prompts to Claude for grounded, context-aware Python code output, eliminating hallucinations in complex nested logic (PERFORM, CALL, XCTL tracking via stack-based reasoning).
 - Built automated data pipelines generating retrievable artifacts (flow diagrams, transcription files, operation mappings) for auditable, systematic GenAI-assisted migrations**, implemented dynamic runtime resolution agents for tricky COBOL patterns (e.g., XX6199XF).
 - Architected **Python FastAPI** backend (35+ REST endpoints, 12 validation modules) where **Claude agents** performed line-by-line retrieval and reasoning to preserve exact rules/error mappings, developed responsive Angular 19 frontend (18 modular components replicating terminal UX) on **AWS serverless (Lambda, API Gateway, S3, CloudFront)**.
 - Set up CI/CD **with GitHub Actions + Terraform IaC** for zero-downtime multi-account deploys (Dev/Test/Prod), mentored juniors on agentic workflows, prompt engineering, and AI-assisted development accelerated team adoption and reduced manual effort significantly.
 - Served as AI Solutions Architect for RFPs: crafted GenAI architectures with Bedrock LLMs, RAG patterns, resilient recovery (RPO/RTO), and cost-optimized AWS models using Pricing Calculator projections.**
- AI Software Engineer – (GenAI, Agentic Systems, Python, React, AWS Bedrock)** **June 2024 – September 2025**
Baltimore, Maryland, USA
Maryland Total Human-services Integrated Network (MD Think/MD Benefits)
 - Led 7-developer team as **primary AI engineer modernizing** legacy Child Support Management System (COBOL on IBM 3270/DB2) into scalable cloud-native app with GenAI features aligned to DHS goals.
 - Engineered high-performance **React.js frontend (React Query + Tailwind CSS)** for responsive UI, static deploys to **AWS S3/CloudFront** for global low-latency access integrated with real-time AI case management.
 - Designed **serverless Python FastAPI backend** on **AWS Lambda/API Gateway**, deployed **custom fine-tuned Claude LLM on Bedrock** for agentic, retrieval-augmented features, reduced user interaction latency by 20% through context-grounded reasoning.
 - Led zero-loss data migration pipeline from **DB2 to AWS RDS (MSSQL)**, **built GraphQL/REST APIs** for integration, achieving <2 hour cutover downtime with GenAI-accelerated validation.
 - Accelerated refactoring via **prompt-engineered Claude agents (Claude Max) for COBOL analysis** and automated translation to clean Python, implemented agentic workflows with retrieval from legacy context to preserve business logic.

- Optimized **MSSQL multi-table schemas (indexing/partitioning)** informed by GenAI insights; delivered 30% faster load times in case registration via efficient data pipelines and retrieval patterns.
- Collaborated on GitHub Actions CI/CD with DevOps engineers to create **Terraform (IaC)** resources as required for a one shot deployment. (branch protection/approvals) reducing deploy time 30%, leveraged CSPO to prioritize stories, drive intuitive UX, and mentor on agent orchestration/GenAI practices.

3) **AI Product Engineer – GPT/AWS**

February 2024 – June 2024

Arizona State University – EPICS

Tempe, Arizona, USA

- Developed a secure **GPT-4 assistant** for “School Fuel,” fine-tuned for resource constrained environments, integrating **AWS Lambda** triggers and **AWS S3** storage, boosting student analytics processing speed by 20% as a teaching assistant bot for educators.
- **Coded** the email-driven teaching bot to analyze educator emails with **GPT-4 API**, generating personalized insights and improvement strategies from student scores, improving recommendation relevance by 15% and aiding student progress. Delivering tailored recommendations from GPT.
- Designed **SQL pipelines** for data retrieval, optimizing with **Pandas** and **NumPy** preprocessing to cut analytics runtime by 25%, enhancing **model accuracy** for reliable educator reports.
- **Deployed** the **GPT-4 assistant** on **AWS EC2** with **Docker**, managing **AWS S3** for scalable log storage, increasing deployment efficiency by 30% across environments. **Secured** student data with **AWS IAM, VPC configurations**, achieving **FERPA** compliance

4) **Data/Security Analyst Intern**

May 2023 – August 2023

Securin Inc.

Albuquerque, New Mexico, USA

- Developed custom **Python scripts** with **Selenium** to scrape package metadata from **NPM, Go Modules**, and **PyPI**, automating extraction for 100,000+ packages with a rate limit of 2000 package’s/API hit.
- **CRON jobs** for local scheduling, Automating, reducing manual effort by 35%. Metadata in **AWS Redshift** for large data warehousing.
- Revolutionized metadata extraction by coding **Python scripts** to process **Selenium**-scraped data, implementing regex parsing, deduplication, and normalization to eliminate 20% noise.
- Compared every package against the **National Vulnerability Database (NVD)** in real time, flagging 15% more outdated dependencies, visualized findings in **Tableau** for actionable security insights. automated vulnerability scans with **OpenVAS**, increasing scan coverage by 30%.

5) **Software Developer**

November 2020 – May 2021

Info Career Pvt. Ltd.

Chennai, Tamil Nadu, India

- **Developed** multiple components of a wholesale e-commerce platform, coding **Angular** for a dynamic **front-end** and **Node.js** for a robust **back end**, handled **WebSocket’s** for real-time order updates, boosting transaction reliability by 25%.
- **Designed PostgreSQL** schemas with optimized indexing for group purchasing, cutting query times by 20% and enabling seamless discount features.
- **Integrated Stripe APIs** for payments and **JWT/OAuth2** for secure authentication. **Deployed** back-end on **AWS EC2** with **AWS S3** for front end static assets, improving load times by 15%. **Collaborated** in an **Agile** team, using **JIRA** and **Bitbucket** to streamline sprints and version control.
- **Collaborated with Product manager and UI/UX Engineers** to deliver a smooth experience for users, making users tasks are aligned with the Business Objectives and KPI’s

OVERALL TECHNICAL SKILLS

Category	Skills
AI / GenAI / Machine Learning	Large Language Models (LLMs: Claude via AWS Bedrock) Retrieval-Augmented Generation (RAG) Pipelines Agentic AI / Multi-Agent Systems & Orchestration Prompt Engineering & Zero-Hallucination Reasoning LLM Integration, Fine-Tuning & Production Deployment (Bedrock) GenAI Microservices & Real-Time AI Features Data Pipelines for AI (Ingestion, Chunking, Embedding, Retrieval)
Programming Languages	JavaScript, TypeScript, Python, SQL
Frameworks / Libraries	React.js, Next.js, Tailwind CSS, React Query, LangChain / LangChain.js, Node.js, Angular.js (Angular 19), FastAPI

Cloud	AWS Services (RDS PostgreSQL, S3, EC2, IAM, Lambda, API Gateway, VPC, EBS, CloudFront) AWS Bedrock (LLM Hosting & Deployment)
Databases	Microsoft SQL Server (MSSQL), RDS PostgreSQL Vector Storage & Embeddings (RAG Knowledge Bases)
API Technologies	REST, FastAPI, GraphQL
DevOps / CI/CD	GitHub Actions, Terraform (IaC), Jenkins
Security	JWT, OAuth2, AWS IAM
Methodologies	Agile, Scrum (CSPO Certified), Test-Driven Development (TDD)

EDUCATION

- Master of Science - Information Technology** **January 2022 - December 2023**
Ira. A Fulton Schools of Engineering, Arizona State University | Tempe, AZ, USA **GPA:4.0/4.0**

Course Work: Big Data Analytics, Machine Learning, Data Visualization, Advanced Database Management Systems (ADBMS), Artificial Intelligence in Cybersecurity, Natural Language Processing (NLP), Fine-Tuning of Large Language Models (LLMs), Gen AI.

- Bachelor of Engineering - Computer Science** **August 2016 - August 2020**
Sri Venkateswara College of Engineering, Anna University | Chennai, TN, India **GPA:3.0 /4.0**

Course Work: Java, Python, and C++, with expertise in advanced database management systems (DBMS), SQL, and software architecture design. Skilled in version control systems, operating systems, network protocols, and distributed systems.

PROJECTS / POC's

- AI-Powered Stock Market Prediction Platform – [Clarividex](#)** **January 2026**
 - Built a real-time data aggregation engine that pulls from 12+ financial APIs (SEC EDGAR, Yahoo Finance, Finviz, options flow) and processes 250+ data points per prediction using **Python/Fast API** backend with async data pipelines.
 - Developed a weighted probability scoring algorithm with transparent reasoning, combining technical indicators (RSI, MACD, Bollinger Bands), sentiment analysis, and options flow data into quantified predictions with confidence intervals. Used an in-house Prediction engine to run multiple algorithms for accurate prediction percentage.
 - Served via **Next.js/React frontend** with responsive UI.
- Cloud Elevate – Infrastructure as Code (IaC) Automation with Terraform** **January 2025**
 - Built an automated cloud deployment platform using **Terraform** and **AWS** microservices, Simplifies infrastructure setup across **AWS using Terraform** modules. Mainly used as an in-house tool to automate and deploy pipelines and resources using terraform
 - Enables scalable and secure deployments for diverse applications. Reducing deployment time, number of AWS Cloud Engineers reduced as well to run a lean team.
- Fine-tuned Cyber LLM** **September 2024**
 - Developed a custom **T5-based LLM** (later upgraded to Gemini 1.5) fine-tuned on NIST data, UI created using **Streamlit**.
 - Provides cybersecurity and Audit teams with tailored risk management guidance, enhancing decision-making with **torch** and **transformers** to produce Data Breaches, Logs, Potential Vulnerabilities, Audit Logging via our AI model.
- Sentiment Analysis of Amazon Product Reviews** **March 2024**
 - Performed sentiment analysis of Amazon product reviews using NLP tools (NLTK, VADER), achieving 80.4% classification accuracy.
 - Utilized Python libraries (Pandas) for data processing and Seaborn for visualization, enhancing data analysis accuracy.
 - Demo of a strong correlation between sentiment scores/product ratings, offering quantifiable insights into consumer opinions.

CERTIFICATIONS

- AWS Certified Cloud Practitioner (CLF-C02) - [Validation URL](#)**
- DASM (Disciplined Agile Scrum Master) Certification**
- Salesforce Certified Administrator – (Cred ID: 4544151)**
- Certified Scrum Product Owner (CSPO)**