



Modernization Success Story

Building Telecom Operational Intelligence on Palantir Foundry and AIP

Performed by: Cadmus (delivered by Ventera, acquired and fully integrated into Cadmus in 2024)

Period of Performance: 2020-Present

Client Organization: Large Telecom Provider Construction & Engineering (C&E) Team

Executive Summary

A large U.S. telecom provider's Construction & Engineering (C&E) organization launched a multi-year initiative to modernize its analytics ecosystem and shift from on-prem, warehouse-centric reporting to a cloud-based, Data Mesh-aligned model. **Cadmus** played a critical role in migrating legacy data (10+ Petabytes) and processes; re-engineering complex datastores to Snowflake; ETL from SQL Server/SSIS to scalable Palantir Foundry pipelines; and transforming the user experience from static "canned" reports to near-real-time, self-service analytics.

The program combined deep engineering (distributed transformations, orchestration, performance tuning, and operational hardening) with change management (training, adoption enablement, and a pragmatic reporting transition). Most recently, we extended the platform with AI-powered capabilities, including a GenAI-enabled conversational analytics chatbot using Palantir AIP, now in production and expanding into operational workflows.



Background & Starting Point

The C&E organization's legacy environment had grown over time and included:

- **120+ SQL Server-based legacy databases** and data marts
- **Cognos** for standardized reporting and “pixel-perfect” outputs
- **SSIS** and SQL-driven transformations (approximately **100+ SSIS jobs**)
- A reporting culture centered on **once-a-day refresh cycles** and pre-defined outputs

This model constrained agility, increased operational burden, and limited the organization's ability to ask new questions quickly. The target state was a modern platform that could support scalable processing, improved governance, decentralized ownership, and self-service insights, moving teams from “request a report” to “answer it yourself,” with fresher data.

“ This model constrained agility, increased operational burden, and limited the organization's ability to ask new questions quickly.

Goals

Cadmus aligned delivery to the C&E organization's modernization goals:

- Defined **data governance standards** (quality, ownership, lineage, and consistency)
- Modernize data transformation design and build using **PySpark, Snowflake, Snowpark,** and **SnowSQL**

“ A defining constraint early in the program was the lack of meaningful vendor support during implementation.

- **Democratize access** to data and analytics tools
- Improve data quality through **decentralized data ownership** aligned to Data Mesh principles
- Enable collaboration and **self-serve analytics** through structured training and documentation
- Move from **static, once-a-day reporting to near real-time, self-service analytics**

A Major Challenge We Overcame: Minimal Vendor Support During Implementation

A defining constraint early in the program was the lack of meaningful vendor support during implementation. When we started in 2020, there was limited structured training and limited hands-on guidance for how to translate a complex legacy data environment into Palantir Foundry at scale.

We closed this gap through deliberate internal capability building:

- **POC-led upskilling:** targeted proofs-of-concept to validate ingestion patterns, transformation approaches, orchestration, and performance strategies before scaling
- **Self-training and internal playbooks:** codified repeatable engineering standards, templates, and runbooks
- **Pattern-driven migration:** standardized

“ This approach reduced vendor dependency and created a sustainable internal delivery capability for the client.

approaches for pipeline design, testing, partitioning, and production hardening

- **Enablement at scale:** “build-with-us” sessions to transition users from dependence to independence

This approach reduced vendor dependency and created a sustainable internal delivery capability for the client.

Technical Solution and Delivery Approach

1) Ingestion and Data Foundation

We migrated application data from SQL Server to Snowflake, so our ingestion approach used a Snowflake-replicated copy of the application database as the primary source. Snowflake served as an intermediary/source system for multiple consumers, while Foundry became the core platform for processing and self-service analytics enablement.

Key actions included:

- Ingesting Snowflake and other sources into the platform with consistent dataset contracts
- Curating governed datasets suitable for downstream analytics and operational use
- Establishing governance-aligned dataset standards (schema discipline, ownership, documentation, and lineage)

2) ETL Modernization: SSIS/SQL to Spark-Based Pipelines

The largest technical lift was re-engineering legacy

transformation logic originally implemented in SSIS and SQL into a cloud-scale execution model. We rebuilt transformations primarily using **PySpark**, and leveraged **Snowpark/SnowSQL** where Snowflake-native processing was the best fit.

Core engineering outcomes:

- Re-implemented hundreds of complex **SSIS jobs** and related SQL logic as scalable pipelines
- Rebuilt orchestration and scheduling with dependency management, automated runs, and production monitoring
- Designed pipelines for reliability: idempotent runs, modular stages, and consistent failure handling

This was not a lift-and-shift. It was a re-platforming effort that changed how transformations executed, scaled, and were operated.

3) Data Mesh Architecture

With the investment in modern data architecture, the client upgraded data operations standards. A **Data Mesh** approach provided the framework to deliver features quickly at scale by treating curated datasets as governed, reusable **data products** owned by domain teams.

As part of the Data Mesh modernization, Cadmus built foundational capabilities that made the platform scalable, trustworthy, and easier to adopt:

- **Data Contracts:** Defined dataset contracts so consumers had a consistent view of schema, definitions, and expected refresh frequency.

“ This was not a lift-and-shift. It was a re-platforming effort that changed how transformations executed, scaled, and were operated.

- **Access Control:** Implemented role-based and attribute-level access so consumers could access only the data and fields they needed.
- **Health Checks:** Strengthened data quality controls and observability across critical pipelines through automated checks and actionable signals.
- **Discoverability:** Enabled consumers to easily discover the feeds and attributes they needed, reducing reliance on tribal knowledge and accelerating onboarding.

These capabilities reduced operational drag and made self-service analytics viable at scale.

4) Near Real-Time Enablement and Operational Readiness

A key objective was moving from once-per-day reporting to near real-time analytics. We designed ingestion and transformation pipelines to support faster refresh cycles and operational use cases, including:

- **Near real-time data loading** patterns for timely insights
- **Operational alerts** and exception monitoring
- Foundations for **predictive modeling**
- A **GenAI-powered chatbot** that reduces friction for common analytic questions

To accelerate delivery and iteration speed, we leaned on Foundry's developer tooling:

- **Code Workbooks** for rapid development and testing of transformations
- **Repositories** for versioned, collaborative

engineering workflows

- **DIY Syncs** to quickly synchronize, validate, and evolve datasets across environments

To improve confidence and operational maturity, we leveraged:

- **Lineage** to provide transparent, auditable dependency tracking end-to-end
- **Health Checks** to strengthen data quality controls and observability across critical pipelines

5) Performance Engineering and Production Hardening

Once end-to-end flows were established, we focused heavily on performance and stability:

- Partitioning and layout optimization for large bucket-file datasets
- Join/shuffle tuning and refactoring "hot path" transforms
- Monitoring, operational runbooks, and predictable release processes

In a high-volume telecom environment, even small delays or instability quickly erode confidence. We focused on performance and reliability, so the platform stayed fast, dependable, and trusted for day-to-day decision making.

Self-Service Analytics Transformation and Change Management

1) Contour Enablement: From Canned Reports to Self-Service

The client's intent was not only data processing modernization, but also self-service analytics adoption. Foundry's **Contour Dashboards** were central to shifting behavior.

Our adoption approach:

- Built an initial set of Contour assets aligned to high-value reporting needs

“ In a high-volume telecom environment, even small delays or instability quickly erode confidence.

“ Users gained control without requiring developers for every new insight, and teams moved closer to real-time, decision-support analytics.

- Conducted working sessions and hands-on training so users learned by building
- Provided time-boxed support for teams as they transitioned to self-service
- Over time, the environment grew to **hundreds of Contour lineages**, reflecting broad user-driven analytics creation and reduced dependency on developers

2) Power BI as a Practical Bridge to Overcome Resistance

Contour adoption faced early resistance, especially from teams used to familiar reporting tools. We avoided disruption by introducing **Power BI as a transitional layer**, providing continuity for key stakeholders while training and adoption matured. As self-service usage increased, reliance on Power BI declined and the organization shifted to Foundry-native analytics.

3) Quiver for Flexible, Intuitive Exploration

To improve usability and analytic flexibility, we also used **Quiver** alongside Contour. This strengthened self-service by giving users more intuitive exploration and interaction patterns while maintaining governance.

Result: users gained control without requiring developers for every new insight, and teams moved closer to real-time, decision-support analytics.

Lightweight Operational Use Cases in Foundry

In addition to analytics, the client used Foundry for small-scale operational workflows that benefit

from governance and traceability but are not high-concurrency transactional systems, including:

- **Reference data management**
- **Metadata management** and controlled edit workflows

These solutions fed downstream systems while maintaining auditability and controlled updates.

AI/GenAI Innovation: Conversational Analytics with AIP

Most recently, we implemented conversational analytics using **Palantir AIP**:

- Delivered as a POC and moved into production
- **Live in production ~3 months**
- Enables users to ask questions in natural language and receive governed, data-backed analytic responses grounded in enterprise datasets and semantics

We are actively expanding AIP to integrate conversational analytics and predictive insights into operational workflows, especially around **project health monitoring**, so insights show up where work happens rather than staying trapped in dashboards.

Outcomes and Impact

This engagement delivered transformation across platform, operations, and user behavior:



Explore Our People-Driven, AI-Empowered Approach.

Learn how we help unlock the power of your data, providing intelligent solutions that drive innovation and operational excellence.

Engineering and Platform Outcomes

- Modernized a complex legacy footprint including **120+ databases** and **100+ SSIS jobs** into scalable, production-grade pipelines
- Improved operational trust through lineage transparency, health checks, and observability-driven support practices
- Enabled near real-time patterns for analytics and operational readiness (alerts, predictive foundations, AI enablement)

User and Adoption Outcomes

- Shifted teams from stale daily reports toward **near real-time, self-service analytics**

- Reduced dependence on centralized reporting development through Contour/Quiver enablement and training
- Used Power BI strategically as a bridge to maintain continuity and accelerate adoption rather than forcing disruption

Innovation Outcomes

- Delivered a GenAI-enabled conversational analytics capability using AIP (production and expanding)
- Established a foundation for embedding predictive and conversational insights into operational workflows (project health monitoring and beyond)

Why This Past Performance Stands Out

This program reflects the full arc of enterprise modernization done the hard way—and done right:

- Deep technical migration (SSIS/SQL – PySpark/Spark and cloud-native processing)
- Data Mesh-aligned governance and decentralized ownership enablement
- Near real-time and operational analytics readiness (alerts, predictive foundations)
- Real change management at scale (Contour adoption, Power BI bridge, training and documentation)
- AI/GenAI enablement in production (AIP conversational analytics)
- Execution resilience (minimal vendor implementation support, capability built through POCs and self-training)

We're here to help you succeed. Cadmus provides government, commercial, and other private organizations worldwide with technology-empowered advisory and implementation services. We help our clients achieve their goals and drive lasting, impactful change by leveraging transformative digital solutions and unparalleled expertise across domains. Together, we are strengthening society and the natural world.

For more information, visit cadmusgroup.com.