Teck

INDUSTRY
Mining

TECHNOLOGIES
Google Cloud Platform, Google Cloud Storage, Google BigQuery, Google Datalab, Google Dataprep, Google Compute Engine, Google Cloud Endpoints, Google Apps Engine, Scikit Learn, TensorFlow, Apache NiFi, Apache Airflow

BUSINESS NEED
Teck, a major Canadian mining company needed to optimize haul truck operations by predicting failures. Having access to vast amount of IoT data from operating machines, Teck turned to AI and Machine Learning to produce unique actionable insights.

SOLUTION
Pythian leveraged raw telemetry data to produce machine learning models and built a production solution on Google Cloud Platform that's embedded into the company's existing on-premises systems.

RESULT
Teck now has an AI solution to predict failures and plan maintenance, avoiding costly downtime.

DATA SCIENCE KEEPS WORLD'S BIGGEST TRUCKS ROLLING

Teck Resources, one of Canada's largest diversified energy and mining companies with business units focused on gold, metallurgical coal, zinc, copper, and energy needed to optimize the efficiency of its mining haul trucks. These trucks are some of the biggest and most costly vehicles to run. Because operating these gigantic machines represents 40 percent of mining site costs—and each minute of productive operation generates revenue—the trucks must keep moving day and night.

A modern mining haul truck is a million-pound IoT device on wheels producing 2GB of raw data per day and this presents a major opportunity to use AI for optimizing performance. The company needed to harness the power of the sensor data to predict costly issues such as electrical failures and suspension degradation before they happened. This would involve analyzing terabytes of raw operational sensor and alerts data, coupled with maintenance, scheduling, and other truck lifecycle records to produce actionable insights.

SOLUTION

Teck chose to partner with Pythian because of their proven expertise in data science and for their deep knowledge and experience implementing innovative solutions on Google Cloud Platform (GCP). By following Pythian's iterative data science framework Teck could make well-educated investment decisions at each step of the project, while carefully managing the uncertainties typically found in data science projects.

Collaborating with Teck's subject matter experts, Pythian's Enterprise Data Science team assessed opportunities to optimize the operational capacity of the machinery by applying Machine Learning solutions to

Teck

IoT data flowing from the haul trucks in the field. After one month of assessment, Pythian and Teck jointly selected the top use case based on a risk/reward analysis. Pythian then ingested all historical data into GCP and within two months built a proof of concept (POC) solution for the top use cases. As the POC demonstrated good predictive performance, it was followed by two months of implementing the haul truck failure prediction and maintenance recommendation solution in production on GCP while also fine-tuning the models to be more robust. Pythian embedded the resulting insights into the company's on-premises systems and made available for end users in their already-familiar tools to help them make timely maintenance decisions. Since the production launch, Pythian has been helping the company maintain the operational state of the predictive system closing the full cycle of AIOps.

The use of GCP scaled as the project evolved, while Teck's engineering team continuously acquired deeper and deeper experience with it. In the initial stages, Pythian made use of such products as Google Cloud Storage, Google BigQuery, Google Datalab, Google Dataprep and Google Compute Engine. For machine learning modeling, Pythian used well-recognized open tools such as Scikit Learn and TensorFlow. Google Cloud Bigtable has been used as a wide columnar sparse data store for thousands of sparse pre-aggregated features that allowed Pythian to iterate rapidly through large-scale feature engineering as part of modeling iterations. During production implementation Pythian deployed a custom API layer on Google Apps Engine, implementing unique business logic around machine learning models and exposing a business-specific predictive API via Google Cloud Endpoints for secure and authenticated consumption of resulted predictive microservices. Pythian also leveraged Apache NiFi to handle data ingestion from the company's on-premise systems and Apache Airflow as job orchestration layer. (The latter is now available as a fully managed service in Google Cloud Platform - Google Cloud Composer.) Thanks to Google Stackdriver and appropriate instrumentation of each component, the state of all components of the predictive system are easily monitored.

Teck

Figure 1: Google Stackdriver AIOps summary dashboard

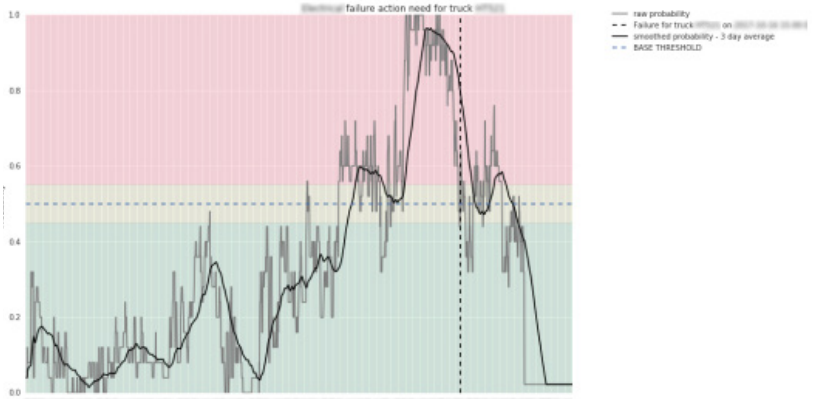

Figure 2: Failure predictions on the timeline indicating increasing probability of failure before the actual failure event.

Pythian
love your data®

Teck

Notable aspects of the project include:

- Available maintenance events were carefully analyzed to produce expanded labeled training data on time series data.
- Pythian performed feature engineering from time-series data with a focus on dimensionality reduction. Given the low number of training samples and a high degree of data complexity, standard high-powered time-series approaches (RNNs, variant CNNs, etc.) were not feasible.
- Pythian designed features and selected models based on low-complexity approximations of high complexity models that preserved relevant time-series information and avoided overfitting.
- Pythian combined generic time-series data aggregations to handle a large variety of sensors and alerts at scale with domain-specific feature engineering.

The collaboration between Pythian and Google has positioned Teck to unlock further value from its unique data and to continue to leverage advanced data analysis and AI capabilities of GCP. These new capabilities are turning Teck into a leading data-driven organization. The solution built by Pythian on GCP is easily expandable - beyond working with end users and enhancements to existing models, the next steps anticipated are growing the number of use cases in optimizing industrial assets maintenance and other optimization of mining operations as well as expanding into other business domains at the company.

RESULT

Pythian used IoT data from truck sensors and other sources to produce a predictive solution that helps minimize costly downtime by predicting failures and recommending maintenance. The resulting solution was a machine learning application built on GCP with integrated IoT data ingestion and predictive microservices embedded into Teck's on-premises products via REST APIs. The ability to anticipate events such as catastrophic electrical failures and suspension performance degradation allows the company to optimize maintenance schedules and reduce trucks downtime.

ABOUT PYTHIAN

Pythian excels at helping businesses around the world use data and the cloud to transform how they compete and win in the data economy. From cloud automation to machine learning, Pythian leads the industry with proven innovative technologies and deep data expertise. For more than 20 years Pythian has built its reputation by delivering solutions to the toughest data challenges faster and better than anyone else.

WORLDWIDE OFFICES

Ottawa, Canada
New York City, USA
London, England
Hyderabad, India

© The Pythian Group Inc., 2019

Pythian®
love your data®