
AWS
re:Invent

ANT201

Data-driven development at Volkswagen Commercial Vehicles

John Fisher

Principal Automotive Business Development Manager

AWS

Dr. Michael Nolting

Director of Digital Services and Data Analytics

Volkswagen Commercial Vehicles

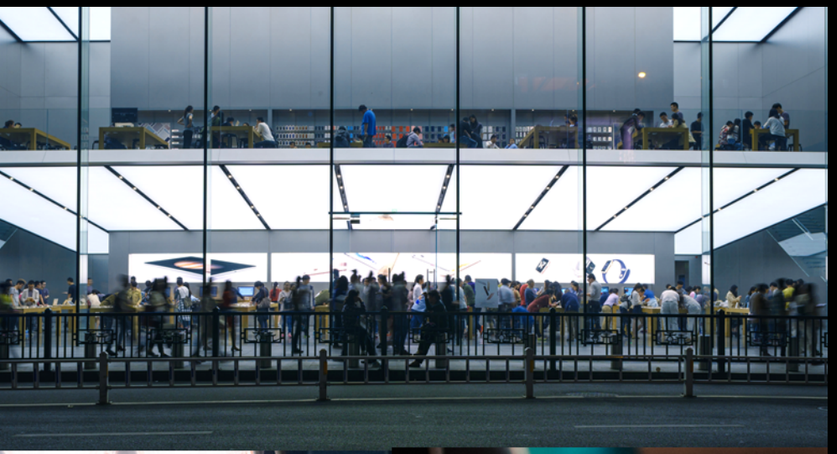
Dr. Mohammad Al-Rifai

Lead Data Scientist

Volkswagen Commercial Vehicles

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The automotive industry is changing…









"The global automotive industry is about to enter a period of wide-ranging and transformative change, as sales continue to shift and environmental regulations tighten."

McKinsey & Company

AWS in automotive

Select AWS automotive customers

OEM's



Tiered Suppliers



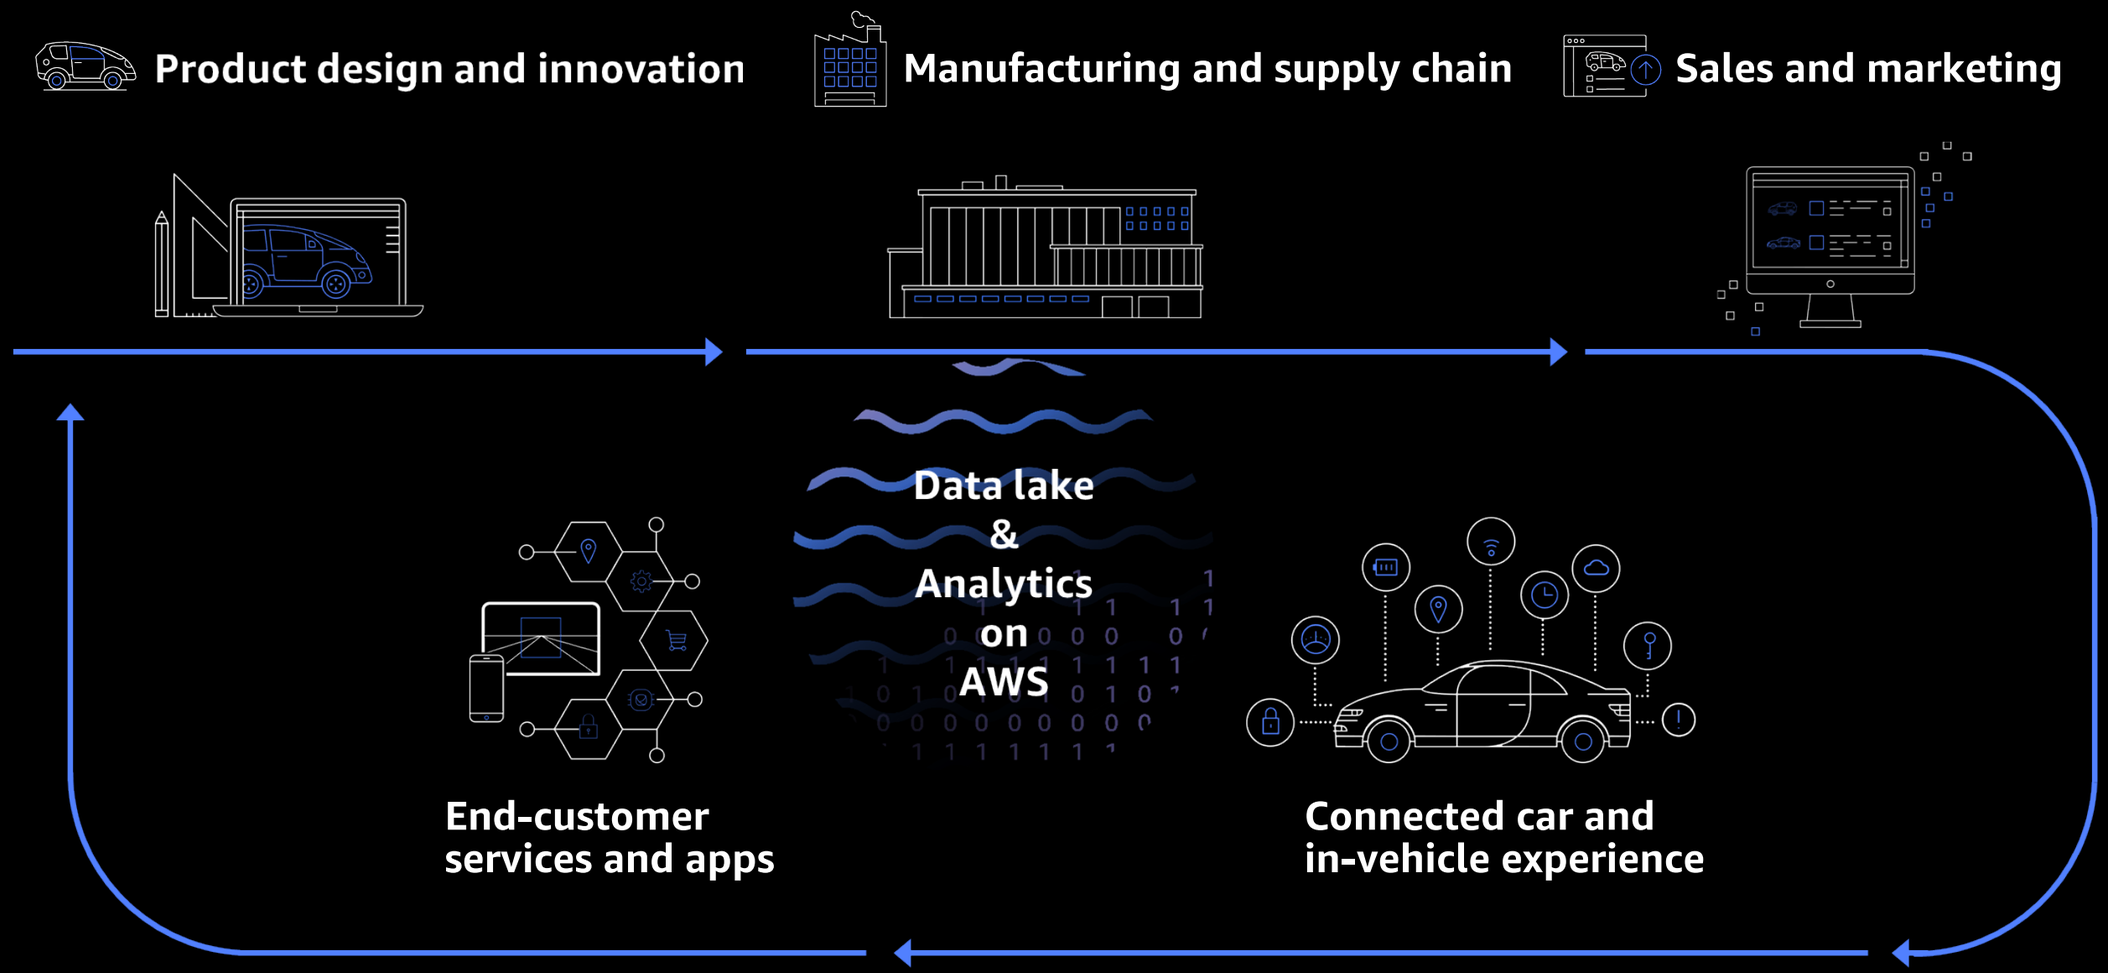
Mobility Providers



Auto Tech



AWS enables the data-driven automotive value chain



Agenda

Why data-driven development of commercial vehicles?

Example use cases

Project SNOWPARK: vehicle analytics at scale



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Volkswagen Commercial Vehicles

WE ARE THE VW BRAND FOCUSING ON BUSINESS CUSTOMERS



Diverse customer groups

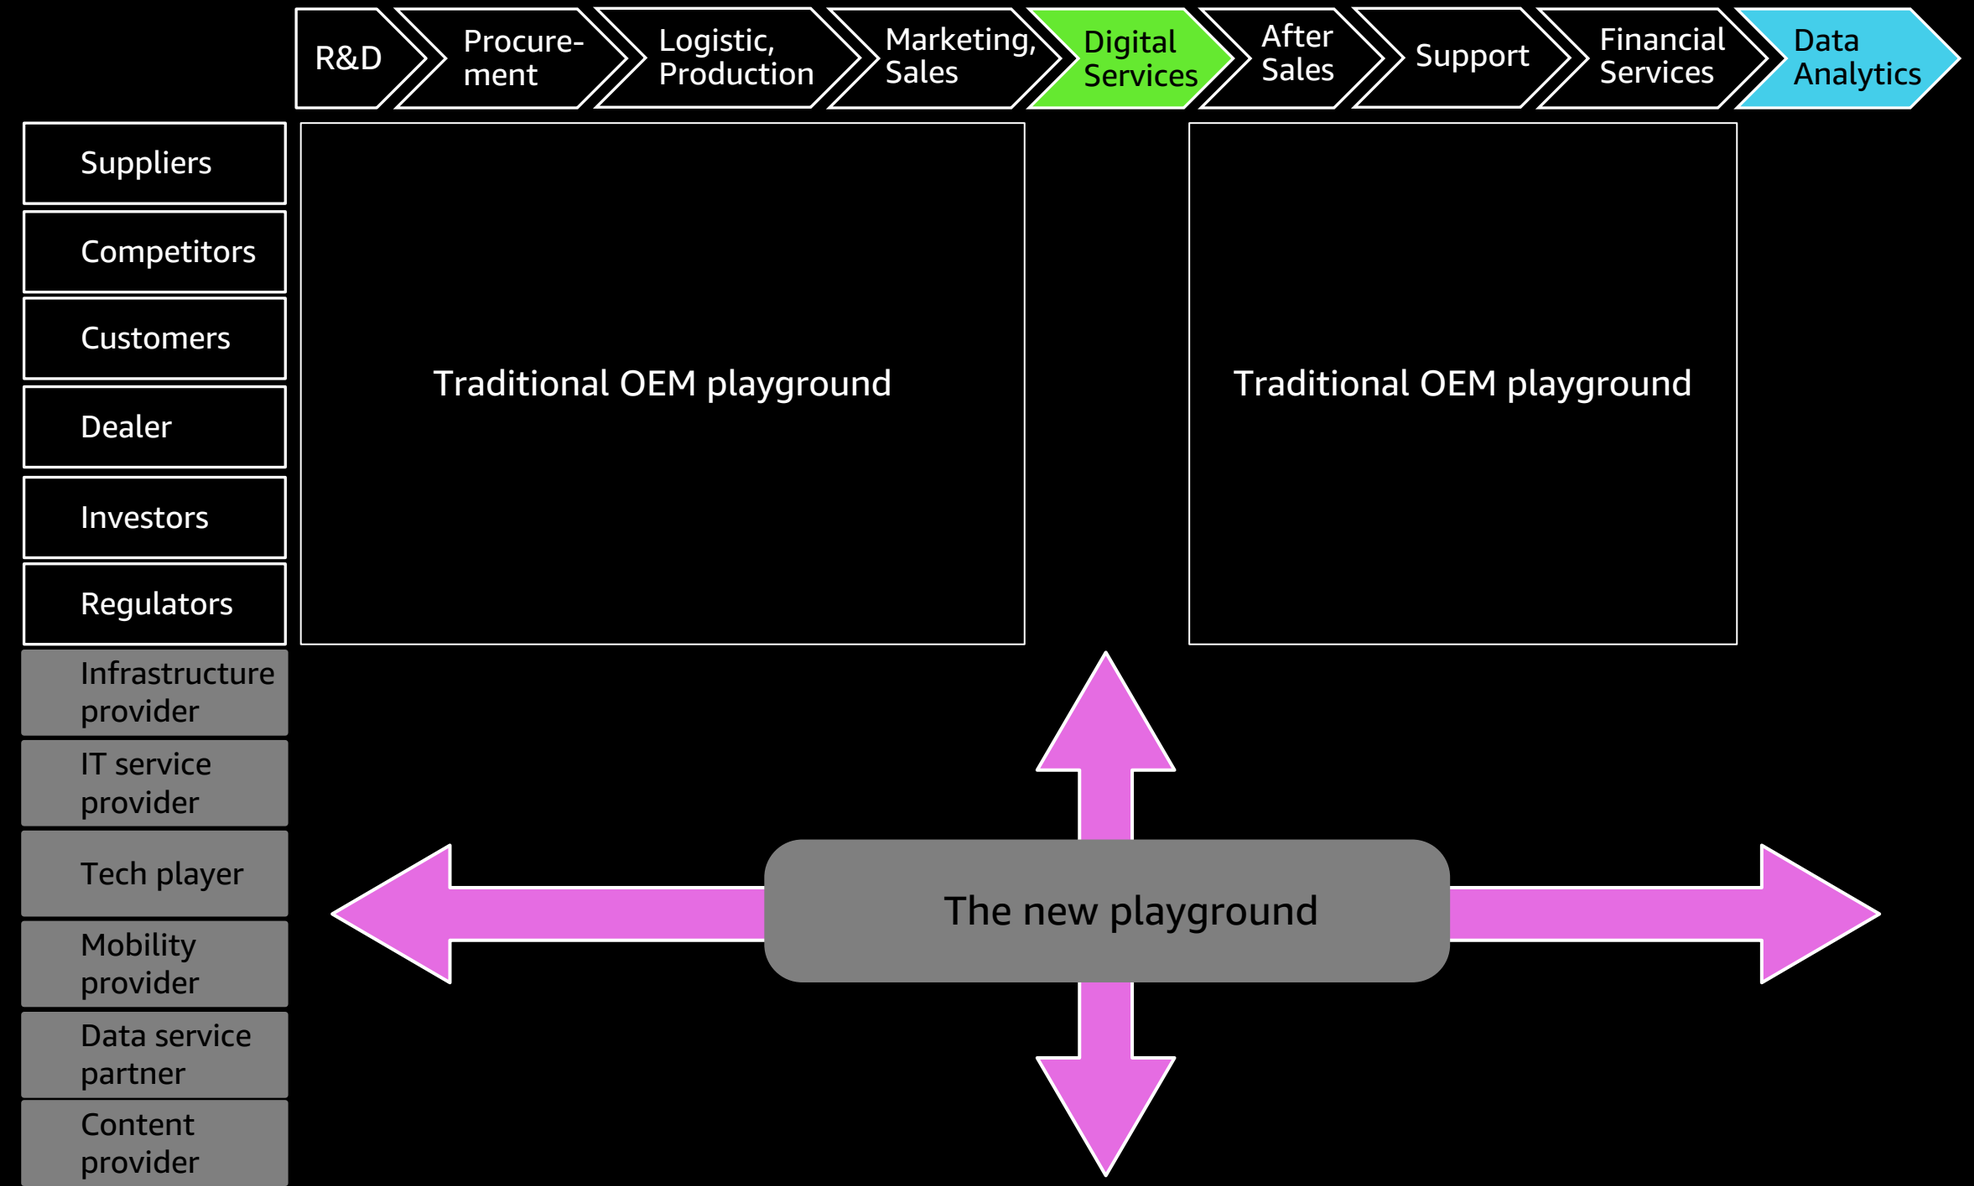








The new automotive value chain



Example use cases



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

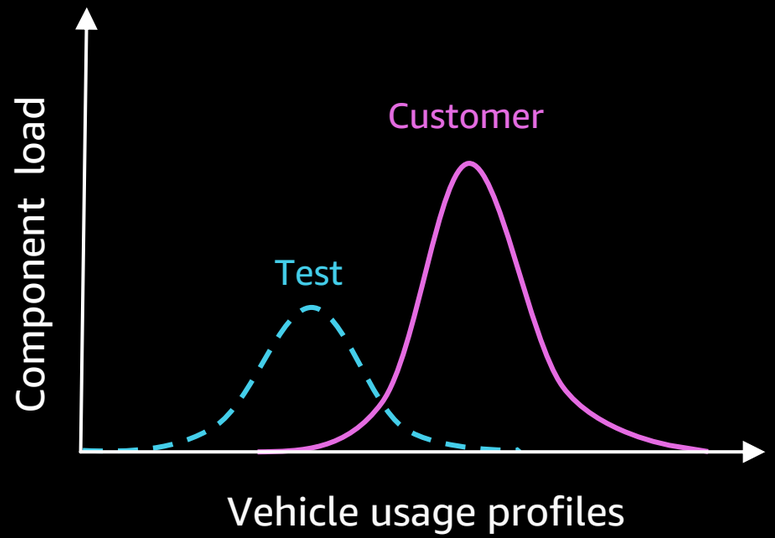
Use case 1: Vehicle design optimization

DESIGN VS. ACTUAL USAGE

- Data from own test fleet and pilot customers
- Load measurement (e.g., brake disks or gears)
- Customer usage profiles vs. test profiles

- Outcome 1: vehicles fit customer needs
- Outcome 2: optimized testing procedures

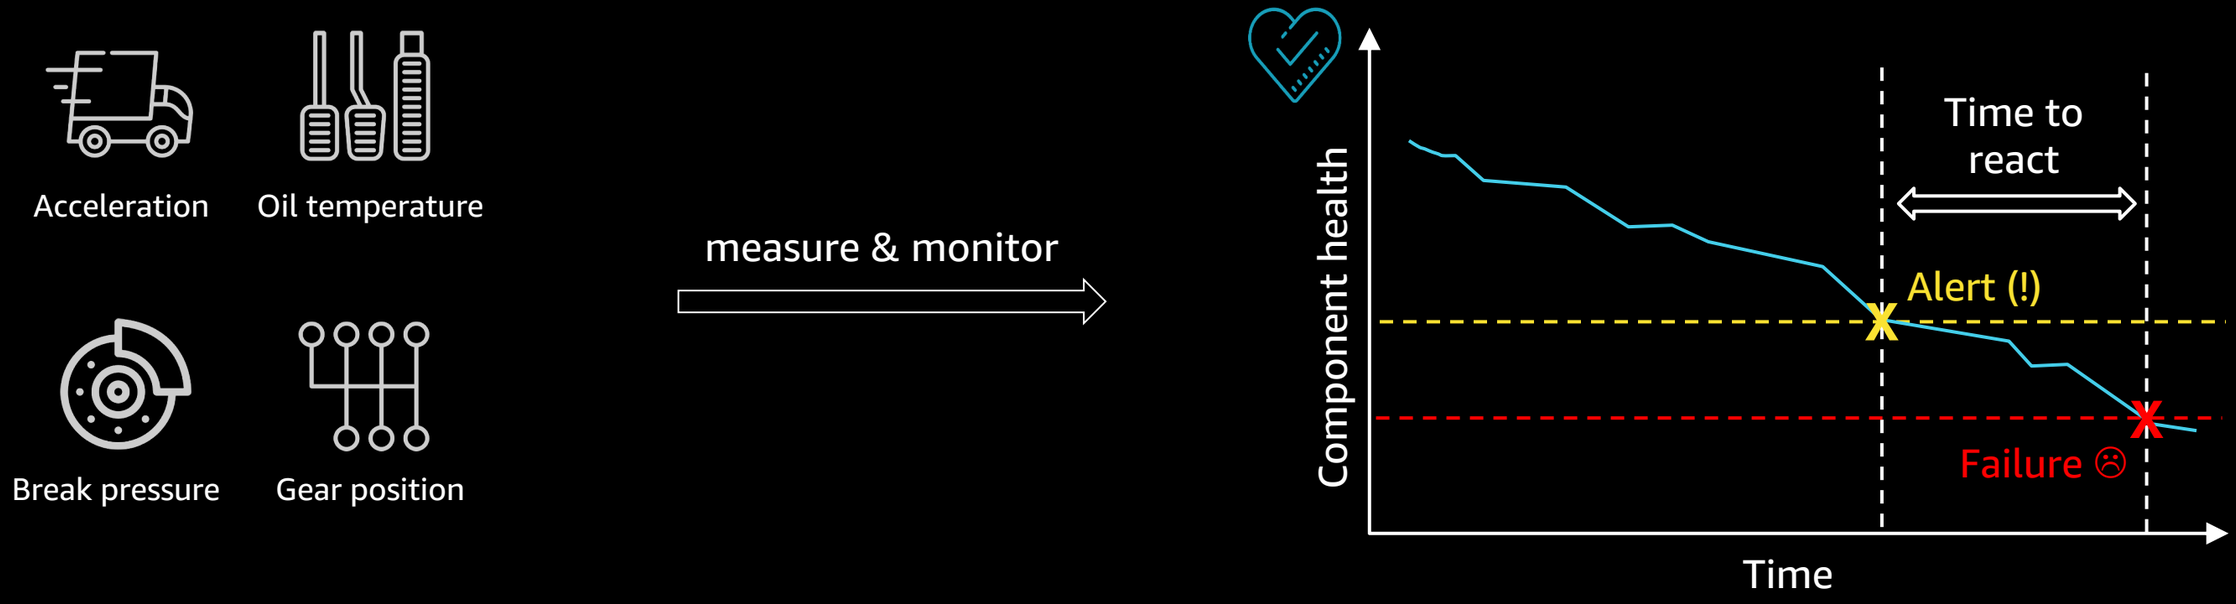




Use case 2: 100% up-time for customer's fleet

PREDICTIVE MAINTENANCE

- Component health monitoring and failure prediction
- Outcome: downtime reduction for fleet vehicles



Project SNOWPARK: Vehicle data analytics at scale



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

SNOWPARK: Challenges

Vast amount of data

Heterogeneous data

Legacy systems on-premises

1000s vehicles

100s Terabyte of data

10 000 000s driving kilometers

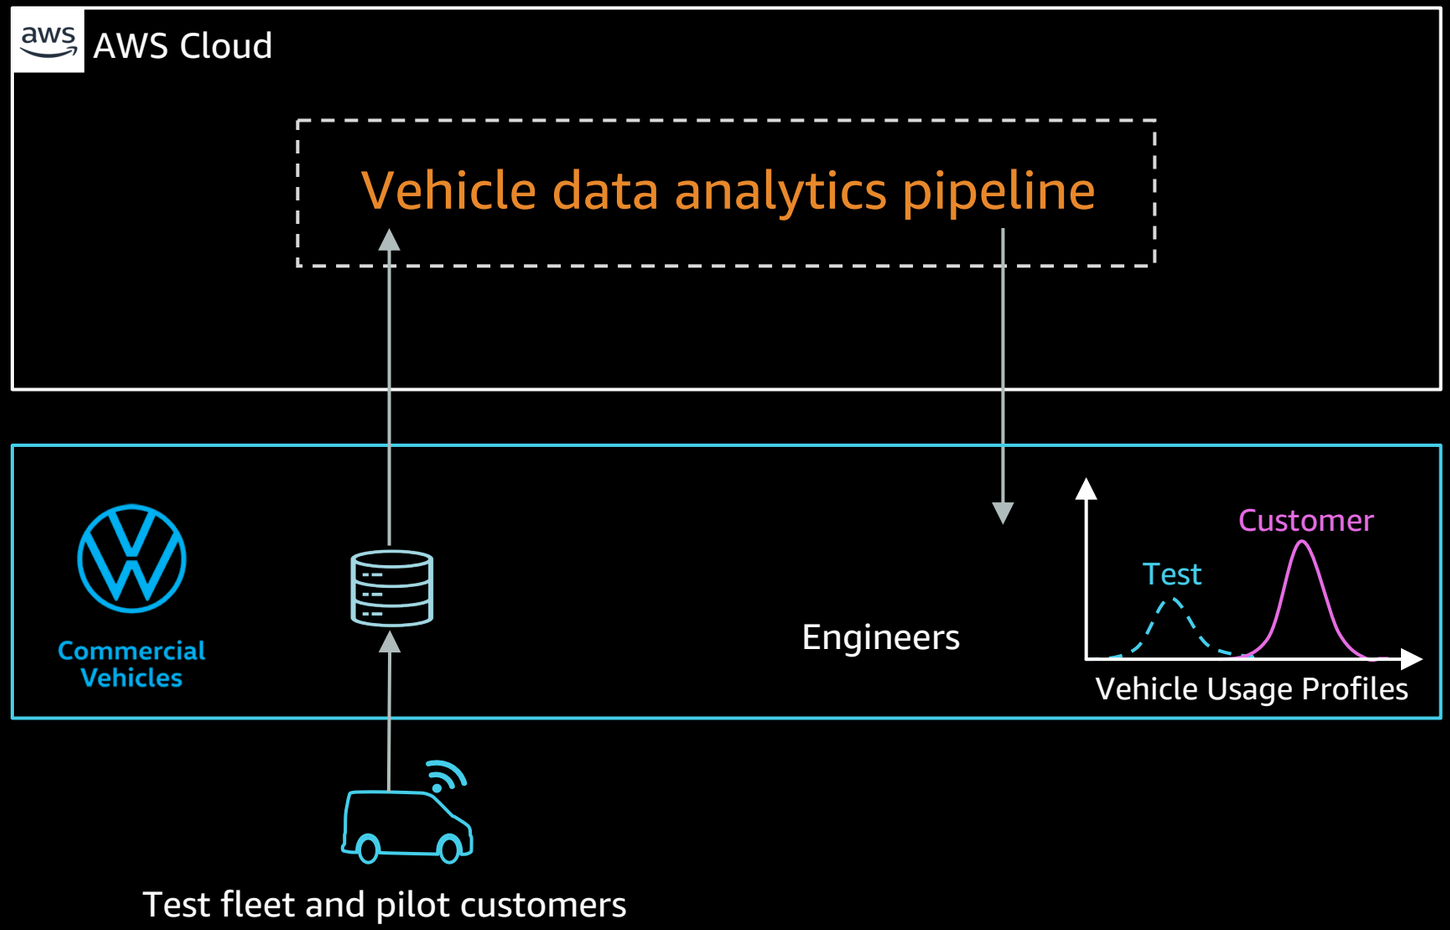
Structured data (measurements) and unstructured data (text)

Vehicle configuration, component specific information etc.

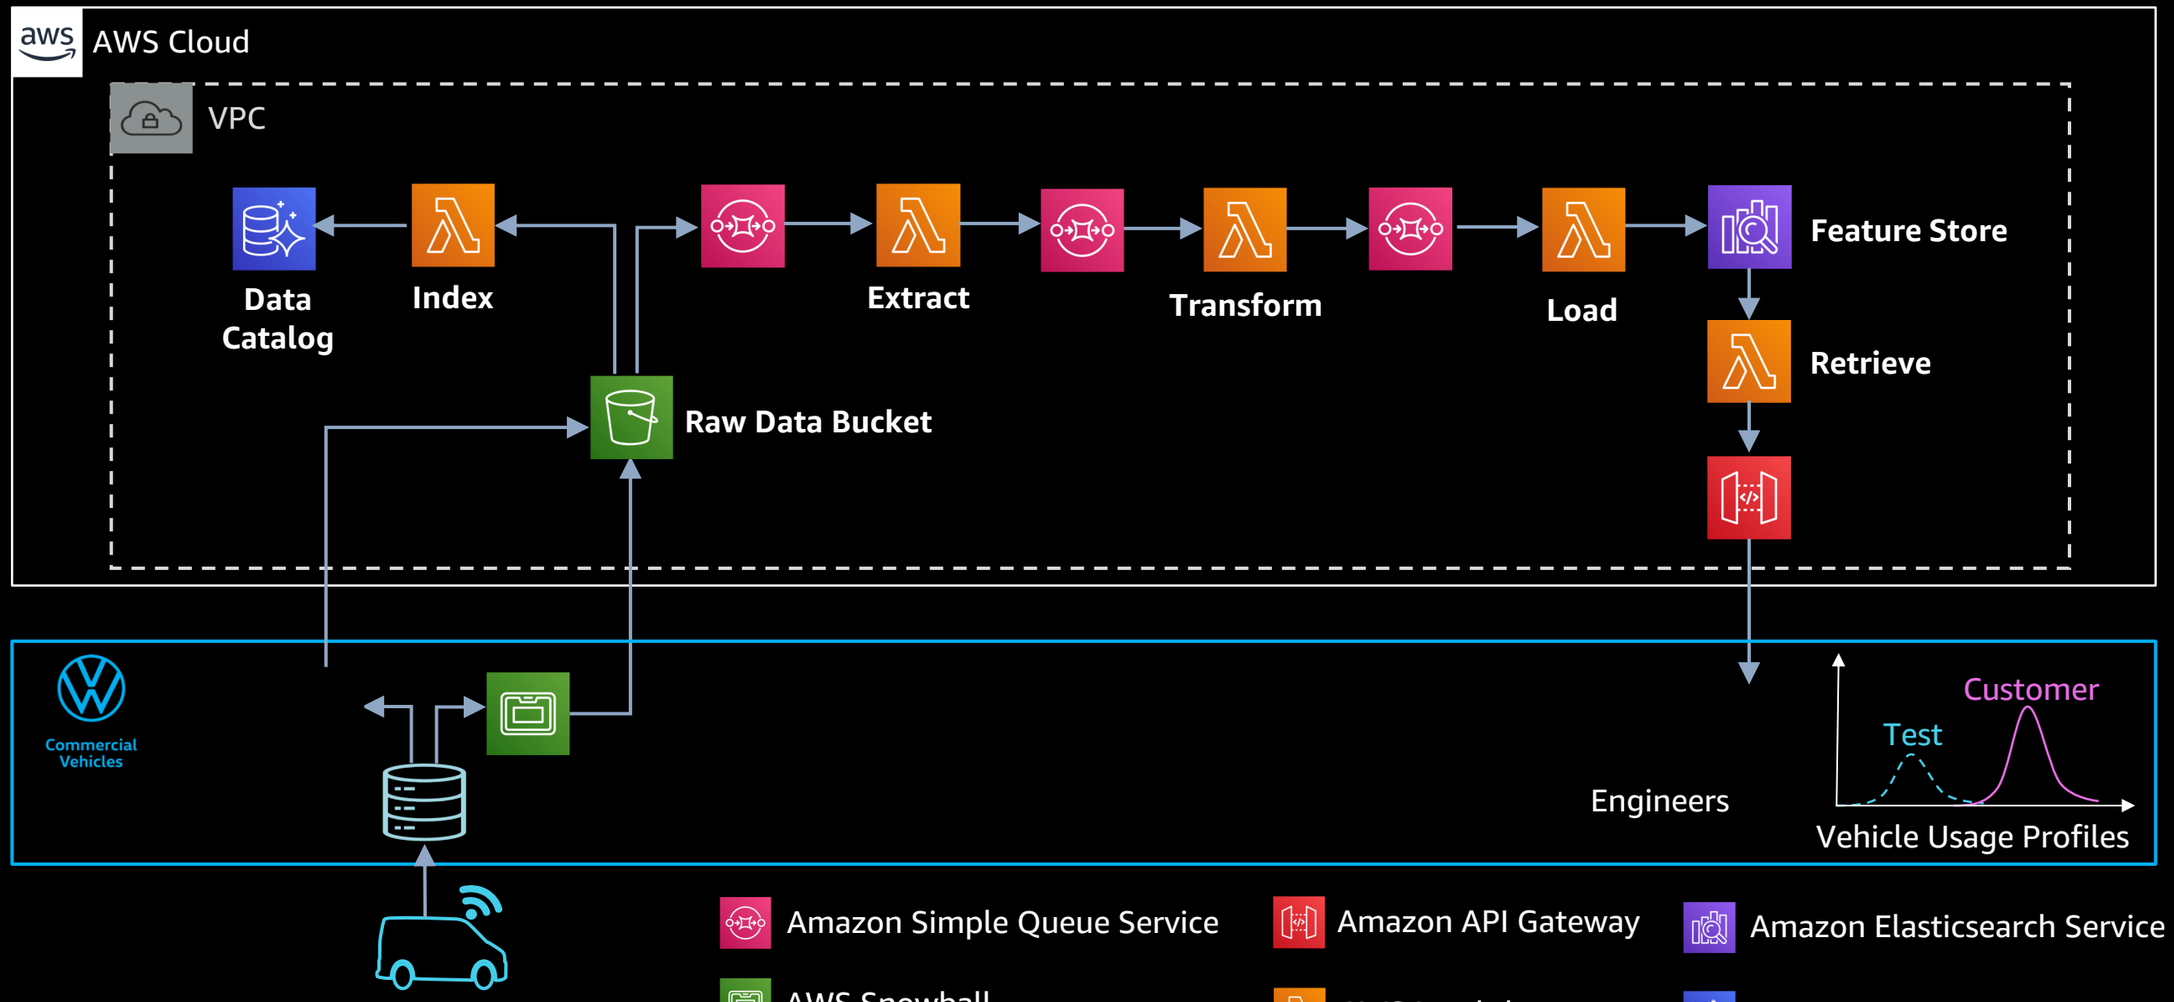
SNOWPARK: 10,000 ft. view

- Scalable data storage, data processing and generation of vehicle usage profiles on AWS

- Analysis of usage profiles using self-service BI tools on-premises

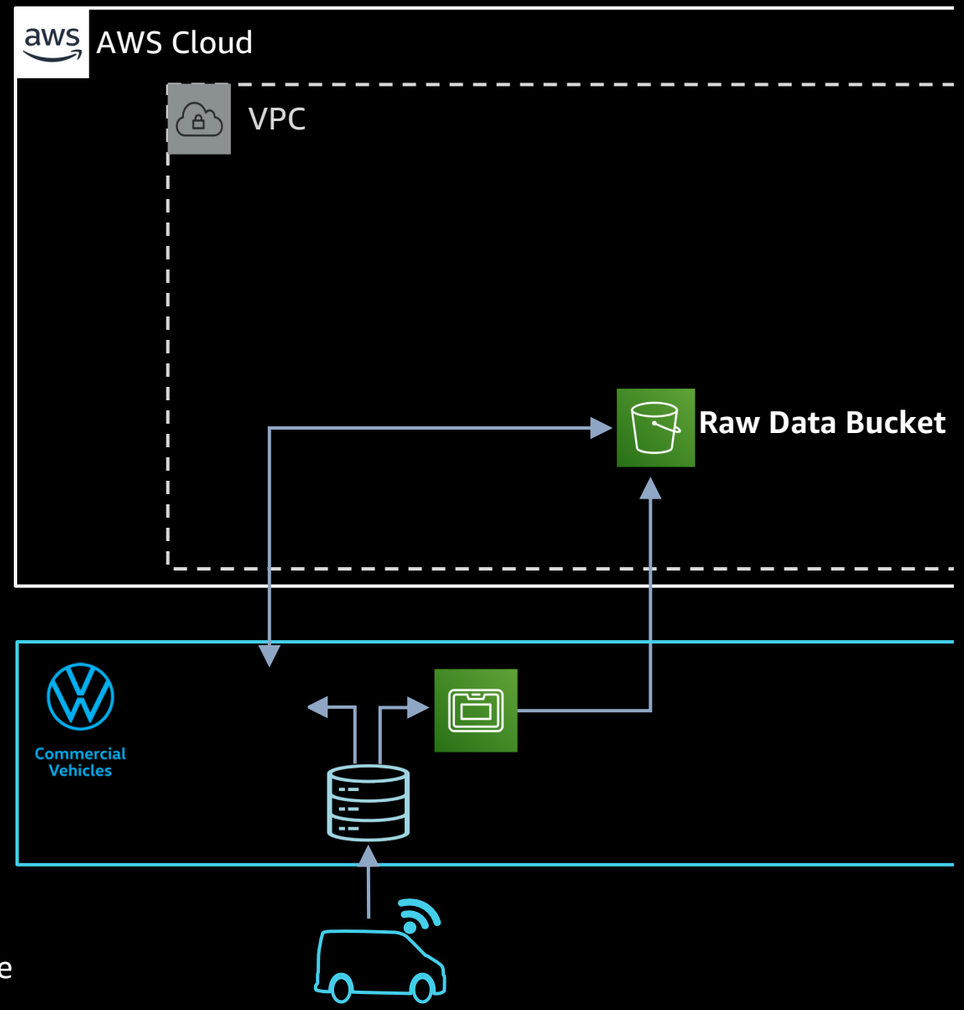


Scalable and serverless architecture



Data import

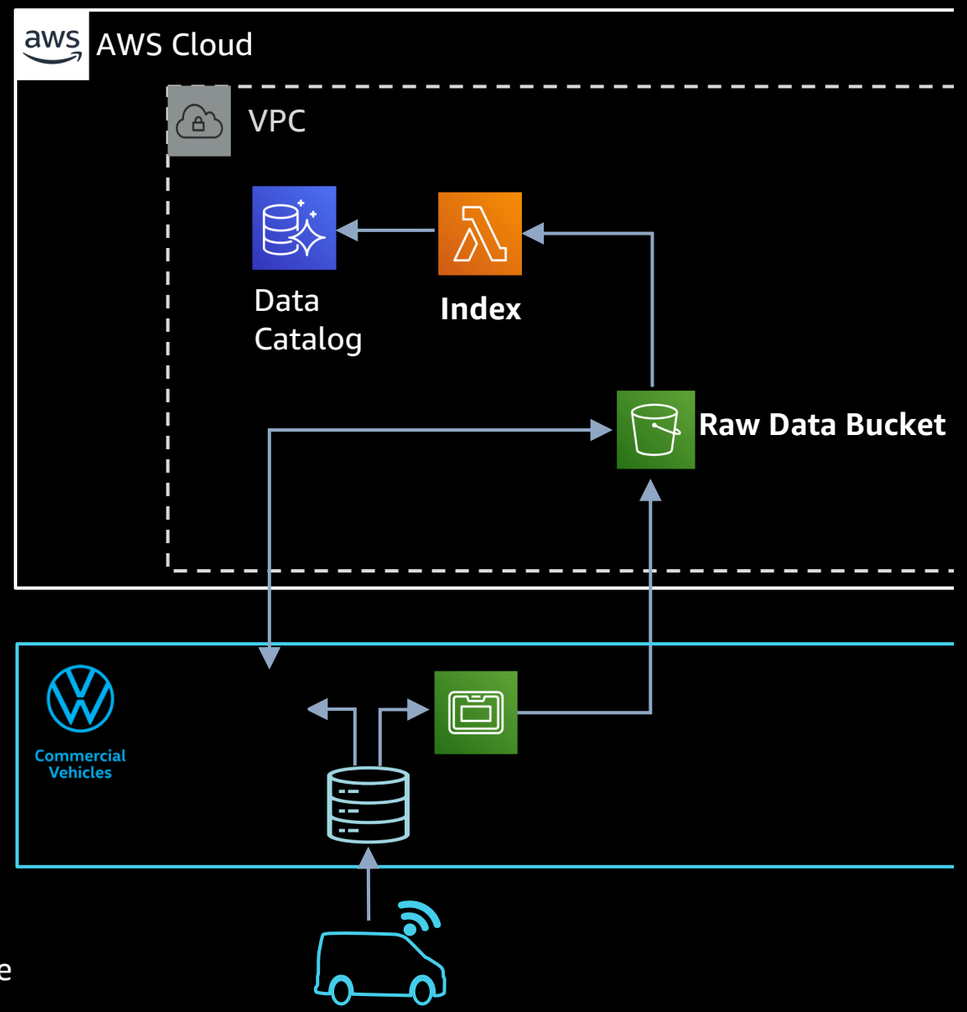
- Historical data
  - > 100 terabytes of sensor data
  - Amazon Simple Storage Service (Amazon S3)
  - Migration using AWS Snowball
- New data
  - ~50 GB per day
  - Own desktop tool for uploading and downloading raw data using AWS Software Developer Kits



Amazon Aurora
Amazon Simple Storage Service (Amazon S3)
AWS Snowball
AWS Lambda

Data import

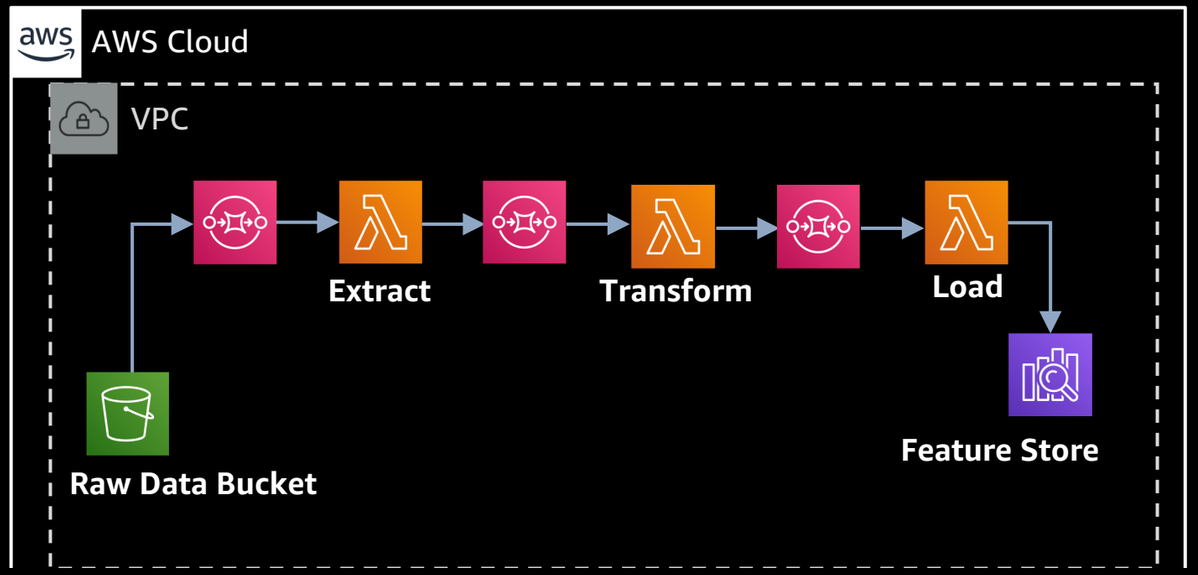
- Data catalog
  - Index of metadata for efficient search
  - Overview of available data and vehicles
  - AWS Lambda for metadata extraction
  - Serverless Amazon Aurora with auto-scaling



Amazon Aurora
Amazon Simple Storage Service (Amazon S3)
AWS Snowball
AWS Lambda

Usage profile extraction

- Profile = set of features, like:
  - average speed
  - number of starts and stops per 100 km
  - number of doors opening per 100 km
  - percentage of driving hours on highway vs. city
  - Fuel consumption
  - component load



AWS Lambda

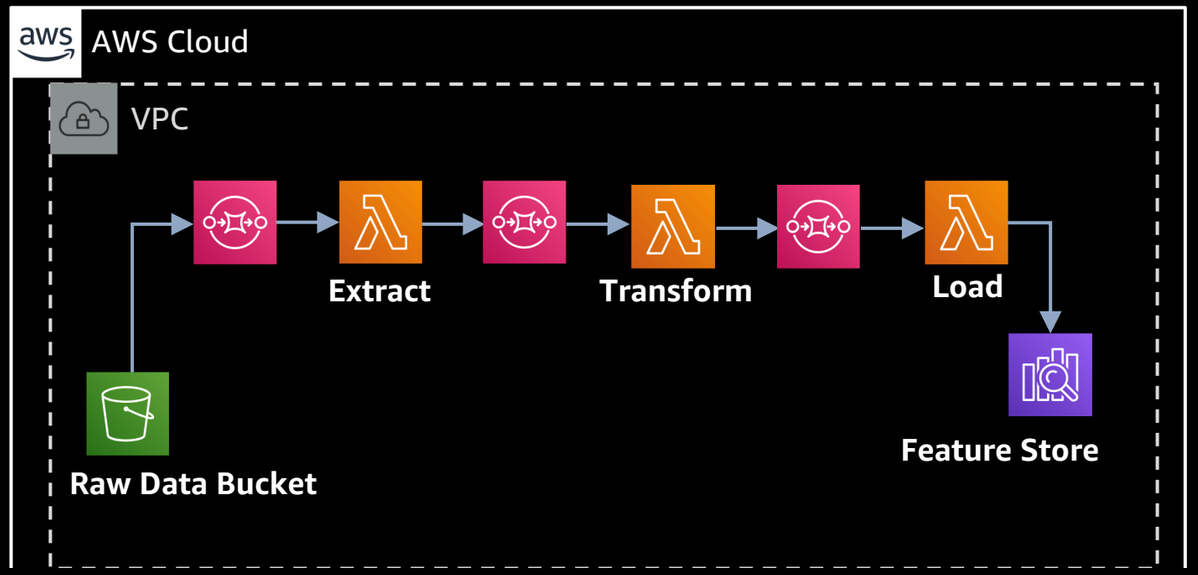
Amazon Simple Queue Service (SQS)

Amazon Simple Storage Service (Amazon S3)

Amazon Elasticsearch Service

Usage profile extraction

- Dynamic set of features
  - New feature definitions
  - New data deliveries
- Feature store
  - No fixed schema for profiles
  - Support for textual information
  - Managed Amazon Elasticsearch Service to reduce complexity



AWS Lambda

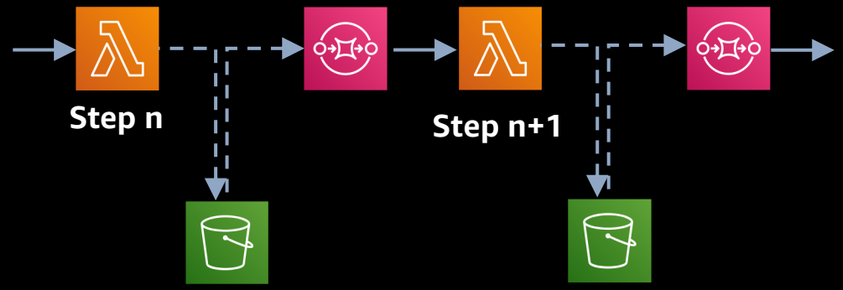
Amazon Simple Queue Service (SQS)

Amazon Simple Storage Service (Amazon S3)

Amazon Elasticsearch Service

Processing pipeline

- Data processing is divided into steps:
  - Preprocessing, feature extraction, ingestion
  - First step is triggered by data upload
  - Manual trigger for reprocessing
  - Each step creates new artifacts in S3 buckets
  - Creation of artifacts triggers the next step
  - Final step stores results in the feature store



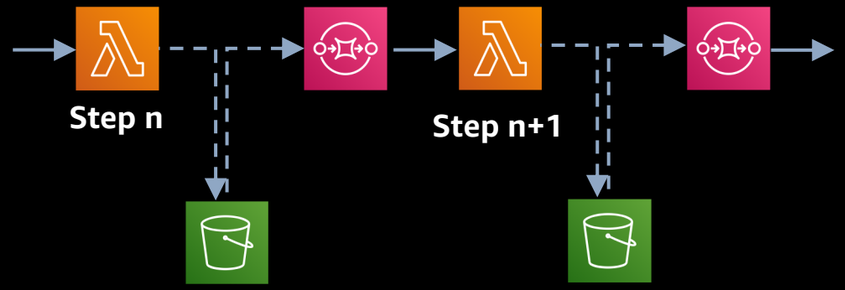
Amazon Simple Storage Service (Amazon S3)

AWS Lambda

Amazon Simple Queue Service (SQS)

Processing pipeline

- Design pattern:
  - Bucket → Queue → Function → Bucket
  - Objects are stored in S3 Buckets
  - Object creation events are sent to SQS
  - AWS Lambda fetches events from SQS
  - AWS Lambda stores the results in S3 buckets
  - and so on …



To ensure that an event is processed only once:
SQS.visibility_timeout > Lambda.timeout

To control the concurrency level of the pipeline, set:
Lambda.reserved_concurrent_executions

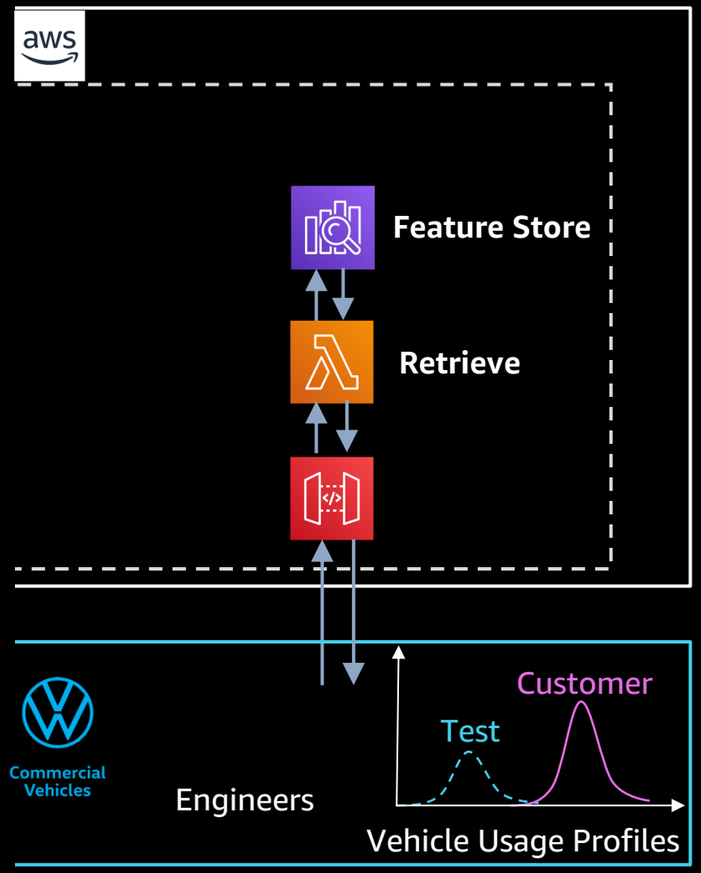
Amazon Simple Storage Service (Amazon S3)

AWS Lambda

Amazon Simple Queue Service (SQS)

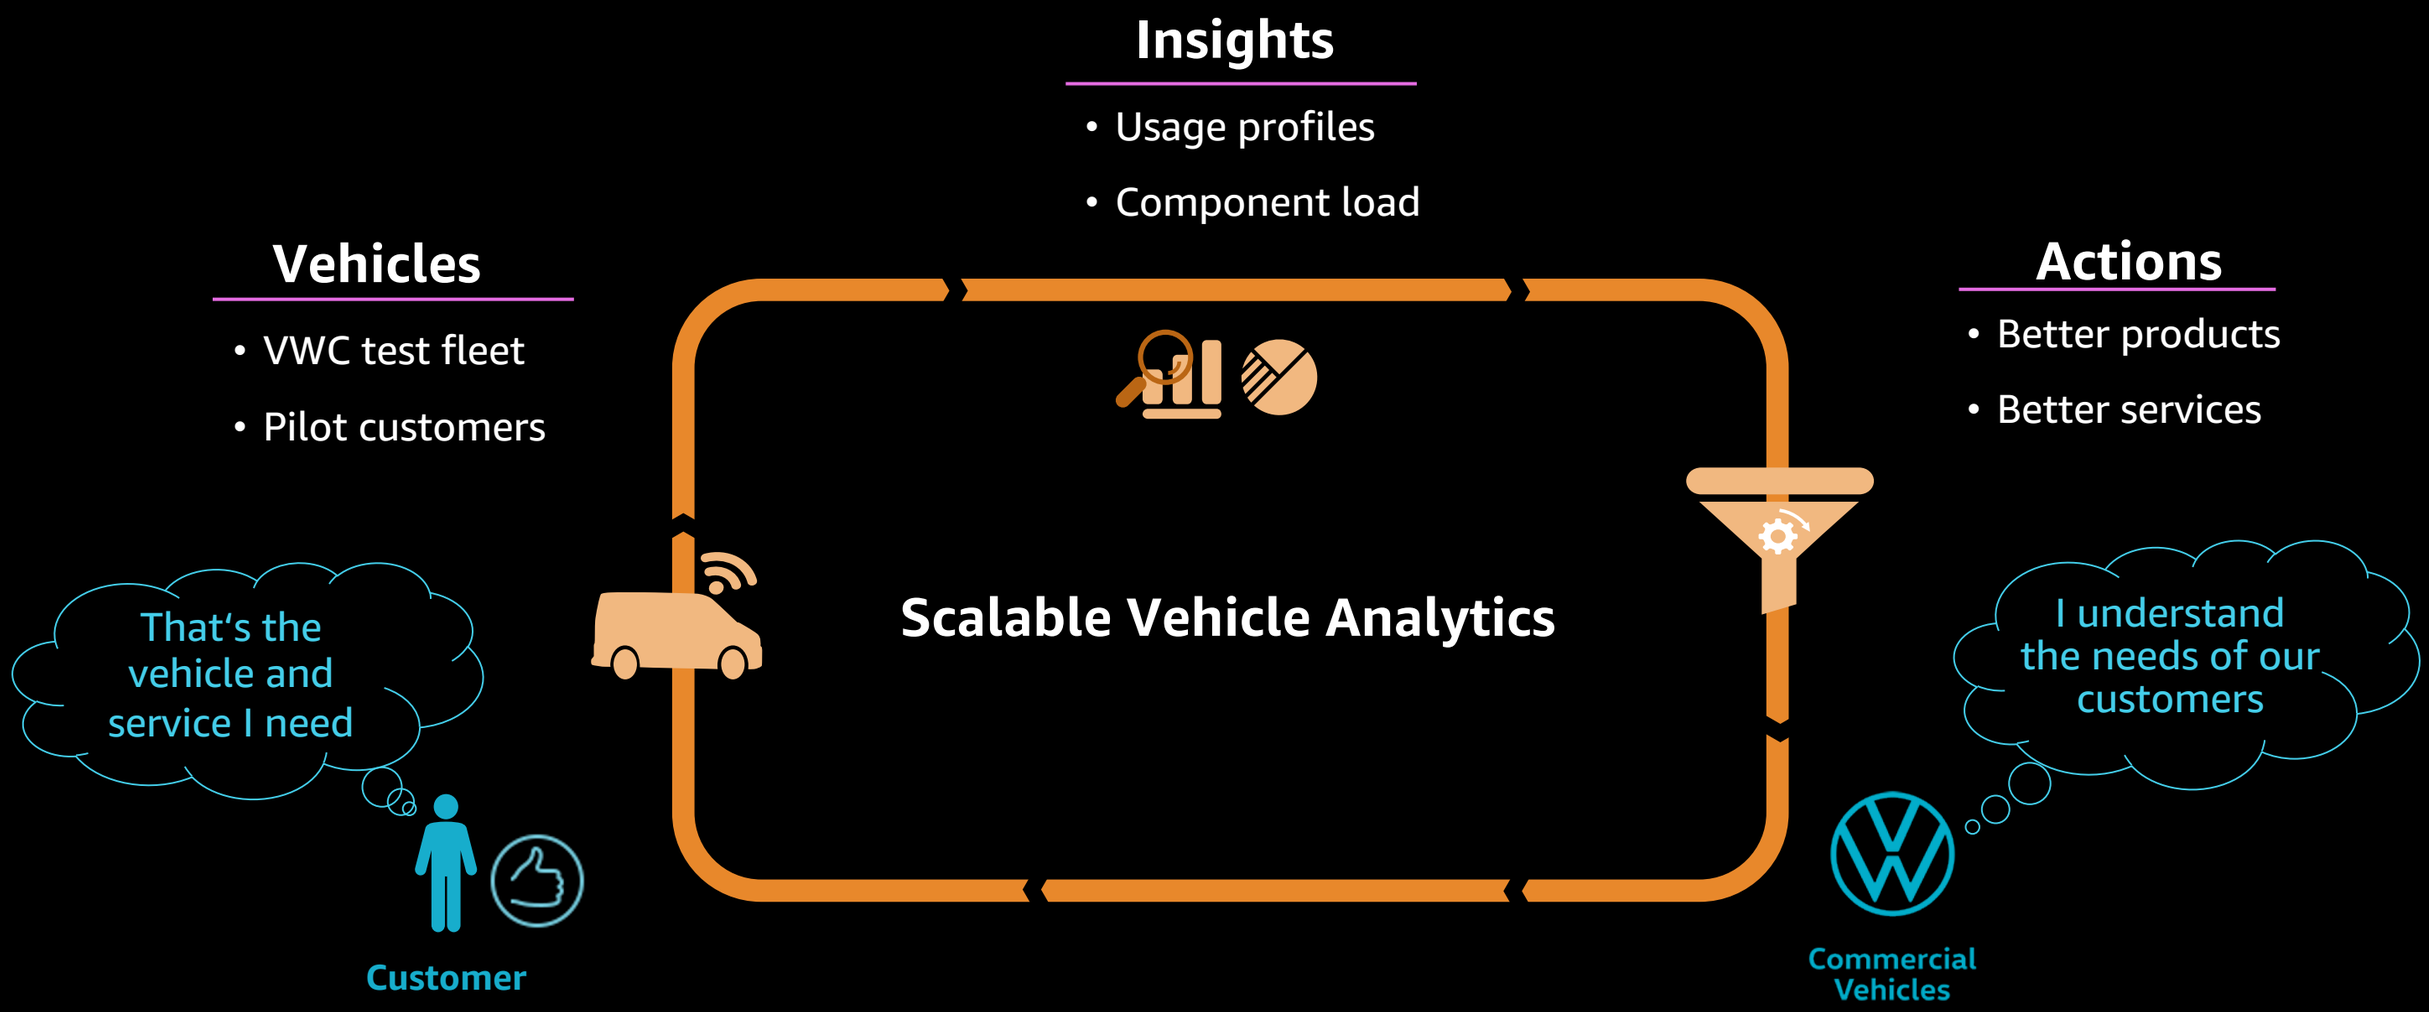
Data export

- Feature store interface:
  - Amazon API Gateway with authentication
  - AWS Lambda as proxy for Elasticsearch
- Download of usage profiles
  - Integration in local data warehouse on premises
  - Engineers using self-service BI tools
  - Comparison of customer usage with test profiles



Amazon Elasticsearch Service AWS Lambda Amazon API Gateway

Data leads the way to customer understanding



Thank you!

John Fisher
Principal Automotive Business Development Manager
AWS

Dr. Michael Nolting
Director of Digital Services and Data Analytics
Volkswagen Commercial Vehicles

Dr. Mohammad Al-Rifai
Lead Data Scientist
Volkswagen Commercial Vehicles

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.