



Real-Time Telematic Vehicle Data Aggregation

We developed a high-performance, multi-cloud data aggregation backbone using AWS Kinesis and Kafka to enable real-time remote diagnostics and proactive service preparation for a leading German automotive manufacturer.

Overview

A leading German automotive R&D center required a specialized backend system—a Vehicle Data Aggregator—to process real-time telemetry data from vehicles for remote diagnostics and preemptive service preparation.

- Designed a scalable data pipeline leveraging Apache Kafka and Amazon Kinesis to process large-scale, asynchronous data streams in real-time.
- Developed the aggregator application in Golang and deployed it on a managed Kubernetes cluster in the AWS cloud, ensuring high performance and resilience.
- Enabled early fault detection and proactive service by transforming raw diagnostic data into actionable messages for downstream systems that handle follow-up actions (alerting users or directing vehicles to prepared workshops).
- Facilitated remote diagnostics, accelerated the detection of faults, and ensured workshops have corrective actions ready-to-go, significantly improving service efficiency.



Client Profile

Headquartered in Germany, our client is the research and development center for the world's largest manufacturer of premium and commercial vehicles. The center focuses on research, IT engineering, and product development.

Challenges: High-Volume, Real-Time Diagnostic Data

The primary challenge was managing the complexity and volume of real-time telematic data streams:

- **Real-Time Requirement:** Diagnostic data (ECU dumps, warning messages) needed to be processed and available live to the organization immediately upon occurrence.

- **Asynchronous Data Handling:** The system needed to reliably ingest and process high-volume, asynchronous data streams from various internal software systems across a multi-cloud network.
- **Diagnostic Complexity:** The aggregator had to transform raw sensor data into specific diagnostic reports (e.g., Quick OBD reports, Drive Control reports) based on different remote use cases.
- **Seamless Routing:** The processed data needed to be seamlessly and reliably routed to different downstream systems responsible for follow-up actions (like alerting the user or preparing a workshop).

QBurst Solution: High-Performance Data Aggregator

A high-performance Vehicle Data Aggregator backend system built on a modern streaming architecture that acts as the mediator between vehicle data sources and diagnostic follow-up systems.

- **Multi-Cloud Streaming Architecture:** The solution was deployed within the client's private multi-cloud network. It listened to existing data streams flowing from other software systems (originating in Azure) and concentrated the diagnostic-related streams within the AWS cloud environment.
- **Real-Time Data Processing:** Amazon Kinesis was utilized for processing and analyzing the large-scale, real-time streaming data. The application specifically subscribed to various Kafka topics within Kinesis to collect the telematic data.
- **High-Performance Application:** The aggregator application was built using Golang for its concurrency and performance advantages, ensuring low latency processing of diagnostic data.
- **Kubernetes Deployment:** The application was deployed onto an organization-level managed Kubernetes cluster in the AWS cloud, guaranteeing scalability, resilience, and operational consistency.
- **Intelligent Routing:** After processing the raw data into use-case specific formats, the application published the transformed messages to new Kafka topics. Downstream systems then subscribed to these topics to initiate follow-up actions (e.g., generating alerts, directing the vehicle, or updating workshop reports).

Technical Highlights

The solution focused on speed, concurrency, and scalable data streaming:

- **Streaming Engine:** AWS Kinesis and Apache Kafka provided the robust, high-throughput backbone for real-time data ingestion and distribution.
- **Core Application Language:** Golang ensured the aggregator application had the speed and concurrency necessary for real-time, low-latency processing of high-volume diagnostic data.
- **Container Orchestration:** Deployed on Kubernetes for automatic scaling, high availability, and efficient resource management.
- **Architectural Integration:** Seamlessly listened to and published data streams across a complex private multi-cloud network (AWS and Azure).

Impact

The implementation of the Vehicle Diagnostic Data Aggregator delivered essential capabilities for proactive service management:

- **Enabled Remote Diagnostics:** Facilitated the immediate collection and processing of critical data (ECU dumps, battery health) live as it happened.
- **Accelerated Fault Detection:** The real-time nature of the processing allows for early detection of faults and potential issues in the vehicle.
- **Improved Service Efficiency:** Downstream systems can now prepare workshops with corrective actions read-to-go, streamlining the service process.
- **Enhanced Scalability:** The use of Kafka and Kubernetes ensured the platform could scale reliably with increasing telematic data volume from a growing fleet.