

UJWAL KANDI

📞 (737) 326-8598 📍 Austin, TX ✉ ujwal.kandi@utexas.edu 🔗 linkedin.com/in/ujwalkandi 🌐 ujwalkandi.in/projects

Professional Summary

Data Operations and Observability Engineer with **4+ years managing production data pipelines and monitoring infrastructure** for Fortune 500 clients including Disney, Bank of America, and Marriott across 16+ enterprise environments. Reduced **Mean Time to Resolution (MTTR) by 75%** and maintained **99.8% SLA compliance** through structured incident management, RCA workflows, and automated observability tooling. M.S. in IT Management, UT Austin (May 2024).

Technical Skills

Languages: Python, Bash, SQL, YAML, XML, Linux

Monitoring: Grafana, Dynatrace, AWS CloudWatch, Azure Monitor, Prometheus, ELK Stack (Elasticsearch, Logstash, Kibana)

Pipelines & Data: Airflow, Matillion, SnapLogic, Kafka, Spark, PySpark, Databricks, PostgreSQL, MySQL, Redshift, SQLite

Cloud: AWS (EC2, S3, Lambda, Redshift, Glue, CloudWatch), Azure (ADLS, Log Analytics, APIM, Azure Monitor)

Infrastructure: Terraform, CloudFormation, Ansible, Docker, Kubernetes, Git, GitHub, Jenkins, GitLab CI/CD

Networking: TCP/IP, HTTP/S, DNS, REST APIs, SNMP, OSI Model

Tools: ServiceNow, PagerDuty, Jira, Confluence, Tableau, Power BI

Experience

Apple Inc. (Contract via Welo Data)

Jan 2025 – Present

Metadata Specialist (Data and Pipeline Operations)

Austin, TX

- Managed metadata ingestion and quality validation pipelines for Apple TV's global content catalog across **100+** international partner feeds, maintaining **sub-2s latency** and **99.9%** uptime through automated XML validation and rollback procedures
- Built **Grafana** dashboards and automated alerting for pipeline health monitoring, enabling early detection of data quality issues; reduced feed failures by **30%** and cut mean time to detection (**MTTD**) from hours to under 15 minutes
- Standardized **Root Cause Analysis (RCA)** frameworks with reusable debugging runbooks and Linux/Bash diagnostic scripts for recurring batch-job failures, trimming resolution time by **25%** and eliminating recurring failure patterns to ensure data availability for downstream AI-driven search and recommendation models
- Designed failover and rollback procedures for critical metadata systems, ensuring uninterrupted content availability in compliance with Apple's **data governance standards** and service-level agreement (**SLA**) targets

Environment: Grafana, Apache Airflow, XML, AWS S3, Python, Bash, Jira, Confluence

Sports Excitement (Sports data analytics startup)

Aug 2024 – Jan 2025

Data Operations Engineer

Austin, TX

- Built and maintained **50+ daily Airflow ETL pipelines** with DAG-level retry logic, SLA breach alerts, and automated failure recovery, achieving **99.8%** job reliability and ensuring consistent data availability for business reporting
- Optimized Azure Data Lake Storage partition schemas and storage tier configurations, reducing query latency from approximately **1.2s to 550ms** and directly improving downstream ML analytics throughput for data-driven decisions
- Automated incident alerting via Airflow trigger jobs integrated with **PagerDuty**, cutting mean response time to under **5 minutes** against SLA targets and reducing operational overhead for the on-call team
- Implemented pre-deployment data quality validation and load-testing frameworks, reducing post-release pipeline failures by **28%** and improving data reliability for end consumers

Environment: Apache Airflow, Azure ADLS, Azure Monitor, PagerDuty, Python, SQL, GitLab CI/CD, Power BI

Dover Fueling Solutions (Contract)

Jun 2024 – Aug 2024

Observability and Site Reliability Engineer

Austin, TX

- Deployed a **3-pillar observability stack** combining Grafana dashboards, Azure APIM log analytics, and Kafka metrics exporters to enable real-time visibility across fuel station infrastructure; reduced issue detection from **15 min to 3 min** and **MTTR by 40%**
- Built an edge diagnostics system for resource-constrained IoT field devices, enabling technicians to autonomously identify and resolve critical faults; reduced support tickets by **35%** and brought fault resolution time under **10 minutes** without remote escalation
- Tuned alert thresholds and implemented intelligent escalation workflows, cutting alert fatigue by **45%** while maintaining **100%** detection coverage for P1 and P2 severity incidents to protect operational continuity

Environment: Grafana, Azure APIM, Apache Kafka, Python, SQLite, Docker, ServiceNow, Jira

Epsilon (Publicis Groupe)

Aug 2021 – Jul 2023

Production Support Engineer (promoted from Data Programmer I, Apr 2022)

Bengaluru, India

Production Support Engineer, Apr 2022 – Jul 2023

- Managed end-to-end **incident management** and response lifecycle covering detection, triage, resolution, **RCA**, and prevention across **16** enterprise client environments including Disney, Bank of America, Unilever, and Marriott, achieving **99.8% SLA compliance** and reducing **MTTR from 45 min to 12 min**
- Built observability infrastructure integrating **AWS CloudWatch**, Kibana, and Dynatrace for metrics collection, log analysis, and distributed tracing; set alert baselines and anomaly detection thresholds that prevented **60%** of incidents before customer impact
- Used Linux command-line tools to inspect logs, trace running processes over SSH, and isolate root causes during triage; wrote Bash and Python automation scripts for scheduled data-quality checks, log rotation, and pipeline health tests that reduced manual monitoring across on-call rotations
- Authored **20+ runbooks**, SOPs, and postmortem documents implementing a blameless RCA culture; eliminated repeat incident categories by **50%** through systematic root cause prevention and knowledge transfer

- Automated production pipeline failover and recovery with Bash scripts on Linux hosts and AWS CloudFormation, reducing **RTO from 2 hrs to 30 min** and sustaining **99.9% uptime** through critical processing windows
 - Built and tuned **SnapLogic** and **Matillion** data integration pipelines and optimized ETL job scheduling across SLA windows and retry logic, resolving mapping and connectivity issues during incidents and **improving batch completion 32%**
- Data Programmer I, Aug 2021 – Apr 2022*
- Engineered and maintained ETL pipelines for **16** global enterprise clients, reducing batch failure frequency by **45%** and ensuring consistent data availability for mission-critical analytics platforms
 - Deployed **CloudWatch**, **Dynatrace**, and **Kibana** observability stack and performed rapid data quality checks and emergency hotfixes, maintaining a **93%** client satisfaction rate across enterprise accounts
 - Wrote Bash and Python automation scripts on Linux for scheduled data-quality checks, log rotation, and pipeline health tests, scheduled via cron, reducing manual monitoring effort across on-call rotations
- Environment:** Linux, AWS CloudWatch, Dynatrace, Kibana, ELK Stack, SnapLogic, Matillion, Apache Airflow, AWS S3, Redshift, CloudFormation, Python, Bash, SQL, Jenkins, Confluence

Isthriwala (E-commerce operations startup) **Jun 2020 – Jul 2021**
Data Management Analyst *Vijayawada, India*

- Optimized SQL database operations for a platform serving **35,000+** customers and **100+ B2B clients** through index restructuring and query execution plan analysis, improving query response time by **~40%** and reducing load on critical transactional systems
- Built real-time KPI dashboards in **Tableau** tracking inventory, sales, and operational metrics; adopted as the primary management reporting tool and contributed to a **15%** improvement in monthly performance targets
- Identified growth opportunities and optimized customer retention strategies, contributing to a **10%** increase in monthly subscriptions

Environment: MySQL, Tableau, Python, Excel

Projects

A.I.D.E. – Observability and Automated RCA Engine — *Prometheus, Grafana, Ollama, Python, LlamaStack* **Apr 2024**

- Built a full-stack observability system combining real-time metrics analysis, log anomaly detection, and automated **RCA** workflows for API and infrastructure health monitoring, enabling autonomous troubleshooting
- Integrated **Prometheus** and **Grafana** for live dashboarding; deployed on-device Llama 3 inference via Ollama under a **2GB** memory constraint, achieving sub-second response and reducing simulated **MTTR by 60%** (Meta ATX Llama Hackathon)

FlipFinder Pro – Real-Time Data Pipeline and Anomaly Detection — *Databricks, PySpark, Python* **2025**

- Built a real-time market data ingestion and ML-scoring pipeline using **Databricks** for distributed processing and feature engineering, enabling statistical outlier identification across a high-volume live dataset
- Achieved **89%+** price-prediction accuracy and surfaced **\$20M+** in detected anomalies through an interactive analytics dashboard with profit, risk, and confidence scoring (Databricks Free Edition Hackathon)

Azure Document Intelligence Pipeline — *Azure AI Studio, Python (Capstone, Stronghold Investment Management)* **Jan–May 2024**

- Built an automated document processing pipeline integrating OCR and Mixtral LLM to extract structured ownership data from unstructured inputs and route verified records through downstream delivery workflows
- Reduced document processing costs by **20%** through end-to-end pipeline automation; benchmarked accuracy, coherence, and latency across **4** Azure-hosted model configurations to select the optimal production model

Leadership and Achievements

ElevenLabs Worldwide Hackathon – First Place and Best Product Built with n8n **Dec 2025**
Global Hackathon (1,000+ participants) *Austin, TX*

- Won **First Place** at a hackathon of **1,000+** participants, building ARIA, a voice-based retail intelligence agent using multi-agent LLM orchestration, ElevenLabs TTS, and n8n-powered automation for personalized recommendations and post-purchase workflows

Seed Studio Ranger – AI/IoT Developer Program **Oct 2024 – Present**
AIoT Developer Educator Program *Austin, TX*

- Designed and delivered **10+** hands-on technical workshops on edge computing, IoT systems, and embedded AI to **100+** university students, building practical skills in AI-driven operational monitoring and edge diagnostics

Dover Fueling Solutions Hackathon – First Place **Feb 2024**
Annual Internal Hackathon *Austin, TX*

- Won **First Place** designing an edge AI diagnostics system for fuel station operations; presented the solution to DFS executive leadership and industry stakeholders at the DFS User Conference 2024

HACKSciCom 2022 – Finalist **May 2022**
Annual Hackathon, MIT Hacking Medicine co-hosted with Roche *Remote*

- Reached the finals building end-to-end **data pipelines** with a globally distributed team at one of healthcare’s premier hackathons, co-hosted by **MIT** and **Roche**

Education

The University of Texas at Austin **Aug 2023 – May 2024**
M.S. in Information Technology and Management, McCombs School of Business *Austin, TX*

KL University **May 2020**
B.Tech in Computer Engineering, Specialization in Data Analytics *Guntur, India*

CGPA: 8.49/10.0 | Department Topper, Dean’s List

Certifications

Google Cloud Associate Cloud Engineer (ACE) | Google Cloud | 2026 | Verify