Jorge Navarro

Senior Data Engineer | MLOps Engineer | Cloud & AI/ML Data Specialist

- jorge@jorgenavarro.dev
- Hemet, CA (Open to Remote in US | Hybrid in Los Angeles Metropolitan)
- LinkedIn.com/in/jorge-raymundo-navarro
- GitHub.com/syntax0111

Summary

Highly accomplished Senior Data Engineer with 10+ years of progressive experience in designing, developing, and optimizing scalable **cloud-native data solutions and MLOps pipelines**. Proven expertise in building robust data pipelines with **Apache Airflow, FME, and Palantir Foundry**, and leading **observability strategies** (Prometheus, Grafana, Loki, OpenTelemetry) that significantly improve system reliability and incident response. Adept at full-stack development, **cloud deployments (AWS, Azure, Kubernetes)**, and leveraging various DBMS (PostgreSQL, Snowflake, Oracle) to drive data health and empower data-driven decision-making. Completed **Cloud DevOps Engineer** and **Data Engineering with AWS** Nanodegrees from Udacity, complemented by **Coursera's Machine Learning Specialization**, strategically positioning for advanced roles in **Artificial Intelligence, Machine Learning, and Robotics data infrastructure.**

Skills

Data & AI/ML Tools	Database Systems	Cloud & DevOps	Core Languages & Frameworks	Observability
Airflow, FME, Palantir Foundry, Spark, Pandas, Jupyter, TensorFlow, scikit-learn	PostgreSQL, Snowflake, MySQL, Oracle, SQLite, Redshift, Cassandra	AWS (S3, EC2, Lambda, Glue, Athena, Redshift, Route 53, Auto Scaling, Load Balancing), Azure, Kubernetes (AKS), Docker, IaC (CloudFormation, Terraform), CI/CD, Git, Linux	Python, SQL, Shell, Java, C/C++, HTML, CSS, JavaScript, Flask, FastAPI, Typer, Bootstrap	Prometheus, Grafana, Loki, OpenTelemetry, OnCall

Professional Experience

Professional Member of Technical Staff, AT&T Inc., CA, JUL 2023 – APR 2024

Focus: Observability Engineering & Incident Response for Critical Systems

- **Spearheaded the development and implementation of a comprehensive observability engineering strategy** utilizing Grafana, OnCall, Prometheus, and Loki, with exploration into OpenTelemetry, significantly improving system reliability and reducing downtime.

- **Engineered an end-to-end incident response pipeline** integrating Grafana alert rules, OnCall's alert groups, and escalation chains, leading to a **quantifiable reduction in incident resolution time, improved system uptime, and faster stakeholder notification.**
- **Administered and contributed to the core Grafana management codebase** (Python, Flask), enhancing email notification functionality and streamlining administrative tasks within Kubernetes (AKS) deployments.
- **Designed and developed advanced, dynamic Grafana dashboards** for data, application, and infrastructure monitoring, providing improved data visibility and **actionable insights for 5+ teams**, accelerating issue identification and response.

Professional-Software Engineer, AT&T Inc., CA, JUL 2019 – JUL 2023

Focus: Scalable Data Pipeline Engineering & Microservices Development

- **Engineered and optimized critical end-to-end data pipelines** using **Apache Airflow, Feature Manipulation Engine (FME), and Palantir's Foundry**, ensuring data health and comprehensive observability for high-volume mobility transport data.
- **Developed multiple complex ETL pipelines with FME** and managed a significantly sized ELT pipeline within Palantir's Foundry, **resulting in improved data accuracy, reduced processing time, and robust support for numerous critical business operations.**
- **Led the design and development of a standardized, deployable Flask-based ReST API template**, addressing cross-cutting concerns (authentication, authorization, logging, OpenAPI Spec) and including **Kubernetes (AKS) deployment artifacts.**
- **Leveraged the API template to launch numerous scalable microservices** supporting critical network elements (e.g., anti-spam, mobility switch orchestration) and internal tools, **accelerating API development and deployment while enhancing service reliability.**
- **Centralized API documentation** through the deployment of a Swagger UI instance, improving developer experience and API discoverability.

Professional-Network Planning Engineer+, AT&T Inc., CA, JUL 2017 – JUL 2019

Focus: Enterprise System Maintenance & Feature Development

- Collaborated on the full software development lifecycle of the **Engineering Rules Database (ERD)**, a legacy N-tier system (Ext JS GUI, Java business logic, Oracle DB), committing **enhancements, fixes, and refactorings** to improve stability and usability.
- Implemented **end-to-end CRUD functionality** for network element capacity planners, directly supporting **15+ planners and streamlining hundreds of engineering rule updates**, enhancing data consistency and availability.

Sr Specialist-Network Planning Engineer, AT&T Inc., CA, NOV 2013 – JUL 2017

Focus: Data-Driven Performance Reporting & Automation

- **Designed, developed, and implemented Key Performance Indicators (KPIs) and Key Capacity Indicators (KCIs)** in utilization and performance reports using Logi Analytics, providing critical insights for proactive network adjustments and improved decision-making.
- **Identified and implemented mechanization opportunities**, developing tools and an **Engineering Rules Database (ERD) web service** to automate repetitive tasks, resulting in a **significant reduction**

in manual effort and improved data accuracy.

Capacity Engineer, AT&T Inc., CA, JAN 2012 – NOV 2013

Focus: Network Capacity Management & AutoCAD Automation

- Managed physical capacity for networks and network elements, including **physical space utilization and heat dissipation analysis** in AT&T central offices.
- Contributed to **C#-based automation solutions for AutoCAD**, streamlining design workflows and increasing efficiency and consistency.

Education

- **Masters of Science in Computer Science**
  - Georgia Institute of Technology, Atlanta GA. *expected 2027*
  - Current GPA: 3.85 (*4.00 scale*)
- **Bachelor of Science in Electrical Engineering**
  - California State Polytechnic University, Pomona, CA. June 2011

Projects & Certifications

- **Cloud DevOps Engineer Nanodegree**, Udacity (Completed)
  - **Key Skills/Projects:** Cloud deployments (AWS, Kubernetes), Infrastructure as Code (CloudFormation), CI/CD pipelines, automated testing, monitoring (Prometheus, Grafana), logging, configuration management (Ansible). Focused on building robust, scalable cloud infrastructure and deploying microservices.
- **Data Engineering with AWS Nanodegree**, Udacity (74% Complete)
  - **Key Skills/Projects:** Data Modeling (Relational, NoSQL; Cassandra), Cloud Data Warehousing (AWS Redshift), Data Lakes (S3, Glue, Athena), Big Data Processing (Spark, PySpark), Data Pipelines Automation (Apache Airflow). Includes building a data Lakehouse for ML model training.
- **Machine Learning Specialization**, Coursera (Completed)
  - **Key Skills/Topics:** Supervised Learning (Linear/Logistic Regression, Neural Networks), Unsupervised Learning (Clustering, PCA), Decision Trees, Recommender Systems, Reinforcement Learning, Practical aspects of ML development. Utilized Python (NumPy, scikit-learn, TensorFlow).