Case Study: Smart Factories

Solving Real-time Data Challenges for Predictive Maintenance

Background

Manufacturing around the world is currently undergoing the next phase of digital transformation. With sophisticated supply chains and the introduction of "just-in-time" processes, they've moved from mass production to mass customization in an effort to meet increasingly high customer expectations. These efforts are not simply a by-product of operational efficiency and improved practices. Instead, they result from a digitization effort in manufacturing that increasingly connects devices and platforms to automate efficiency gains. But while companies strive to create a competitive edge through investments in IoT, AI, and the Cloud, success really relies on their transformation towards becoming software-enabled organizations.

In general, manufacturing processes typically involve a significant amount of planning, scheduling, and tracking. Establishing high-performance operations that yield a competitive advantage usually requires a thorough historical data-based analysis. There's a huge amount of data already generated and stored on the factory floor. Recent efforts in digital transformation tend to focus on architectures where edge gateways periodically collect data, transmit that data in batch format to a cloud-based data lake, analyze the more recent data against historical trends, and finally provide visual summaries through business intelligence tools. While this approach may improve long-term planning, it adds an increasingly complex data infrastructure that remains reactive and not proactive.

The future of manufacturing needs to adapt to the increasingly dynamic needs of the customer. Data generated by most factories will be too voluminous to send to the cloud. Batch processes will add too much latency for making real-time decisions. And adding data scientists and algorithms will be too complex for many organizations, much less relying on a 3rd party without the institutional knowledge of the factory. A different approach is needed – one that drives real-time insights – as factory operators and line engineers often make complex decisions each and every day.

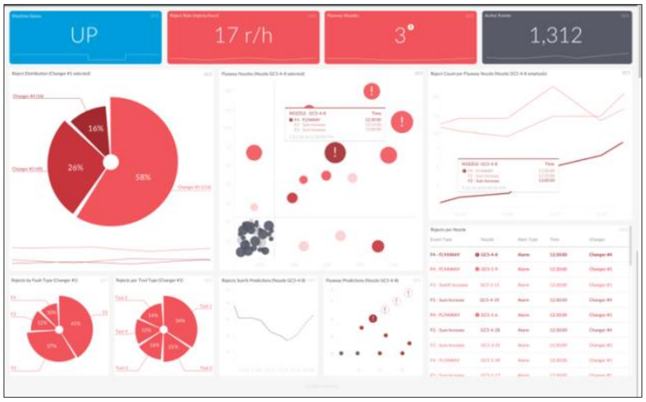

Figure 1. User interface of the Swim-enabled real-time monitoring and predictive maintenance application.

Overview

An innovative company with a long tradition in industrial automation recognized the opportunity to incorporate their operational knowledge and expertise into software applications for creating "Smart Machines." As some factories invest in equipment designed to operate for decades, they needed to create flexible solutions that handle large volumes of data locally while delivering real-time performance, and ultimately help create a more pro-active culture that focuses on operators and not data scientists. Updating any of the existing equipment on the factory floor was not an option, especially for facilities that ran 24/7 operations, as it would cause machine down-time and impact production. The alternative of collecting, transmitting, and storing machine data with a traditional database-first architecture would incur significant network and infrastructure costs, require a complex design for handling the large volumes of data, and inevitably adds latency. Finally, users would need to identify the data to which they needed access. Without the clear definition of an application, traditional business intelligence tools might lend a user towards a data deluge and "analysis paralysis." A focused application framework that delivers the right insights at the right time was necessary to create transformation in a cost-effective timeframe.

The company evaluated several alternatives in the market and decided on Swim for a solution.

Solution

Designed from first principles, Swim is the world's first application platform for building stateful, real-time, distributed applications. It provides a single, unifying code base that can run locally at the edge or in the cloud to analyze, visualize, collaborate, and act on streaming data as it's created.

The creation of an application in Swim starts with building Web Agents – self-contained objects that manage their own data streams, possess individual logic, and continuously communicate their state to any subscriber. These Web Agents compose naturally and allow the simple creation of complex applications that can be easily distributed across infrastructure. The design pattern provides the simplicity and flexibility for a repeatable process while drastically reducing the design time and implementation costs from complex data modeling.

For this application, a Web Agent was built to replicate the production run for a given machine type on an electronic assembly manufacturing (EAM) line. This included the ability to distinguish the different data elements and continuously interpret performance during runtime. For example, some of these data elements included the identification of faults in the assembly process, either collected from a machine identified error or a visual indication from automated cameras in the quality assurance process. These faults were continuously processed by the Web Agent, significantly reducing data volumes by transforming raw data into pertinent information.

Web Agent enabled the application of Machine Learning (ML) algorithms and operational logic to interpret and act on results. When the Web Agent identified a machine was trending outside of acceptable limits, it would categorize the anomaly and generate an alarm, after removing false positives and ensuring the anomaly met the applicable threshold criteria. Each data element managed by a Web Agent is immediately available through an individually addressable URL to any authorized subscriber.

Using WARP, a multiplex streaming upgrade to HTTP, Swim is able to create a mesh of application services which continuously share data streaming insights with each other in network real-time.

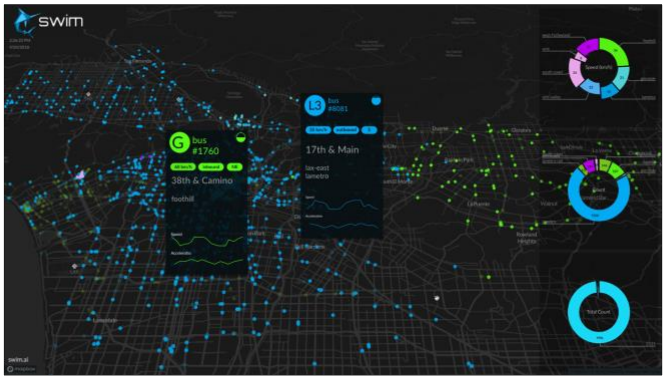

Figure 2. Swim can monitor assets whether they're in a single location or on-the-move and uses machine learning to identify real-time insights.

This composability allows the machines on a production line to share their data models and analytic information with each other. It also enables the creation of live user interfaces with charts and graphs that summarize the analysis and make it available in real-time. The entire Swim framework enables a more dynamic manufacturing environment, where line engineers receive immediate notification of potential issues and have the ability to diagnose machine behavior patterns on the factory floor.

Swim leverages any available compute on the factory floor and provides a variety of deployment options that can be optimized to meet their customer's needs. Each machine connected to the application is connected through a distributed cloud infrastructure that enables the factory to perform detailed analytics on real-time production information. Swim is the only application platform designed to scale this way on distributed compute and it helped the company avoid adding expensive infrastructure on existing equipment and the factory floor.

Conclusion

Our partner's application has been deployed in U.S. factories and seamlessly analyzes over a 100 million data points across thousands of production runs. Becoming a software-enabled organization helped the company create "Smart Machines," with the ability to automate failure detection and prevent downtime, improving quality and reducing waste.

Learn More

You can learn more about this application and more by visiting **www.swim.ai**