Case Study: Data-as-a-Service

Creating a Data Exchange for Connected Traffic Infrastructure

Background

The transportation industry is undergoing a major digital transformation. In an effort to support Smart City and autonomous vehicle applications, transportation service providers are embracing vehicle-to-infrastructure (V2I) communications, or the sharing of information between vehicles and infrastructure in real-time. Streaming data from connected infrastructure can provide additional context to improve passenger safety and reduce congestion and emissions through efficient routing. While investments in new technologies such as 5G networks, LIDAR sensors, and GPU compute will provide better infrastructure for cities and vehicles, the volumes of data generated and communicated by connected traffic infrastructure makes it difficult to act in a timely manner.

As an example, the data created by traffic infrastructure in Palo Alto, CA can generate more than 30,000 data points per second from streaming traffic lights, vehicles and pedestrians. With just over 100 connected traffic intersections, Palo Alto's traffic infrastructure will generate more data points per second than all of Twitter on any given day.

There are more than 300,000 intersections and 270 million vehicles in the U.S. alone, so it's simply not cost effective to process all this streaming data using traditional software architectures. It's also impractical for government agencies, logistics companies, routing companies, and insurance agencies to be able to consume and process that volume of data for public or private use. While improving our transportation infrastructure is a move in the right direction, it won't be enough to make effective use of the data generated. A different approach is needed – one with extremely low latency – as drivers and autonomous vehicles often make decisions on a sub-second basis for their safety.

Figure 1. User interface of the Swim-enabled real-time traffic & infrastructure monitoring application.

Overview

An innovative company specializing in traffic infrastructure wanted to transform raw data from their equipment into a real-time streaming API for connected and autonomous vehicles. They faced several challenges in building a solution that could be cost effective, deliver real-time performance, and scale across distributed infrastructure. While the company already owned compute infrastructure across agencies around the world, none were suited to run a complex application at the scale required by the data. Even more, traditional architectures incur latency, reducing the effectiveness (and value) of the data. Finally, users of the data needed to pull discrete pieces of the data available, as opposed to consuming a "firehose" of all available data. Any vehicle traveling through a particular city may have a unique destination and is only interested in subsets of the traffic data available from particular intersections along their route. While weather API's solve this addressability with individual city requests (a many-to-one relationship), connected and autonomous vehicles pose the challenge of requiring several requests to multiple discrete sources (a many-to-many relationship).

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, unified 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.

Swim applications start with the building of 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 provide a simple way to create complex applications that can be easily distributed across infrastructure. This design pattern provides the simplicity and flexibility for a repeatable process while drastically reducing design time and implementation costs for complex data modeling.

For this application, Web Agents were configured to replicate each traffic intersection. This included the ability to convert voltage readings from traffic lights into a logical progression related to green, yellow, and red states. It also included a data element that tracks the detection of vehicles at a light and counts of each vehicle that passes, along with the detection of each pedestrian that presses a crosswalk button. Latest state is processed and updated locally by the Web Agent as the data is generated and received. By maintaining the latest state and only transmitting updates to data subscribers, Swim drastically reduces the volume of data that sensors or other systems pass over a communication network.

Data streams processed by Web Agents are made immediately available via URI to any authorized subscriber. As data changes, subscribers receive state changes through the WARP protocol. WARP was developed by SWIM.AI as a multiplexed streaming upgrade to HTTP and enables the creation of bidirectional streaming links between distributed services. Using WARP, Swim is able to create a mesh of application services which continuously share data streaming insights with each other at network real-time.

This directly solved the company's challenge in delivering real-time data without incurring unnecessary latency. It also solved the challenge of enabling a range of unique data subscriptions to live data streams through individually addressable APIs.

Finally, the solution leverages available compute on the existing infrastructure of each city to run a distributed application. A simple software update is all that is required for Swim to instantiate Web Agents for each intersection. Every city with the application is then aggregated in the cloud to enable the company to manage and publish real-time streaming traffic information to their subscribers. Swim is the only application platform designed to scale this way on distributed compute and it helped the company avoid adding expensive infrastructure for each city.

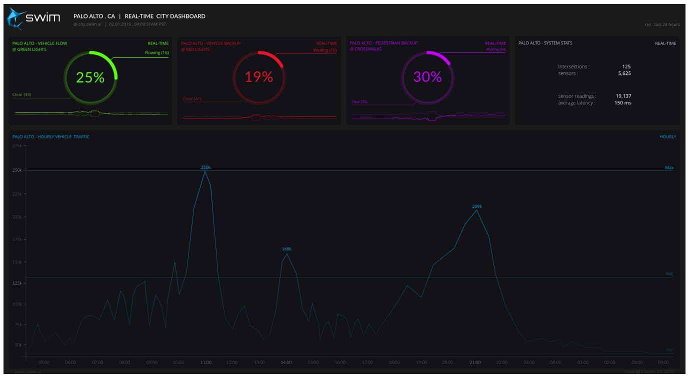

Figure 2. Swim enables the monitoring and introspection of connected traffic infrastructure in real-time.

Conclusion

Swim enabled our partner to successfully launch a real-time, streaming traffic information service with unmatched performance and scalability. The application has been deployed across a number of U.S. cities. Companies that subscribe to the service now receive updates with under 150 millisecond latency. Furthermore, Swim reduced the volume of data passing over the network by 99.2%.

Learn More

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