

Foundry Streaming

Use real-time data to drive real-time decisions — with the full context of your organization.

palantir.com

Copyright © 2023
Palantir Technologies Inc.

All Rights Reserved



Foundry Streaming Overview

- 01 How it Works
 - 02 Impact
 - 03 Reference Architecture
-



Introduction

Across businesses and industries, the ability to make critical decisions in near real-time — based on real-time data — has become a requirement. From asset monitoring in manufacturing to flight swaps in aviation, missed signals and momentary delays can have negative immediate consequences and destructive ramifications.

To keep up, many organizations have invested in complex streaming infrastructure, stitched together patchworks of services, and hired skilled software engineers to manage an unwieldy landscape. However, even with a solid technical foundation, real-time action depends on more than just capturing the right information.

Organizations also need the ability to contextualize real-time data alongside operational data and processes. Only then can they surface, understand, and take action on thousands of transient, ephemeral events in near real-time.

That is the power of the Foundry Ontology.

The [Foundry Ontology](#) drives responsive operations with object-aware alerting, analytics, and action-centered applications, unlocking a unified end-to-end streaming experience with:

Automated state monitoring

The sheer volume of streaming data precludes manual identification of risks and opportunities. On top of this, the events captured by such data streams are inherently fleeting, and easily lost in a sea of noise.

With the Foundry Ontology, organizations can harness streaming data to automatically generate alerts and escalate them to human operators or algorithms for analysis and action in near real-time.

This means the right data gets to the right place at the right time.

Data in context, always

Just as readings from a temperature sensor on a machine are all but meaningless without an understanding of maintenance history, operator notes, and equipment specs, data from even the lowest-latency pipelines must be put in context to properly inform action.

The Foundry Ontology is built to fully contextualize low-latency data alongside other relevant batch data (e.g., ERP and MES systems), and to enrich it with AI/ML models for advanced analytics.

By binding streaming data to their ontology's objects — whether a manufacturing robot, a wind turbine, or a hospital bed — organizations have a real-time, all-encompassing view of the entities, relationships, and events that make up their business.

This complete picture of past and present intelligently informs the future.

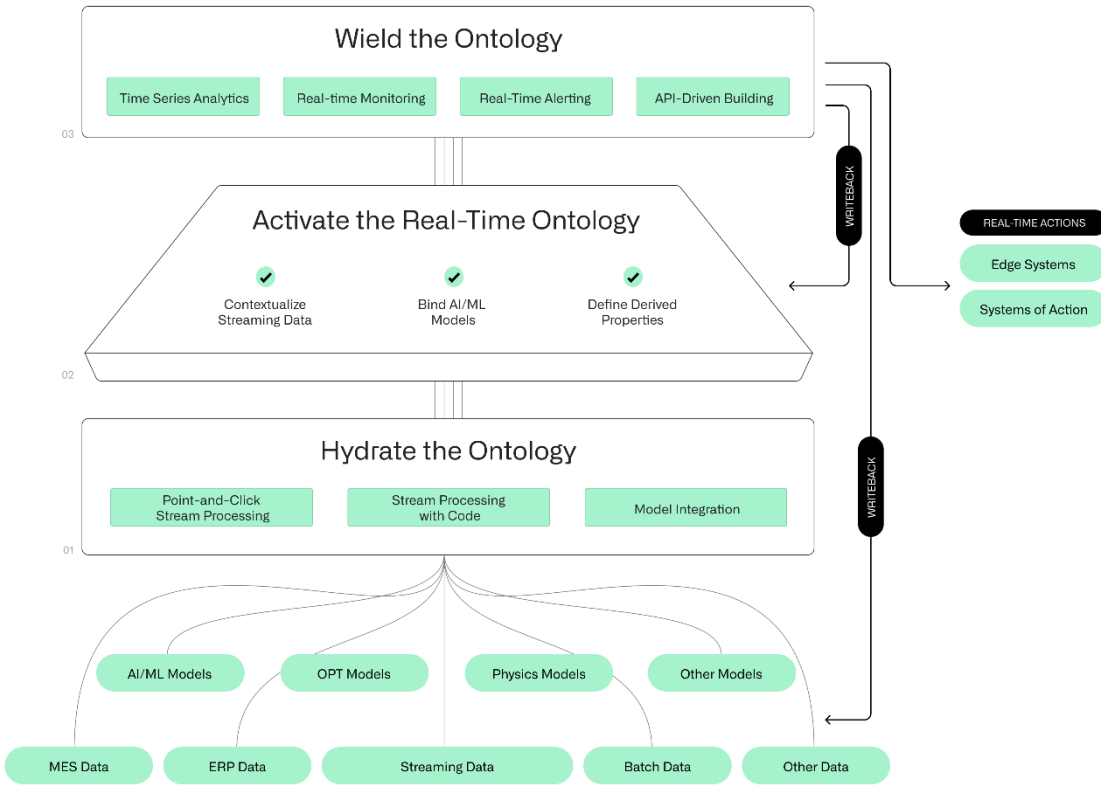
Closed-loop action

To take data-driven action at speed (and on a repeatable basis), organizations need seamless connections between real-time data, real-time analytics, and real-time operations.

The Foundry Ontology not only delivers object-aware workflows that allow operators to take immediate action based on alerts and associated context, but also writes back to systems of action and the organization's ontology, generating compounding value over time.

How it Works

- 01 How it Works
- 02 Impact
- 03 Reference Architecture



01 Hydrate the Ontology

Foundry's out-of-the-box connectors make it easy to connect with systems that generate relevant data and models.

02 Activate the Ontology

Bind streaming data to your ontology's objects and actions — delivering complete context to streaming data and creating the foundation for closed-loop action.

03 Wield the Ontology

Enable your operators in the field to identify risks and opportunities — and act in real-time. The Foundry Ontology orchestrates writeback to operational systems, fueling continuous improvement over time.

01 Hydrate the Ontology

Foundry's out-of-the-box connectors make it easy to connect with systems that generate relevant data and models.

Integrate streaming data sources with Foundry's out-of-the-box connectors and stream creation wizard.

A guided UI walks users through integrating existing streaming sources such as Apache Kafka, Google Pub/Sub, OSI PI, Azure Event Hub, and Amazon Kinesis. With [out-of-the-box connectors](#), you can start pulling historic and livestreaming data into your ontology in minutes.

A stream creation wizard provides standard REST API endpoints to push streaming data directly into your ontology, programming language-agnostic. This means users can customize their IoT architecture while delivering secure, scalable, and performant streaming capabilities.

Transform and process streaming data at scale while continuously ensuring data quality.

[Pipeline Builder](#) combines the ease of low-code with rigorous change management and health monitoring. It allows you to filter, join, and aggregate data streams, and land the outputs directly in your ontology.

Users can conduct Flink-based transforms, creating custom, user-defined transformations using open standards from a full code-based stream processing workbench.

[Data Health](#) is a Foundry service that monitors for issues across datasets; in the event of a failed check, Data Health sends in-platform notifications and emails with alerts of the failure. Foundry allows users to set mechanisms for handling data quality issues, which are then applied to all data in the stream.

Integrate AI/ML models with built-in connectors.

With a [suite of interoperable connectors](#), you can bring externally developed models into your ontology, making them easily consumable by other users, pipelines, and operational applications.

02 Activate the Ontology

Bind streaming data to your ontology's objects and actions — delivering complete context to streaming data and creating the foundation for closed-loop action.

Connect data streams to Ontology objects.

The Foundry Ontology is built to contextualize all streaming and low-latency data with other relevant information, such as structured, unstructured, geospatial, and media properties.

Bind AI/ML models to objects.


Bind process optimization algorithms that run on streaming data to your ontology's objects and processes to generate predictions, detect patterns and anomalies, and perform other advanced analytics tasks.

Set a cadence for ML models to retrain on new data or decisions captured, empowering data science teams to visualize and analyze the performance of models over time.

Perform real-time operations.

Perform complex calculations — for example, combining and interpolating streams or applying regression analyses on your streaming data — to derive real-time properties such as machine performance KPIs (e.g., OEE) and expected yield.

These derived properties are fed into your ontology, where they are displayed alongside other object properties and can be leveraged by any application to power intelligent decision-making.



03 Wield the Ontology

Enable your operators in the field to identify risks and opportunities — and act in real time. The Foundry Ontology orchestrates writeback to operational systems, fueling continuous improvement over time.

Analyze real-time data.

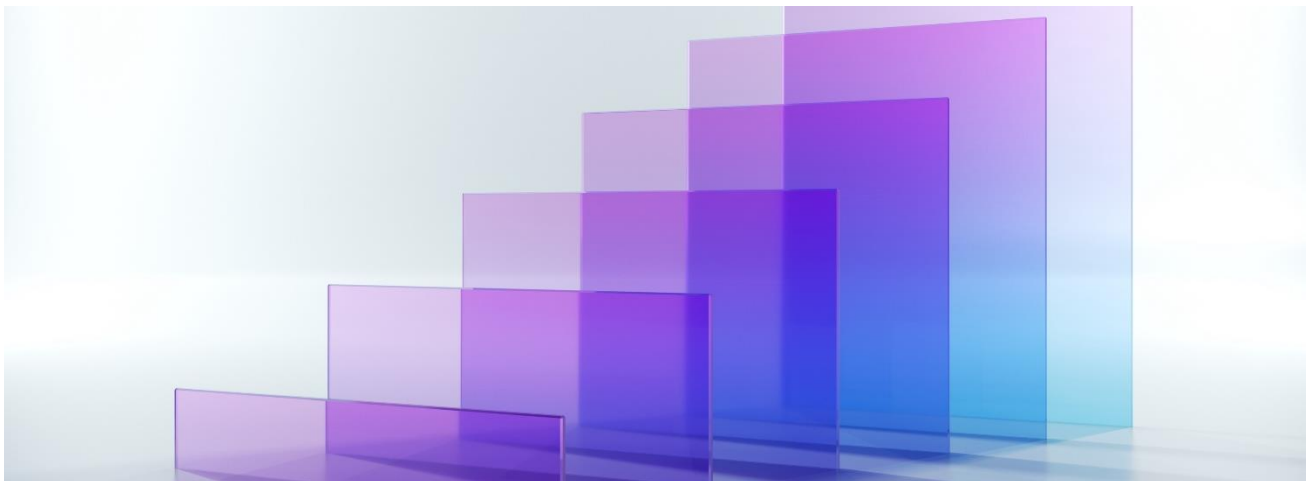
Visualize and analyze time series and object data in a point-and-click interface and charting library using Foundry's analytical and dashboarding suite, [Quiver](#). Apply object-based transformations such as batch analysis or derive properties on thousands of objects. Build and publish analyses into interactive dashboards that embed in operational applications.

Compose real-time alerts.

Proactively monitor objects and create alerts with [Foundry Rules & Real-Time Alerting](#). Set criteria for real-time alerts, route alerts to downstream applications, and configure notifications via email, in-platform, or through third-party integrations for custom notification services. Full compatibility with Quiver means operators can incorporate rich conditions into their rules.

Power closed-loop actions.

Build rich Ontology-powered applications in [Workshop](#). Operators use these applications to evaluate alerts from Foundry Rules & Real-Time Alerting, analyze published views from Quiver, and take action. The Foundry Ontology coordinates actions from Workshop applications to underlying systems of record, as well as writeback to your ontology's objects and to your source systems, ensuring your entire enterprise remains in sync.



Impact

Jacobs + Palantir

- 01 How it Works
- 02 Impact
- 03 Reference Architecture

Blower Optimization at Agua Nueva

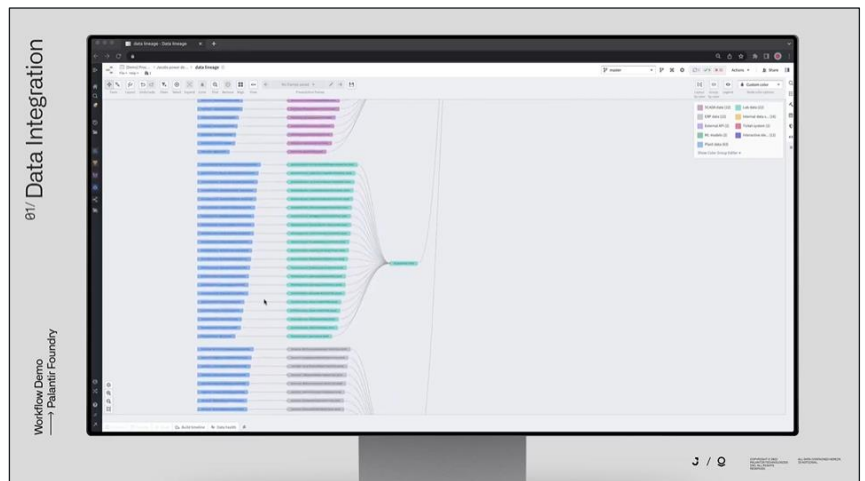
~20%

Reduction in electricity usage plant-wide

Global solutions firm Jacobs deployed Foundry at Agua Nueva Water Reclamation Site in Pima County, Arizona to accelerate the site’s digital transformation and implement predictive operations across the plant — a proof of concept that could scale across more Jacobs sites, and to their clients.

At Agua Nueva, [Jacobs used Foundry](#) and their Smart Algorithms to build a solution that optimized blowers (which consume the majority of electricity at water treatment plants) across the site, directly resulting in a **~20% reduction in electricity usage plant-wide.**

The ability to harness sensor data – as well as the connectivity between data, analytics, and operations enabled by the Foundry Ontology – was critical to these savings.



→ To begin, Jacobs used Foundry to integrate real-time sensor data, lab data, CMMS data, Oracle Financials, and more. This data is fed into their ontology, meaning that Jacobs can see a complete picture of operator usage, electricity consumption, model error and drift, and more, all in a centralized platform.

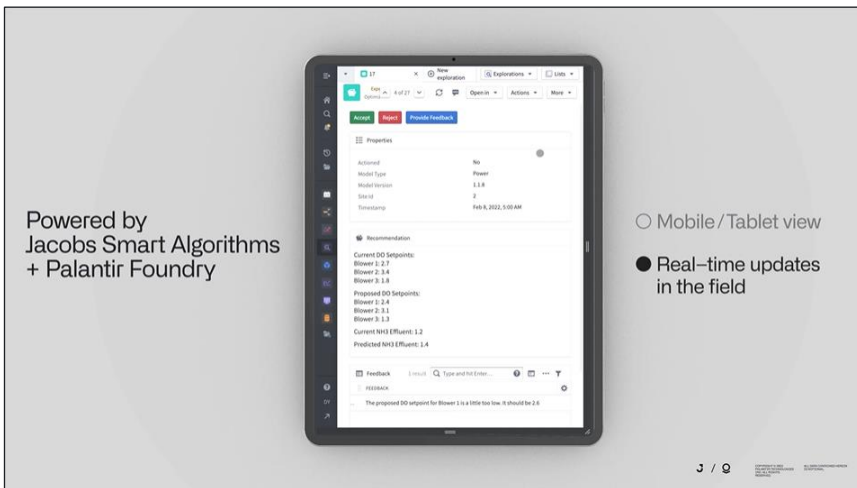
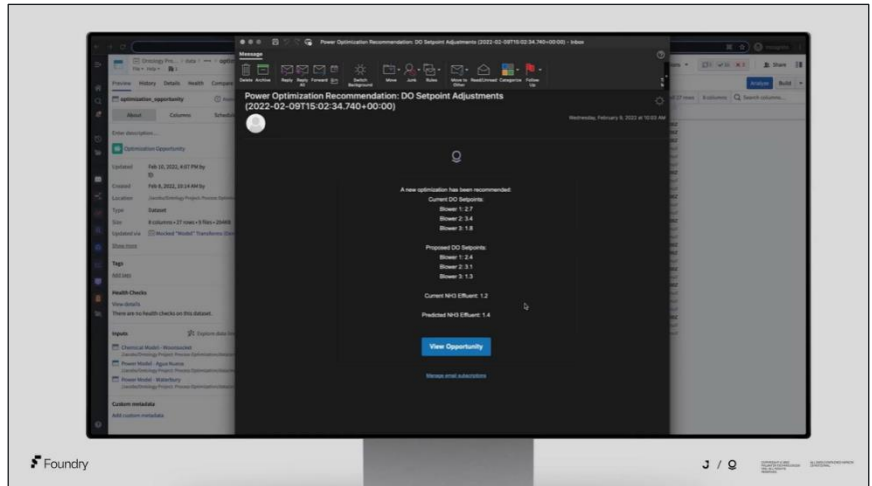
Impact

Jacobs + Palantir

Sensor data is then run through Jacobs' process models, managed by a remote data scientist, to make real-time predictions on ammonia levels in the plant's aeration basins.

These predictions are then compared against required compliance levels and mapped into simple, human-readable emails recommending updated set points for the blowers.

Operators receive these email notifications twice daily, timed to their planning meetings, so they can easily review and quickly action blower optimization for maximum power savings.



Powered by
Jacobs Smart Algorithms
+ Palantir Foundry

- Mobile/ Tablet view
- Real-time updates in the field

When data scientists deploy models, these models are fed real-time information from the plant, providing them with both current and forward-looking indicators to suggest set points for the blowers.

Should the site be experiencing extenuating circumstances, operators are able to reject notifications and provide rich text feedback to data scientists for review.

This ensures operators remain in ultimate control while directly engaging with the data science insights they need to achieve optimal performance at the plant.

Reference Architecture

- 01 How it Works
- 02 Impact
- 03 Reference Architecture

