INTELLIGENT PREDICTIVE REPAIR USING DIGITAL TWINS

The auto industry is rapidly embracing a digital transformation on multiple fronts. On the one hand, online and direct sales models are trending along with the widening consumer adoption of a digital economy. Manufacturers are experimenting with digital threads to enhance collaborations between the design department and the shop floor, with the aim to find optimal design parameters early in the product cycle, rather than after multiple iterations of costly physical prototyping. During its operational phase, the most important consideration is to keep the vehicle in a healthy functional condition for the longest uptime possible. A breakdown during its ownership ruins the user experience, tarnishes the brand image, drives up the operation cost, and in a commercial fleet setting, directly affects the profitability of the business.

The concept of predictive maintenance has been making a buzz in the industry for quite a while. The benefits of reducing unexpected breakdowns between scheduled maintenance are numerous. This is achieved by leveraging machine learning models which can continuously probe the health state of major car components from a host of sensor readouts. This intelligence based approach supplants the traditional preventive maintenance, which necessarily involves discarding a significant number of perfectly good components based on a mean time to failure parameter, and still having to contend with fairly high percent of unexpected breakdowns.

BIGRIO'S PROVEN EXPERTISE IN PREDICTIVE REPAIR INTELLIGENCE

Boston based BigRio led a pioneering effort in delivering a machine learning solution to a global premium brand car company, tackling a particularly challenging component which is seeing a higher than expected failure rate. This happened at a time when not only is machine learning a recent fad, software capabilities as a whole largely remains a myth to car companies who devoted their entire competitive history almost exclusively on powerful engines and intricate hardware mechanics.

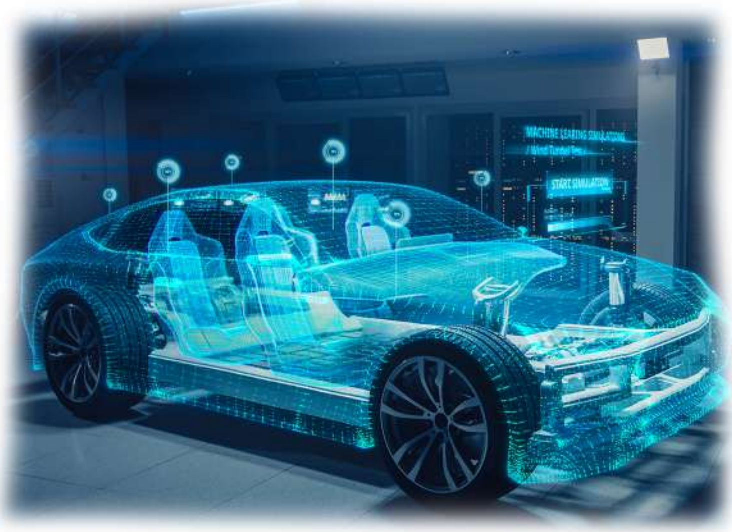

Once our team engaged in the rubber hitting the road moment, the lofty statistical theories of AI/machine learning quickly receded to the background. The real challenge in getting a working model has to do with absorbing the auto domain lingos, out-of-control data anomalies, a sketchy sample size, non-existent computing infrastructure, data security that rivals that of the financial industry, and ever-shifting directives from an executive team who is also learning on the fly what constitutes reasonable deliverables for the project, in spite of the phased and agile mechanisms we put in place.

At the end, we were able to deliver a predictive model that yielded up to 99% accuracy, with contingencies built in for market-dependent adjustments (e.g. North American vs. Far East). Along the way, the client learned their most valuable lessons in, first, the wake-up call that AI is for real and represents a strategic competitive advantage, and, second, the wider implications of digital capabilities, from data lake, cloud computing, to above all, data governance. So much so that the company and its parent organization implemented a top-down reorganization, which called for no less than spinning off a software group to carry out the enterprise wide software/AI initiatives.

THE ANATOMY OF PREDICTIVE REPAIR INTELLIGENCE

Auto predictive repair has a glamourous aspect in its use of real-time IoT data and over-the-air transmission which is built into the new generation of electric cars. The predictive model training is an almost necessary precursor to the actual real-time data collection from the car on the road to the central operational system in the back end. Bandwidth limitation dictates that only essential IoT data be transmitted with high frequency. Each transmitted data point must have an economic justification, which comes from analysis done by the predictive model. The classifier we use can enumerate the input parameters by their relative attribution to the model accuracy.

Scalability and operational efficiency are two primary concerns in deploying predictive repair. While a single model focusing on historical data makes an interesting proof of concept, applying model prediction for all the physical cars on the road means collecting and updating data for sometimes millions of units just for one particular car model. The solution comes in the form of a digital twin.

DIGITAL TWIN SOLUTION

A Digital Twin is a virtual replica of a physical entity. A car is an ideal real world product to model with a twin. A car's twin keeps in sync with the physical instance by means of over-the-air transmission of IoT sensors while the car is speeding along on the road. Every operational parameter from engine health, ambient temperature, battery capacity, road condition, to any mechatronics settings is constantly collected on the backend server. The backend hosts various simulation and machine learning models trained from years of fleet-level data to predict pending failures, and feedback advisories to the driver. In the case of planned trips, the system can test out different driving scenarios and suggest the best route for congestion avoidance and accessibility to charging stations.

A Digital Twin platform serves much more function than the caring of individual vehicles. Intelligent models often require sufficient data samples which are collected only at a fleet level. Fleet operators have the most to gain from this capability. With the volume of data they are able to access, they can create models which minimize repair expenses, forecast technical glitches, manage part inventory, schedule deliveries, while maximizing uptime. Failure modes unique to each operator can be deciphered independent of industry trends. The locations and health states of all vehicles are transparent at any moment, allowing for instant reallocation in the event of an unexpected inventory disruption.

In truth, an enterprise needs multiple versions of digital twins for their different concerns. Manufacturers have a warranty department that is motivated to keep abreast of overall fleet health and recognize faulty repairs. Design engineers want to learn insights from the failure statistics to improve future product designs. Repair shops want to be able to diagnose quickly and pinpoint the failed component when the diagnostic trouble code is ambiguous. Car owners and fleet management want to know about potential breakdowns when planning a critical trip or delivery job. The car brand as a whole wants to maintain its reliability image and be the first to know in the event of any catastrophic systemic failure. Imagine the most severe scenario of a part recall, by having the comprehensive repair and manufacturing information in the twins, a car companyhas a chance to quickly trace the problem to certain shop floor equipment or a particular parts vendor, and thus significantly limit the exposure of the recall.

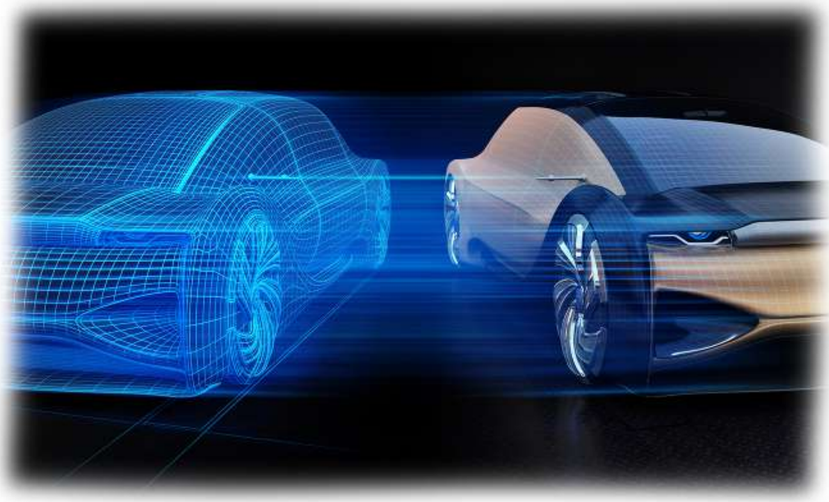

DIGITAL TWIN IMPLEMENTATION USING A CLOUD PLATFORM

Various cloud platforms provide IoT and digital twin capabilities in a hyper scalable way. They provide internet gateways to receive real-time updates from deployed sensors. By spec, an Azure twin server can host up to 1 Million twin instances with each twin carrying 5000 relationships.

For real-time processing, these platforms offer in-memory computing capabilities to analyze incoming data streams, before storing the data. This allows dashboard applications to follow a huge population of cars and perform custom analytics on the fly.

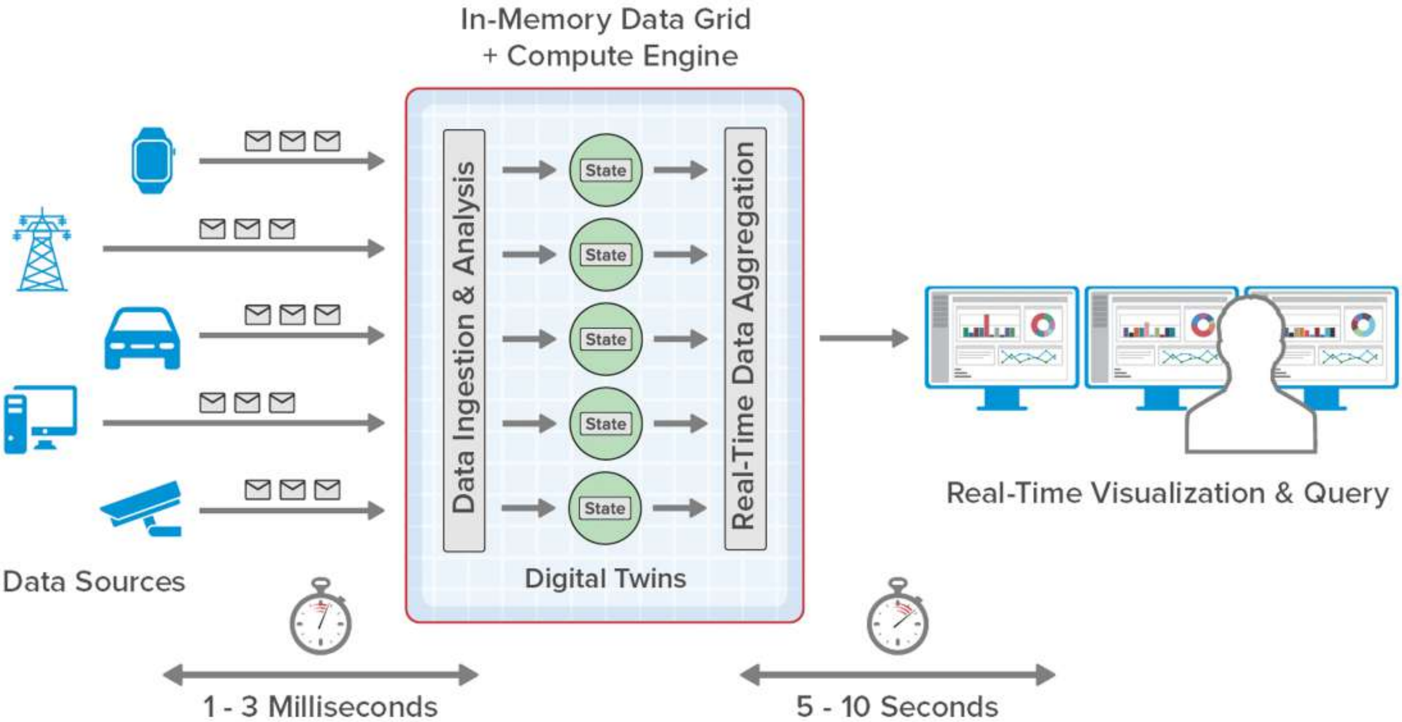

Twins tend to grow into huge graphs as the taxonomy expands. Graph database and query capabilities are all necessary parts of a digital twin platform. Query languages like the Open Cypher are the best ways for retrieving large twin graphs, with sophisticated features like variable depth search.

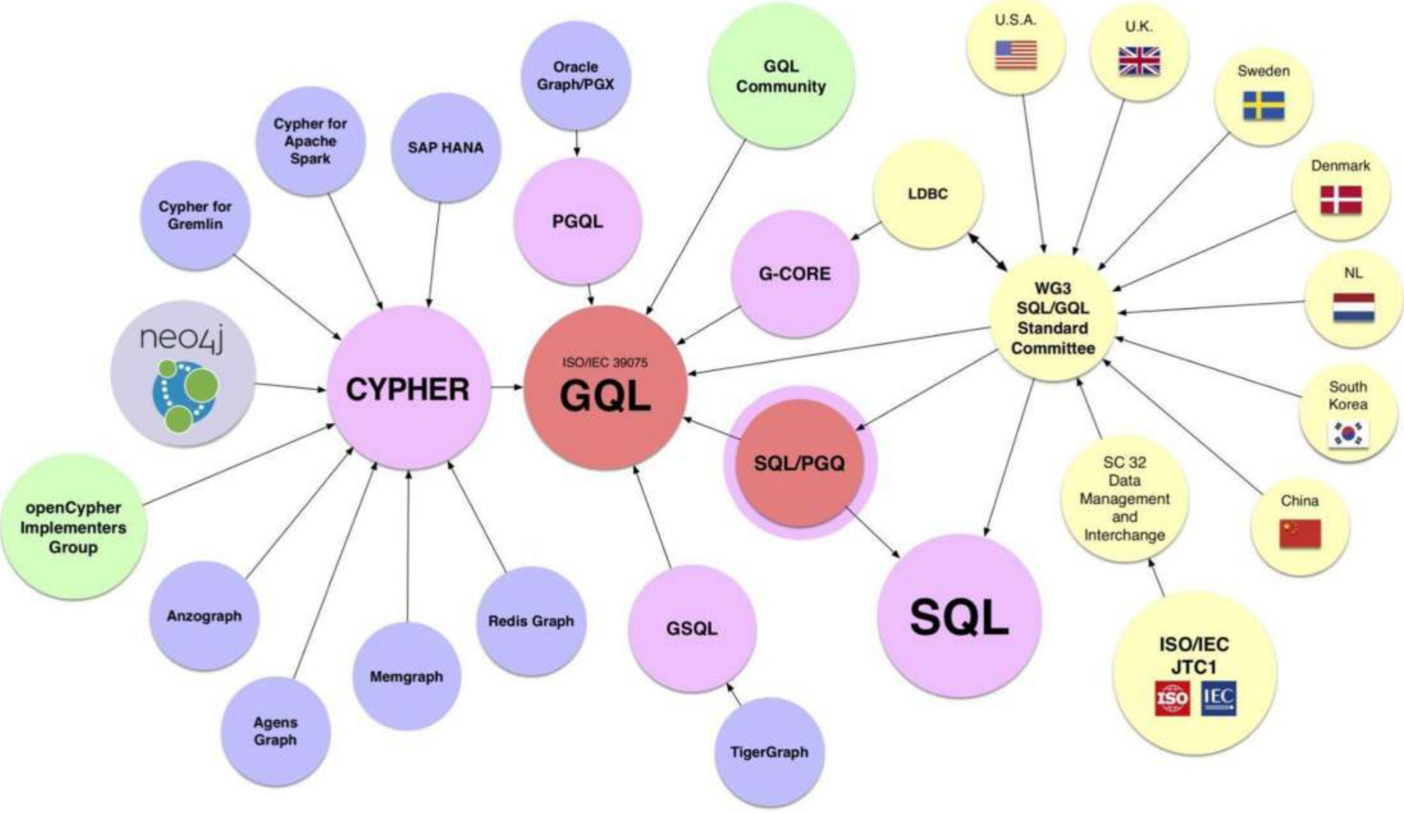

Other areas of concern is visualization of the twin graphs. Here third party visualization tools are great supplementary additions to the Azure cloud offerings.

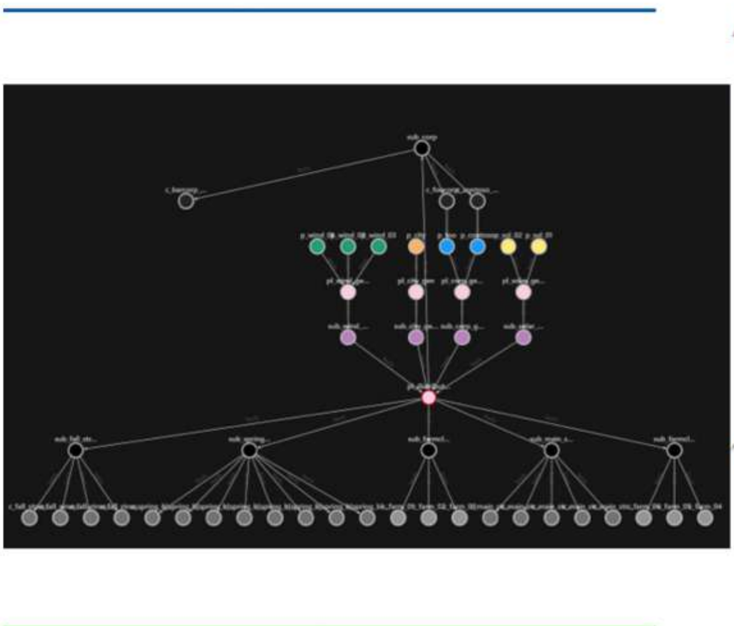

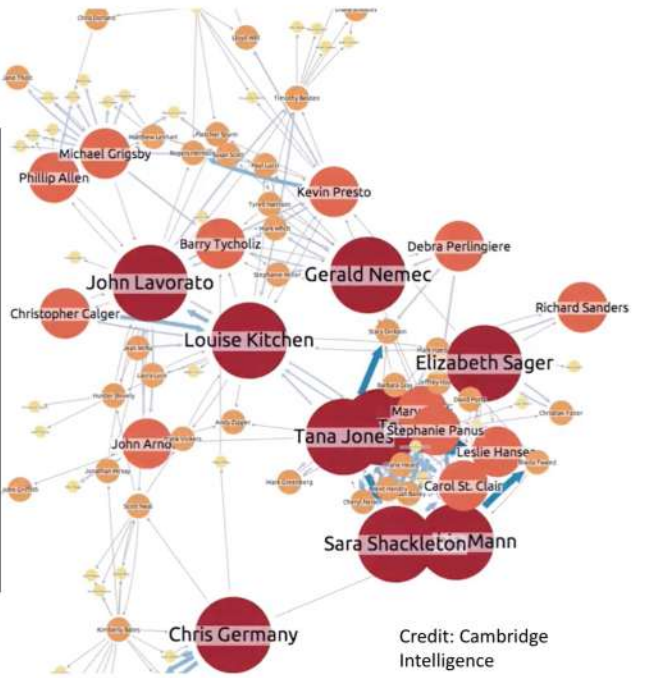

POPULATING A TWIN

A fully developed twin can capture the information on a car down to any degree of detail. While telemetry and sensor data are temporal in nature, there is much more data that are either static or semi-temporal. The component and software versions are unique to each particular car instance. Each car has a distinct history. It was assembled on a particular factory floor machinery, by a particular technician, using parts from a specific vendor shipped in an identifiable batch. At the end, a twin design may incorporate information from many departments: design, factory, sourcing, logistics, warranty, even sales.

To populate a twin from all these data sources, what is needed is the concept of a digital thread coupled with a data architecture called the data mesh. A digital thread interconnects all the digital twin information in a network throughout a car's product life cycle. At its foundation, it is a two way communication framework which facilitates the exchange of a digital product definition and bridges information silos.

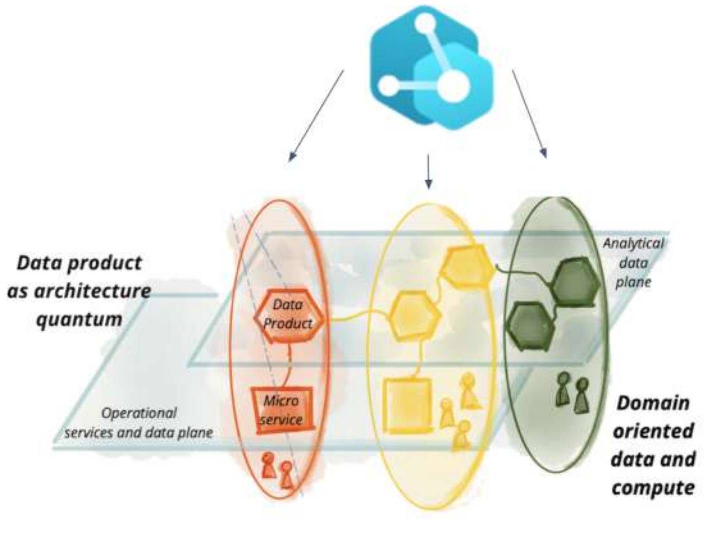

Credit: Zhamak Dehghani

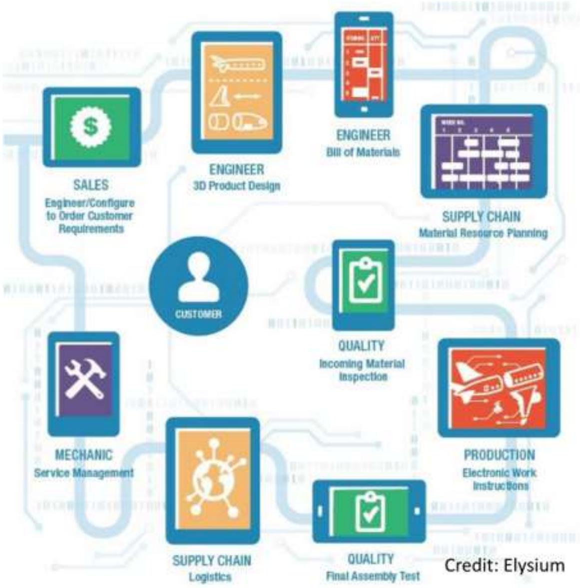

Credit: Elysium

The data mesh architecture arose out of a need to overcome the shortcomings of the traditional operational vs analytic database infrastructure, which is solely based on the concept of an ETL (Extract Transform Load) bridge. Instead of a centralized effort to create a monolithic operational data store, and pipe data to the data warehouse once a specific analytic requirement has been identified, the Data Mesh concept calls for each department to own their data and publish selected data as a product to other departments. Using the data plane idea, the consumer entity can stitch together data from different sources for a particular use case. This configuration is a perfect match for a digital thread implementation.

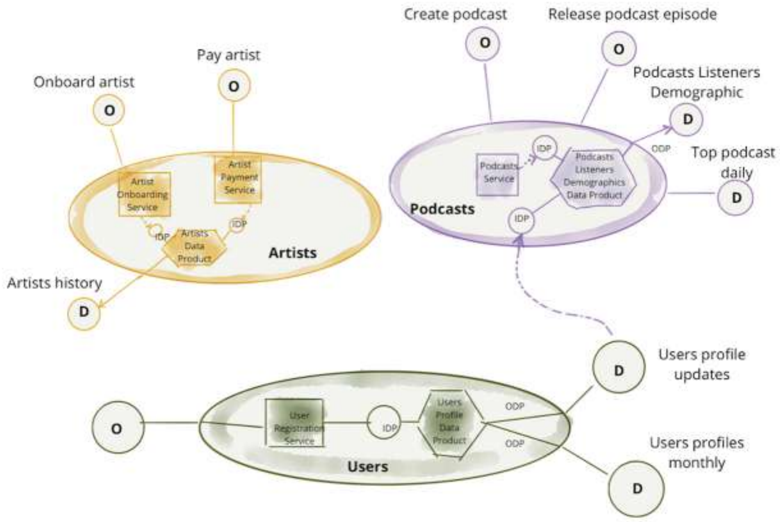

MOVING TOWARDS A FLEET LEVEL PREDICTIVE REPAIR INTELLIGENCE

Intelligent predictive repair is a very important undertaking for optimizing fleet operations. To create such a system, a fleet operator must build a digital infrastructure for collecting IoT data and capture the unique configurations of every car. BigRio is a specialty AI consultancy with proven expertise in predictive repair and data engineering. Our data engineering/cloud computing AI team can implement a complete digital twin system for fleet operators and auto manufacturers to monitor and optimize fleet operations, increase profitability, and improve business reputation. Please contact us to examine your particular needs and illustrate the vision of a digitally transformed company.

Illustrations Credit:
martinfowler.com/articles/data-mesh-principles