



# Khalid Rizvi

## Generative AI, Cloud-Native & DevOps Specialist

khalid.rizvi@icloud.com

Vienna, VA - 703-656-6394

[linkedin.com/in/khalidrizvi](https://www.linkedin.com/in/khalidrizvi)

[khalidrizvi.com](https://khalidrizvi.com)

Innovative technology leader with 30+ years of experience architecting and delivering cloud-native, AI-powered solutions for global enterprises and the public sector. Specializing in generative AI, large-scale system modernization, and secure, scalable infrastructure across AWS and Azure. Highly skilled in designing high-performance backends with Go and Java, and integrating emerging technologies such as LLMs and GenAI into legacy and greenfield environments. Recognized for aligning technical strategy with business goals, mentoring high-performing teams, and leading digital transformations that accelerate innovation and operational excellence. A hands-on architect with a strong engineering mindset, continuously learning and pushing boundaries to deliver future-ready, intelligent solutions that drive measurable value.

## LEADERSHIP & STRATEGIC VALUE

### Enterprise Architecture Leadership:

*Led major digital transformations across industries by aligning architecture with business goals to deliver scalable, high-impact solutions.*

### Innovation and Digital Transformation:

*Modernized legacy systems into AI-driven, cloud-native solutions—cut mainframe processing by 99% and delivered 99.9% uptime on critical platforms.*

### Emerging Technology Integration:

*Delivered AI/ML, GenAI, and big data solutions to solve complex challenges—built award-winning modernization tools and AI-powered analytics to drive competitive edge.*

### Team Leadership and Mentorship:

*Built and mentored high-performing teams—fostering innovation, growth, and a culture of excellence through hands-on leadership and collaboration.*

### Strategic and Crisis Management:

*Proven ability to lead high-pressure projects—delivered urgent solutions like COVID-19 systems and federal cloud migrations with speed, quality, and security.*

## EXPERIENCE

*Solution architect overseeing cloud innovation, AI/ML solution delivery, and digital transformation initiatives for major clients in telecom, finance, public sector, and transportation.*

- **Nokia Networks – KPI Analytics & AI Platform (Dec 2024 – Jun 2025):** Spearheaded the design of a next-generation network KPI analytics platform for a telecom client, leveraging intelligent automation and AI-driven insights to transform performance monitoring. Architected a high-throughput data processing pipeline in Go, deployed on Azure Container Apps (ACA) using advanced Terraform automation (including Azure AQL packaging). Achieved a **95%** reduction in manual data handling and ensured near real-time availability of critical network metrics. Integrated advanced analytics features, including dynamic Power BI dashboards and a conversational (chatbot) AI interface (LangChain and OpenAI GPT) for natural language data queries. These innovations improved data reliability by **40%**, sped up information retrieval by **60%**, and cut stakeholder decision-making time from days to minutes.
- **Network of Giving – Real-Time Donations Microservices (2023 – 2024):** Architected and delivered an enterprise-grade charitable donations platform that processes millions of transactions in real time with **99.9%** uptime. Designed a multi-tier security architecture (OAuth 2.0, mTLS, JWT) to meet stringent financial industry standards, ensuring end-to-end encryption and compliance while maintaining sub-second transaction response times. Built a fault-tolerant, event-driven microservices backend on AWS (Java/Spring Boot and Go on ECS Fargate, integrated with RabbitMQ and DynamoDB) capable of seamlessly handling traffic surges (300% spike) during peak campaigns.
  - Leveraged AWS Bedrock (Claude v2 and Amazon Titan) to integrate generative AI features for personalized donor experiences and automated campaign operations:
    - **Donor-Cause Matching Engine:** Employed Claude v2 to analyze donor histories and generate tailored charity recommendations, significantly improving donor conversion and retention.
    - **Donor Assistant Chatbot:** Deployed a natural language assistant (using Claude v2) to help donors explore causes, retrieve giving summaries, and receive campaign updates with safe, conversational AI responses.
    - **Automated Content Generation:** Enabled nonprofits to auto-generate campaign descriptions and donor outreach content using Amazon Titan (Text G1) from brief prompts, greatly reducing manual workload for fundraising teams.
  - Established comprehensive monitoring and alerting with Amazon CloudWatch, cutting incident resolution time by **70%**. Mentored the engineering team in cloud-native best practices and DevOps, improving development velocity by **50%** and instilling a culture of reliability and performance.
- **Grants.gov Modernization – Cloud Migration & Security (2022 – 2023):** Led a full modernization of a legacy federal grants management system, migrating two decades of data and applications from on-premises Sun servers to a scalable AWS cloud environment. Implemented infrastructure as code using **AWS CloudFormation** (later refactored to Terraform) to automate provisioning of VPCs, subnets, IAM roles, and S3 storage. Employed **Ansible** for OS-level configuration and AMI hardening, establishing a robust hybrid IaC approach. Automated end-to-end deployment and data onboarding with custom Bash scripts orchestrating resource setup, application build, and database initialization—eliminating manual errors and speeding up environment provisioning. Re-architected the monolithic application

into secure microservices, introducing **JWT-based authentication** and an **API Gateway** with Lambda authorizer for centralized access control. Defined serverless components using AWS SAM templates for versioned, repeatable deployments. Enabled cross-environment connectivity via VPC peering and fine-grained IAM policies, ensuring minimal latency and maximum security between components. Introduced centralized logging and tracing (CloudWatch Logs, AWS X-Ray), improving system resiliency by **45%** and providing real-time performance insights for a platform handling millions of grant applications annually.

- **Los Angeles Metro – Transit Systems Integration (2020 – 2022):** Developed an advanced integration solution for LA Metro to consolidate disparate transit and fleet management systems using Dell Boomi and custom Java services. Integrated GPS telematics data from 2,000+ buses and trains into a unified fleet monitoring platform, enabling real-time tracking and route optimization that reduced fuel consumption by 15% (approximately **\$2M** in annual savings). Linked maintenance depots, inventory systems, and external logistics partners to establish end-to-end supply chain visibility, cutting parts shortages and delivery delays by **30%**. Unified scheduling, ticketing, and live vehicle location data to power rider-facing applications with up-to-the-minute transit information—boosting mobile app engagement by **25%** and markedly improving the passenger experience. Implemented predictive analytics on vehicle sensor data for proactive maintenance, reducing equipment downtime and service disruptions.
- **NY Metropolitan Transportation Authority – Enterprise Architecture & Systems Modernization (2018 – 2020):** Served as Enterprise Architect across a portfolio of modernization initiatives for NYMTA, providing architectural leadership in data integration, infrastructure automation, and asset management systems.
  - **Electric Bus Predictive Maintenance & Energy Optimization:** Architected a real-time analytics platform for NYMTA's electric bus fleet to predict maintenance needs and optimize energy usage. Ingested high-volume charging station and telemetry data (from 500+ buses) into the SPEAR analytics system. Trained regression models in AWS SageMaker to predict battery degradation and schedule preventive maintenance, using historical charging and environmental data. Employed XGBoost for enhanced interpretability and validated models with MAE/RMSE metrics. Deployed SageMaker endpoints to serve real-time predictions to a dashboard, resulting in a **20%** reduction in battery failures and a **35%** improvement in maintenance scheduling accuracy, aligning with sustainability goals.
  - **COVID-19 Sanitation Compliance Platform:** Designed and delivered an enterprise application to automate pandemic-related sanitation workflows for Metro-North Railroad. Transformed train consist data into scheduled cleaning work orders in Infor EAM using Java and Apache Camel for high-throughput, parallel processing. Ensured end-to-end auditability and real-time status tracking with persistent data storage and automated notifications. This platform ensured 100% regulatory compliance, eliminated manual data entry, and provided complete visibility into sanitation operations across the fleet.
  - **Asset Management Integration Framework:** Led integration of Infor EAM with Bentley AssetWise to unify asset lifecycle data. Architected a microservices solution on Red Hat OpenShift with Redis for caching and custom schema mediation services to translate between MIMOSA and OAGIS data standards. The unified integration fabric combined asset condition monitoring, maintenance history, and IoT sensor data, laying the groundwork for predictive maintenance across NYMTA's infrastructure. This initiative reduced asset downtime by **40%** and improved data quality for long-term capital planning.

- **Mainframe Batch Modernization – Daito Corp (2017 – 2018):** Transformed a critical batch processing system for a Japanese insurance client, migrating legacy COBOL and CL programs on AS/400 to a modern Java Spring Boot framework. Optimized an overnight batch job (formerly 23 hours) to run in just **13 minutes** (a **99%** improvement) by streamlining legacy logic and introducing caching (Redis) and efficient state management patterns. Led the end-to-end migration and data validation, delivering a cloud-ready solution that preserved all business logic and sequencing while dramatically improving performance and maintainability.
- **Mastercard – Global Loyalty Platform Integration (2016 – 2017):** Led the redesign of Mastercard’s Global Customer Loyalty platform to support rapid integration of new partner programs and acquisitions. Architected a flexible, metadata-driven integration engine allowing dynamic adaptation to each partner’s data structures and business rules without extensive recoding. This capability enabled faster onboarding of loyalty partners and was central to Mastercard’s growth strategy. Rebuilt core components using Spring Boot microservices, improving system throughput and batch processing efficiency by **60%**. The enhanced platform provided seamless interoperability with external partner systems via REST APIs and scalable batch processes, directly contributing to revenue growth through expanded loyalty offerings.

## Senior Solutions Architect

11/2006 - 07/2015

*RIZ Consulting*

- Led high-impact federal IT modernization projects, designing and delivering advanced solutions for agencies like EPA, VA, USDA, and Grants.gov.
- Engineered automated rules engines and modern web applications that increased operational efficiency by up to 65% and improved data accuracy, compliance, and user satisfaction for millions of users.

## Senior Architect

07/1997 - 11/2006

*CSC Consulting*

- Led multi-million dollar software modernization and integration projects for Fortune 500 clients, accelerating delivery by developing reusable frameworks and automation tools.
- Architected secure, high-availability platforms in finance, manufacturing, education, and healthcare, reducing development timelines by 70% and ensuring zero production downtime.

## EDUCATION

### Master of Science in Computer Science

01/1995 - 07/1997

*California State University, Sacramento*

### Master of Science in Mechanical Engineering

01/1986 - 01/1988

*NED University of Engineering & Technology*

### Bachelor of Science in Mechanical Engineering

01/1982 - 01/1986

*NED University of Engineering & Technology*

## MODERN TECHNICAL SKILLS

### Programming & Scripting:

*Python, Go, Java, Javascript, Typescript, Bash, SQL*

### Cloud Platforms:

AWS (EC2, S3, RDS, Lambda, API Gateway, SageMaker, Bedrock, Amazon Q), Azure (Azure ML, Container Apps), Google Cloud

**Machine Learning & AI:**

Skilled in building and tuning machine learning models like regression, clustering, neural networks, and NLP using transformers such as BERT and RoBERTa.

**Generative AI Frameworks & Tools:**

Experienced with tools like NeMo, Whisper, LangChain, and Streamlit to build secure and interactive generative AI apps.

**Frameworks & Libraries:**

TensorFlow, PyTorch, Keras, scikit-learn, Pandas, NumPy

**Big Data & Analytics:**

Apache Spark, Apache Kafka – for large-scale data processing; Tableau, Power BI – for data visualization and dashboards

**Databases:**

SQL (MySQL, PostgreSQL, Oracle) and NoSQL (MongoDB, DynamoDB, Elasticsearch, Redis)

**DevOps & MLOps:**

Docker, Kubernetes, Jenkins, CI/CD pipelines; Infrastructure as Code (AWS CloudFormation, Terraform, AWS SAM); Configuration management (Ansible). Skilled in end-to-end MLOps lifecycle management from model experimentation and versioning to production deployment and monitoring.

**Security & Compliance:**

OAuth 2.0, mTLS encryption, JWT authentication

CERTIFICATIONS

MCP Servers Made Easy with Python and OpenAI Agents - 07/2025

AWS Certified Cloud Practitioner - 08/2024