Mihir Patel

Lead / Staff Software Engineer

fullstackfusions@gmail.com | linkedin.com/in/fullstackfusions | github.com/fullstackfusions | fullstackfusions.com

PROFESSIONAL SUMMARY

Staff Software Engineer with 7+ years building Python backends, microservices, and AI-driven automation at scale. I design secure APIs and distributed systems, engineer high-performance data pipelines with Kafka and Airflow, and integrate LLMs — RAG, agents, and function calling — into enterprise workflows that measurably cut costs and accelerate decisions. I lead cross-functional delivery and translate business requirements into scalable, maintainable architectures.

PROFESSIONAL EXPERIENCE

Lead / Staff Engineer — RBC
May 2025 – Present

Leading network automation initiatives to transform manual validation and certification processes into intelligent, scalable platforms.

- **Network Validation Automation:** Spearheaded development of pre/post change validation system for Cisco NX-OS, IOS and IOS-XE devices, reducing validation time from weeks to 10 minutes while preventing "quiet failures" often missed by manual processes.
- **Device Certification Platform:** Built automated certification workflow that parses PDF requirements documents using RAG techniques, generates pytest suites, and produces versioned Allure reports — eliminating contractor costs and accelerating onboarding from weeks to minutes.
- **Topology & Impact Analysis:** Extended network visibility capabilities with scheduled data collection and Neo4j graph-backed mapping to support real-time impact analysis during change management and incident response.
- **Enterprise Integration:** Embedded platforms into change management workflows, collaborated with network engineering teams to drive adoption, and delivered board-ready compliance reports with full audit traceability.

Senior Developer — RBC
Aug 2024 – May 2025

Architected and delivered enterprise data platform MVP, establishing scalable ETL pipelines and real-time analytics capabilities.

- **Data Platform MVP:** Spearheaded development of comprehensive ETL pipelines extracting and transforming data from Grafana, PostgreSQL, Elasticsearch, and Dynatrace into analytics-friendly S3 data lake, enabling enterprise-wide data insights.
- **Streaming & Orchestration:** Implemented streaming solutions using Kafka and automated data workflows with Airflow, ensuring timely and reliable data delivery to downstream consumers with 99.9% uptime.
- **Microservices Architecture:** Designed scalable microservices using FastAPI and deployed on OpenShift/Kubernetes (EKS) with comprehensive observability, monitoring through Grafana dashboards, and operational guardrails.
- **Real-time Systems:** Developed event-driven systems using webhooks and WebSockets, optimizing responsiveness for enterprise applications and improving user experience across internal tools.
- **Data Governance & Leadership:** Researched and implemented AWS Glue Data Catalog and Lake Formation for data governance, mentored cross-functional teams, and influenced architecture decisions through code reviews and best practices.

Software Engineer — RBC
Jun 2023 – May 2024

Transformed monolithic chatbot architecture into intelligent microservices platform with advanced NLP capabilities and enterprise-grade scalability.

- **AI-Powered NLP Integration:** Spearheaded integration of advanced NLP models to automate chatbot responses, significantly improving response time and accuracy while handling complex domain-specific queries.
- **Microservices Transformation:** Successfully decomposed monolithic chatbot application into 3 scalable microservices, implementing Kafka-backed messaging for reliable communication between multiple frontends and backend services.
- **RAG Pipeline Development:** Built sophisticated Retrieval-Augmented Generation pipeline using LangChain framework and Hugging Face models, merging LLM knowledge with internal data sources to create a more versatile and intelligent chatbot.

- **Performance Optimization:** Introduced effective MongoDB caching mechanisms, significantly reducing response times by serving cached responses and optimizing overall system performance.
- **Cloud-Native Deployment:** Containerized applications using Docker, orchestrated on Kubernetes with autoscaling and health-focused rollouts, and streamlined CI/CD pipelines for fast, consistent, low-risk deployments.

Senior Python Developer — CPPIB
Jul 2022 – May 2023

Developed complex capital markets backend systems and led cloud migration initiatives to modernize trading infrastructure.

- **Capital Markets Backend:** Developed complex backend systems for capital markets using Python frameworks (Flask, Django, SQLAlchemy), designed RESTful APIs and middleware microservices for seamless integration with trading systems.
- **Cloud Migration Leadership:** Led comprehensive cloud database migration from physical servers to Data Fabric, optimizing performance using AWS SDK tools like Boto3 and modernizing data access patterns for enhanced reliability and performance.
- **Full-Stack Development:** Implemented complete development lifecycle from code changes to production deployment, including testing, building pipelines, and performance testing on pre-prod environments before client-facing releases.
- **Enterprise Deployment:** Deployed applications using AWS ECS/EC2 and IBM UrbanCode, managing microservice architectures with focus on high availability, resilience, and smooth rollouts.
- **Automation & Reporting:** Automated critical workflows using Python and JAMS scheduler tools to generate daily/weekly reports, reducing manual effort and improving operational efficiency.

Full Stack Developer — CGI
Mar 2021 – Jul 2022

Built comprehensive data processing pipelines and full-stack applications while implementing infrastructure as code and cloud-native CI/CD practices.

- **Workflow Automation:** Automated various manual workflows using Python scripts and Unix shell scripting, developed change data capture scripts for processing incremental records between new and existing relational database data.
- **Infrastructure as Code:** Implemented Terraform to map complex dependencies, identify network issues, and leverage key features like Infrastructure as Code, execution plans, resource graphs, and change automation.
- **Full-Stack Dashboard Development:** Developed and tested comprehensive dashboard features using Python, React, Bootstrap, CSS, JavaScript, and jQuery, with expertise in scientific computing stack (NumPy, SciPy, pandas, matplotlib).
- **Cloud Data Pipelines:** Built reliable AWS data pipelines using Glue ETL jobs triggered by S3 events, Lambda functions, Step Functions, and scheduled Airflow jobs for data loading and time series manipulation.
- **DevOps & Reporting:** Implemented CI/CD for microservices deployment in Kubernetes clusters on AWS Cloud, generated capacity planning reports using Python visualization packages, and deployed comprehensive AWS stacks including EC2, S3, RDS, and Load Balancing services.

Python Developer — Kroger
Aug 2019 – Feb 2021

Developed patient data management systems with multi-database architecture and established robust CI/CD infrastructure on AWS cloud platform.

- **Healthcare Data Systems:** Implemented comprehensive data tables using PyQt to add, delete, update, and display patient records and policy information, with modules to connect and monitor Apache Cassandra instance status.
- **Multi-Database Architecture:** Designed hybrid data storage solution using Apache CouchDB (NoSQL) on AWS Linux instances in parallel with RDS MySQL, improving data security and report generation efficiency through strategic caching and data reuse.
- **CI/CD Pipeline Implementation:** Built comprehensive end-to-end CI/CD pipelines in Jenkins to retrieve code, compile applications, perform automated tests, and push build artifacts to Nexus repository for continuous integration and delivery.
- **AWS Infrastructure Automation:** Managed and automated all aspects of AWS infrastructure including compute, storage, network, permissions, and cost optimization using configuration management tools like CloudFormation and custom shell scripts.

TECHNICAL SKILLS

Programming & Frameworks	Python: FastAPI, Flask, Django, JavaScript/TypeScript: React, Next.js, Go: Gin, Echo, Fiber
APIs & Messaging	Webhooks, WebSocket, REST, GraphQL, gRPC
Databases	PostgreSQL, MongoDB, Elasticsearch, Redis, SQL Server, Cassandra

Data Engineering	Kafka, Airflow, AWS S3, ETL Pipelines, AWS Glue
GenAI Features	Agentic RAG, fine-tuning, function/tool calling, prompt chaining, orchestration, and re-ranking
Platform & Environment Span	On-prem, cloud-native, hybrid-cloud, multi-cloud, containerized workloads, Kubernetes/OpenShift
DevOps & Cloud	Docker, Kubernetes, Jenkins, GitHub Actions, OpenShift, AWS
Network & Observability	Grafana, Dynatrace, NetBrain, ExtraHop, Corvil
Leadership & Delivery	Architecture reviews, mentoring, shell scripting, CI/CD, and delivery automation

EDUCATION

Sheridan College — Postgraduate Diploma   Sep 2017 – Sep 2019

Gujarat Technological University — B.Tech, Electronics Engineering   May 2011 – Jun 2015

CERTIFICATIONS

- Elasticsearch Engineer
- Certified Kubernetes Administrator (CKA)
- Advanced Python Programming - Coursera
- AWS Certified Solutions Architect – Associate
- AWS Cloud Practitioner (CLF-C02)
- Intermediate Go Language Development – Dreams of Go