

# Autonomous Vehicles

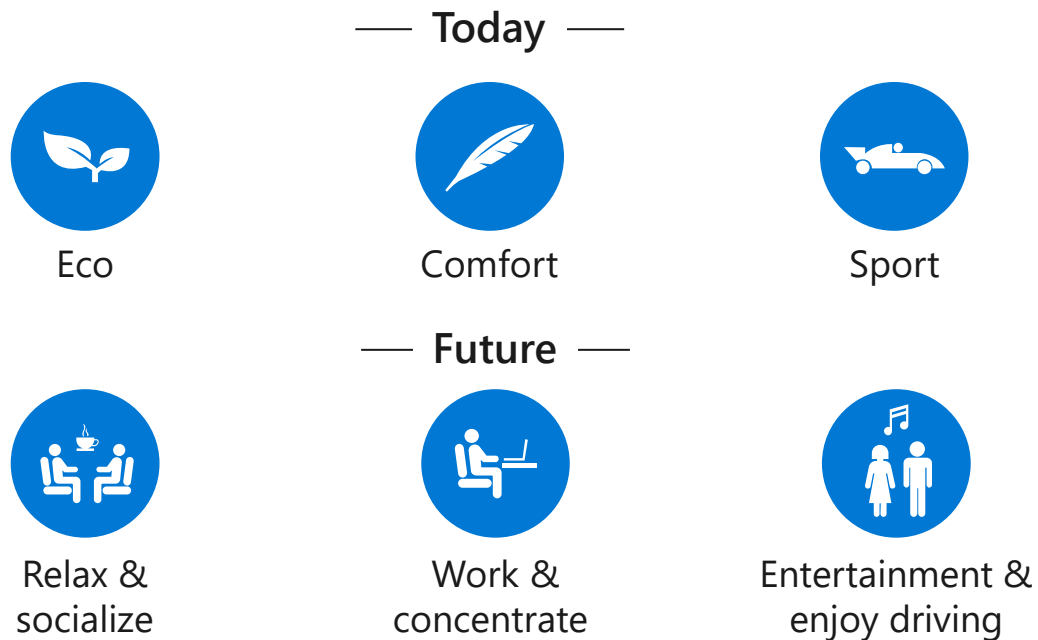
Accelerating development with Intelligent cloud and intelligent edge

Gabriel Sallah EMEA GBB HPC/AI

# The Race to Autonomous Explained

- Since 2015, [44 corporations from automotive & tech](#) have announced large investments or alliances for autonomous vehicles.
- Today, approximately 15% of entire auto OEM R&D budgets are spent on autonomous vehicle research.

## Envisioning drive modes today vs. future



Today's vehicle buying criterion are fuel economy, driving comfort, & drivability.

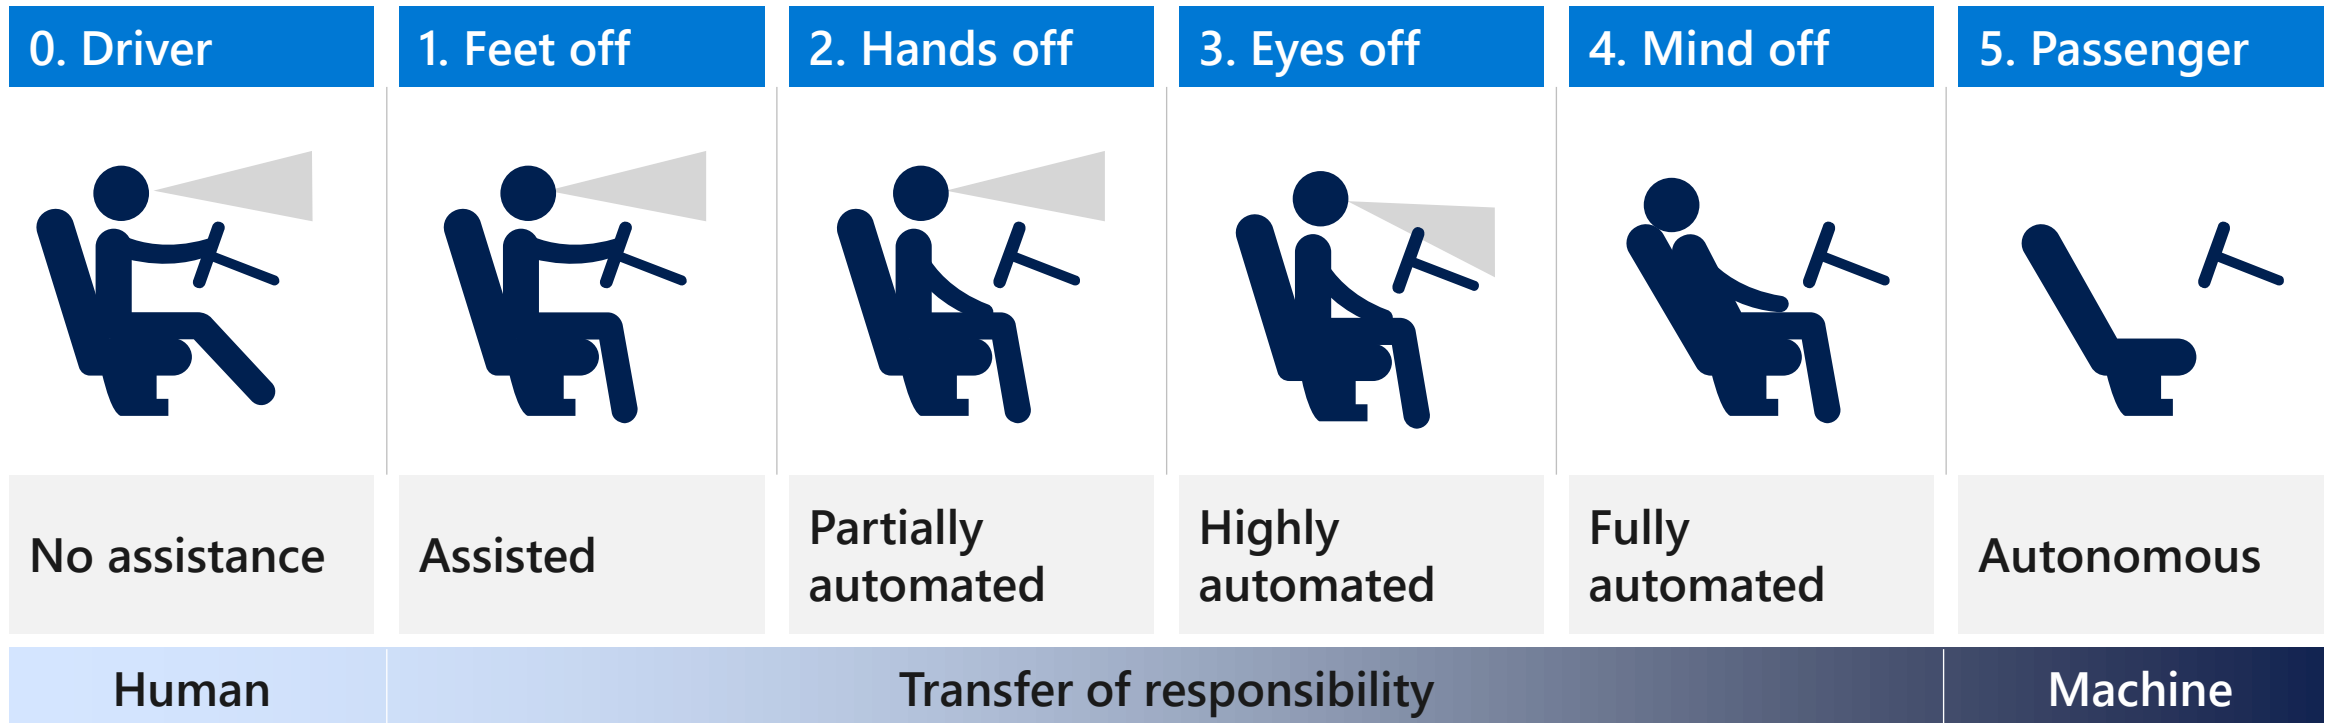
Autonomous "Mobility as a Service" customer decisions will center around productivity & infotainment – all new revenue streams for automotive companies.

Whoever gets to full autonomy first will take share & grab these new sources of revenue.

In this era of disruption and digital transformation, auto OEMs are therefore fighting for survival.

# Definitions: Levels of Autonomy

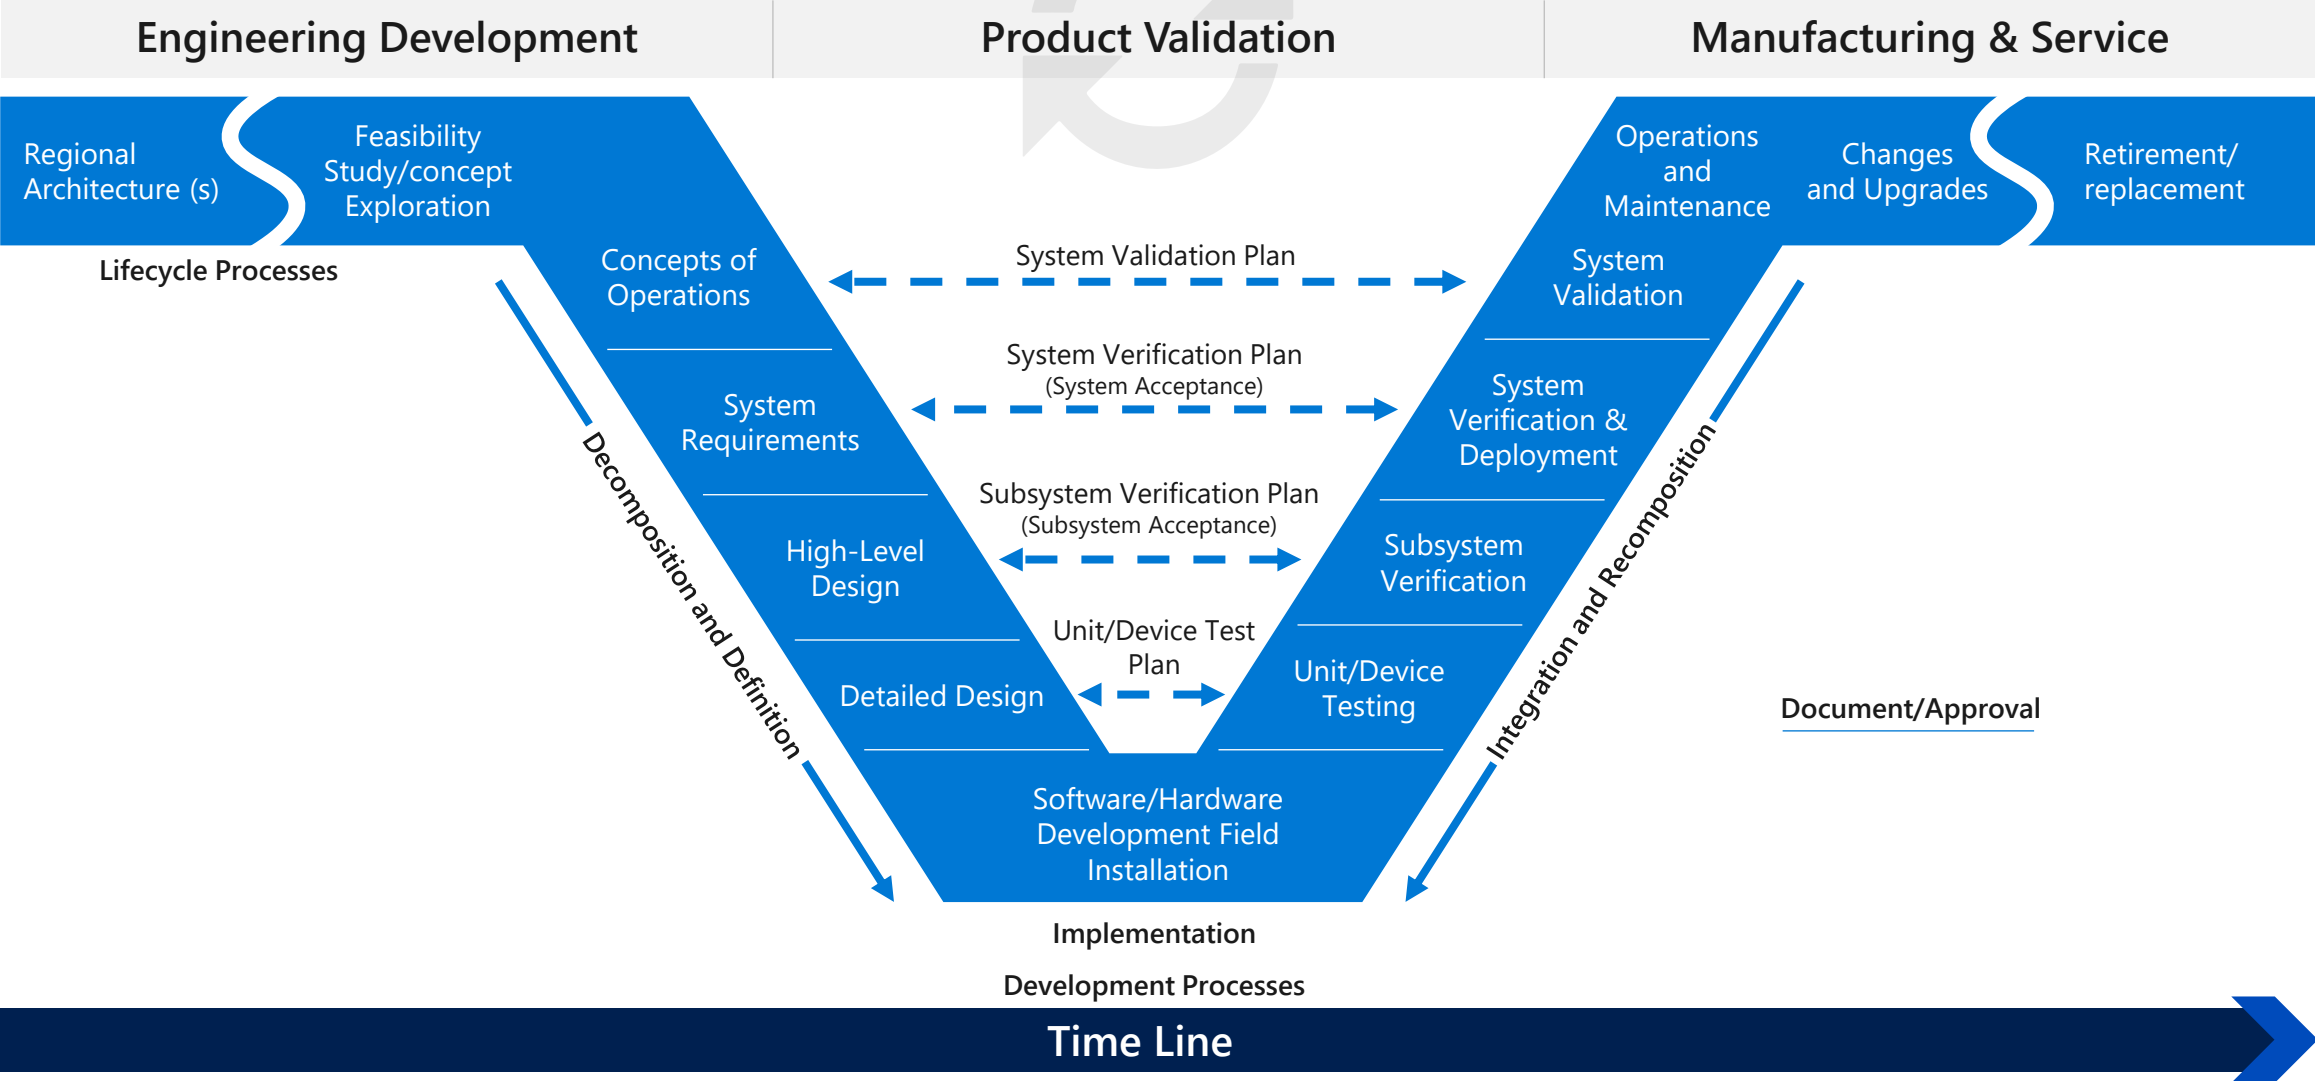
## The five stages of autonomy



Sources: Evercore ISI, SAE International


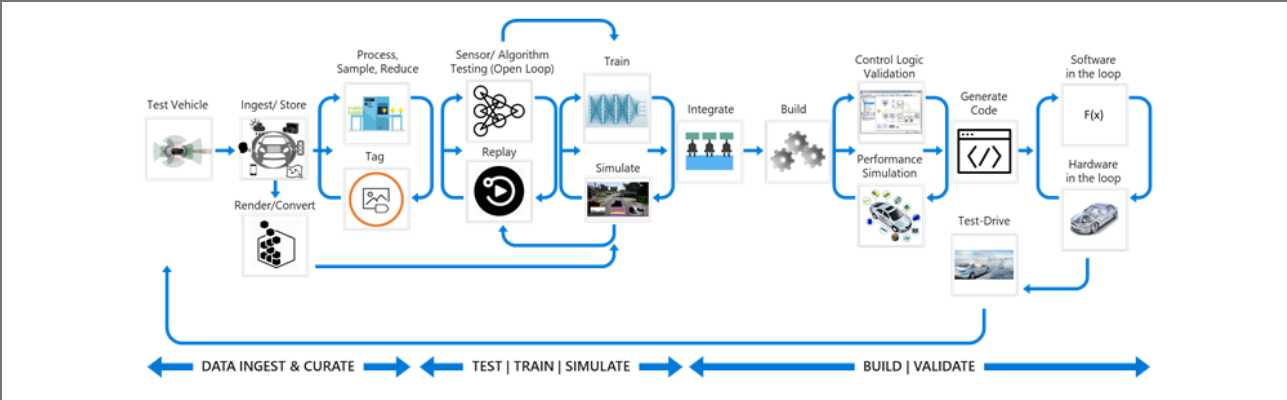


# OEM System Engineering Process

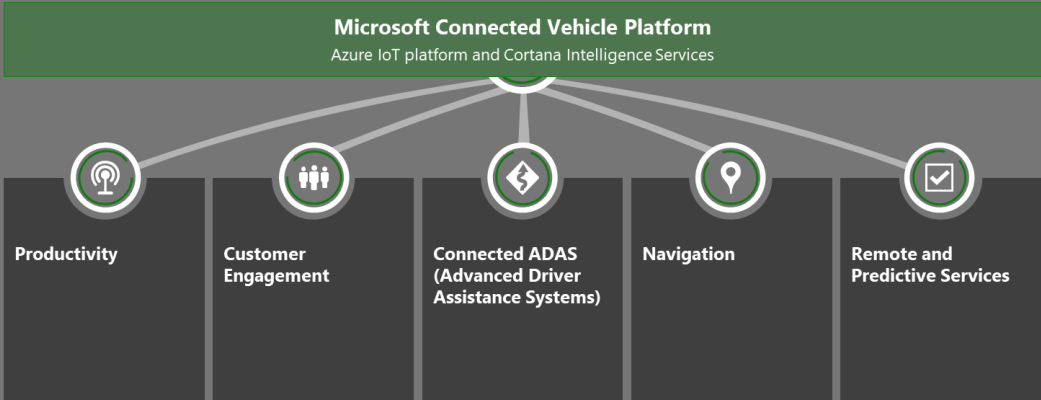


# OEM Product and Services Lifecycle

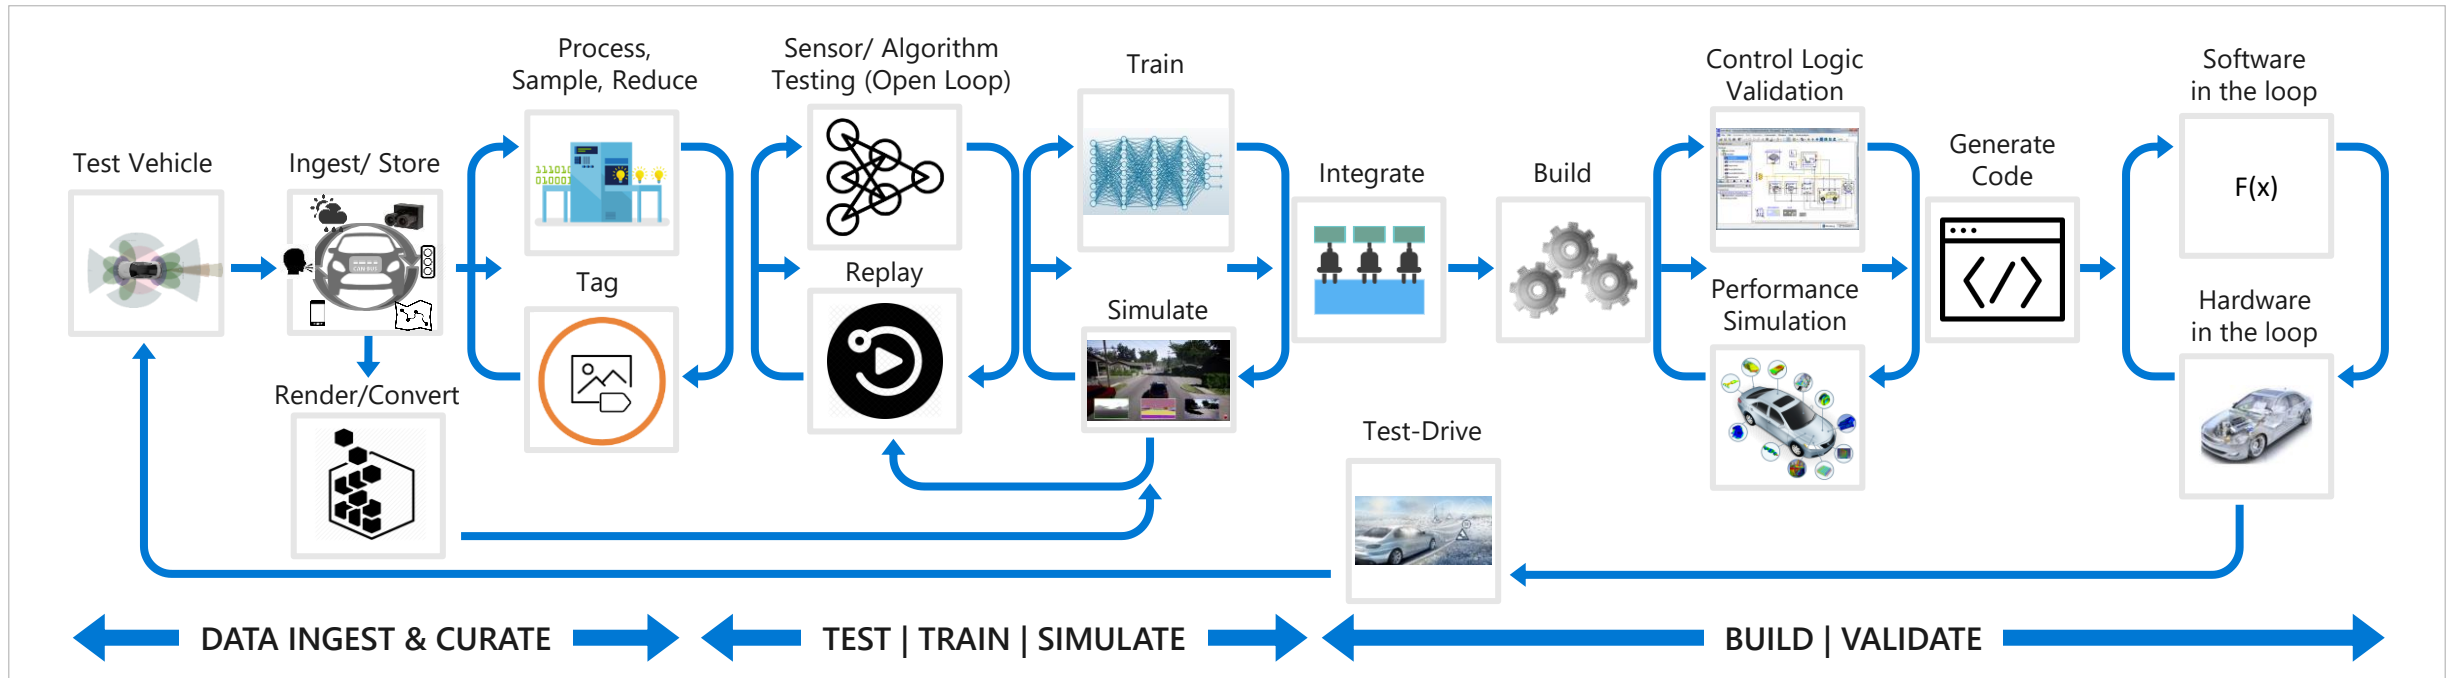
## Autonomous Driving Development



## Mobility Services



# AD Dev/Test: E2E Workflow & Partner Ecosystem



Azure Express Route

Azure Storage, Data Lake Gen2

Azure Data Box

Azure Data Factory

Azure Batch

Cycle Cloud

Databricks & HDInsight

Azure Machine Learning

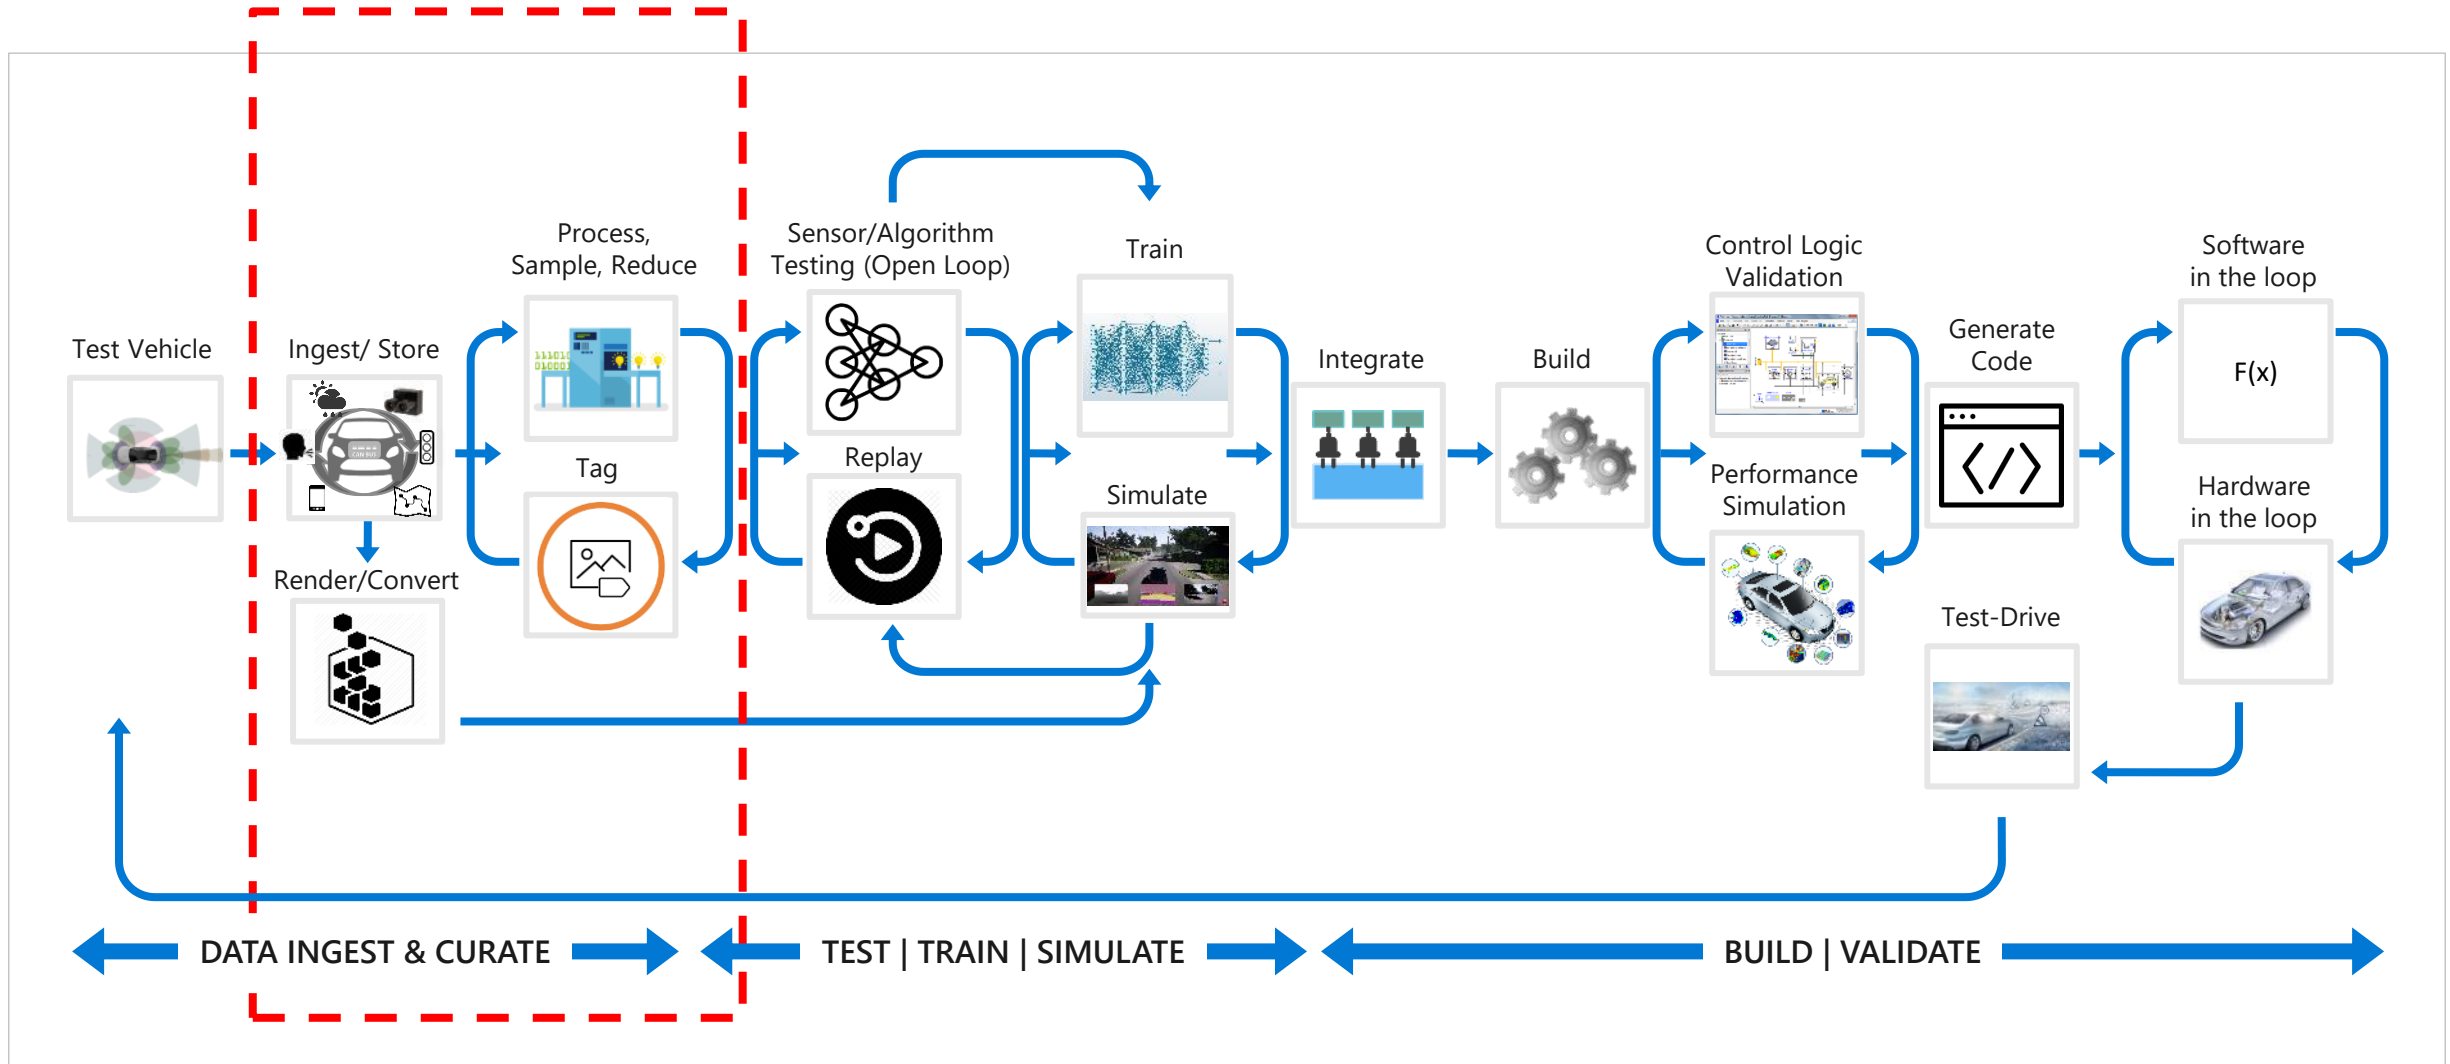
Cognitive Services

Azure IoT Edge

Azure Kubernetes Service

Azure Key Vault, AAD

# AD Dev/Test: End-to-End Processing & Workflow



## Data Ingest

### Express Route

- Private, secure, predictable network
- 100+ carrier partners
- 10+ Gbps

### Data Box



Disk Import



Data Box Disk



Data Box – 100 TB

### Azure Edge

- IoT Edge
- AI Toolkit

Multiple options to filter and ingest PBs data every day regardless of fleet type or location

## Data Storage

### Scalable

- Foundational service for Microsoft\*
- 40 million transactions per second
- Multi-PB accounts

### Performant

- 100 Gbps ingress
- 200 Gbps egress
- Account-scale object throughput

### Secure & Compliant

- Client & Service Encryption
- AAD Integration + ACLs
- Broad & deep compliance portfolio

### Durable

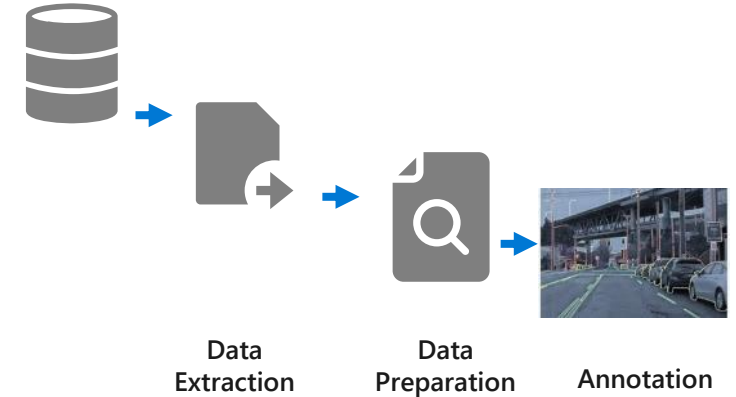
- Multiple redundancy options
- Strong consistency, data integrity
- Policy: Versioning & WORM locks

### Cost Effective


- Interested storage tiers
- Lifecycle management
- Rich Metrics

Massively scalable object storage for unstructured data

## Data Curation



### Microsoft Services



Transform and process, PII data redaction, annotation & training data preparation on both 1<sup>st</sup> and 3<sup>rd</sup> party tools.



# Generating Ground Truth with Labeling?




## Semantic Segmentation

Each Pixel of the image is assigned a category



## Object Detection and Classification

Bounding box drawn around each object of interest



## 3D Point Cloud Labeling

Objects of interest as assigned a category in 3D LIDAR point cloud

“Ground truth” is the accuracy of the training set's classification for supervised learning techniques

Currently done manually

Longer term – auto labelling

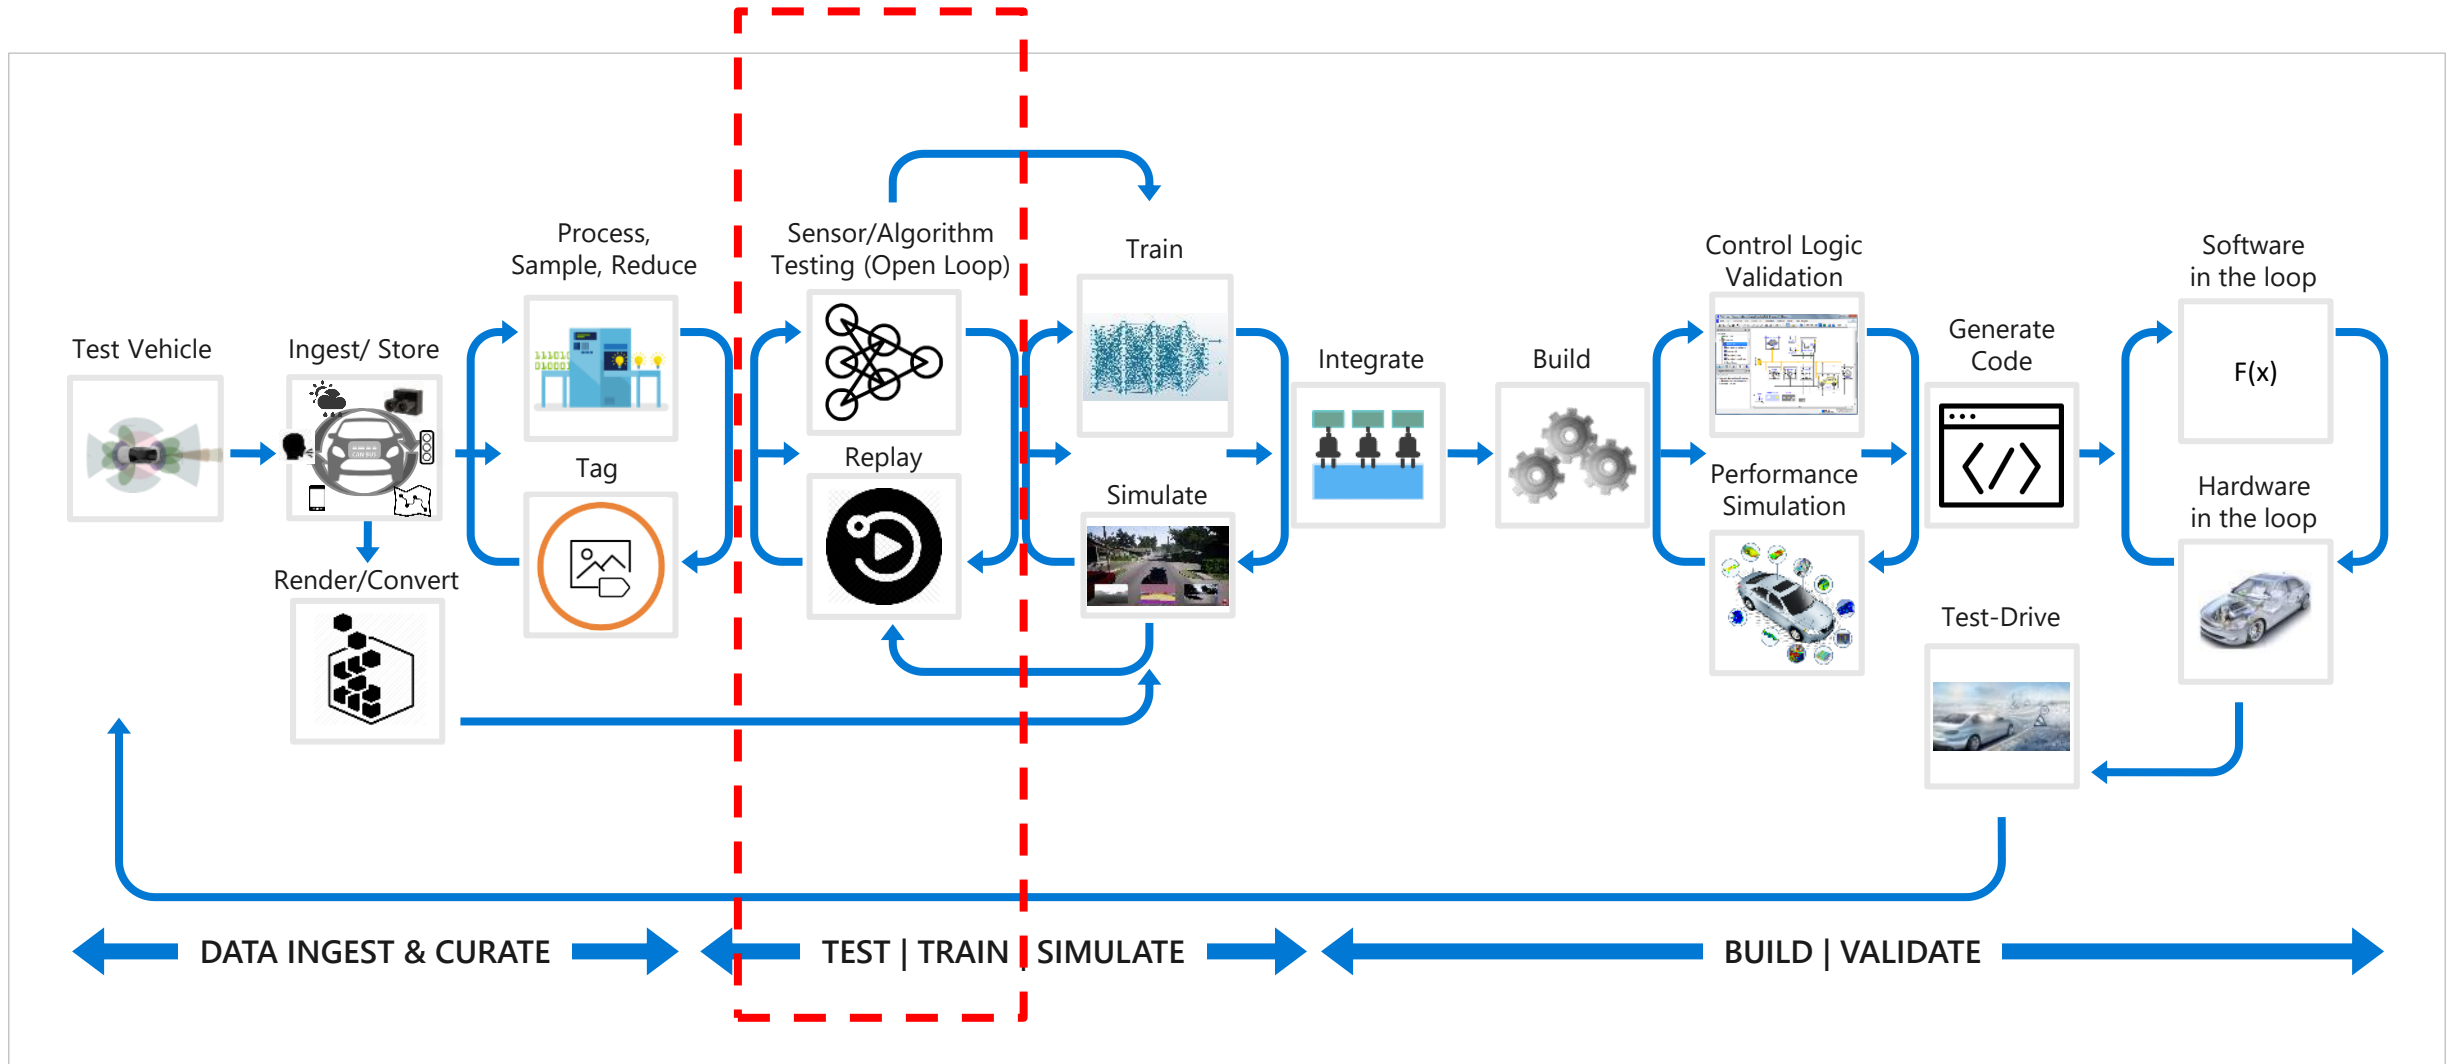
Partners provide

- Results based managed service contracts
- Trained workforce, on demand
- Mature labeling tools



Ground Truth is one of the most critical elements of Machine Learning for Training and Validation

# AD Dev/Test: End-to-End Processing & Workflow



# What is Open Loop Testing?


RGB Camera



Radar



Lidar




Open Loop or “sensor reprocessing” is a large-scale embarrassingly parallel compute job that processes 10~100s PBs of data using tens of thousands of cores and requiring very high I/O throughput of >30GB/sec.

Data sets are fused from multiple sensor types representing a singular view of what the on-vehicle computer vision systems “saw” when navigating the real world. An open loop test is where the performance of the algorithms is tested & validated against ground truth using replay and scoring. The output is used later in the workflow for algorithm training.


Compute VMs



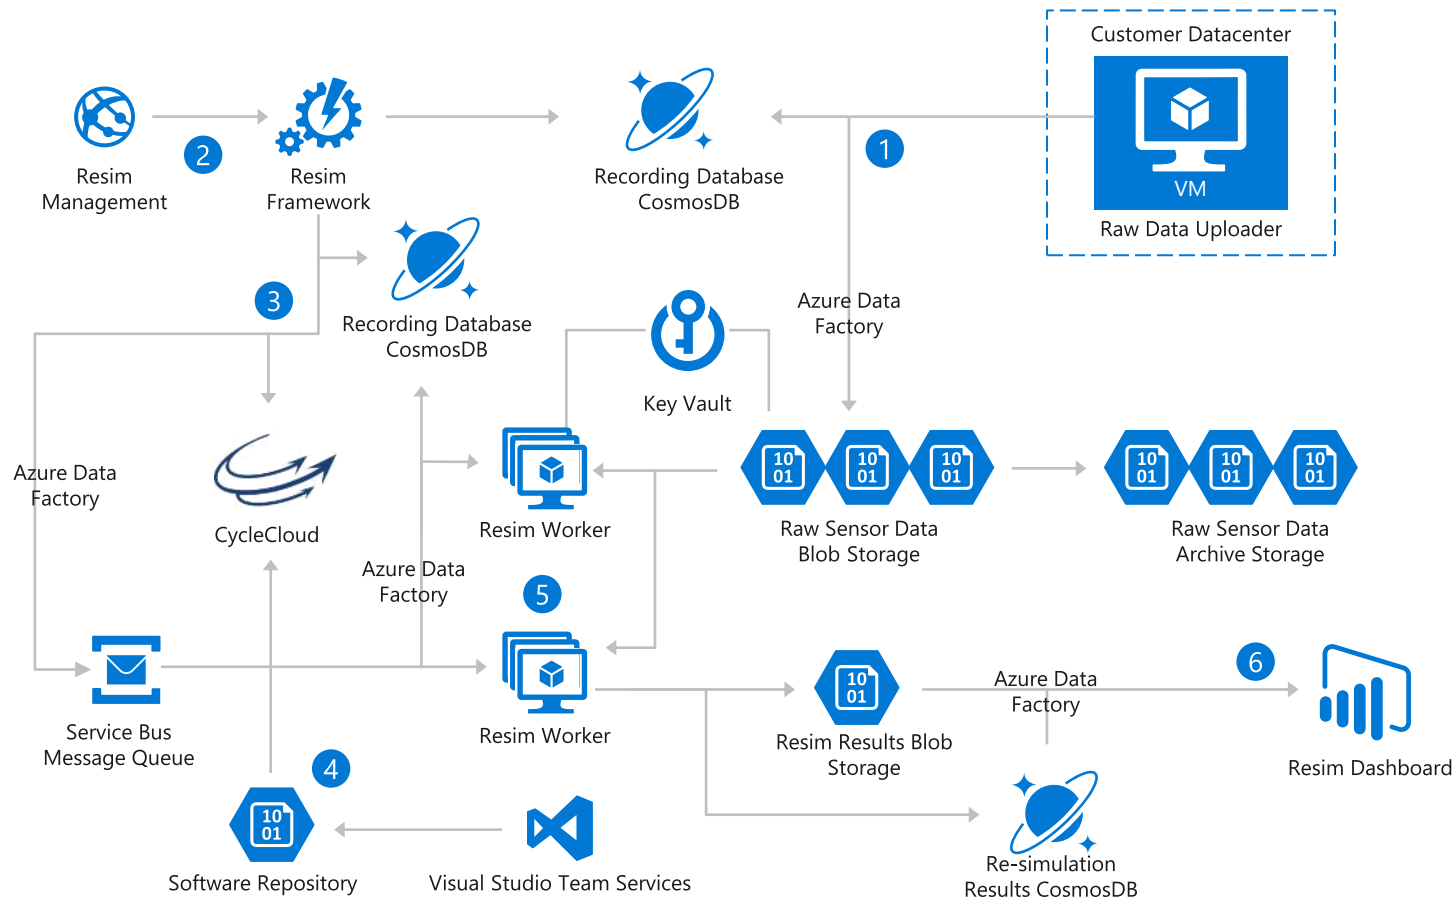
files



Storage



# Algorithm Validation

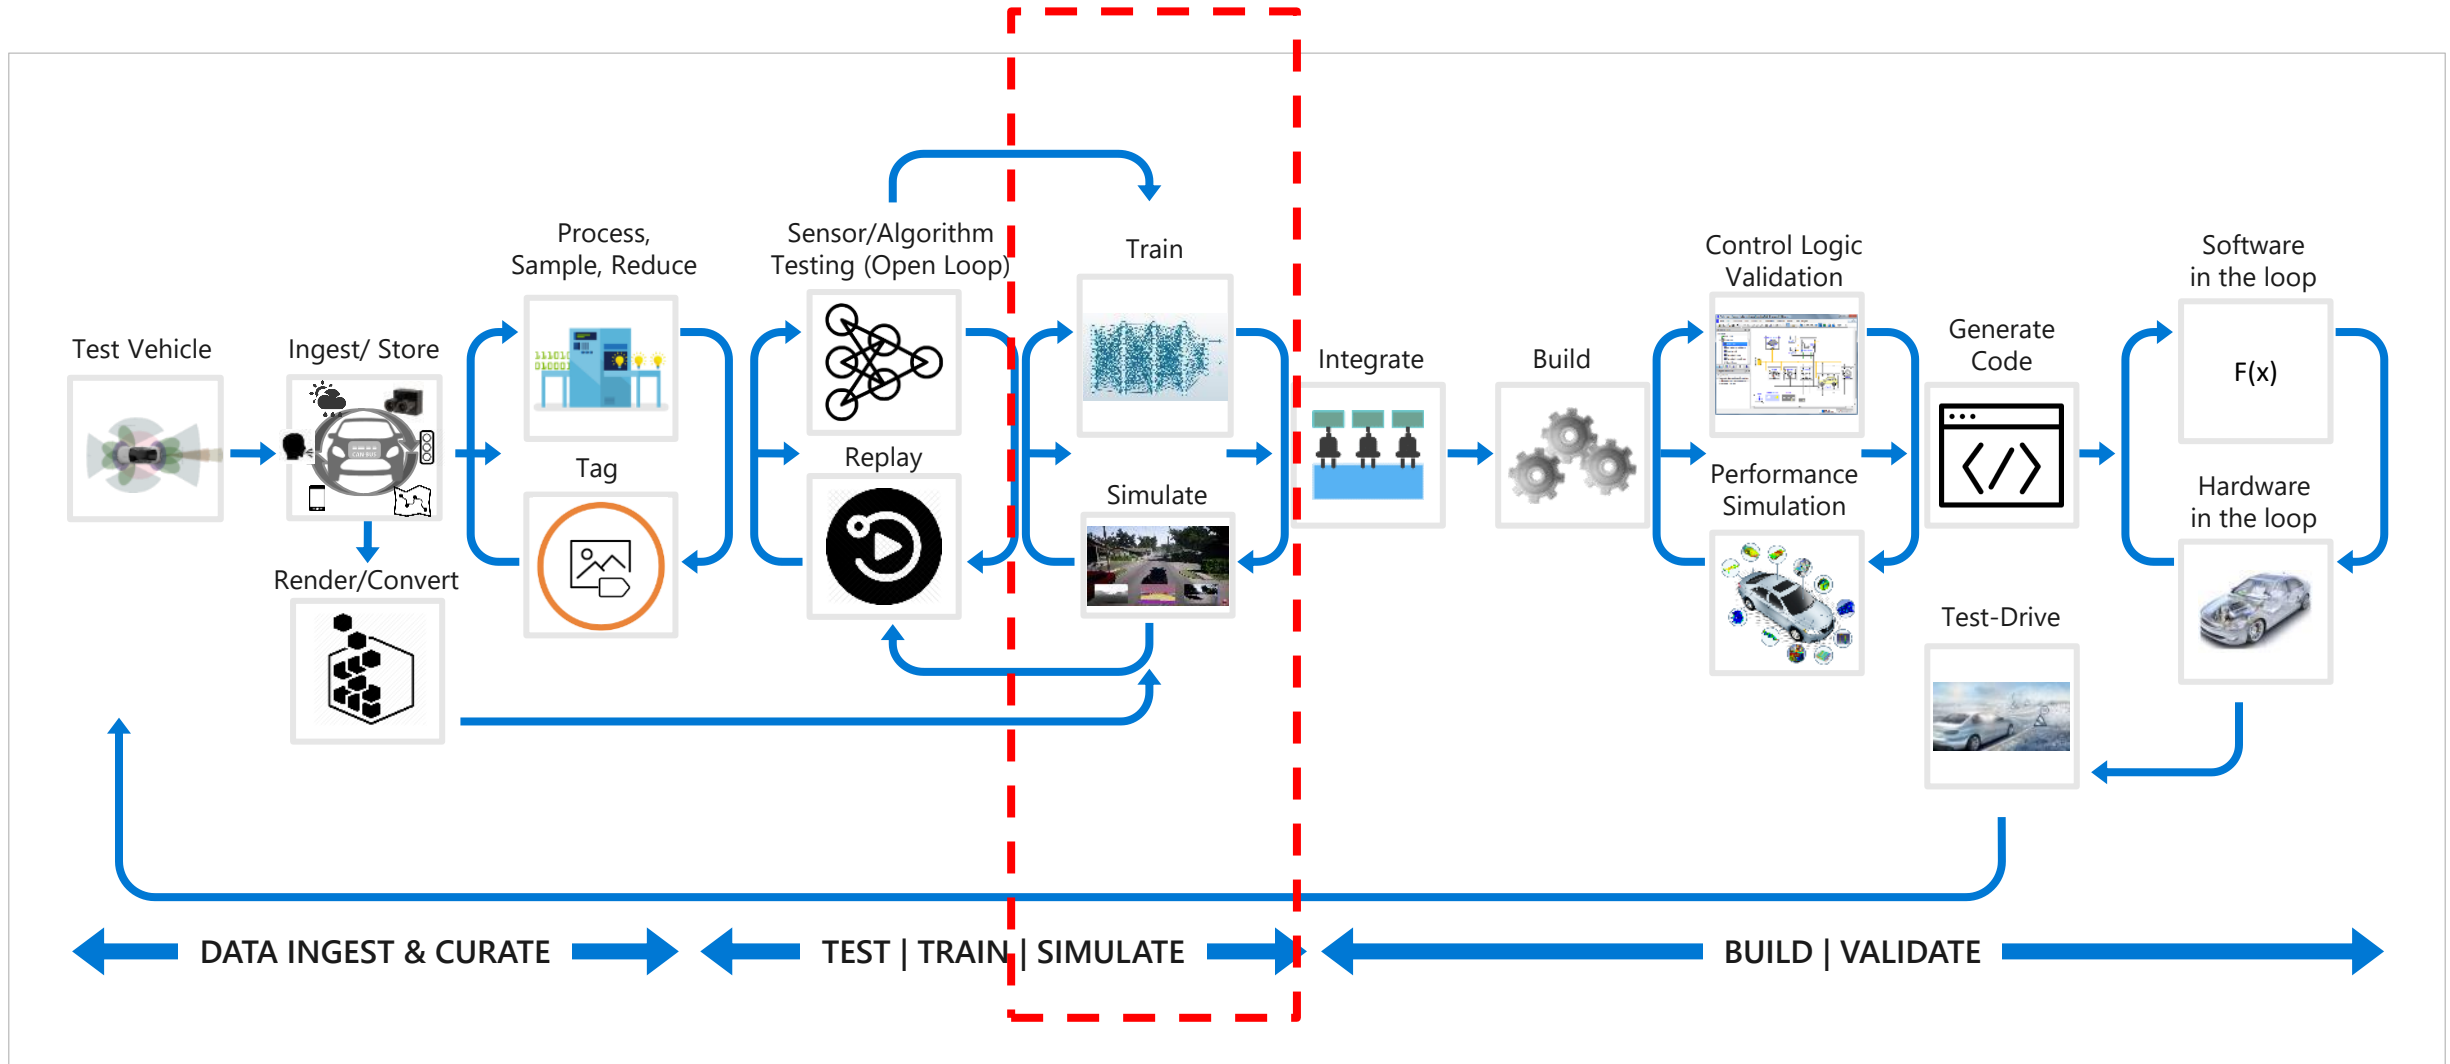


Open loop testing tools  
(ex. ADTF)

Comprehensive Test  
management framework

Verification and validation of training algorithms and sensors via open loop testing

# AD Dev/Test: End-to-End Processing & Workflow



# What is Training?


Leveraging libraries with classified images, numerous training jobs run in parallel – each trained to recognize a specific type of object

## Some examples:

- Traffic signs
- Pedestrians
- Other vehicles
- Road lanes, lines, & edges
- Buildings
- Shrubbery/vegetation.
- Non-vehicle traffic (E.g., cyclists)
- Roadside objects

## Each type of job will then have multiple framework variants

- E.g., for road signs, separate ones for speed limit, warning signs, etc.



# What is Deep Learning?

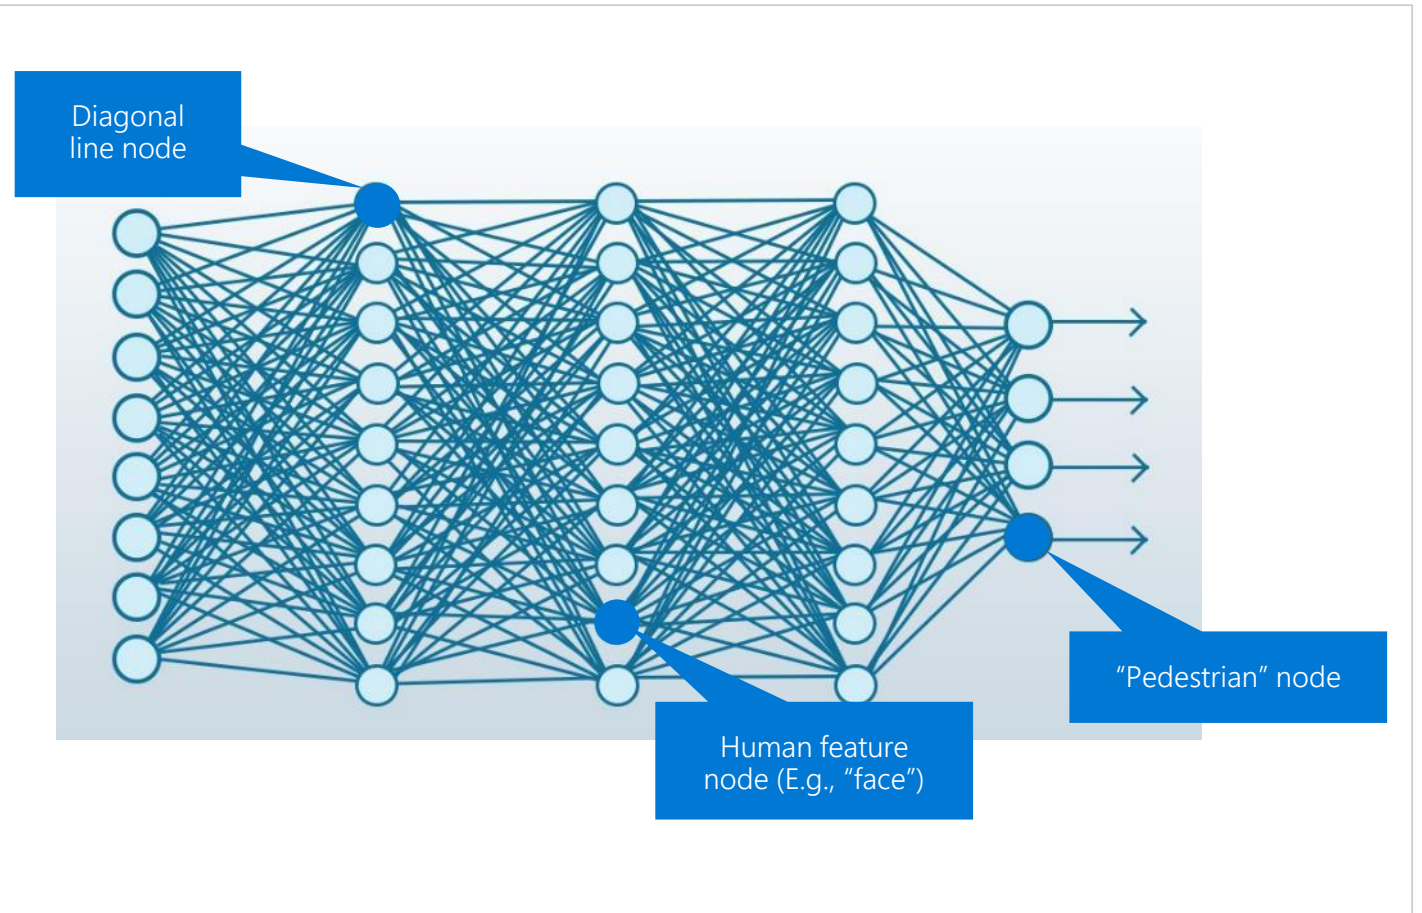
With the data ingested from vehicles, image recognition training using deep learning is a key technical enabler for AD/ADAS development. GPU clusters predominantly used for these jobs.

## How Deep Learning works

It first identifies what are the edges that are most relevant to find out (for example...) a pedestrian or inanimate roadside object

It then builds on this hierarchically to find what combination of shapes and edges we can find. For example, if arms/legs are present, or whether heads/faces are present, etc.

After consecutive hierarchical identification of complex concepts, it then decides which of these features are responsible for finding the answer.



# Training vs Inference

- Two distinctive types of inferencing are characteristic of AD/ADAS algorithm development:
  - Simulation-Driven Inferencing
  - Inferencing on Vehicle
- Simulation-driven Inferencing enables engineers to test training models digitally thereby reducing the number of miles test fleets must drive in order to develop more robust autonomous capabilities.
- On-vehicle inferencing is the run-time implementation that powers individual vehicles' autonomous features

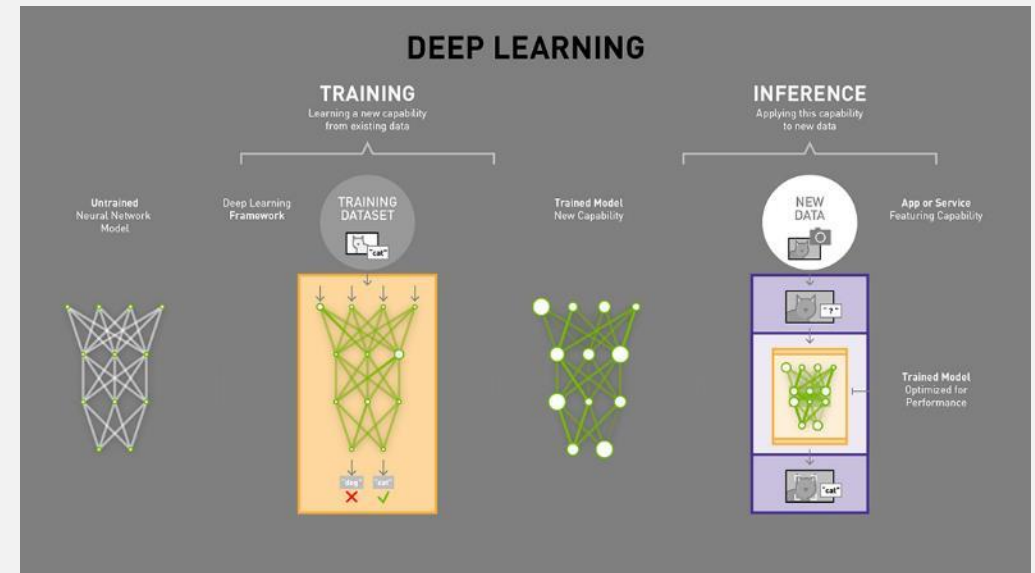


Image Source: [NVIDIA blog](#)




# Training


## Training and Deploy Custom AI End-to-End



Your Data with Any AI Tools

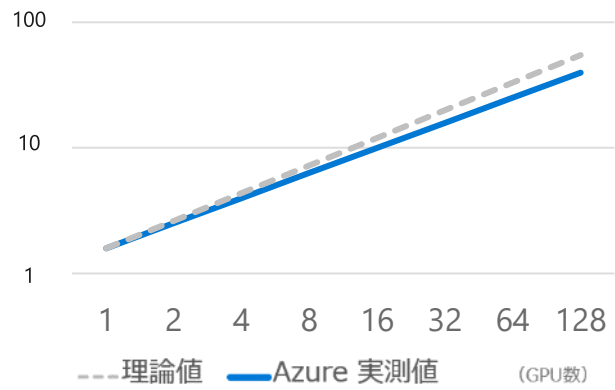


Training with Scale-Out GPU Clusters on Demand



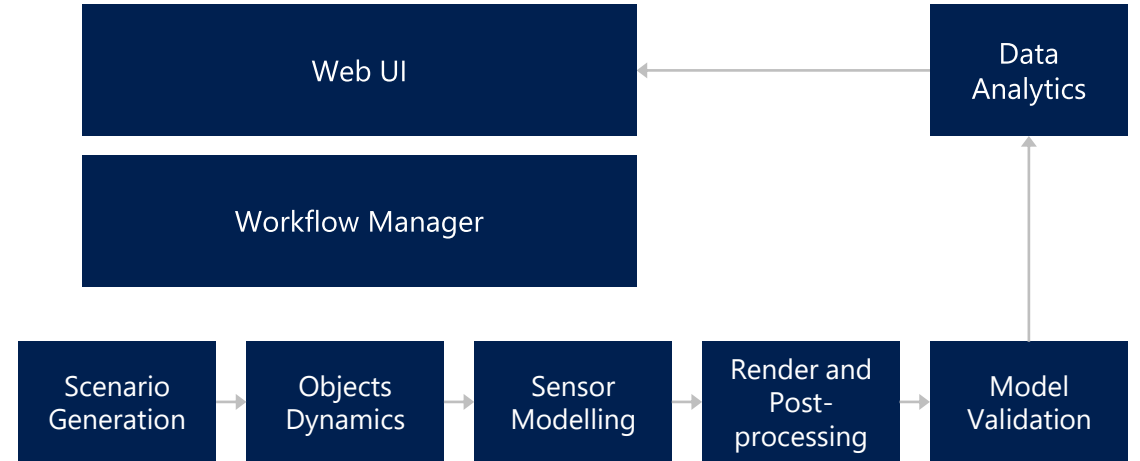
Intelligence In your Apps and Data Services

## Reduce Training Job Run-time with ChainerMN on Azure




With InfiniBand, Scaled to 128 GPUs – job run-time shrunk by a factor of > 100x

# Simulation (Partner based)



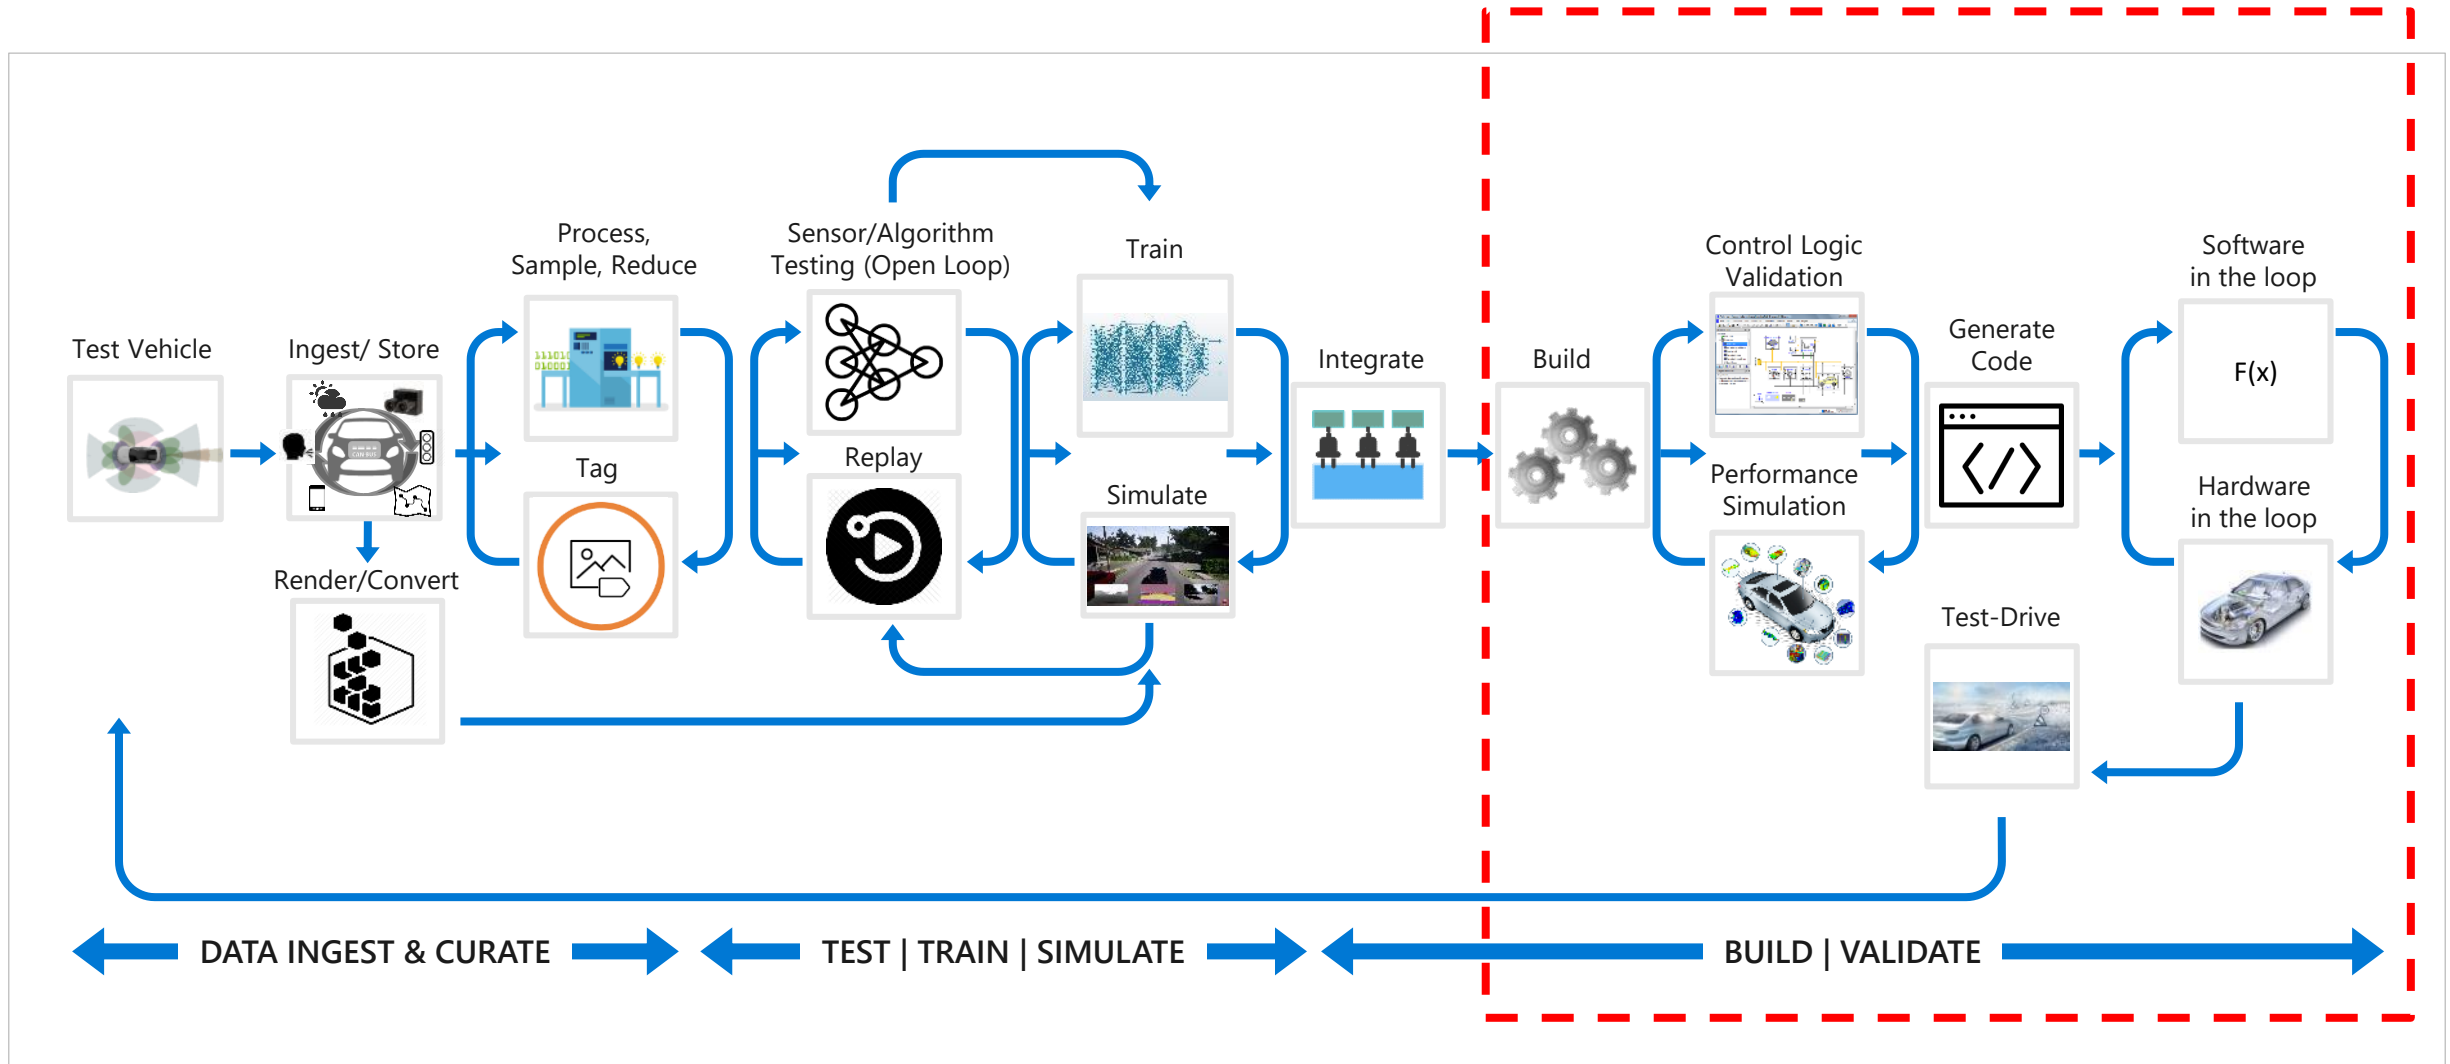
## Simulation As a Service

### Microsoft Services

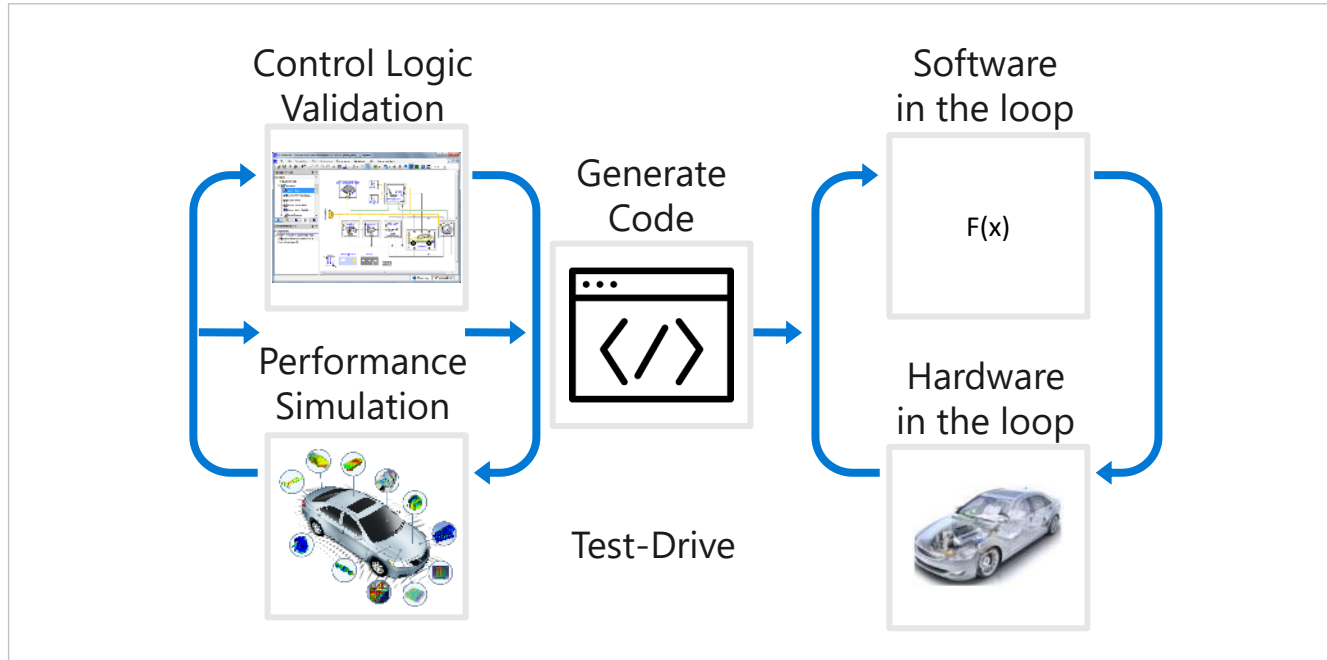


Driving billions of miles necessary for the development of Autonomous vehicles, is only possible by running Simulation at scale on the cloud

# AD Dev/Test: End-to-End Processing & Workflow



# HIL and SIL Validation



Microsoft Services

- Blob storage
- Batch
- GPU VM
- Active Directory
- Container service

Embedded system validation via hardware-in-loop and software-in-loop

## Working with partners on

- Comprehensive test management framework
- HIL solutions
- System validation tools
- Workflow management services
- Managed services

