

Narendranath Edara

+1 (573) 466-6656 | edara.narendranath@gmail.com | [linkedin.com/in/narendranathe](https://www.linkedin.com/in/narendranathe) | narendranathe.github.io

EDUCATION

Missouri University of Science and Technology

Master of Science in Information Science and Technology; GPA: 4.0/4.0

Rolla, MO

Jan. 2022 – Dec. 2023

EXPERIENCE

Data Engineer

July 2024 – Present

ExponentHR

Dallas, TX

- Architected governed semantic layer on Microsoft Fabric with DAX-based metric definitions, establishing a single source of truth for executive dashboards, BI reports, and AI-ready analytics, reducing metric discrepancies to zero and client support volume by 40%
- Built and deployed an Azure DevOps driven Contained AAG automation framework with automated restore, security, recovery model, and CDC setup, saving ~1 hour per request across 20+ daily refreshes.
- Built CI/CD deployment pipelines with Pytest unit tests, integration validation, and automated schema checks, compressing release cycles from 3 months to 14 days (85% faster) with zero production incidents post-deploy
- Redesigned SQL Server OLAP into star schema with partition pruning and file compaction (20% latency reduction); reengineered CDC-based ETL with idempotent loads and backfill support (70% faster, 30% cost savings)
- Designed automation agents with exponential backoff, dead-letter routing, and error handling, integrating Git and Azure DevOps APIs to eliminate 15+ hours/sprint of manual work with real-time observability

Data Engineer

Aug. 2023 – July 2024

Missouri S&T

Rolla, MO

- Engineered data observability pipelines with automated schema validation, freshness monitoring, and volume anomaly detection, achieving 95%+ accuracy and preventing silent data failures in production
- Designed configurable REST APIs with contract-based input validation and schema enforcement, enabling dynamic anomaly detection tuning with zero manual intervention for schema drift, reducing false positives by 40%

Business Analyst

Mar. 2018 – Sept. 2020

Zomato

Hyderabad, India

- Designed real-time analytics platform with governed metric definitions for eCommerce KPIs, enabling data-driven pricing strategies that contributed to 9% market share increase
- Optimized search relevance using ranking models with A/B testing frameworks and automated quality monitoring, improving search-to-conversion rates across millions of daily queries
- Built Elasticsearch-powered search engine indexing 100K+ documents with modular ranking components and automated query resolution, reducing Support Desk workload by 80%

PROJECTS

Real-Time Fraud Detection Platform | *PySpark, Kafka, Airflow, MLflow, Docker, AWS S3*

2026

- Built real-time streaming inference with Airflow DAG orchestration serving 100+ TPS using Kafka consumers, containerized endpoints, dead-letter queues, and automated retry with exponential backoff
- Engineered ML pipeline with Pytest validation suites, data quality gates, LightGBM training with MLflow tracking, and automated schema checks in CI/CD deployment workflow

Automated Portfolio Risk Analytics | *Kafka, Spark Streaming, PySpark, FastAPI, Docker*

2025

- Built Kafka + Spark Streaming pipeline with broadcast joins processing 15K+ records at 47.8 TPS, with backfill-safe idempotent loads and sub-5-second Value at Risk computation
- Designed RESTful APIs with schema enforcement and contract validation, deploying containerized microservices with health probes, self-recovery, and real-time observability dashboards

CERTIFICATIONS & PUBLICATIONS

DP-700: Microsoft Certified Data Engineer Associate | **Publication:** Sentiment Analysis for Visitor Insights – [DOI](#)

TECHNICAL SKILLS

Languages & Testing: Python, SQL, Bash, Pytest, FastAPI, LangChain, REST APIs

Data & ML Engineering: Spark, Kafka, Airflow, Delta Lake, ETL/ELT, Data Modeling, Feature Engineering, MLflow, RAG

Cloud & Infrastructure: Azure (AKS, Fabric), AWS (S3, Kinesis), Databricks, Snowflake, PostgreSQL, Docker, Kubernetes, CI/CD, Git, Power BI