

CHAITANYA VULLAGANTI

Palantir Foundry Engineer · Certified Data Engineer & Application Developer

Ontology-Driven Platforms · Real-Time Data Engineering · Enterprise AI Applications

Jacksonville, FL · (904) 933-6924 · chaitanyavvk98@gmail.com

chaitanyavullaganti.com · [LinkedIn](#) · [GitHub](#)

PROFESSIONAL SUMMARY

Certified Palantir Foundry Data Engineer and Application Developer with 8 years of experience designing enterprise data platforms and operational applications across banking, utilities, healthcare, and financial operations. Specialized in ontology-driven application design, governed batch and streaming pipelines, real-time event-driven architectures, and Workshop/OSDK applications using Palantir Foundry, Python, PySpark, SQL, and TypeScript. Delivered production platforms spanning CRM, AML case management, real-time collections, and cash-flow forecasting — turning complex operational and regulatory requirements into governed, auditable Foundry products.

SELECTED IMPACT

24h → min

Collections time-to-action via event-driven eligibility

19.7M

Invoice records via incremental CDC/batch transforms

26,749

Collection actions in first production sweep

3 platforms

CRM · AML Case Manager · Cash Flow delivered E2E

FLAGSHIP PROJECTS

Cash Flow Forecasting & Financial Operations

Sep 2025 – Jul 2026

Medallion Foundry platform for treasury/FP&A — Bronze/Silver/Gold pipelines, finance Ontology, ML forecasting, Workshop apps, and OSDK React console.

Collections Management — Real-Time Event-Driven Platform

Oct 2024 – Sep 2025

Replaced Oracle BRM batch collections for 921K accounts / 19.7M invoices; dual state machines and Automate lifecycle cut time-to-action from 24 hours to minutes.

CRM Application on Palantir Foundry

Dec 2023 – Jun 2024

Full-stack CRM with Ontology Manager, Pipeline Builder, Workshop, Automate, Vertex, and AIP — LLM synthetic data and TypeScript auto-assignment by state and industry.

CERTIFICATIONS

- Palantir Foundry Data Engineer Professional — Palantir Technologies Jul 2026 – Jul 2028
- Palantir Foundry Application Developer Professional — Palantir Technologies Jul 2026 – Jul 2028
- Palantir Foundry Aware Professional — Palantir Technologies Jul 2026 – Jul 2028
- Foundry & AIP Builder Foundations — Palantir Technologies May 2026
- Data Science 360 — Analytixs Labs Dec 2021

TECHNICAL EXPERTISE

Ontology & Modelling Ontology Manager · Object/Link/Action Types · State Machines · Cascading State · RBAC · Markings · Lineage
Pipelines & Transforms Pipeline Builder · Code Repositories · Python · PySpark · SQL · CDC/Batch · Streaming · Databricks Delta/Iceberg Pushdown
Functions & Automation TypeScript Ontology Functions · Automate Monitors · Event-Driven Orchestration · Rules Engines · AIP / LLM Integration
Applications Workshop · Slate · Contour · Quiver · Vertex · Foundry SDK / OSDK · React · Power BI · Tableau
Data Ecosystem Oracle · Oracle BRM · SQL Server · Snowflake · Databricks · ADF · ADLS · AWS · Azure · GitHub · Splunk · JIRA
Languages Python · TypeScript · SQL · PL/SQL · R · UNIX Shell

PROFESSIONAL EXPERIENCE

Palantir Foundry Engineer

Sep 2025 – Jul 2026

Pacific Gas and Electric Company

Enterprise Cash Flow Forecasting & Financial Operations Platform

- Designed and delivered an end-to-end cash flow forecasting and financial operations platform on Foundry for treasury and FP&A — governed Bronze/Silver/Gold pipelines, a semantic finance Ontology, ML-based monthly forecasting, and an OSDK-powered React console.
- Modeled a finance ontology (CashTransaction, GLEntry, CashFlowForecast, MonthlyCashFlowStatement, ReconciliationBreak) with link types connecting transactions, accounts, forecasts, and exceptions.

- Built ingestion and curation pipelines in Pipeline Builder and Code Repositories for bank transactions, GL entries, chart of accounts, and reference data; tuned PySpark jobs to stabilize month-end workloads.
- Implemented data quality expectations, reconciliation logic, and governed Action Types for break assignment, forecast overrides, and approvals with object-level permissions for treasury roles.
- Delivered Workshop applications (CFO Cash Dashboard, Forecast vs Actuals, Line Item Explorer, Reconciliation Workbench) and a React/TypeScript OSDK application for workflows beyond standard Workshop.

Senior Foundry Engineer

Oct 2024 – Sep 2025

U.S. Bank

Collections Management System — Real-Time, Event-Driven Collections Platform on Foundry

- Designed and deployed an end-to-end collections platform replacing a legacy Oracle BRM batch engine across a production portfolio of 921K billing accounts and 19.7M invoice records — reducing time-to-action from a nightly 24-hour batch cycle to minutes via event-driven eligibility.
- Architected CDC and batch merge pipelines across six Oracle BRM tables using LEFT ANTI JOIN merge with timestamp-based deduplication; built incremental PySpark transforms with derived ageing analytics.
- Designed an eight-object-type ontology with dual state machines (collection record: 5 states / 10 transitions; collection action: 7 states / 12 transitions) and cascading exit logic.
- Authored 15+ production TypeScript Ontology Functions and configured Automate monitors for real-time entry plus full lifecycle automation (scheduling, execution, promise-to-pay, disputes), removing routine manual handling on standard-path cases.
- Led parallel-run migration against legacy batch with rollback capability; delivered initial production sweep creating 1,917 collection records and 26,749 collection actions.

AML Alert Manager & Case Manager — Anti-Money Laundering Workflows

- Developed and delivered end-to-end Foundry applications (Alert Manager, Case Manager) to support AML workflows for a banking client — translating business requirements into production features.
- Built and maintained ingestion pipelines and operational tooling to onboard historical and daily transactional data from legacy systems into Palantir Foundry, ensuring data availability for downstream teams.
- Designed and implemented data transformation pipelines using Pipeline Builder and loaded curated datasets into the Foundry Ontology to standardize entities and relationships for compliance use cases.
- Implemented business rules and validation logic for regulatory workflows, enabling automated alert triage and case creation while preserving auditability and traceability.
- Led feature demos and iterative feedback cycles with business stakeholders and compliance SMEs; incorporated feedback through sprint planning, grooming, and story refinement.
- Collaborated cross-functionally with Data Ingestion, Downstream Processing, and Application Development teams to align data models, resolve ingestion issues, and accelerate time-to-value for regulatory applications.
- Monitored and troubleshooted production data pipelines and ingestion jobs; investigated data quality issues and coordinated remediation with engineering and operations teams.
- Developed user-facing Workshop applications and UI components to support end-user workflows for alert review, investigations, SAR preparation, and FINCEN filing.
- Integrated regulatory filing workflows into the Case Manager application to prepare, submit, and manage Suspicious Activity Reports (SARs) and handle FINCEN feedback and rework within the platform.
- Applied Agile practices (Scrum ceremonies, story estimation, sprint planning, retrospectives) to deliver incremental value and maintain predictable release cadence.
- Documented technical designs, runbooks, and onboarding guides to reduce operational risk and accelerate handoffs between teams.
- Mentored new team members on Foundry best practices, ontology design, and pipeline debugging to raise team capability and consistency.

Palantir Data Engineer

Jun 2024 – Nov 2024

TeksolveIT Solutions LLC · NC, United States

Real-Time Streaming & Batch Operations Platform — Ontology, Workshop & Databricks Pushdown

- Developed and optimized real-time streaming and batch data pipelines in Palantir Foundry, managing multiple development branches and ensuring seamless ingestion, transformation, and streaming of clean, standardized data.
- Monitored pipeline health dashboards continuously, proactively resolving warnings, errors, late data, and schema changes to prevent outages and maintain high data availability.
- Scaled compute resources and fine-tuned cluster configurations and partition keys in response to increasing data volumes, optimizing pipeline performance and cost-efficiency.
- Managed code versioning with pull requests and coordinated release schedules to ensure stable, outage-free production pipeline deployments with complete audit trails.
- Enhanced data correlation models by integrating new data sources and business rules through collaborative planning with product owners and analysts, resulting in improved incident detection and reduced unnecessary field dispatches.
- Designed and implemented data cleansing steps including deduplication, missing value imputation, unit translation, and schema alignment within Pipeline Builder to ensure data quality and consistency.
- Developed and refined Palantir Foundry Workshop self-service applications featuring overlays, widgets, and role-based access controls to empower users with interactive data exploration and secure ticketing management.
- Created and maintained governed Ontology elements (Object Types, Action Types, Link Types) for devices, circuits, and outages, establishing property constraints, lineage rules, and RBAC in collaboration with architecture and governance teams.

- Built and maintained interactive dashboards and reports in Quiver and Contour, visualizing pipeline health, data quality, outage maps, and key operational metrics with drill-downs and alerting features for executive and operational stakeholders.
- Developed scalable PySpark transformations and reusable TypeScript and Spark SQL UDFs within Foundry Code Repositories and Workbooks, ensuring data readiness for real-time pipelines and analytics applications.
- Developed a proof of concept leveraging compute pushdown by integrating Databricks Delta tables and Iceberg Tables with Palantir Foundry via data connections; created virtual datasets in Foundry that utilize Databricks compute for PySpark transformations.
- Saved transformed data back as Delta tables in Databricks, then reaccessed and modeled them as Object Types within Palantir Ontology for seamless downstream analytics.
- Conducted detailed performance benchmarking comparing standard Databricks compute, Photon acceleration, serverless compute, and native Palantir Pipeline Builder, demonstrating efficiency gains and informing compute strategy decisions.

Foundry Application Developer

Dec 2023 – Jun 2024

CRM Application Development (Palantir Foundry)

Enterprise CRM on Foundry — Ontology, Pipeline Builder, Workshop, Automate & AIP

- Developed a comprehensive CRM using Code Workspaces, Ontology Manager, Workshop, Pipeline Builder, Automate, and Vertex — covering client–employee management end to end.
- Built automated Pipeline Builder ETL; performed analysis and cleansing on company/people datasets in Python (Jupyter) and SQL Workbooks.
- Implemented transformations and synthetic data via custom AIP Logic and LLM integration; created Ontology Object Sets and Action Types.
- Developed TypeScript Automate functions including auto-assignment of employees to new clients by state and industry; designed Workshop UI with visualizations, filters, widgets, overlays, and real-time metrics.

Data Scientist

Dec 2022 – Dec 2023

Centene Corporation

Healthcare Data Intelligence & Predictive Analytics Platform

- Built predictive models supporting healthcare operations — risk assessment, population segmentation, and operational forecasting — across structured and unstructured datasets.
- Delivered end-to-end modeling workflows as modular R scripts; applied supervised/unsupervised techniques selected against precision, recall, ROC-AUC, and cross-validation.
- Built NLP workflows for free-text feedback analysis; developed Tableau dashboards communicating model outputs and operational KPIs to leadership.

Business Intelligence Engineer

Jun 2018 – Dec 2022

National Employment Savings Trust (NEST)

Enterprise Business Intelligence & Reporting Platform

- Designed analytical data models and complex SQL (CTEs, window functions, ranking, aggregations) transforming transactional data into executive-ready datasets.
- Built interactive Tableau dashboards with calculated fields, parameters, drill-downs, and dynamic filters; automated recurring reporting to reduce manual effort.

LEADERSHIP & MENTORING

- Mentor engineers and analysts on Foundry best practices — ontology modelling, transform reuse, incremental design, security marking hygiene, and Automate debugging.
- Run peer code reviews and knowledge-sharing on idempotent function design, Spark profiling, partition/join tuning, and lineage analysis.
- Bridge business and engineering by translating requirements into ontology models, pipelines, object relationships, and action types; lead RCA and runbook-driven prevention after incidents.

EDUCATION

Bachelor of Technology, Information Technology

St. Martin's Engineering College, JNTUH — Hyderabad, Telangana, India

2019