References implementation of operational data storage for Škoda Auto client

Our IT solutions will excite you!

Škoda Auto is the leading Czech manufacturer in the automotive sector and is a major customer of GEM System, which supplies IT solutions for the automotive, financial and many other sectors. Thanks to many years of cooperation between Škoda Auto and GEM System, they form an important alliance, creating first-class and smart IT solutions developed in professional teams. The implemented IT solutions embody the core business principles of Škoda Auto, raise the bar for service quality and improve the unrivalled customer experience for countless thousands of people around the world who are users of Škoda Auto vehicles, part of the Volkswagen Group.

Clear evidence of this is the implementation and delivery of solutions, including further support and development. The UVD project is an Operational Data Store for user and vehicle data. The entire UVD project is a key building block in the vast world of Škoda Auto's data ecosystem. It literally serves as the nerve centre of the carmaker, carefully collecting and consolidating data.The data itself represents operational reports or data warehousing, as well as a portfolio of user-focused products, including engaging marketing campaigns, a digital vehicle certificate service, the popular My Škoda app and many other key digital products.

For the implementation of the project, GEM System provided Škoda Auto with a team of software developers, data analysts and system architects with extensive experience in the field of integration. Currently, theUVD product team is involved not only in solving data assurance issues, but also in various data analysis activities:

- DATA CATALOG PREPARATION,
- DATA VISUALIZATION,
- DATA SYSTEM ARRANGEMENT,
- DATA ANOMALY RESEARCH,
- AD-HOC REPORTING AND MORE.

Today, UVD integrates dozens of user and vehicle data sources to provide its users with consolidated, cleaned, and enriched data.

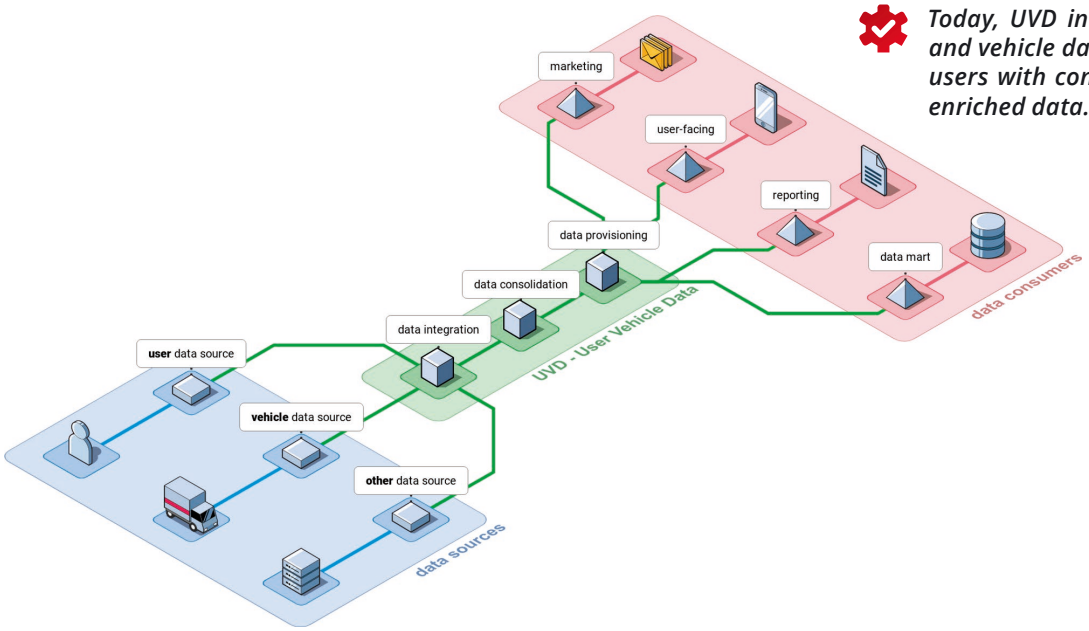

GEM System | Štětkova 1638/18, 140 00 Praha 4 – Nusle, Czech Republic | sales@gemsystem.cz | www.gemsystem.cz

UVD project in numbers:

- Project duration: 3 years (start Q1 2020)..
- Team size: 11 FTE..
- Team competencies: covering the most important parts of the software development lifecycle: architecture, development, quality assurance and operations.
- Logical decomposition of the system: 28 subsystems.
- Physical decomposition of the system: 50+ microservices.
- Storage size: 1TB+ of active data.
- Monthly storage increment: 150 GB of live data.
- Number of unique entities updated per month (batch and stream processing): 100,000,000+.
- Domain coverage: user personal data, vehicle production data, digital service license data, digital service usage data, vehicle telemetry data, vehicle maintenance data, etc.
- Main architectural challenges: availability, resilience, data consistency, data quality anomaly detection, maintenance of integration contracts.

UVD project in detail:

Team growth time

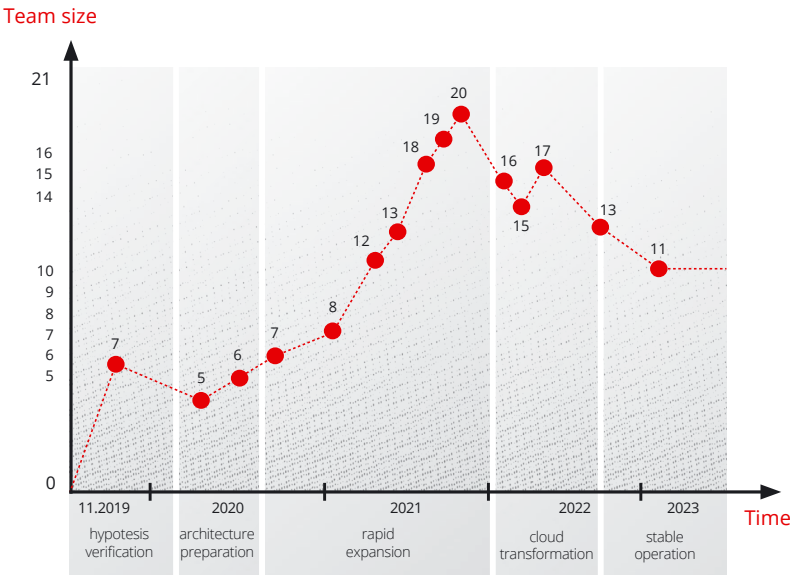

Product growth

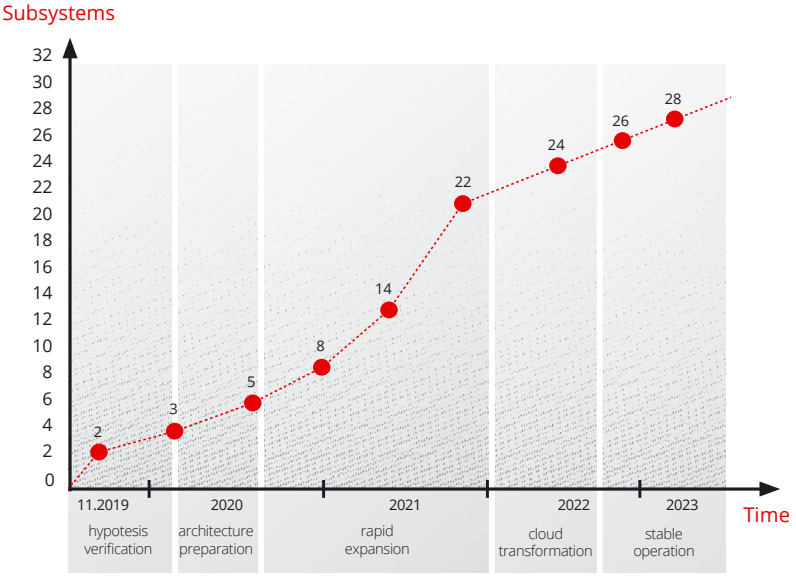

GEM SYSTEM

SKODA

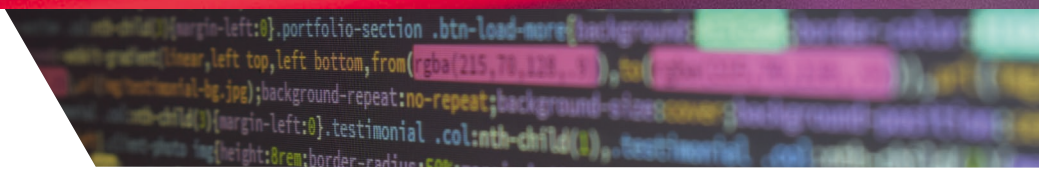

1 The „hypothesis verification“ phase:

The project began by testing the hypothesis that the UVD product is capable of serving as an ODS (Operational Data Storage) in the Škoda Auto ecosystem, which means:

- collect data on users and vehicles from various data sources in the Škoda Auto ecosystem,
- perform consolidation processes on the collected data (data transformation, data cleaning, data quality assessment, etc.),
- providing requested data in real time (or near real time) to consumers in the Skoda Auto ecosystem.

By successfully integrating several data sources and completing proof-of-concept projects in the area of data consolidation and data provision, it has been proven that UVD is capable of serving as a full-fledged Operational Data Store (ODS) solution in the Škoda Auto ecosystem.

2 The „architectural preparation“ phase:

After verifying the hypothesis, the next logical step was to prepare a UVD product for the integration of multiple data sources and data consumers.

Taking into account the fact that one of the main objectives ofODS is „operational decision making“ (as opposed to EnterpriseData Warehouse, whose main objective is „tactical and strategic decision making“), the following characteristic was considered important for UVD:

- performance and scalability - to be ready to process large amounts of data and provide data to consumers in real-time and near real-time,
- consistency - to ensure the highest possible level of data consistency for consumers,
- modifiability and flexibility - to quickly implement changes to the system, both when adding new functionality and when experimenting with data processing algorithms.

Příprava architektury, navržená a implementovaná umožnila exponenciální růst produktu UVD.

3 The „rapid expansion“ phase:

During the rapid expansion phase, the UVD product expanded from 2 to 20 subsystems in the production environment:

- 11 integration subsystems, integrating data from different data sources located in the ŠkodaAuto ecosystem using different enterprise integration patterns such as file transfer, remote procedure call, database sharing and messaging),
- 4 consolidation subsystems, consolidating integrated data in the UVD,
- 5 implementation (or use-case) subsystems, implementing the necessary business logic and providing data to UVD consumers.
- The rapid expansion has significantly increased the value of the entire UVD solution in the Škoda Auto ecosystem. Last but not least, the phase revealed not only the weaknesses and limitations of the technologies that UVD was using, but also some gaps that needed to be fine-tuned in terms of processes and organisation.

The architectural sub-team decided to solve all problems conceptually. Thus, an overarching „improvement mind map" for the problems emerged. The map covered all aspects that required improvement (i.e. CI/CD procedures, product quality control, support and operational cost reduction, team process maturity):

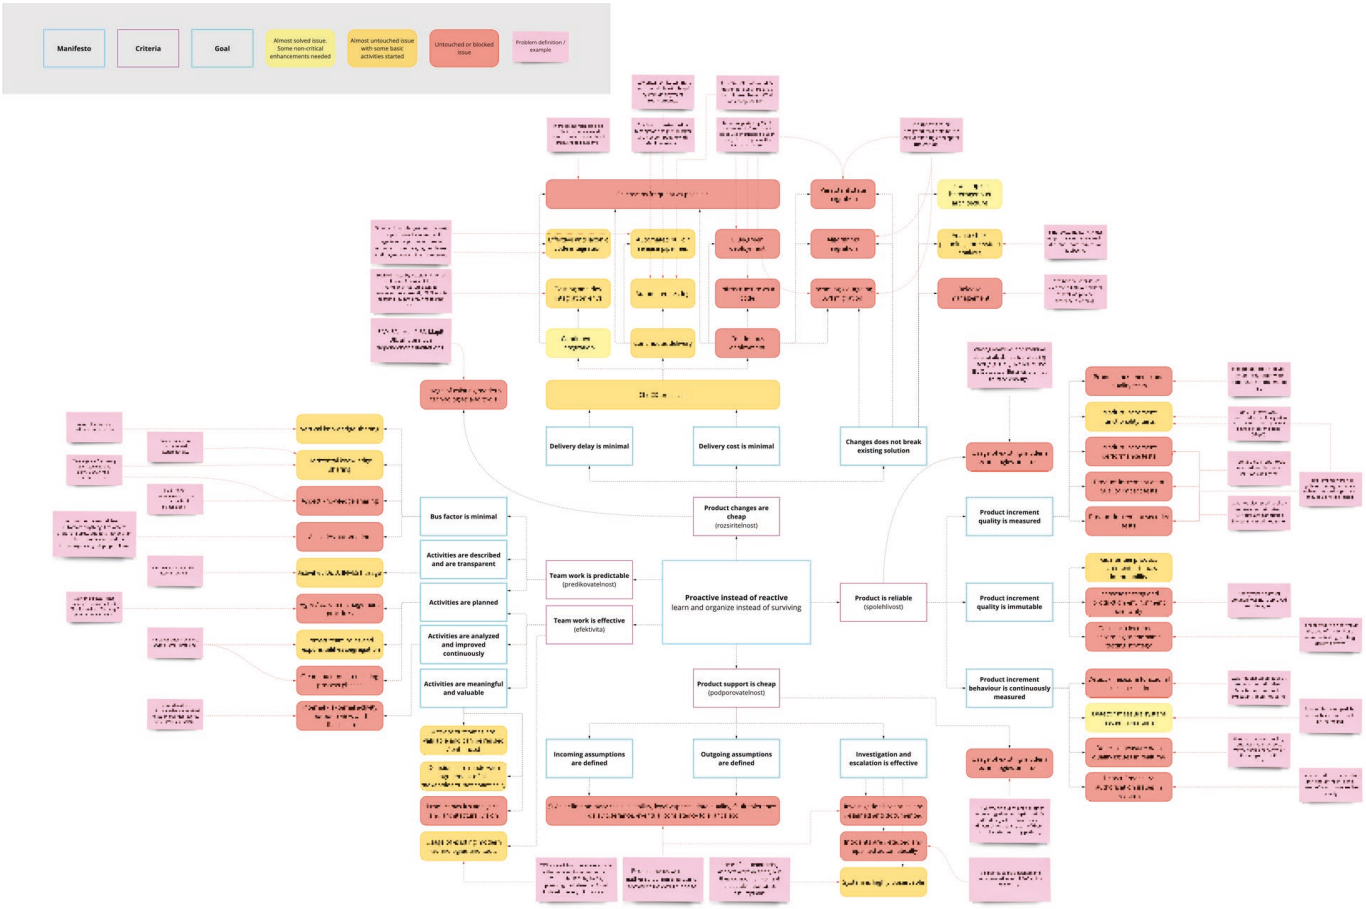

Idea map for improvements

4 The „cloud transformation" phase:

The detailed high-level overview of all events provided by the „Improvement Mind Map" helped the UVD team to prioritise the issues and find appropriate solutions.

The transformation of the UVD product to cloud-based technologies was one of the first solutions to be applied, helping to get rid of the technological constraints that were blocking product development in many areas.

From a technology perspective, the transformation primarily involved replacing the old proprietary „HPE EzmeralData Fabric" (formerly „MapR") solution with a combination of native tech Azure Cloud technologies (Azure Event Hubs, Azure Blob Storage, Azure Service Bus, etc.) and Mongo Atlas in Azure.

The UVD architectural sub-team managed to design, plan and execute the migration to the new technology stacks with zero downtime and zero data loss. As a result, the cloud transformation moved the UVD product forward not only in terms of technology, but also in terms of processes - significantly improving deployability, testability, developer experience and maintainability.

5 „Stable operation“ phase:

After the „Cloud Transformation“ phase, the UVD product has entered a phase of stable operation, which is ongoing.

The benefits of the cloud transformation have saved resources (both human and technological) for setting up use cases and development functions in UVD.

The UVD product is ready for a new phase of expansion. In addition, there is room for further technological improvements, which are clearly defined and mapped.

Thanks to the comprehensive approach of continuous improvement, the architectural team continues to work and the process of technological improvement at UVD continues.

The improvement map provides a high level overview of the improvement activities and the dependencies between them.

The improvement map provides a high level overview of improvement activities and the dependencies between them.