Bhaskar Boruah

boruah.bhaskar@gmail.com | https://www.linkedin.com/in/bhaskar-boruah-b6141218/ |
+1 210 305 6977 | Texas |GitHub - https://github.com/boruahbhaskar

Professional Summary

AI and Cloud Solution Architect with 16+ years of experience building enterprise-scale applications, including 8+ years in financial services (USAA, Bank of America) delivering mission-critical systems for risk management, compliance, and customer experience. Deep hands-on expertise in Python, Java, AWS cloud infrastructure, and full-stack development with specialized skills in ML/AI solutions. Proven ability to rapidly adapt across diverse technology stacks—from building data pipelines and microservices to implementing generative AI platforms. Strong track record delivering production systems at scale using modern practices: CI/CD automation, infrastructure-as-code (Terraform), containerization (Docker/Kubernetes), and Agile methodologies. Proficient in AI-assisted development tools (GitHub Copilot, Gemini) for accelerated delivery.

Technical Skills

- **Programming Languages**: Python, SQL, Java, PL/SQL, JavaScript, TypeScript, Angular JS, Shell Scripting
- **Cloud Platforms & Services**: AWS (EC2, S3, Lambda, SageMaker, Amazon Bedrock,RDS, API Gateway, CloudWatch, IAM, EKS)
- **Infrastructure & DevOps**: Docker, Kubernetes, Jenkins, GitHub Actions, CI/CD Pipelines, Terraform
- **Databases & Processing**: PostgreSQL, MySQL, Oracle, Redshift, Snowflake, Apache Spark, Kafka
- **Backend Development:** RESTful APIs, Microservices, FastAPI, Spring Boot, Event-Driven Architecture
- **Data Engineering & Pipelines:** ETL/ELT, Kafka, AWS Glue, Data Lakes, CDC (Change Data Capture), Data Validation Frameworks
- **Machine Learning & AI**: PyTorch, TensorFlow, Scikit-Learn, RAG Systems, Multi-LLM Architectures, MLOps, NLP, LangChain, Model Routing and Lifecycle Management
- **APIs & Microservices:** RESTful APIs, Microservices Architecture, API Gateway
- **Development Tools:** VS Code, Git, GitHub Copilot, Claude Code, Jira, Confluence
- **Methodologies:** Agile/Scrum, DevOps, Test-Driven Development, Infrastructure as Code

Professional Experience

Technical Consultant/Technical Solution Architect, Toyota Motors North America, Austin, TX (Dec 2024 – Current)

- **Enterprise AI Platform - Toyota GPT | Multi-LLM Orchestration & MCP Design**
  - Architected and delivered production-grade conversational AI platform (Toyota GPT) serving 15,000+ employees, implementing a Model Context Protocol (MCP) orchestration layer that routes queries intelligently across Amazon Bedrock (Claude), Azure OpenAI GPT-4, and Google Gemini — optimizing for cost, latency, and query-type with SSE streaming to a React SPA via CloudFront CDN.
  - Designed model lifecycle management framework covering provider onboarding, versioning, fallback routing, and deprecation — enabling zero-downtime model

upgrades across the multi-LLM fleet; established AI safety guidelines and responsible AI coding standards adopted platform-wide.
    - Built multi-tier RAG pipeline using Amazon OpenSearch (vector search), Airflow-triggered ECS ingestion, SharePoint connectors, and OpenAI embeddings; polyglot data strategy — DynamoDB, RDS , ElastiCache/Redis (context caching), S3 (document store) — driving measurable reduction in knowledge-retrieval latency.
    - Solutioned multi-account AWS architecture (VPC, ECS, ALB, API Gateway, Direct Connect) with defense-in-depth security: Global + Regional WAF, AWS Shield Advanced, Lambda Edge token validation, Azure AD / TMNA SSO, least-privilege IAM, and TLS 1.2 across all boundaries.
    - Built LLM evaluation and monitoring framework tracking output quality, performance SLAs, and model drift across providers via CloudWatch and custom dashboards — enabling data-driven decisions on model selection and cost optimization.
    - Drove CI/CD pipeline using GitHub Actions, Terraform, Docker, and Kubernetes/EKS for automated multi-environment deployments; platform scaled to 10,000+ concurrent users; secrets managed via AWS Secrets Manager and KMS encryption at rest and in transit.
    - Technologies: Python, FastAPI, TypeScript, AWS (ECS, EKS, Lambda, API Gateway, OpenSearch, ElastiCache, DynamoDB, RDS, S3, CloudWatch, KMS), Amazon Bedrock, Azure OpenAI, Google Gemini, Docker, Kubernetes, GitHub Actions, Terraform
- **Agentic AI Data Discovery Platform | NLP-Powered Enterprise Data Catalog**
    - Designed and built agentic AI orchestration layer using Python and AWS Bedrock enabling natural language queries against enterprise data lake metadata — agents autonomously invoke tools (schema lookup, lineage traversal, sample data fetch) through a structured MCP-compatible tool-use protocol.
    - Developed microservices architecture on AWS EKS (Docker containerization) for dataset recommendation engine; implemented RESTful APIs in Python (FastAPI) with JWT auth, rate limiting, and comprehensive error handling.
    - Integrated AWS Bedrock for generative AI capabilities with fine-grained IAM access control; built automated pytest suite with 85% code coverage integrated into Jenkins CI/CD pipeline for continuous quality assurance.
    - Drove AI adoption across 3 engineering squads by embedding GitHub Copilot into standard SDLC workflows; tracked developer productivity metrics (PR velocity, review cycle time, boilerplate reduction) — delivering measurable efficiency improvements reported to engineering leadership.
    - Built visualization connectors for Power BI and Tableau enabling seamless chart generation from conversational queries.
    - Technologies: Python, FastAPI, AWS (EKS, Lambda, Bedrock, S3), Docker, Kubernetes, Jenkins, pytest, GitHub Copilot, Git
- **SAP Data Pipeline to AWS – Full-Stack Data Engineering (Python, AWS, Terraform)**
    - Engineered end-to-end CDC data pipeline in Python replicating 5M+ daily records from SAP Datasphere to AWS S3 data lake, replacing legacy infrastructure
    - Built secure cross-account data transfer service using Python, AWS Transfer Family (SFTP), Lambda, and IAM role assumptions following security best practices

- Developed Infrastructure as Code using Terraform and Terragrunt for provisioning complete AWS stack (S3, Lambda, EKS, IAM, CloudWatch) across multiple environments
- Implemented automated deployment pipeline using Jenkins with multi-stage validation (unit tests, integration tests, security scans)
- Created Python-based monitoring framework with Lambda functions for real-time health checks, alerting, and automated recovery procedures
- Built data quality validation framework using Python and SQL for anomaly detection and data integrity verification
- Technologies: Python, AWS (S3, Lambda, Transfer Family, EKS, IAM), Terraform, Jenkins, SQL, Kafka, VS Code

- **Tax Engine Migration & Integration (Java, Python, Microservices)**
  - Migrated legacy tax system to Vertex Cloud On-Demand using Java and Python with zero downtime cutover
  - Built real-time tax calculation microservices integrated with SAP ERP using RESTful APIs deployed on AWS EKS
  - Implemented resilient messaging using Kafka with retry logic, dead-letter queues, and circuit breakers for guaranteed delivery
  - Developed reconciliation framework in Python with SQL validating 1M+ daily tax calculations against source systems
  - Technologies: Java, Python, Spring Boot, AWS (EKS, Lambda), Kafka, SQL, Docker, Kubernetes
- **SAP Integration Platform Migration – 600+ Interfaces (Java, Python, Kafka)**
  - Led technical implementation migrating 600+ SAP PI/PO interfaces to SAP BTP Integration Suite using Java and Python
  - Developed integration adapters in Java connecting SAP ERP with Ariba, Taulia, and 50+ external systems via RESTful APIs
  - Implemented event-driven architecture using Kafka for real-time data synchronization between SAP and enterprise systems
  - Created CI/CD pipeline using Jenkins and Terraform for automated interface deployment with rollback capabilities
  - Developed monitoring dashboard using Python, CloudWatch, and custom metrics for interface health tracking and SLA compliance
  - Technologies: Java, Python, SAP BTP, Kafka, Jenkins, Docker, Terraform, RESTful APIs, SQL

Technical Solution Architect/Technical Consultant, Estuate Inc, Austin, TX (Apr 2023 – Oct 2024)

Project: Automated IT Risk and Control Testing Assessment Platform

Key Client: Autodesk, Expedia

- Designed and implemented GRC solutions using Python and Java for enterprise risk management, compliance monitoring, and audit management
- Built integration pipelines connecting GRC platforms with ERP, CRM, and financial systems using RESTful APIs, Python, and Kafka
- Developed automation scripts in Python for compliance reporting, policy lifecycle management, and real-time risk dashboards
- Led Agile development teams managing sprints in Jira and implemented CI/CD pipelines using Jenkins and Docker

- Created data transformation services using Python and SQL for cross-system data synchronization and reporting
- Technologies: Python, Java, SQL, Kafka, Jenkins, Docker, AWS (Lambda, RDS, S3), RESTful APIs

Technical Consultant/Principal Technical Architect - GRC Solutions, MetricStream Inc, San Antonio, TX (Jan 2016 – Apr 2023)

Project: GRC Journey - Integrated GRC Solution with AI and ML Enhancements

Client: United States Automobile Association - USAA (Jan 2016 – Apr 2023)

- **Cloud Migration & Platform Modernization (AWS, Terraform, Python)**
  - Led AWS cloud migration of enterprise GRC platform using Terraform and Python automation achieving 30% operational cost reduction
  - Built real-time risk monitoring pipeline using Python, Kafka, and AWS Lambda processing 100K+ events/day with 65% improvement in incident response
  - Implemented comprehensive CI/CD pipeline using Jenkins and Terraform for automated multi-environment deployment
  - Developed AWS infrastructure using Python (Boto3) and Terraform for EC2 auto-scaling, RDS failover, and S3 lifecycle management
  - Technologies: Python, Java, AWS (EKS, Lambda, RDS, S3, EC2), Terraform, Jenkins, Docker, Kubernetes, Kafka
- **Data Integration & API Development (Python, Java, SQL)**
  - Implemented data validation framework in Python ensuring data quality across 50+ integrated systems
  - Created reporting APIs using Python (Flask) serving 5000+ business users with sub-second query performance
  - Built batch ETL processes using Python and SQL for nightly data synchronization and reconciliation
  - Technologies: Python, Java, Spring Boot, Flask, SQL, Oracle, PostgreSQL, RESTful APIs, SOAP Web Services

Technical Lead/Senior Member Technical Staff, MetricStream, Bangalore, India (Aug 2009 – Jan 2016)

- Developed GRC solutions for Bank of America using Java, Python, and SQL serving 10,000+ users across risk and compliance functions
- Implemented web applications using Java (Spring MVC), JavaScript, and Oracle database with responsive UI design
- Developed reporting modules using SQL and Python generating regulatory compliance reports for banking audits
- Technologies: Java, Python, SQL, Oracle, Spring Framework, RESTful APIs, JavaScript, HTML/CSS

ML Projects Portfolio(Personal Learning)

- **RayScale ML – Distributed Machine Learning Platform –** GitHub Link
- **BinSense AI - Automated Bin Item Verification System -** GitHub Link
- **AI-Assisted Learning for NVIDIA SDKs and Toolkits -** GitHub Link
- **Graph Analysis using Machine Learning for Clustering Ego Graphs -** GitHub Link

Education

- **Master of Computer Application (MCA)**, Assam Engineering College, Gauhati University (2009)
- **Bachelor of Science (Physics Hons)**, Jorhat Science College, Dibrugarh University (2002-2005)

Certifications

- AWS Certified Solution Architect – Associate (Oct 2022)