

Pranav Modh

Professional Summary

Data Engineer with 5+ years of experience specializing in **Palantir Foundry platform development**, data pipeline engineering, and custom application delivery across healthcare, semiconductor, and supply chain domains. Expert in building production-grade Ontologies, transforms, and React/TypeScript-based custom widgets within Foundry. Proven track record of reducing processing times by 50%, cutting cloud costs by 30%, and delivering end-to-end solutions that drive measurable business impact. Strong background in PySpark, AWS integration, and client-facing solution architecture.

Professional Experience

Data Engineer (FDE Contract)

04/2026 – Present

- Own the **Foundry reporting layer end-to-end** — building Python/PySpark transform pipelines that consolidate 10+ upstream sources into 25+ curated KPI datasets, refreshed on schedule with automated data-quality checks and full lineage.
- Modeled **12+ Ontology object types and link types** (suppliers, KPIs, metric snapshots, portal links) backing the reporting layer, and moved heavy builds to incremental transforms, cutting pipeline runtime by ~45%.
- Built a **reusable TypeScript custom-functions library** that generates KPI content elements (cards, trend charts, thresholds, variance flags), standardizing 40+ KPI definitions and reducing new-KPI turnaround from 3 days to under 2 hours.
- Shipped a **React application deployed via Foundry Developer Console** that consumes these Functions to render 30+ interactive charts and KPI cards, giving 300+ business users self-serve reporting and eliminating ~20 hours/week of manual Excel reporting.

Data Engineer (FDE)

12/2024 – 03/2026

- Architected and deployed **Palantir Foundry-based inventory management platform** integrating 15+ data sources (ERP systems, Excel, flat files, APIs) into unified Ontology, enabling real-time supply chain visibility for 200+ global users across 12 manufacturing facilities.
- Developed **custom React/TypeScript widgets** within Foundry Workshop for interactive inventory alerts, PO generation workflows, and supplier performance dashboards, resulting in **90% reduction in manual review effort** and **99.2% order accuracy improvement**.
- Built end-to-end data pipelines using **Foundry Pipeline Builder and Code Repositories (Python/PySpark)**, processing 500GB+ daily data volume with automated quality checks, data lineage tracking, and governance controls.
- Created **Workshop applications with drill-down analytics** connecting Ontology objects to executive dashboards, reducing decision-making time by 40% through one-click PO summary reports and threshold-based alerting system.
- Prevented **12+ stockout incidents** in first 6 months through predictive alerting system built using **Foundry Model Studio** and Trigger-based monitoring.

Database Developer

06/2024 – 12/2024

- Designed and implemented **multi-source data ingestion pipeline** from Excel, Microsoft Access, and legacy systems into AWS S3 using Python and Airflow, enabling centralized data access for global analytics teams.
- Built **interactive Power BI dashboards** with 25+ visualizations connected to Redshift, providing real-time insights on operational KPIs and reducing manual **reporting time by 15 hours/week**.

Data Engineer

03/2023 – 12/2023

- Streamlined **end-to-end semiconductor raw-material data flows** in Palantir Foundry by integrating supplier feeds, ERP exports, and fab consumption data into a unified **Ontology**, creating a single source of truth for supply and demand alignment across suppliers and fabs.
- Built **Foundry transforms and PySpark pipelines** in Code Repositories to standardize lead times, safety stock, and consumption patterns, improving planning accuracy and reducing manual data reconciliation for supply-chain and operations teams.
- Developed interactive supply-chain analytics using **Quiver dashboards embedded in Workshop apps**, connecting supplier delivery performance, fab inventory positions, and material risk scores so planners could proactively identify shortages and re-route supply.

- Development of **multi-source social media data aggregation pipeline** using Facebook Graph API, AWS S3, **Apache Kafka**, and **Apache Airflow**, reducing data inconsistency by 40% and enabling real-time analytics on 5M+ daily events.
- Implemented **automated ETL processes with Snowflake** and Unix/Linux scripting, enhancing data processing speed and accuracy by 60% through parallelized workflows and incremental loading strategies.

- Developed and deployed **microservices for data processing** using Python, FastAPI, and MongoDB, integrated with Docker Compose and Kubernetes, reducing system downtime by 40%.
- Built **automated web scraping solution using Python and Selenium**, extracting structured data from 15+ international freight websites into AWS RDS, improving data acquisition efficiency by 25% and access speed by 30%.

Technical Skills

Palantir Foundry: Ontology Design & Modeling, Pipeline Builder, Code Repositories, Custom Transforms, Workshop Applications, Quiver, Custom Widget And App Development (React/TypeScript), Functions, OSDK, Developer Console, Data Connection, AIP Integration, Contour, WorkFlow

Programming & Scripting: Python, TypeScript, JavaScript, SQL, Bash

Data Engineering & Big Data: PySpark, Airflow, DBT, Redshift, BigQuery, ETL/ELT Pipelines

Cloud & Infrastructure: AWS (EC2, S3, Lambda, SQS), Google Cloud (Cloud Functions, Pub/Sub, Compute Engine)

Frontend & Frameworks: React, TypeScript, Flask, FastAPI, Django, D3.js, Vega-Lite

Databases: PostgreSQL, MongoDB, MySQL, Oracle, DynamoDB, AWS RDS

Machine Learning Certification: <https://www.coursera.org/account/accomplishments/certificate/M2G78BN8P9NX>

Projects

AI-Powered Inventory Demand Forecaster & Optimizer

Palantir Foundry (Ontology, Model Studio, Data Pipeline, Code Repository, Workshop), PySpark, AutoGluon, AWS

- Consolidated fragmented supply-chain data for thousands of SKUs across regions by joining Foundry data cubes (SKU master, historical orders, open orders, quotes, sales, cost metrics) with an external oil price index into a single, model-ready dataset.
- Partitioned SKUs into **regular, intermittent, and sparse** demand profiles using ADI and CV² metrics, then trained three dedicated forecasting pipelines in **Foundry Model Studio** with time-series models and AutoGluon to predict 6-month demand per SKU.
- Achieved strong forecast quality (MASE $\approx$ 0.78, log-scale error $\approx$ 0.65) and built validation and production forecast pipelines whose outputs fed **Workshop dashboards** showing predicted min/max levels and six-month forecasts, with approved demand signals flowing into the main planning channel to drive production decisions.

Self-Serve KPI Reporting App on Foundry

Palantir Foundry (Ontology, Code Repositories, Functions, OSDK, Developer Console), PySpark, TypeScript, React

- Designed a **three-repository architecture** — a transforms repo for the reporting layer, a TypeScript library for KPI content elements, and a Foundry Functions repo — decoupling data, business logic, and UI so each layer ships independently.
- Built a **config-driven KPI engine** where each KPI (formula, thresholds, trend window, chart type) is defined once and rendered consistently as cards, line/bar charts, and tables, scaling from 8 to 40+ KPIs with zero front-end code changes.
- Deployed the **React front end through Foundry Developer Console** using OSDK-generated typed clients and Ontology-enforced row-level security, delivering sub-2-second page loads and role-specific views for every business unit.