AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

AUT201

AWS for autonomous driving with Aurora Innovation

Tim Kelton
Director Software Engineering
Aurora Innovation Inc.

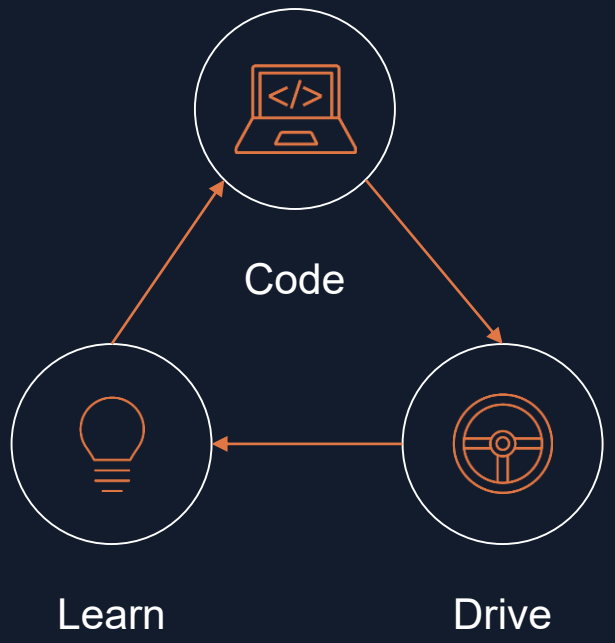
Stefano Marzani
WW Tech Leader, Autonomous Driving
Amazon Web Services

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

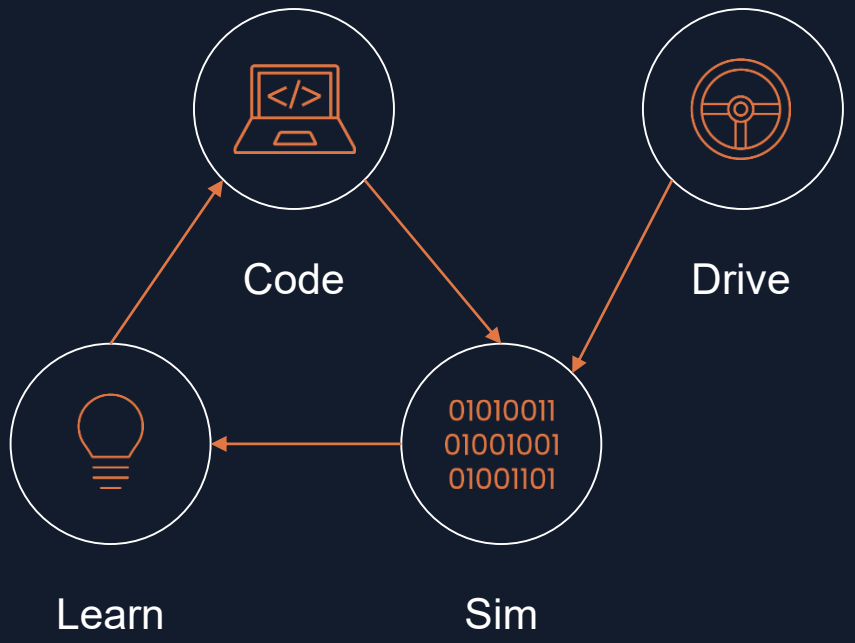


Aurora

©2021 | Aurora Proprietar



SLOW



FAST!

Aurora

©2021 | Aurora Proprietary

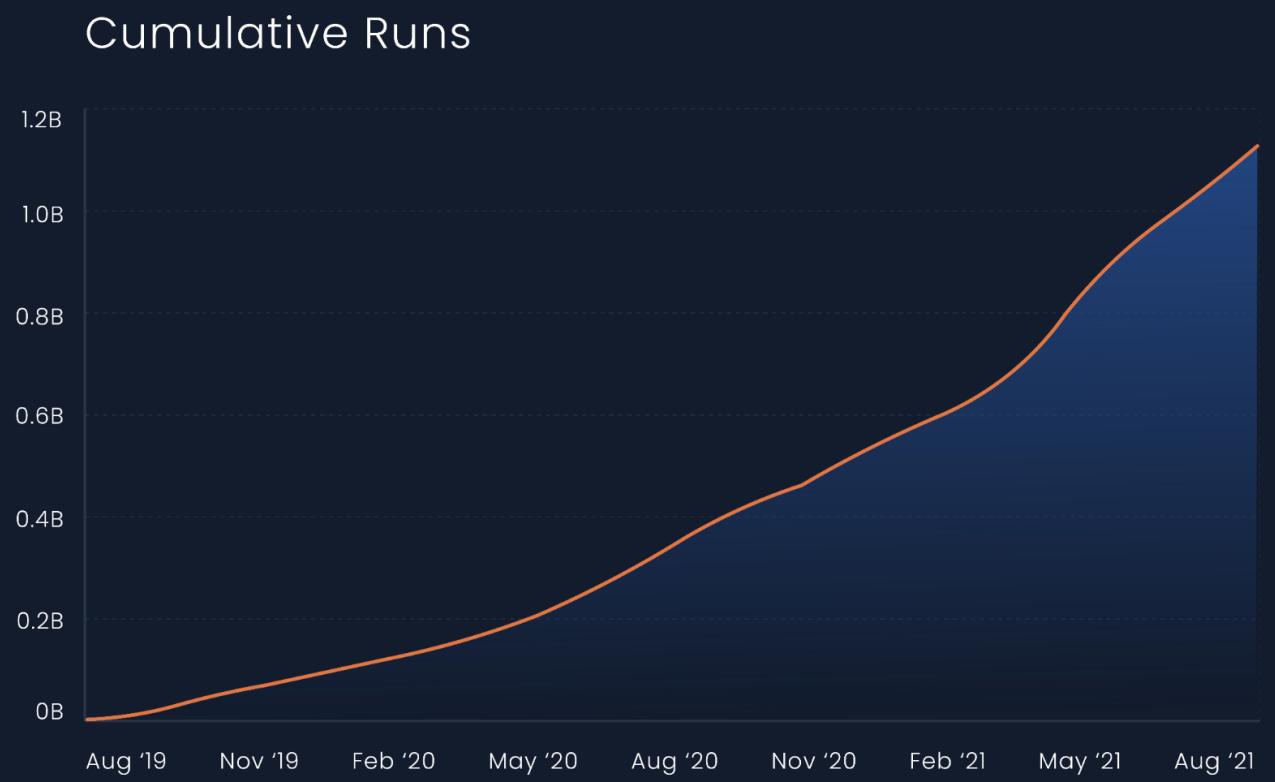
Simulation Stats

5 million Sims run each day

60,000 CPUs utilized each hour

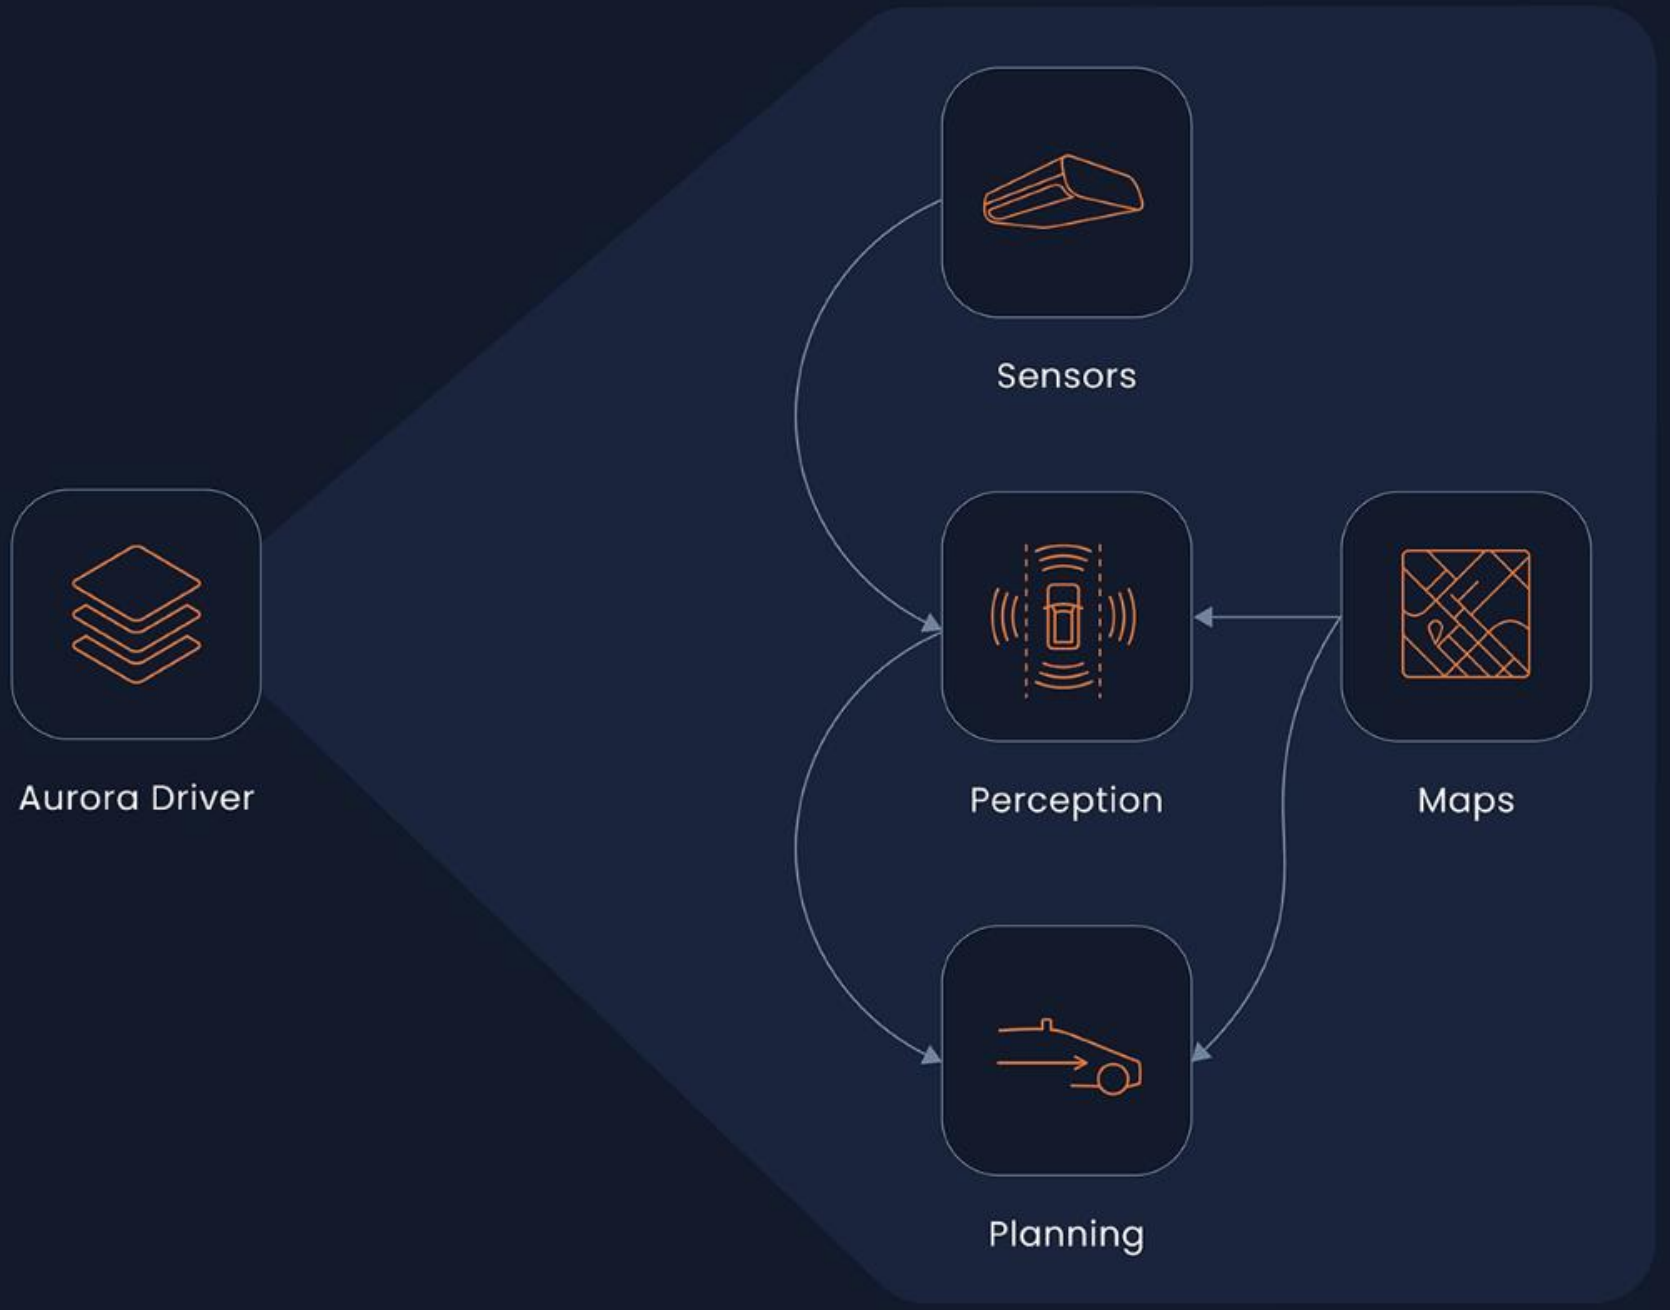
>5,000 Compute nodes parallelized

6 billion Equivalent miles driven



Aurora

©2021 | Aurora Proprietary

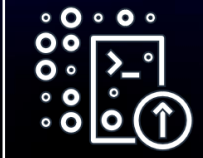


©2021 | Aurora Proprietary

AWS for autonomous driving



Key industry challenges



HYPERSCALE

Thousands of cores of compute for development and validation. TBs of data to collect, ingest, and store every day translates into a PB-scale data processing, storage, and management problem.



AGILITY & SPEED

Optimized software engineering to reduce development and validation costs and enable faster time to market. Future-proof R&D cycles. Integrated and agile to rapidly innovate.



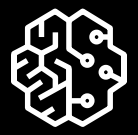
COST

PB-scale data storage and large-scale compute costs, managing fleet operations, significant capex of on-premises compute, lack of AV expertise requires significant human investment



SAFETY

Safety of the passengers and surrounding environment are top of the mind for all of our customers and their end customers as the decisions are moving to vehicles from humans



ECOSYSTEM PLAY

Interoperability and seamless integration of multiple first-party and third-party workload-specific tools

0101
0110

SW DEFINED

More software, from different vendors, additional non-functional requirements, integration issues, testing coverage



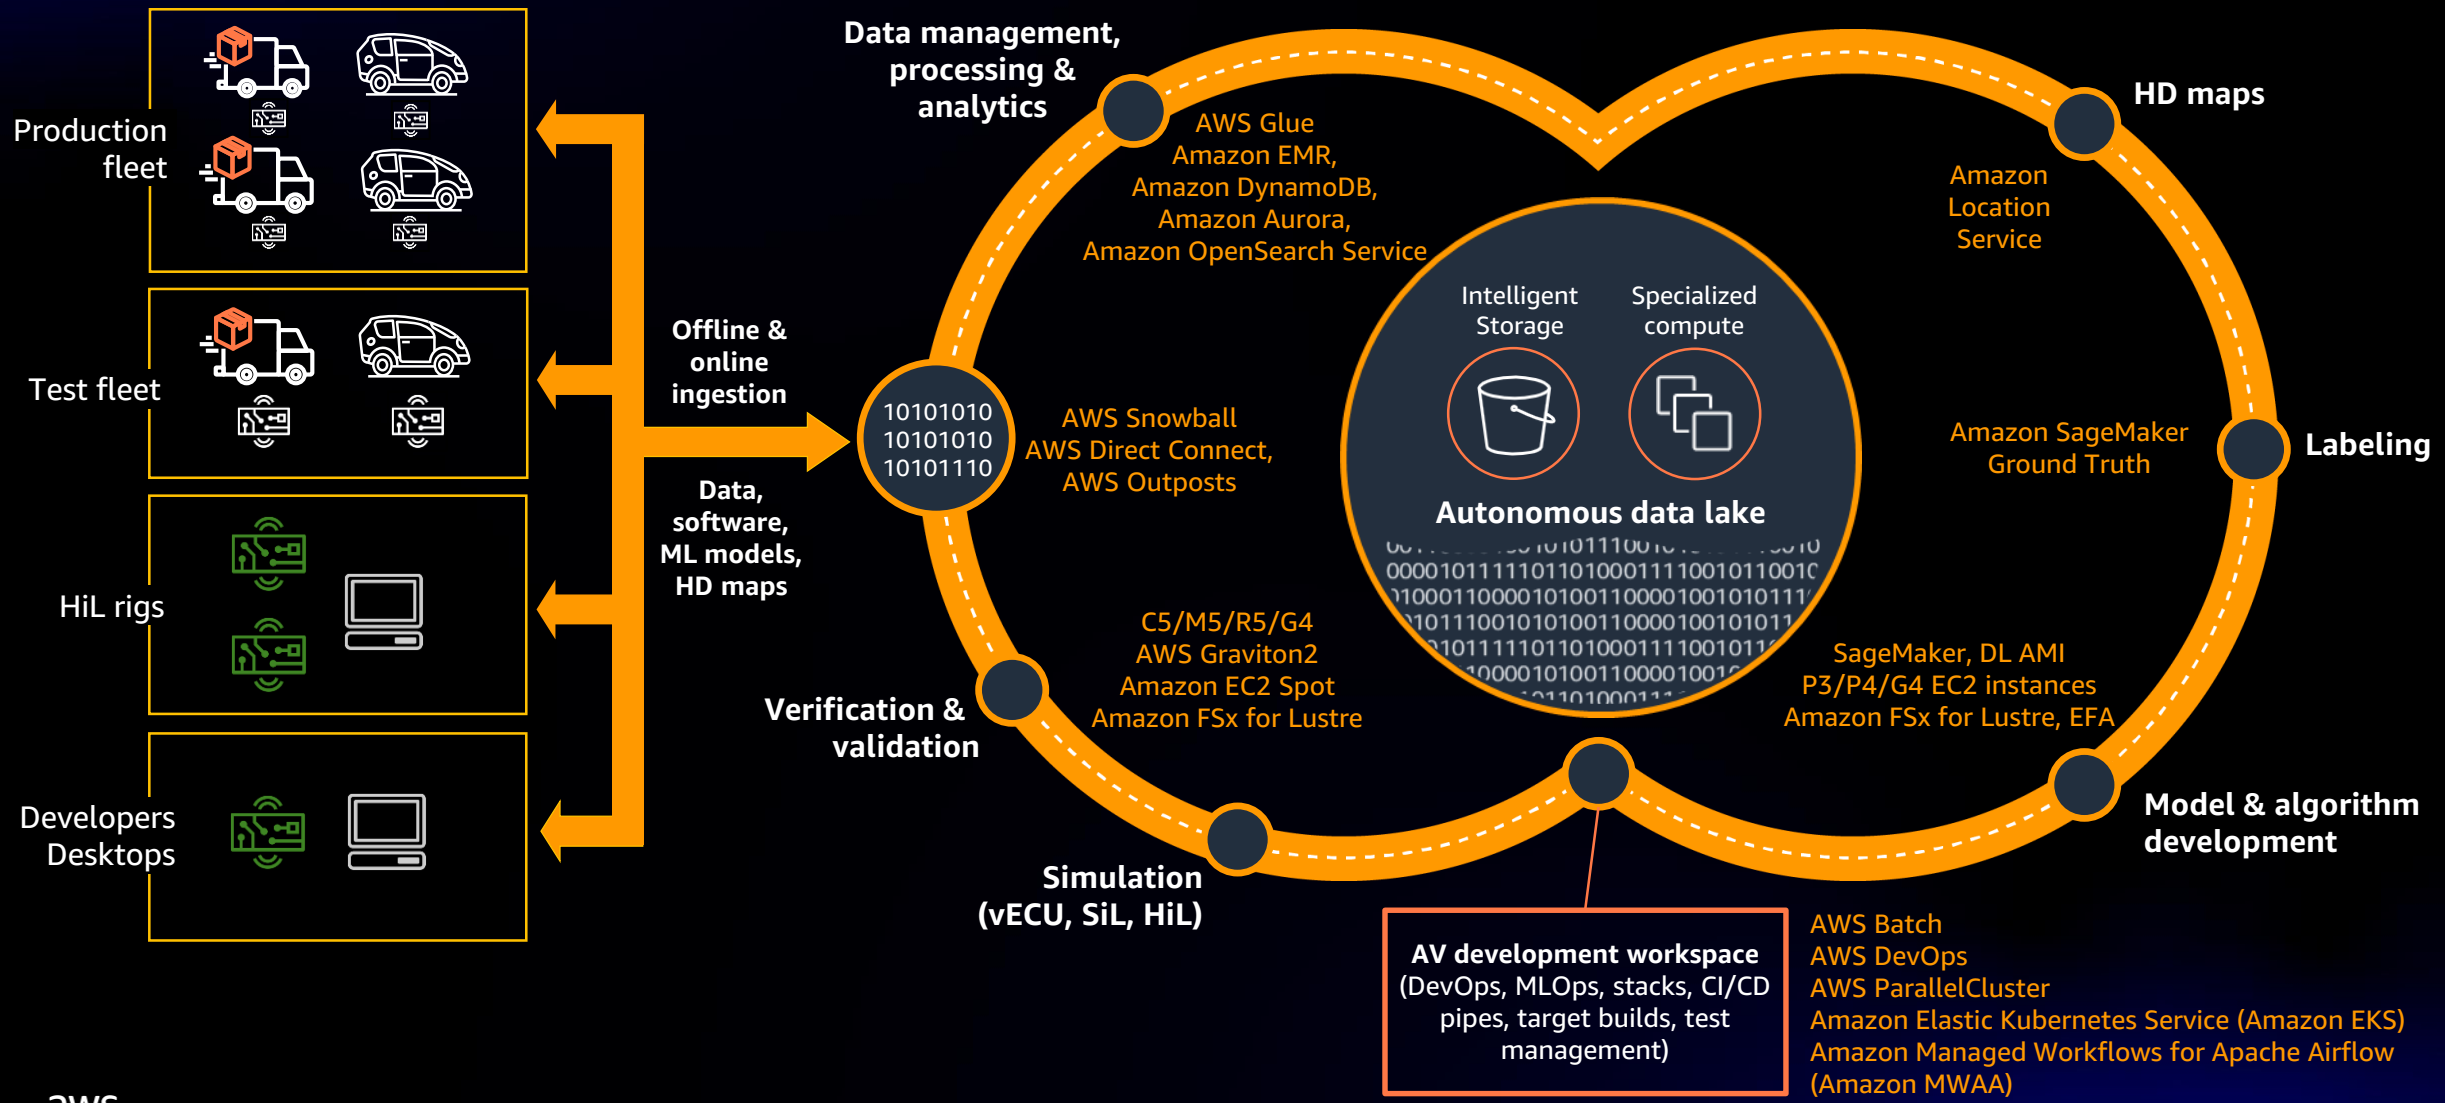
GLOBAL/SECURITY/ DATA PRIVACY

Global fleet requires managed service for complex operations, attaining data, and security compliance across the globe

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Autonomy “infinity” workflow



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Evolving the connected vehicle for ADAS/AV

SNAPSHOT, SHADOW MODE



Autonomy ECU

Autonomy software [test in shadow]

Autonomy software [prod]

OTA agent

Edge intelligence agent

SW

DATA

10101010
10101010
10101110

Data management, processing & analytics

HD maps

Intelligent storage

Specialized compute

Autonomous data lake

Labeling

Verification & validation

Model & algorithm development

Simulation (vECU, SiL, HiL)

AV development workspace







Sensors

Radars

Cameras

Lidar

...

Actuators

Steer-by-wire

Brake-by-wire

...

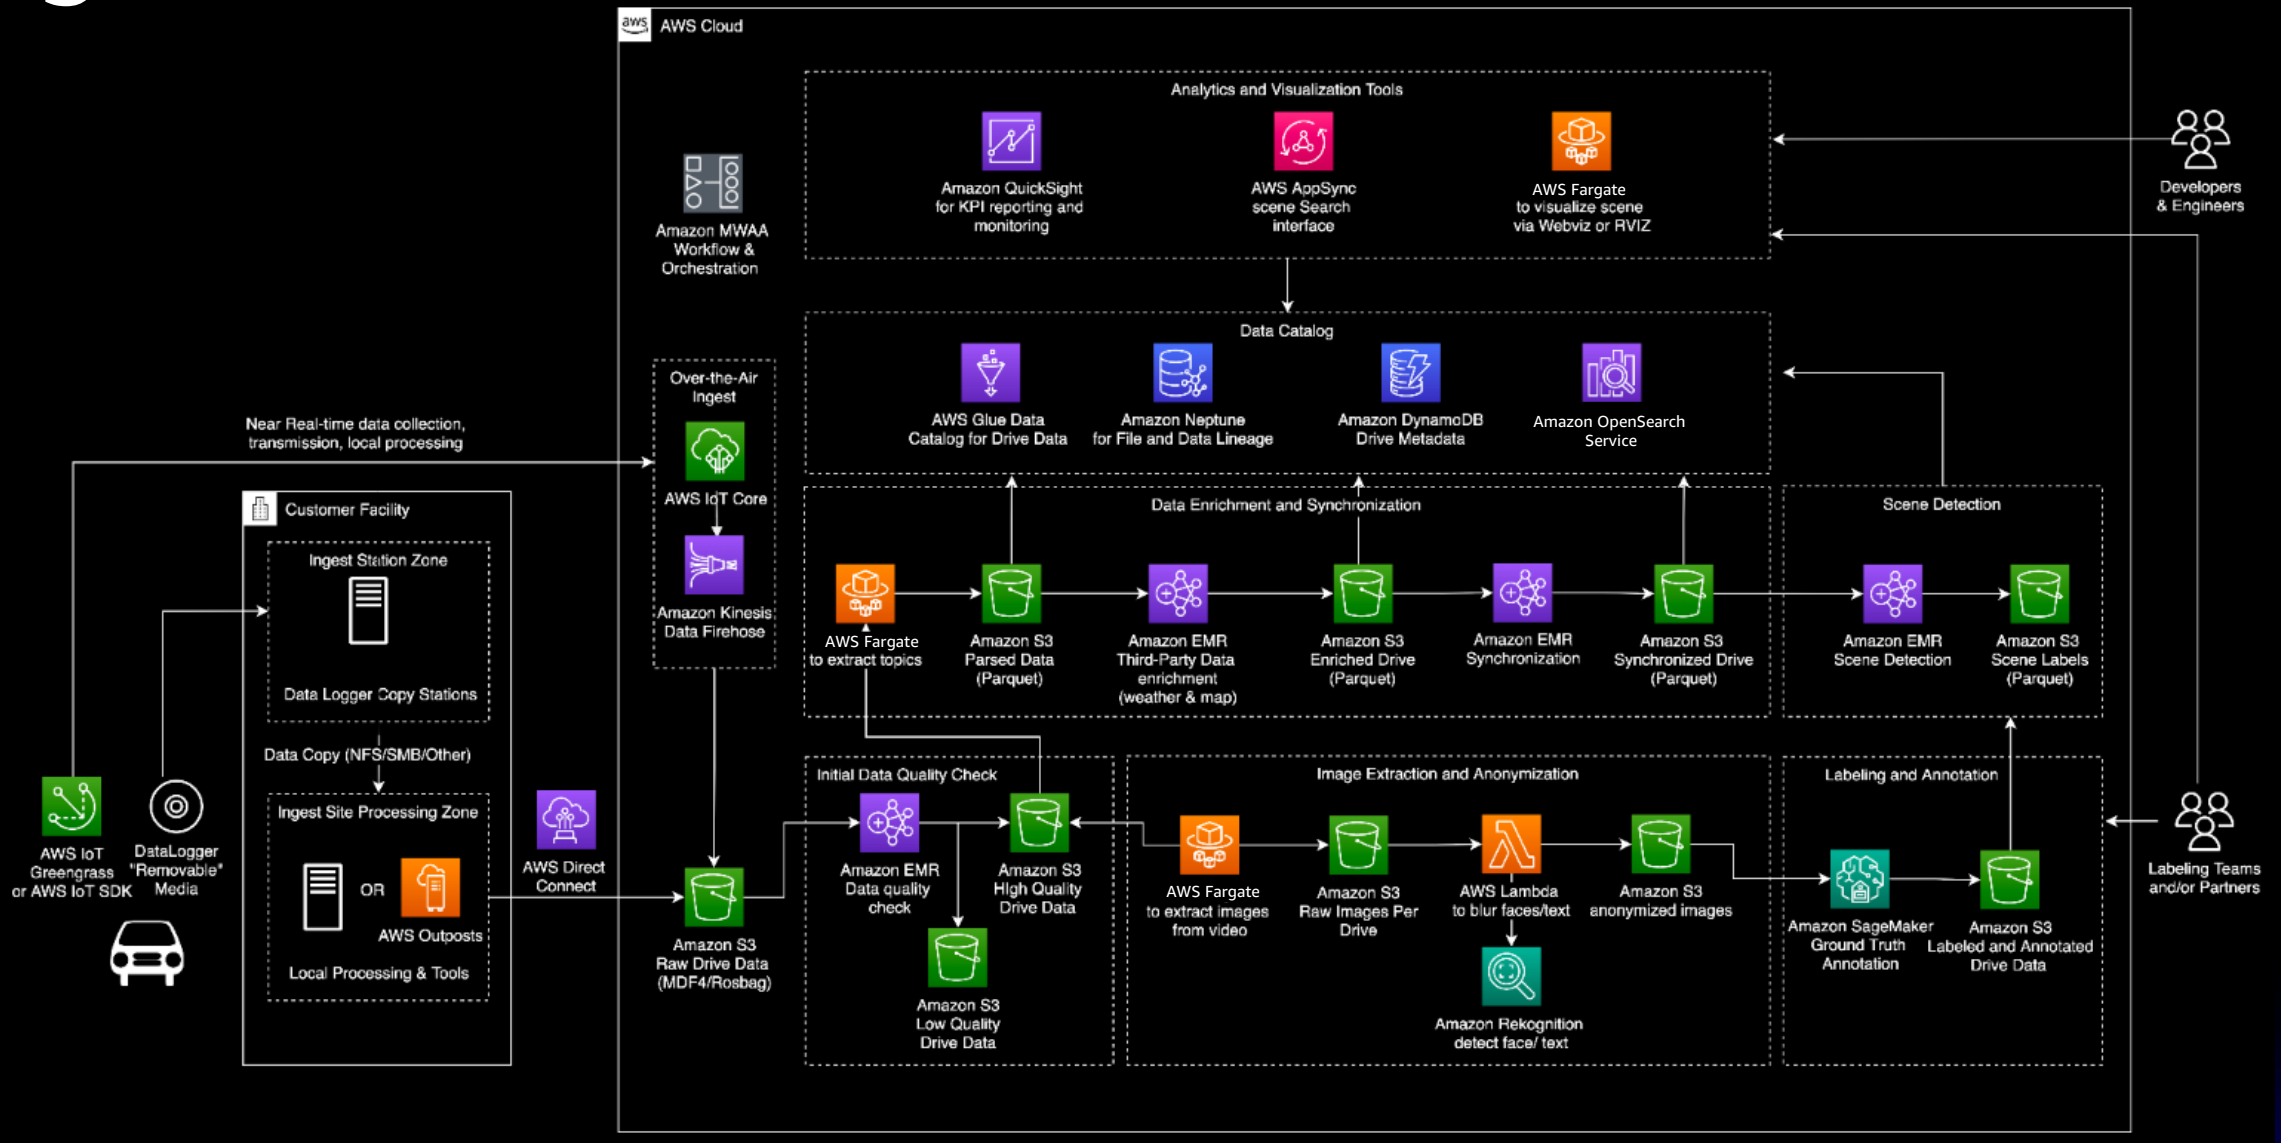


Driver

aws

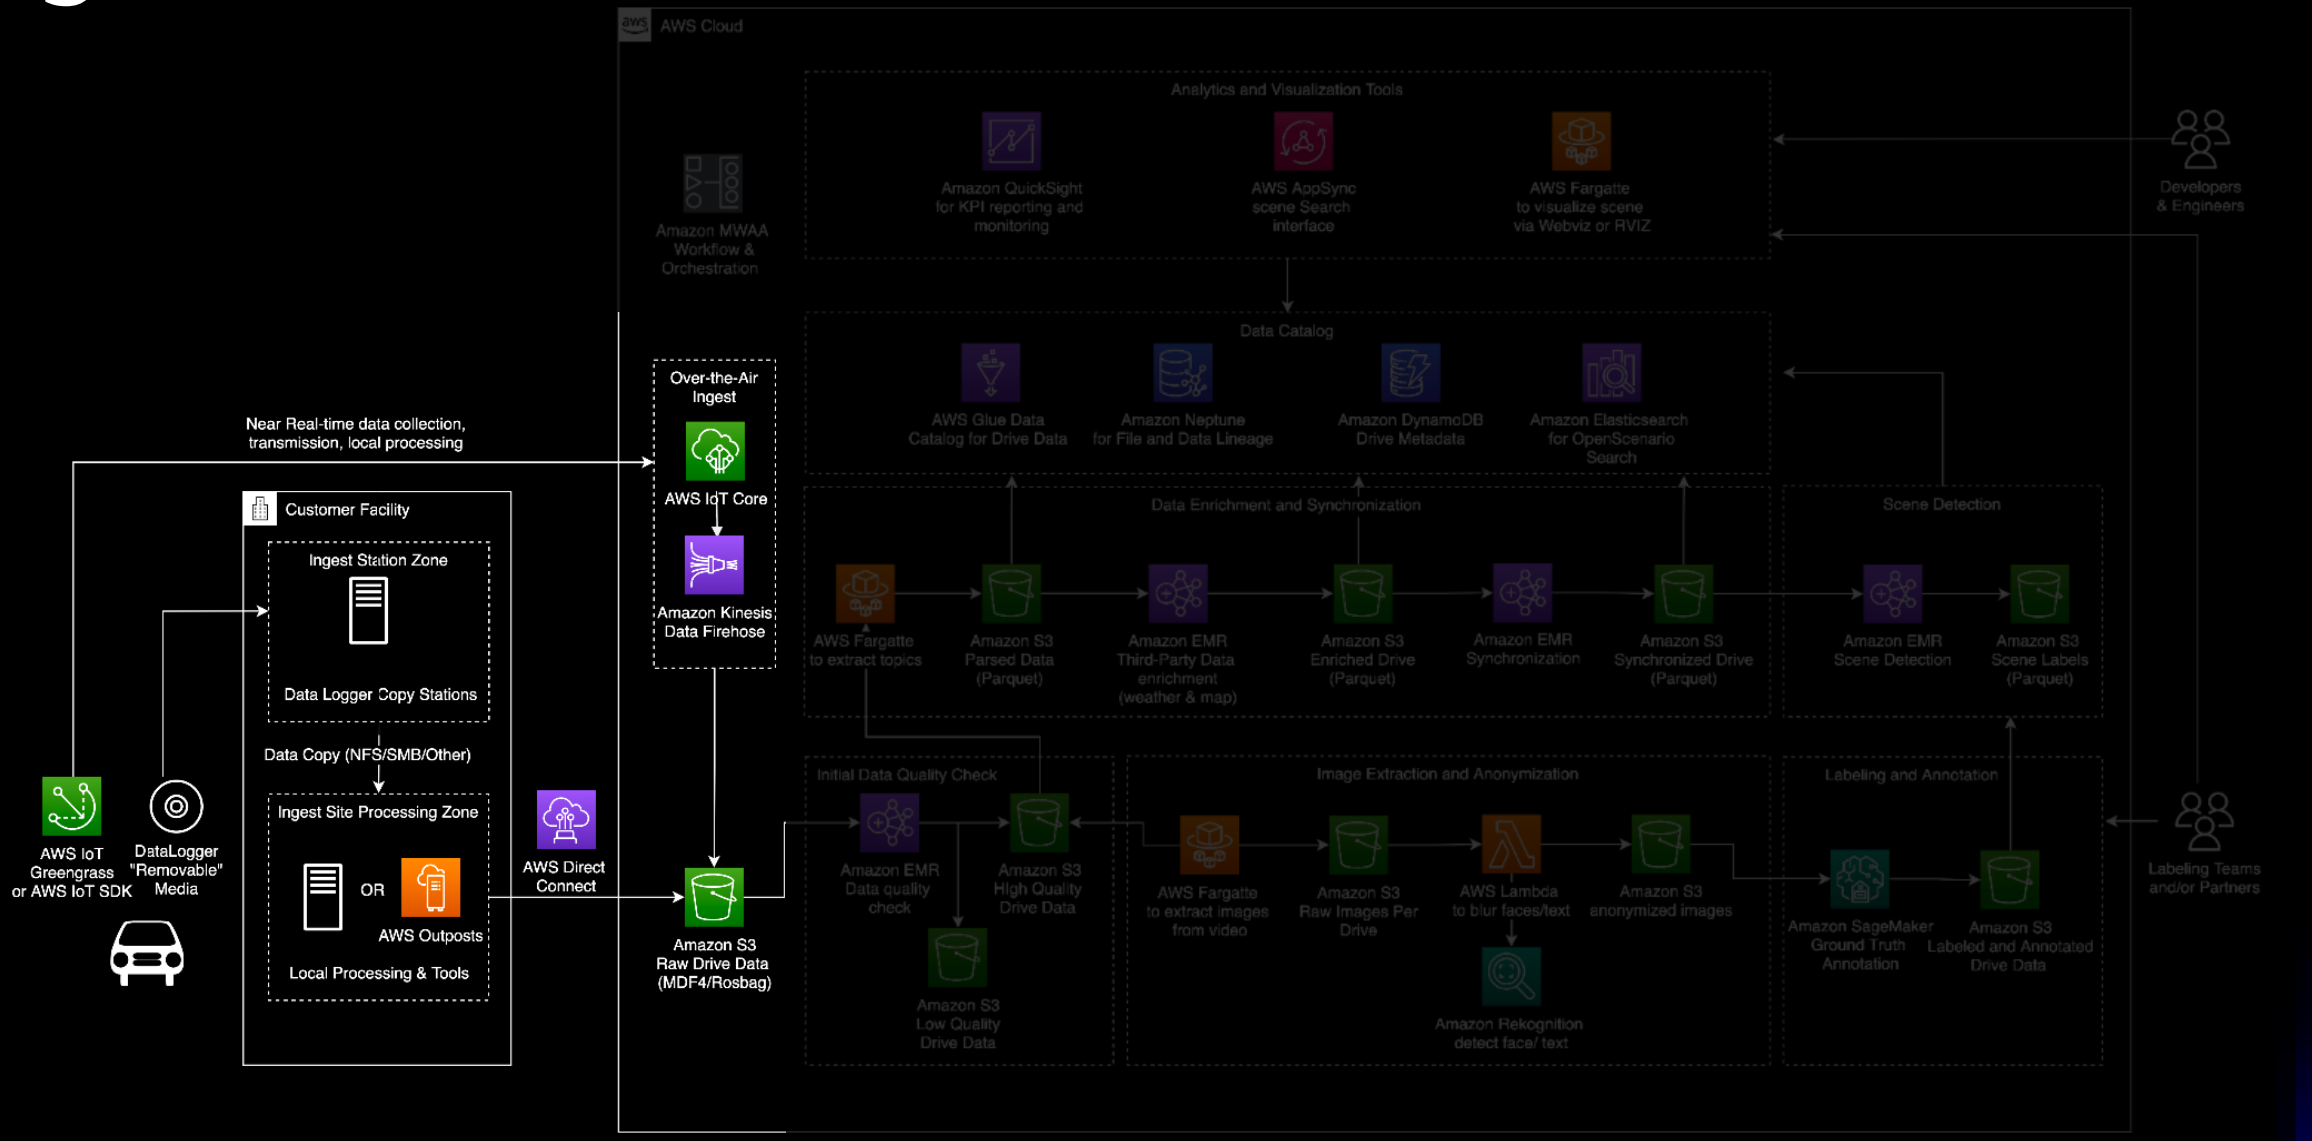
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ingestion and data lake



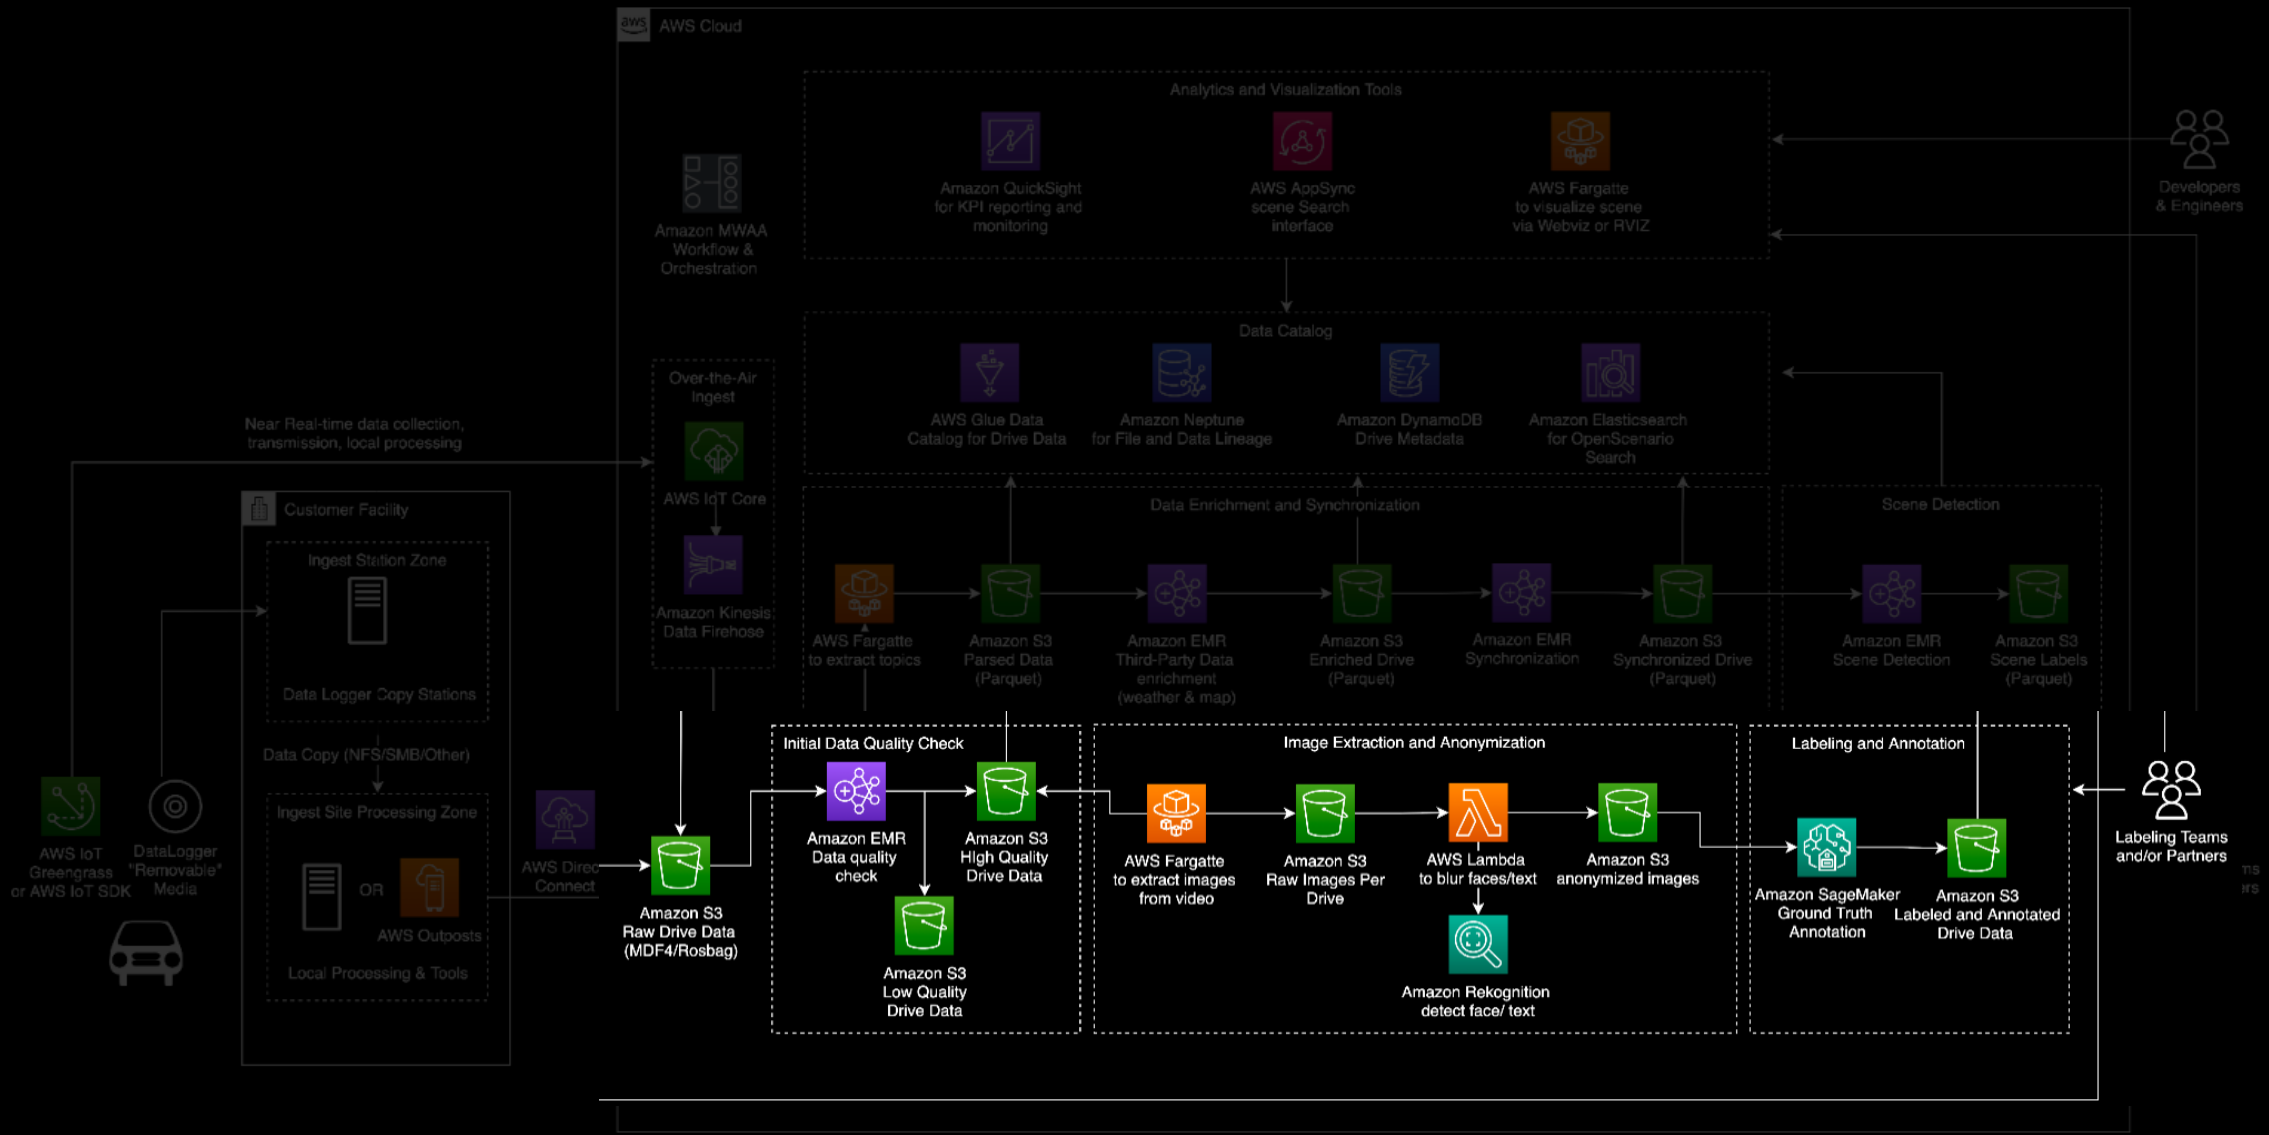
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ingestion and data lake



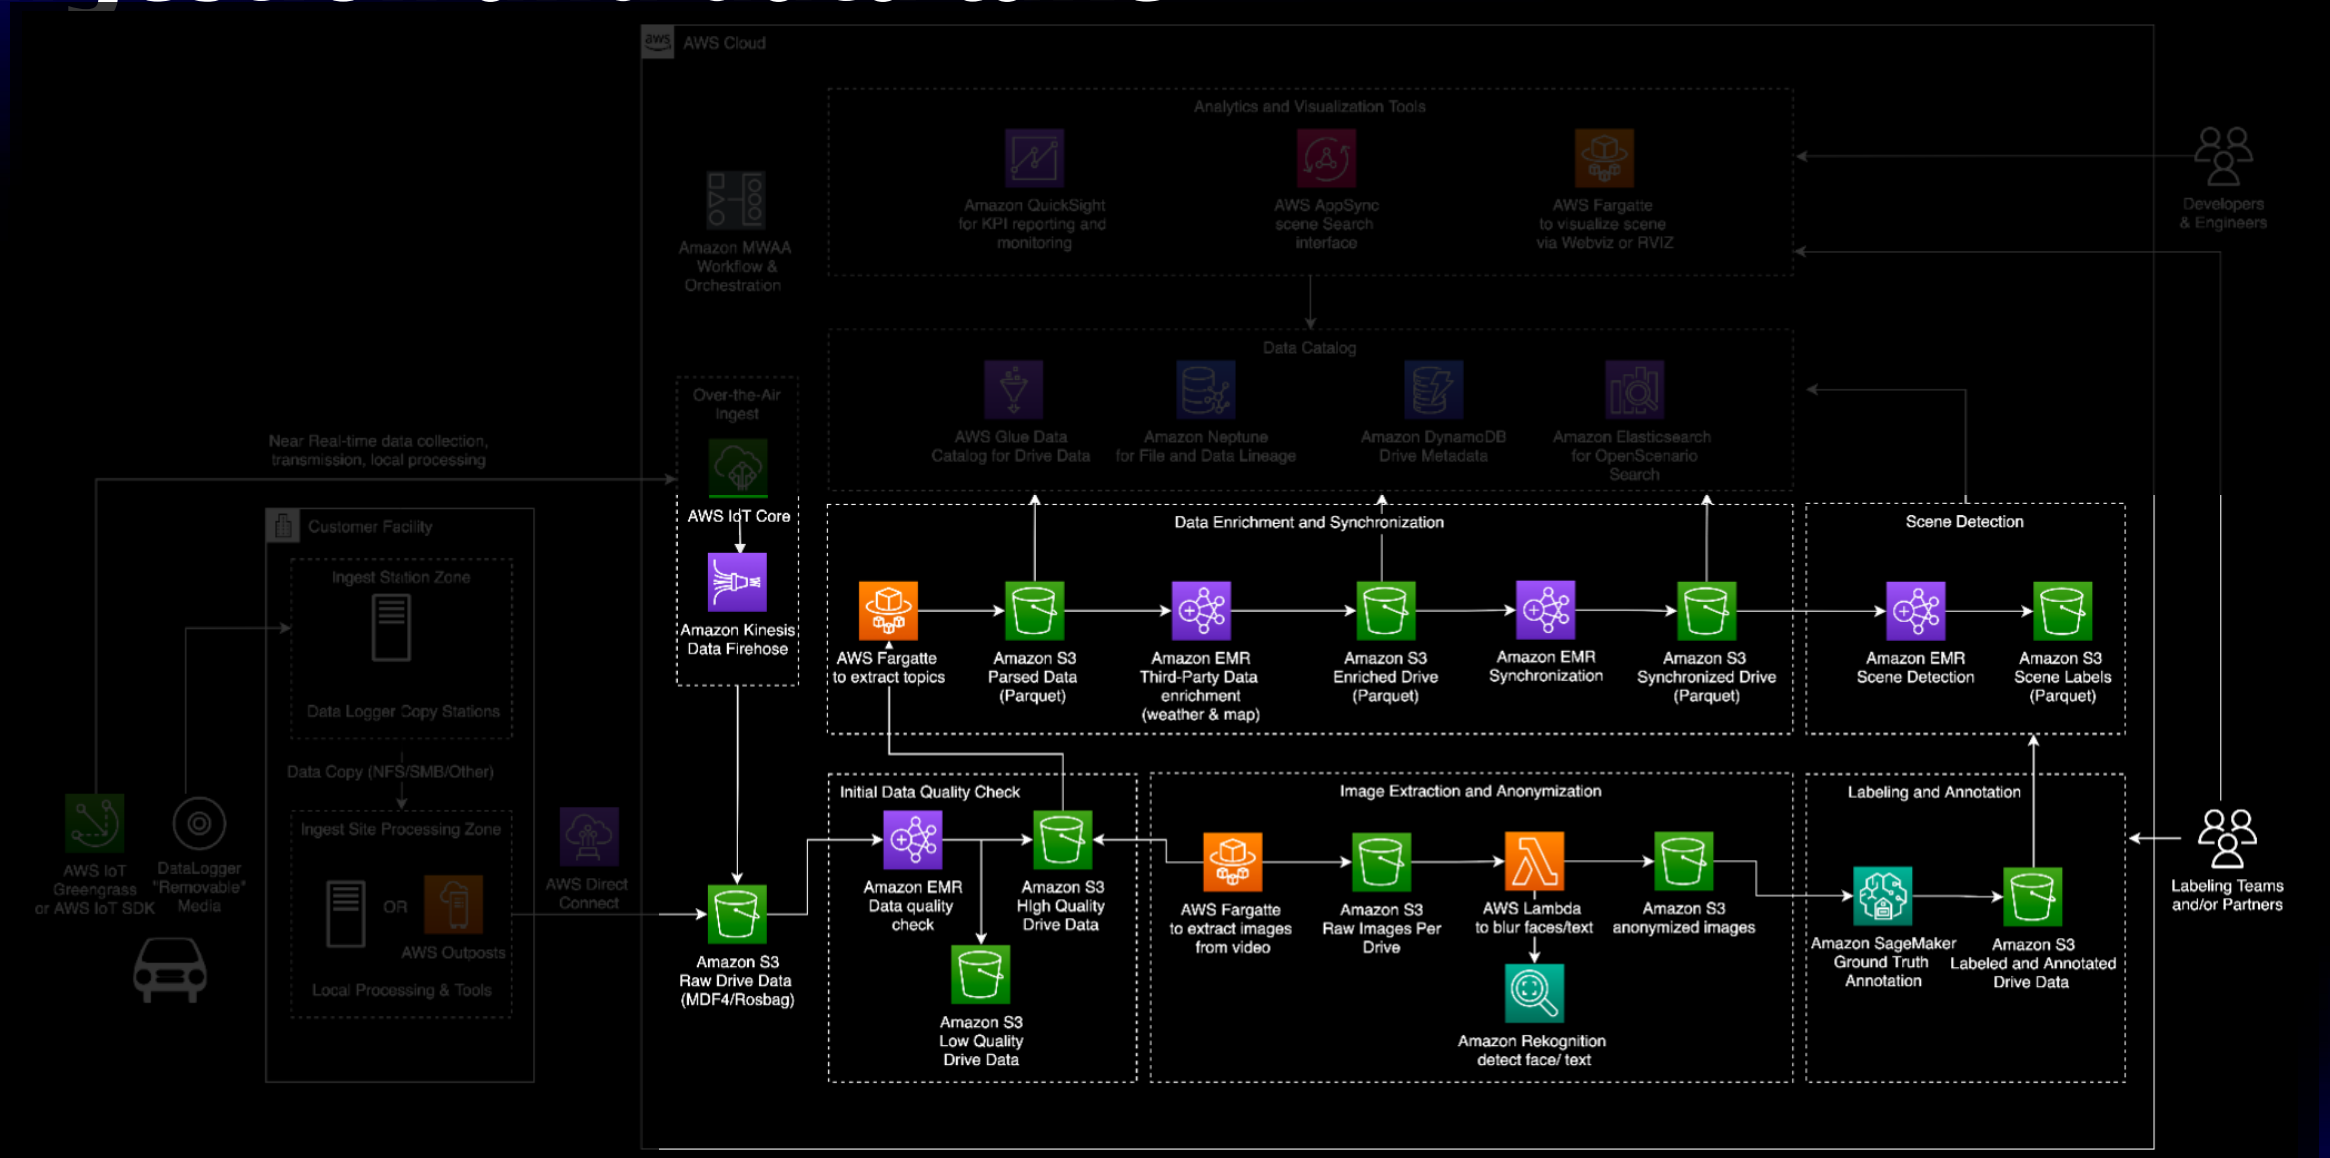
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ingestion and data lake



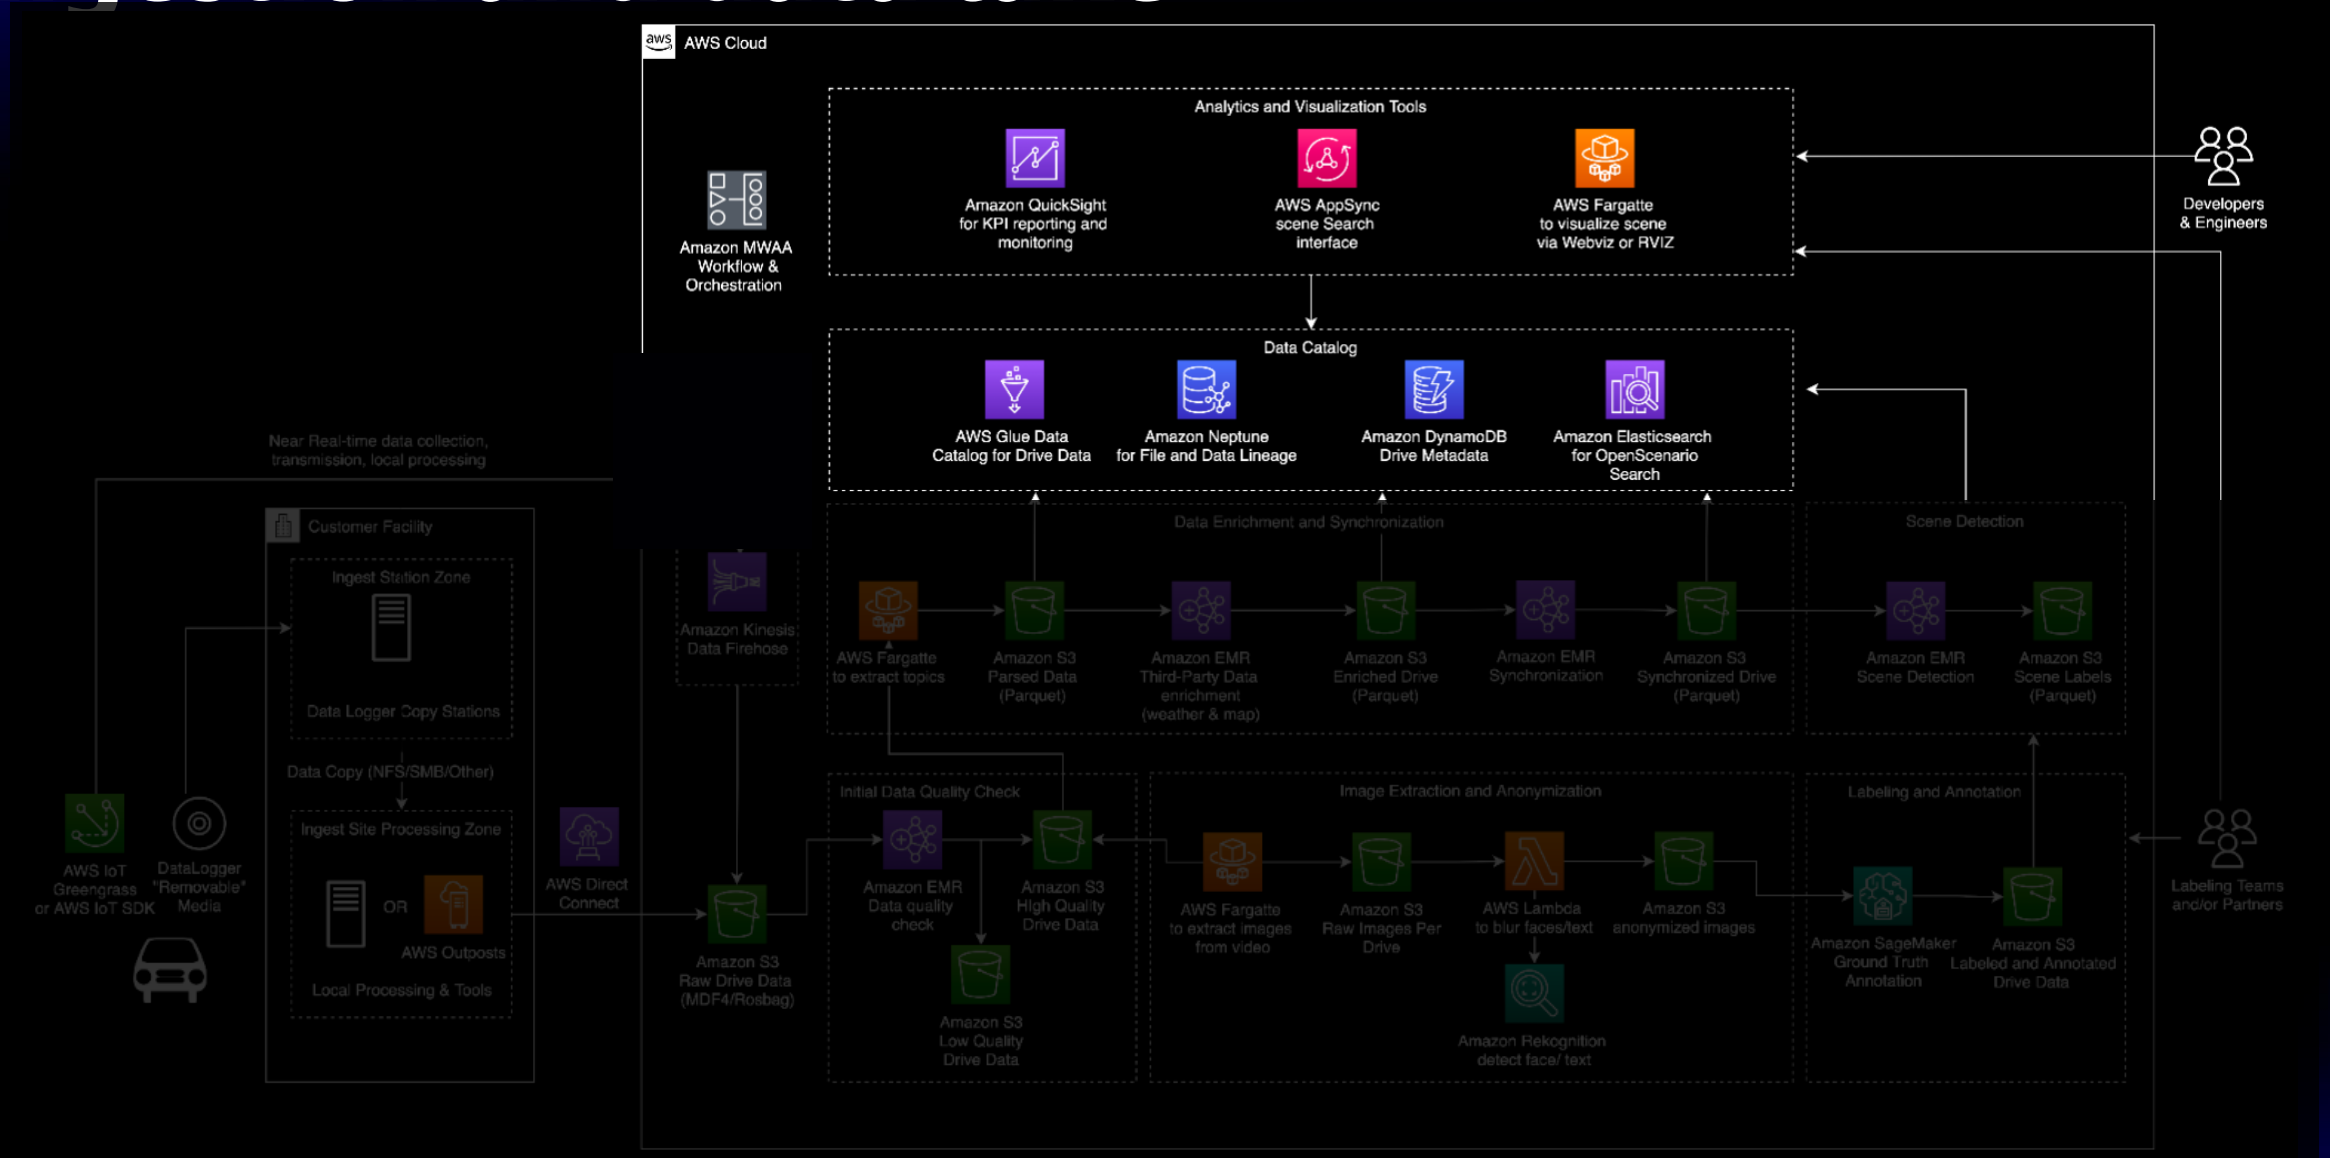
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ingestion and data lake



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Ingestion and data lake



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The AWS ML stack

BROADEST AND DEEPEST SET OF MACHINE LEARNING CAPABILITIES

AI SERVICES

VISION	SPEECH		TEXT			SEARCH	CHATBOTS	PERSONALIZATION	FORECASTING	FRAUD	DEVELOPMENT	CONTACT CENTERS
Amazon Rekognition	Amazon Polly	Amazon Transcribe	Amazon Comprehend	Amazon Translate	Amazon Textract	Amazon Kendra	Amazon Lex	Amazon Personalize	Amazon Forecast	Amazon Fraud Detector	Amazon CodeGuru	Contact Lens

ML SERVICES

Amazon SageMaker

Ground Truth

AWS Marketplace for ML

SageMaker Studio IDE

Built-in algorithms | Notebooks | Experiments | Processing | Model training & tuning | Debugger | Autopilot | Model hosting | Model Monitor

Neo

Augmented AI

ML FRAMEWORKS & INFRASTRUCTURE

TensorFlow | PyTorch | mxnet

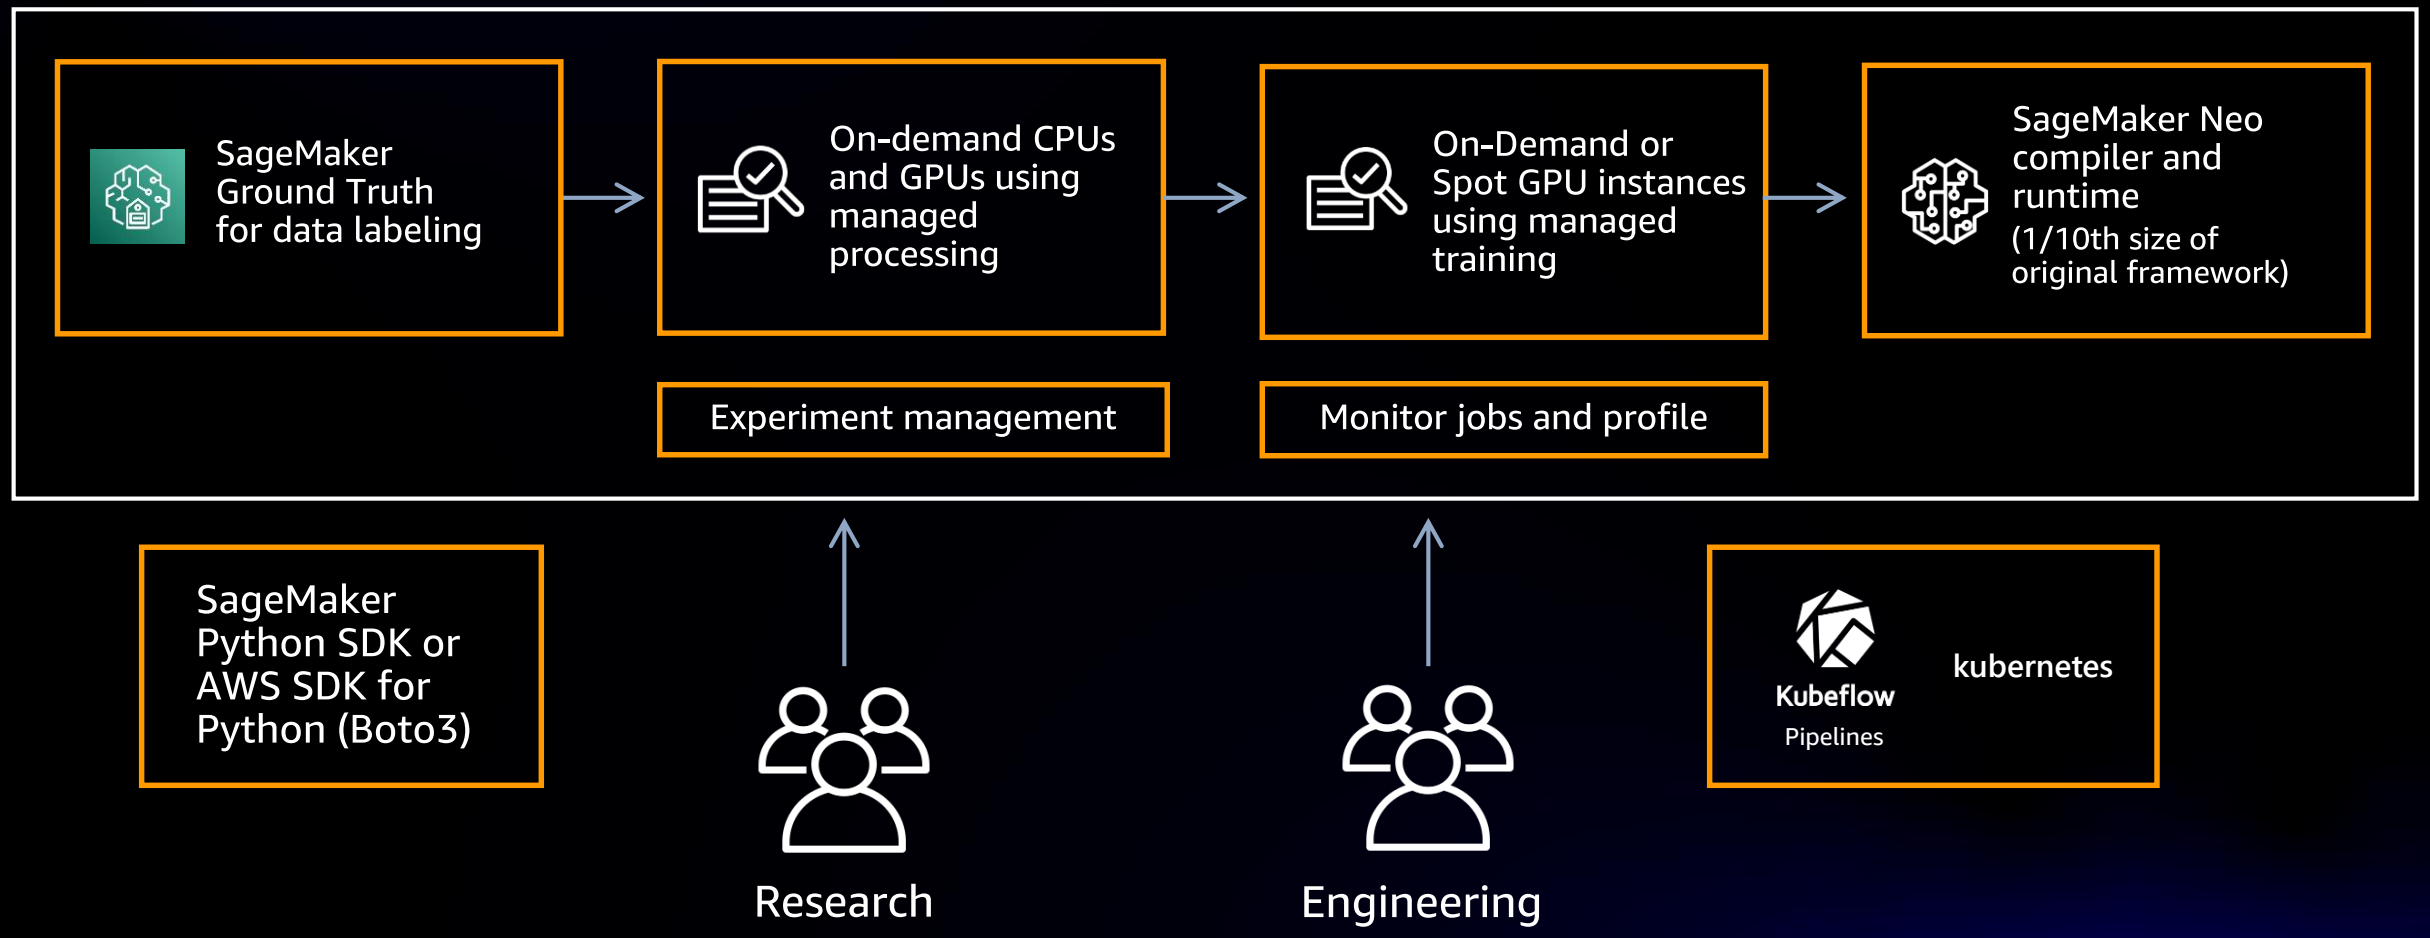
GLUON | HOROVOD | DeepGraphLibrary | Keras | scikit-learn | Deep Java Library

Deep learning AMIs & containers | GPUs & CPUs | Amazon Elastic Inference | Inferentia | FPGA

aws

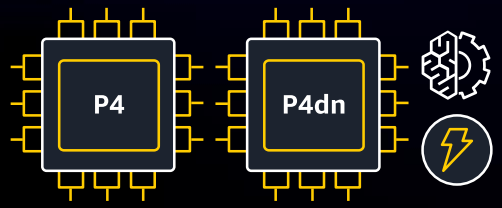
PyTorch, the PyTorch logo and any related marks are trademarks of Facebook, Inc.
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Using SageMaker to manage ML development for AV



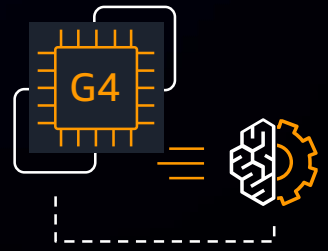
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon EC2 accelerated computing instances



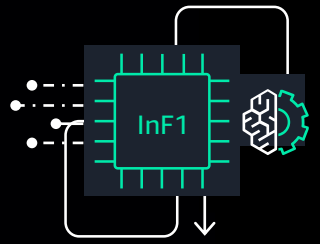
P4/P3/P3dn GPU compute instance

- Up to 1 petaFLOP of compute with 8x NVIDIA V100 GPUs
- Up to 256 GB of GPU memory and up to 100 Gbps of networking
- Designed for HPC and to handle large distributed machine learning training jobs



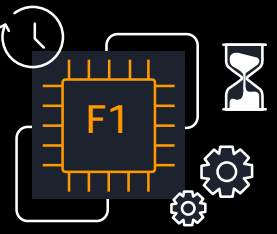
G4 GPU compute instance

- Up to 520 teraFLOPs of compute with 8x NVIDIA T4 GPUs
- Up to 1.8 TB of local NVMe storage and up to 100 Gbps of networking throughput
- Designed for cost-effective machine learning inference and graphics-intensive applications



Inf1 Inferentia instance

- Up to 2000 TOPs with 16x AWS-designed Inferentia accelerators
- Featuring AWS Inferentia, the first custom ML chip designed by AWS
- Designed for high-throughput and low-latency machine learning inference



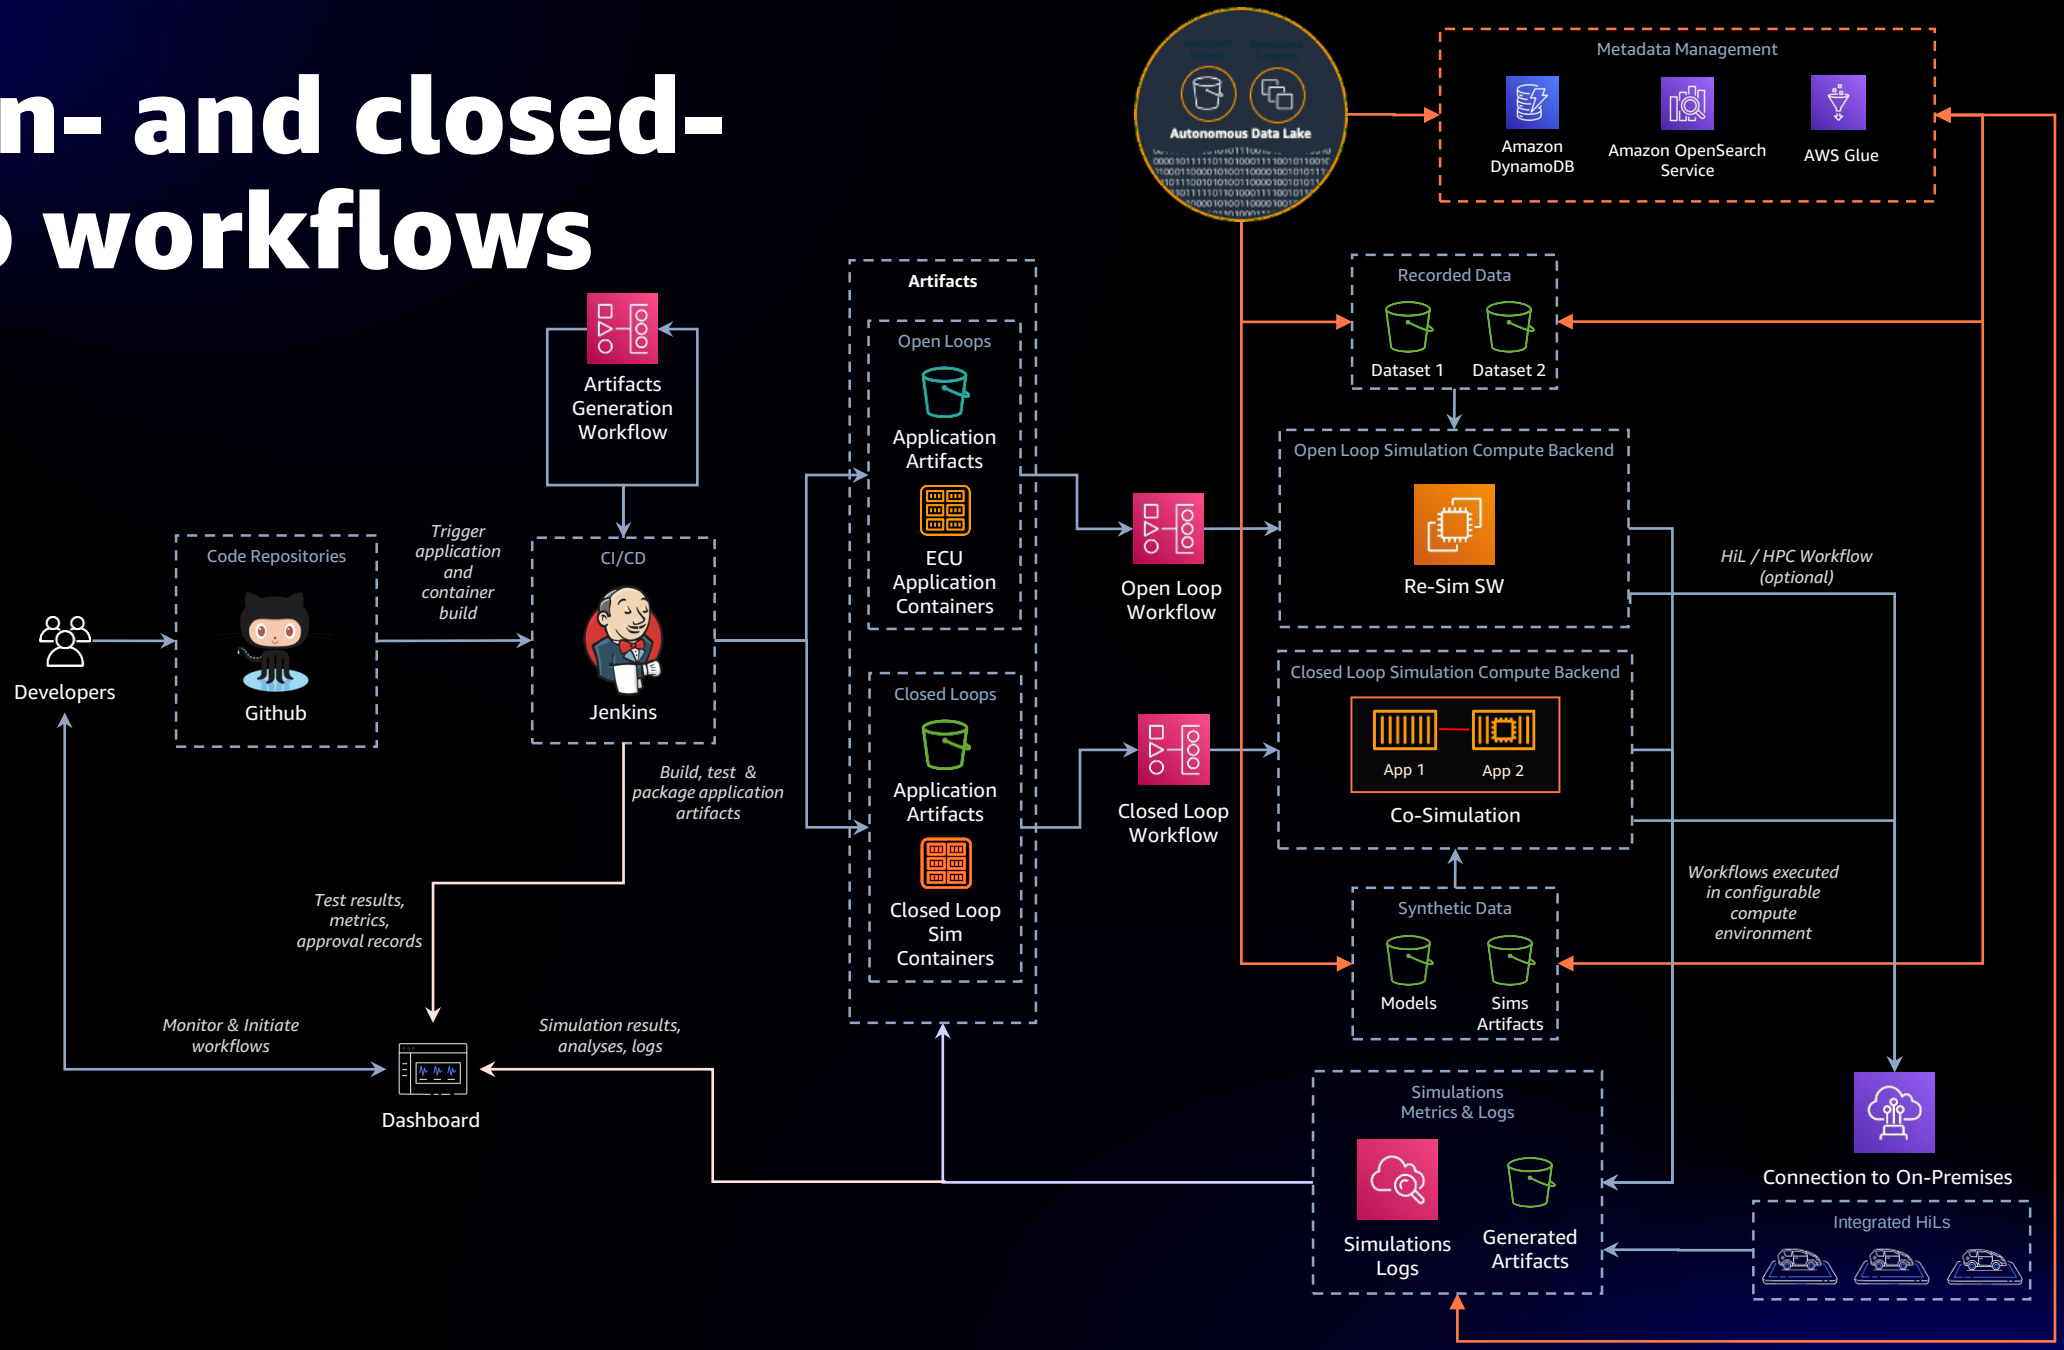
F1 FPGA instance

- Up to 100x acceleration compared to CPUs for a diverse set of compute-bound applications
- Up to 8 Xilinx Virtex UltraScale+ VU9P FPGAs in a single instance
- Designed for hardware-accelerated applications including financial computing, genomics, accelerated search, and image processing

aws

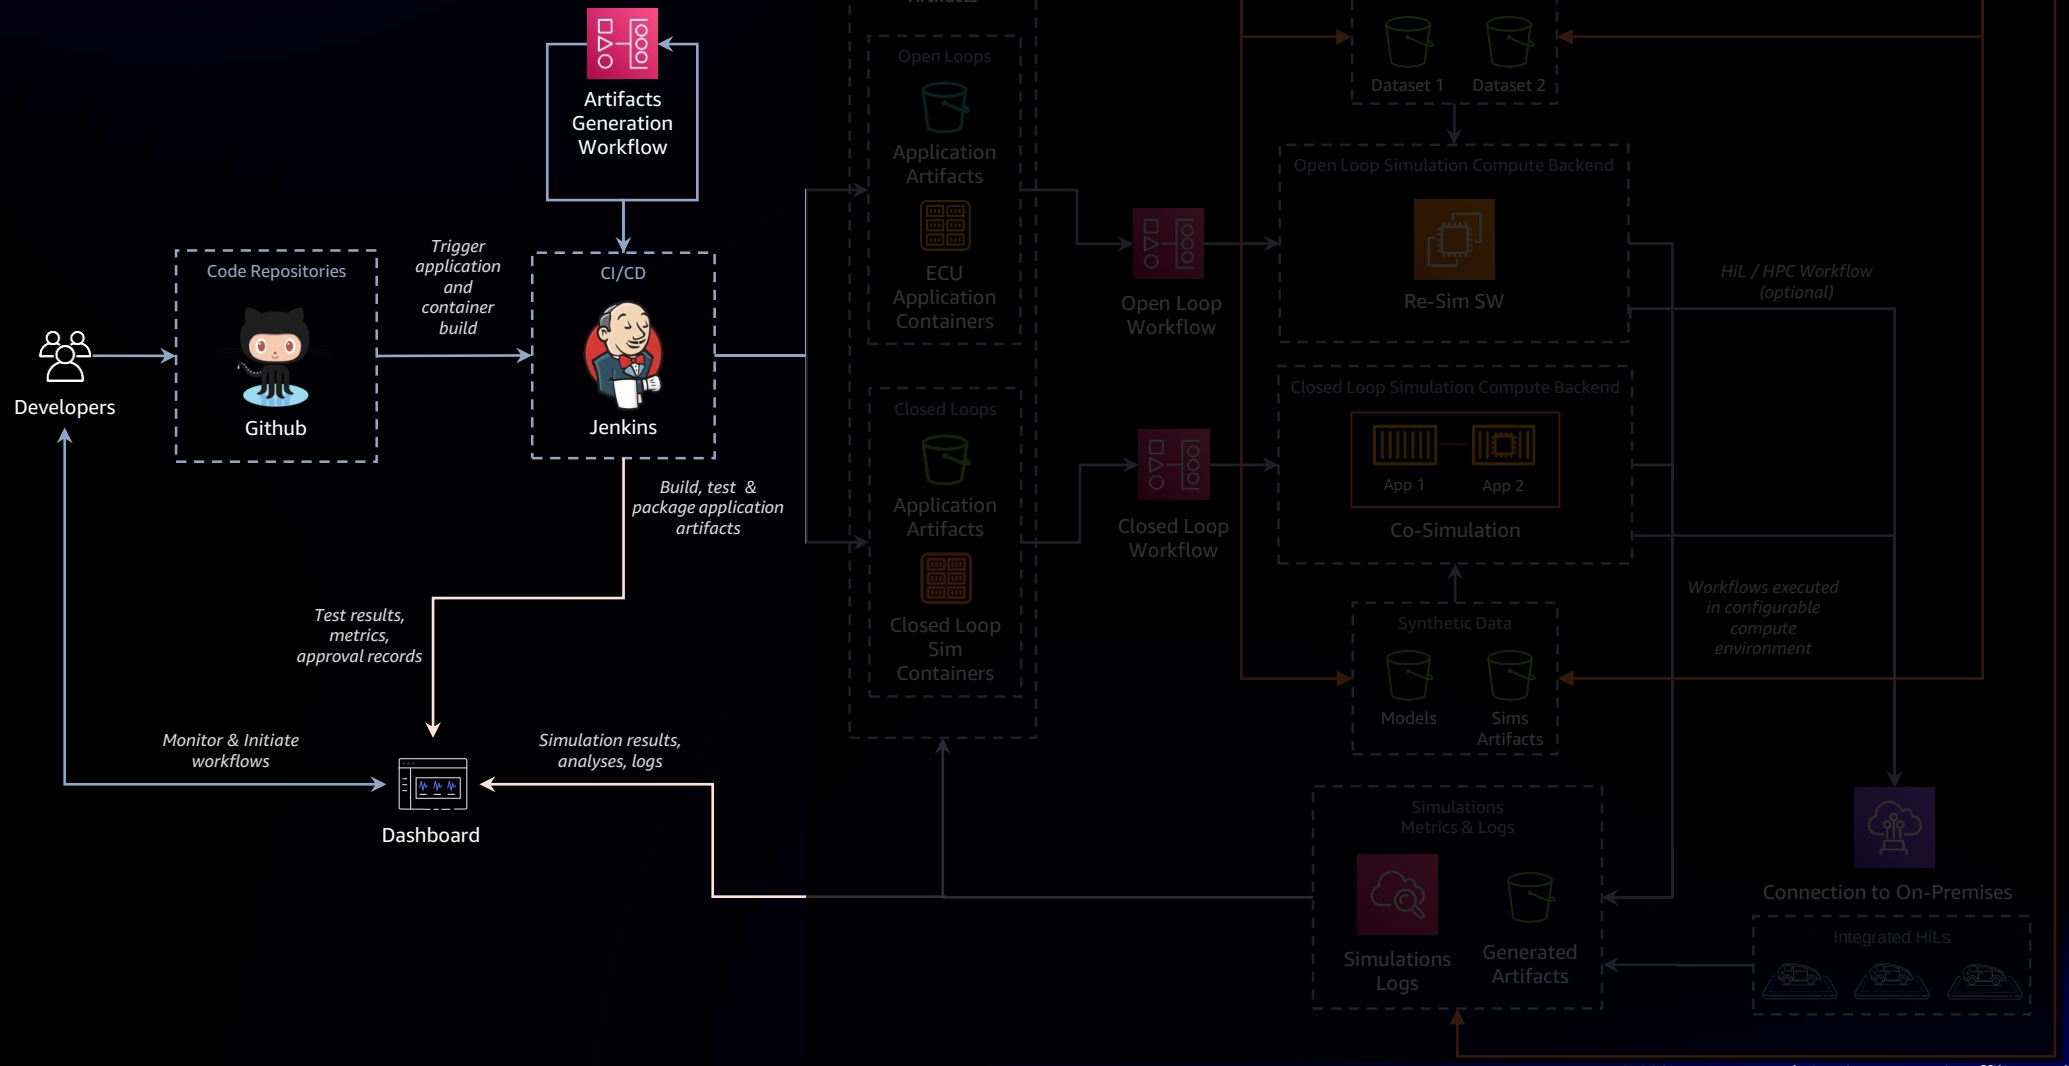
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Open- and closed-loop workflows



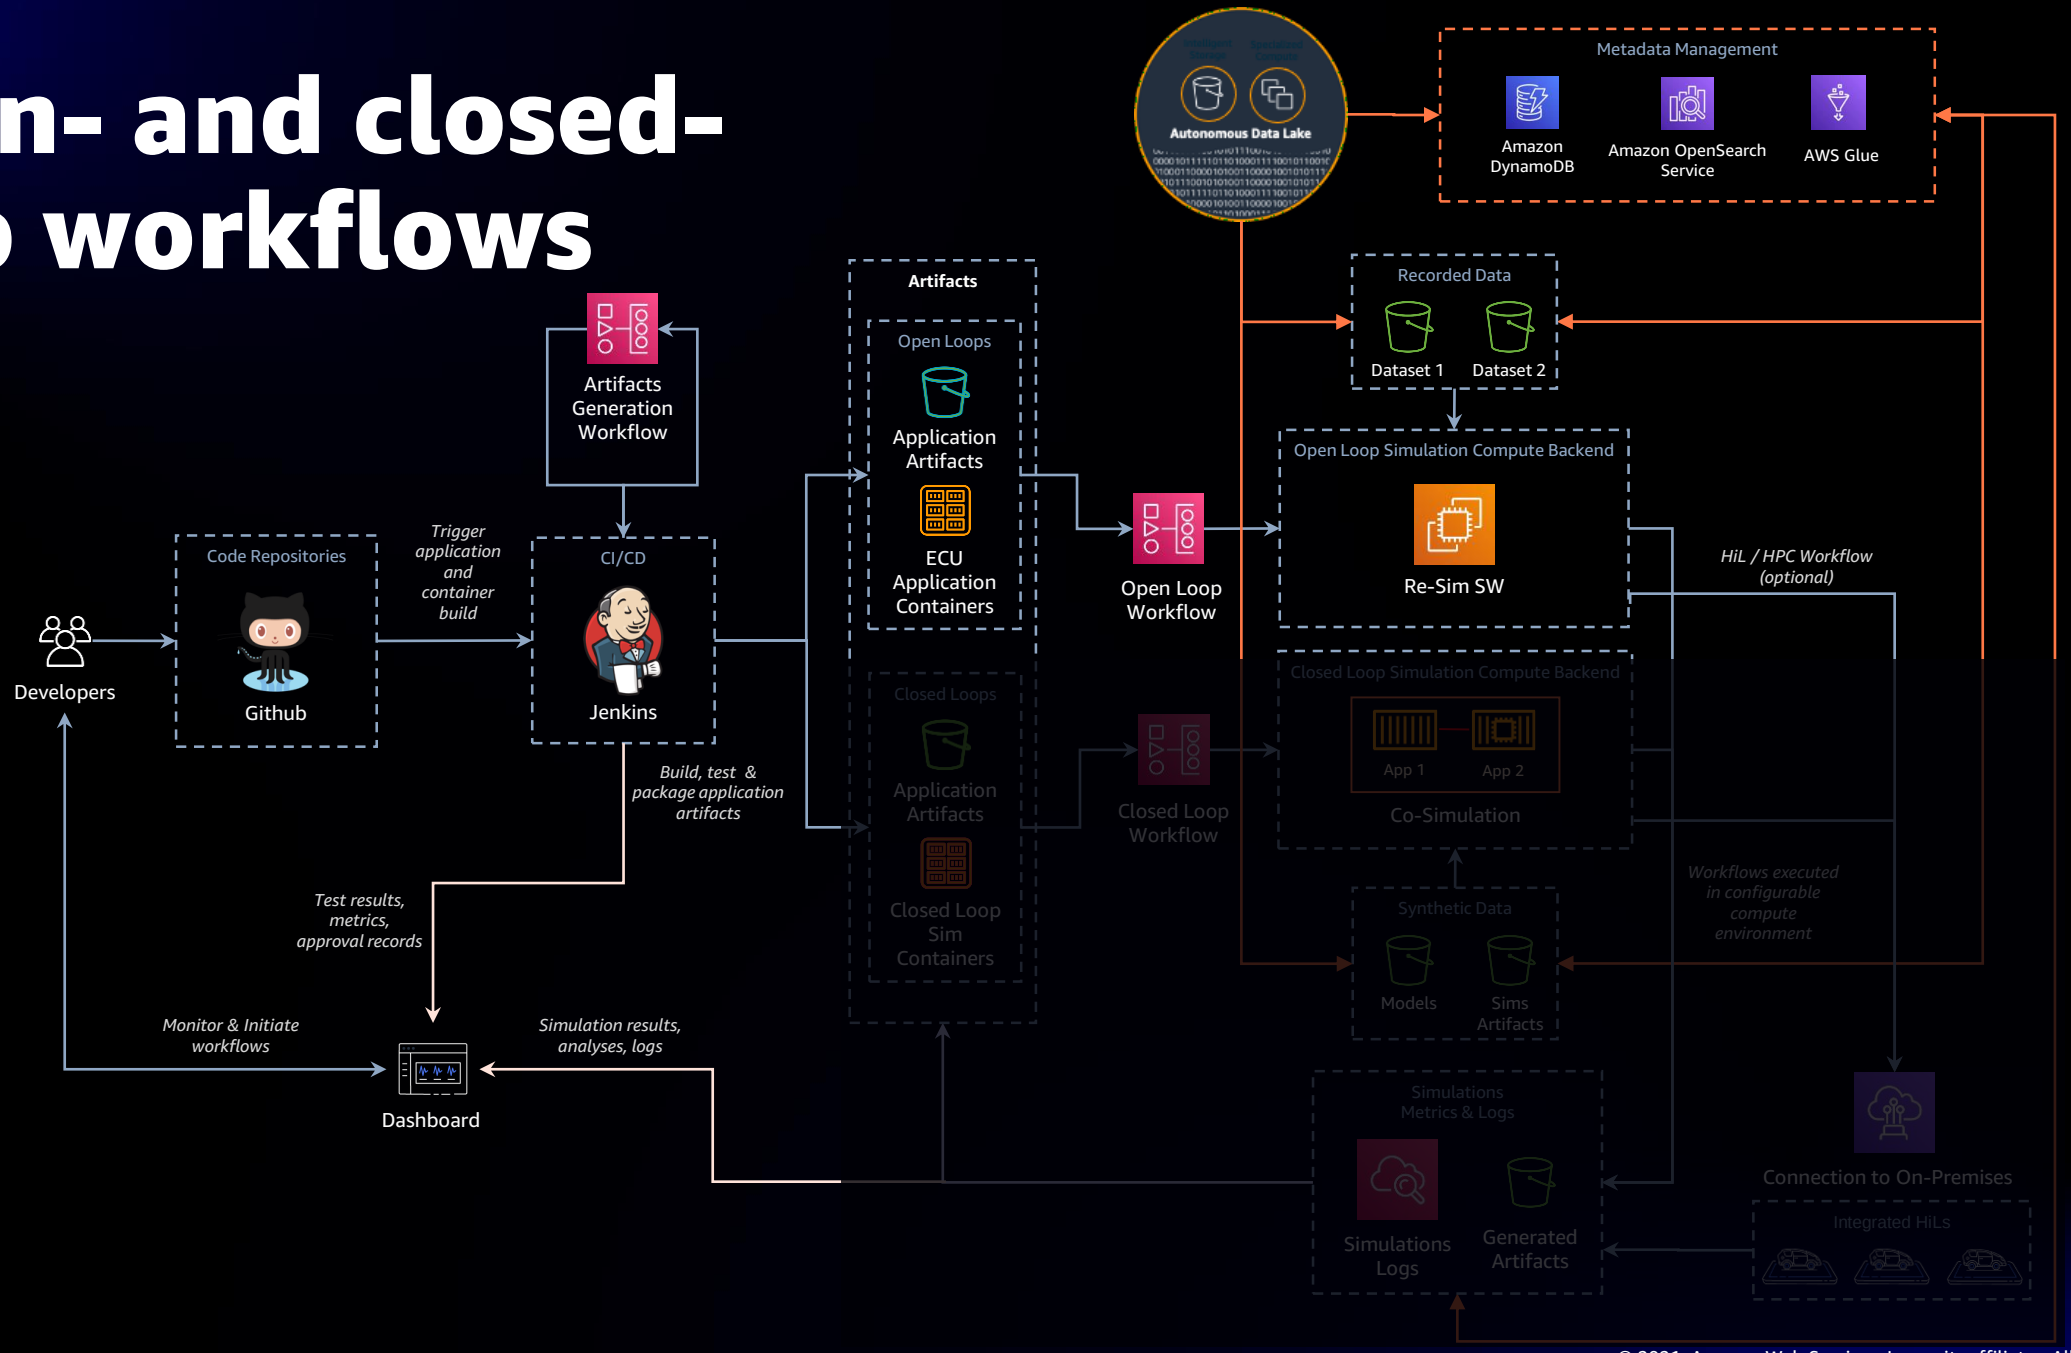
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Open- and closed-loop workflows



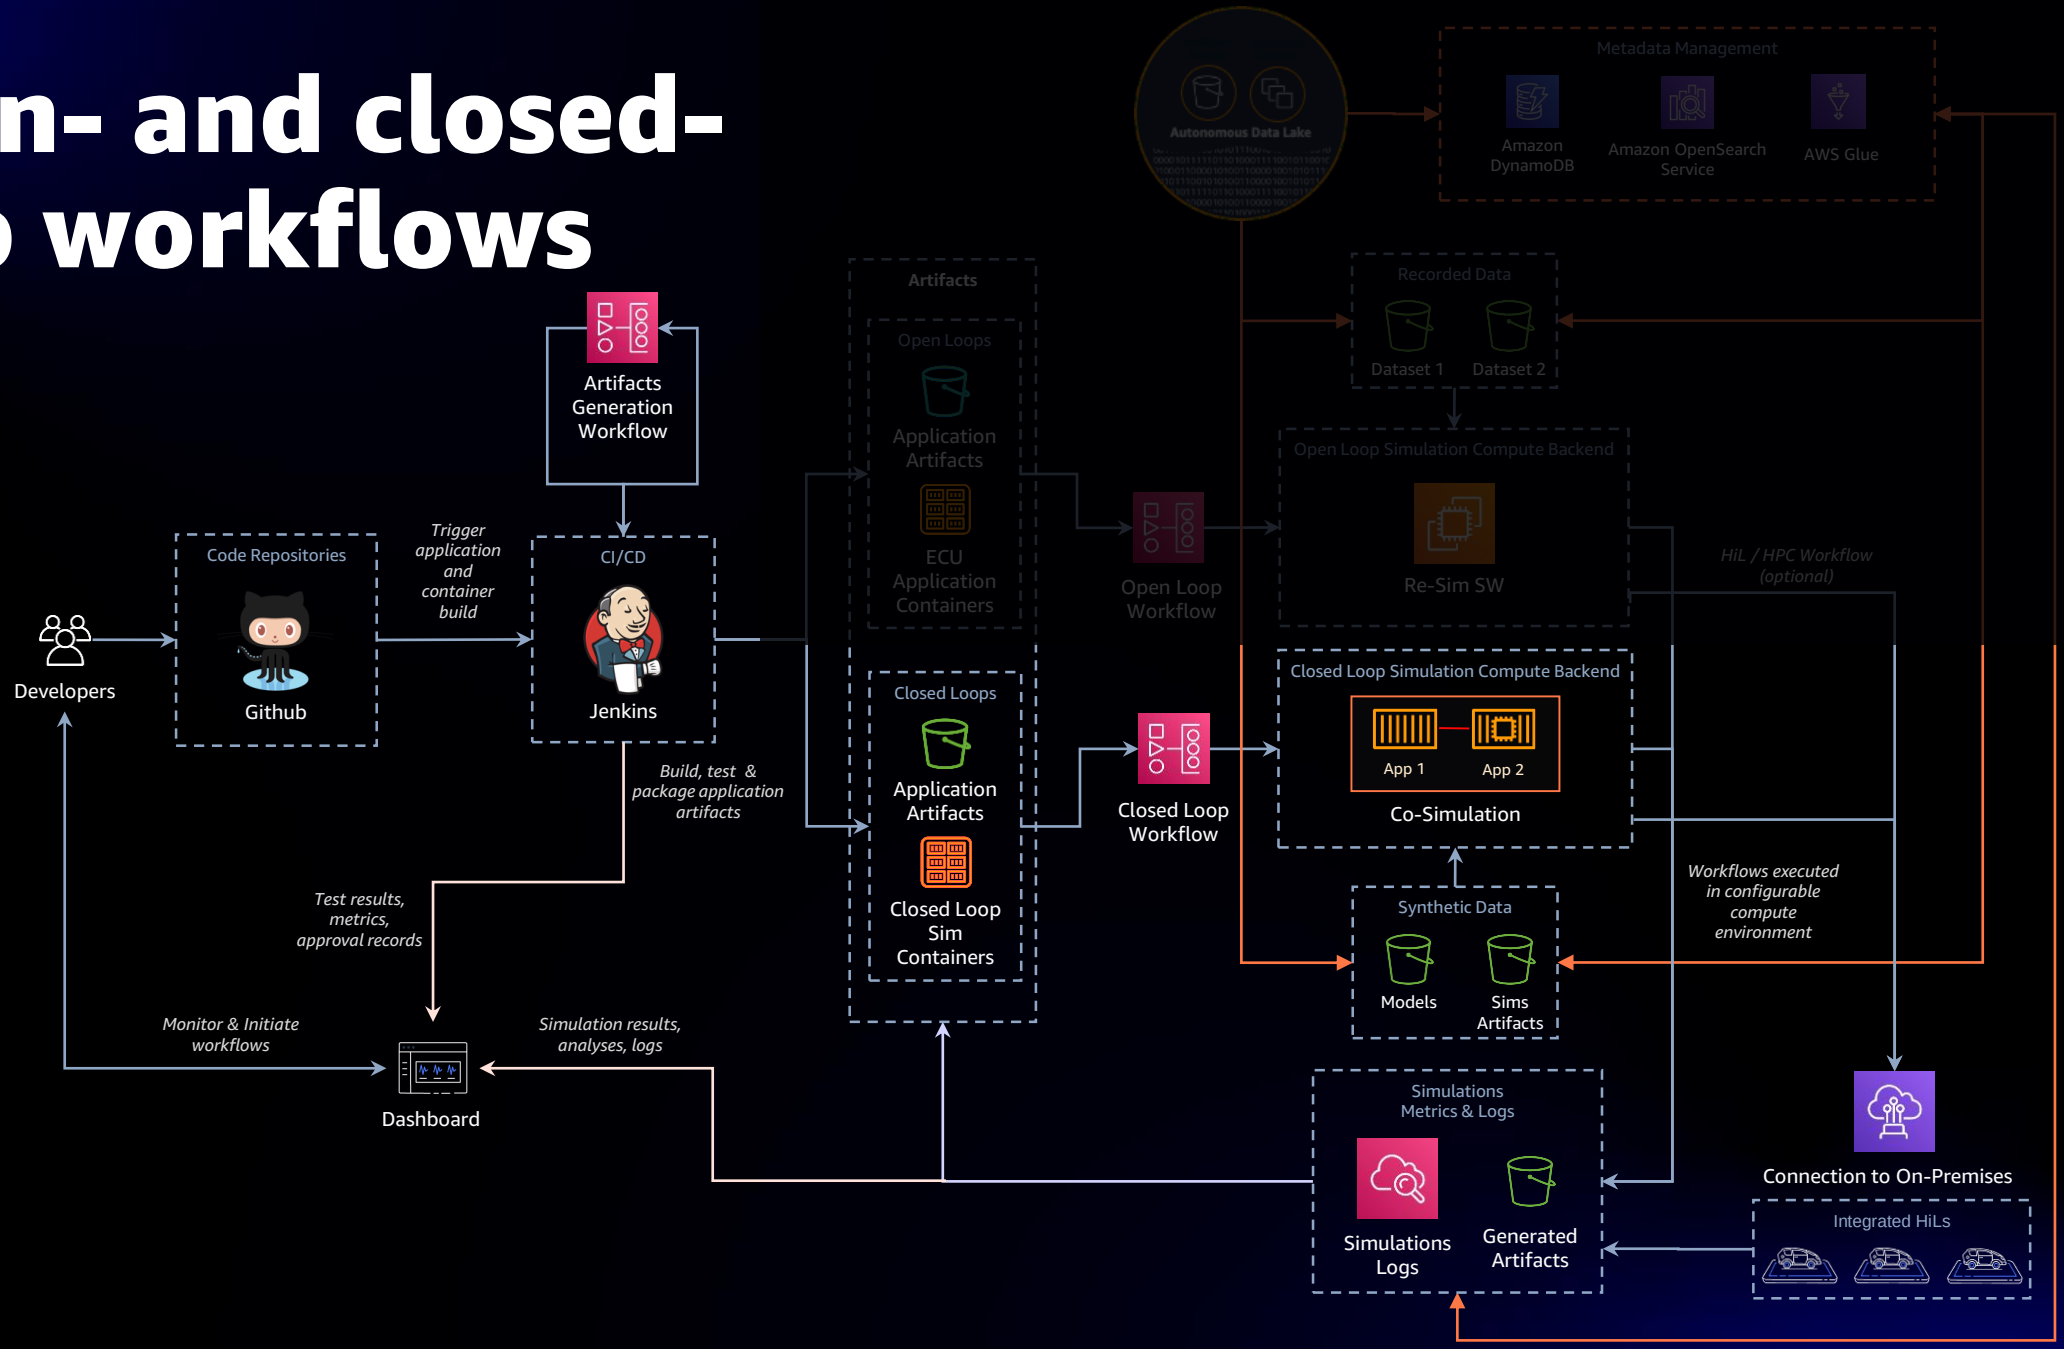
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Open- and closed-loop workflows



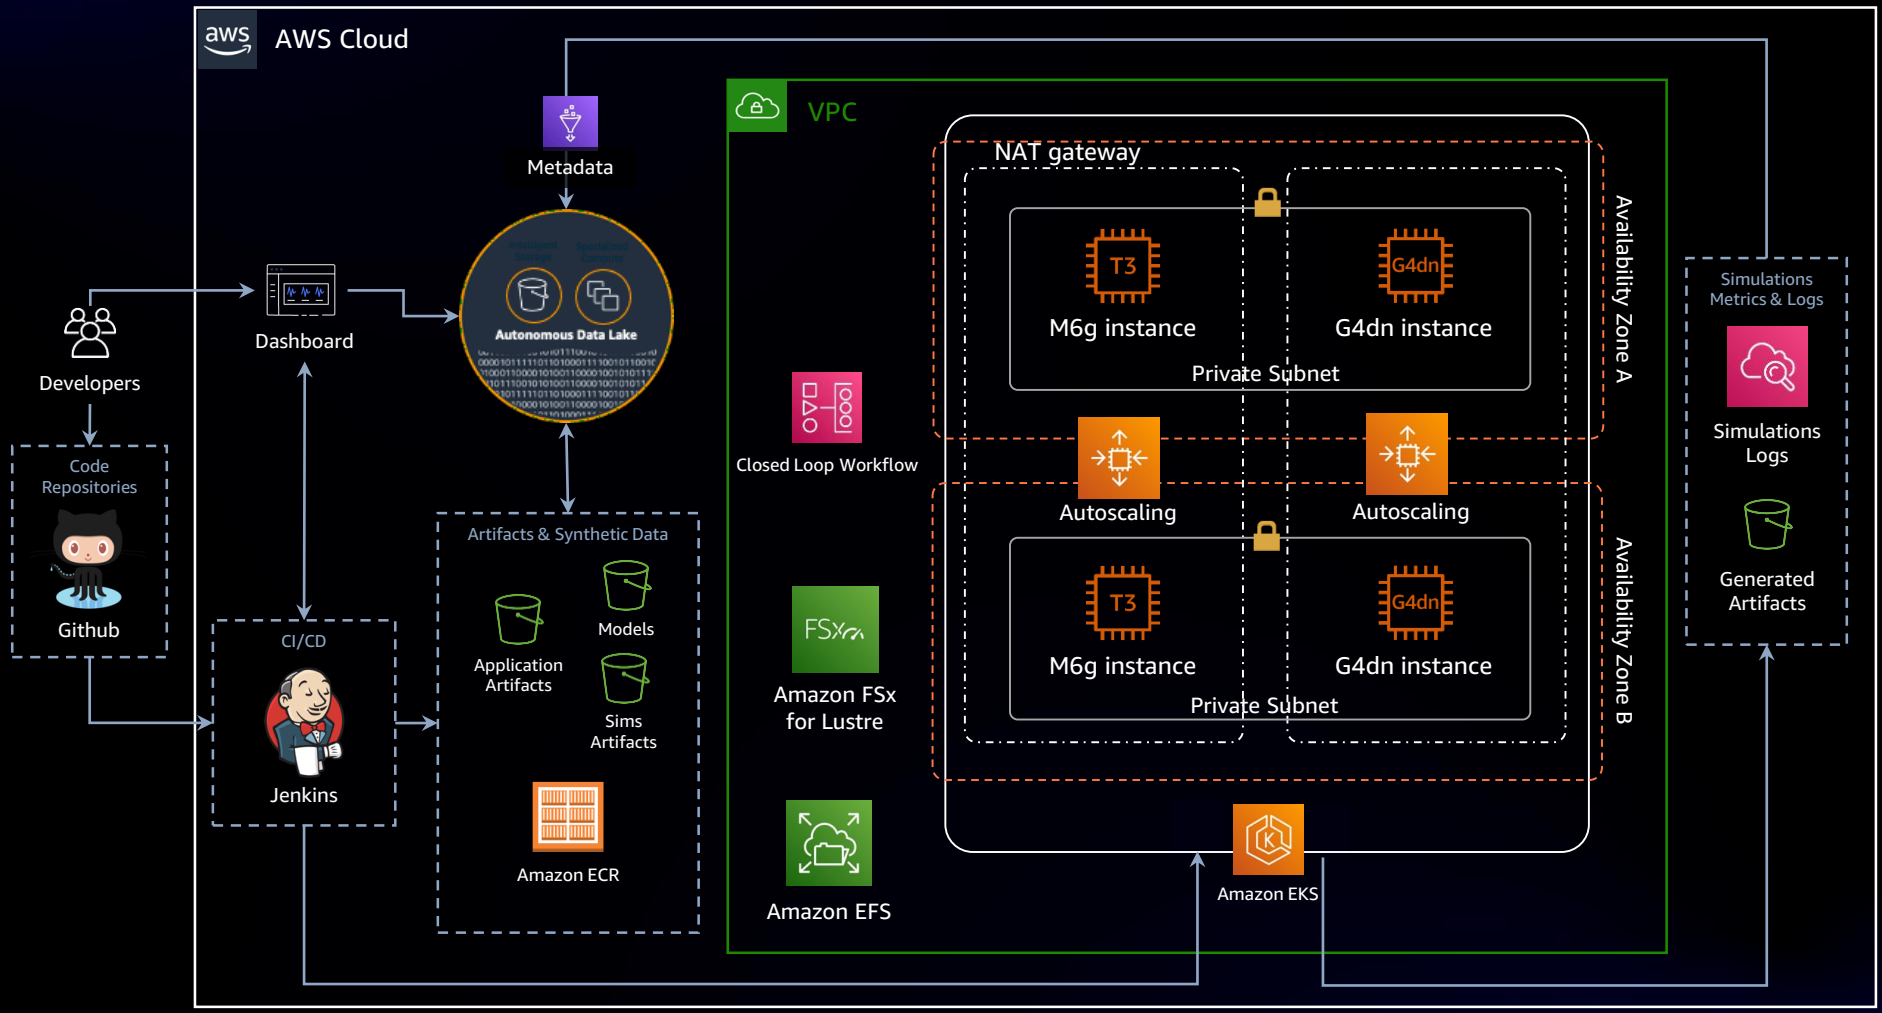
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Open- and closed-loop workflows



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Closed-loop co-simulation with proper instances



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AV Partner ecosystem

Data management, processing & analytics

wejo
Elektrobit
otonomo
heex
siasearch
RIDECELL

Labeling

scale
dSPACE
iMerit

Model & algorithm development

MathWorks

Verification & validation

SYNOPSYS
Elektrobit
ALTAIR
dSPACE
foretellix
d(risk)
VIRES

HD map development

here
MOBILEYE
mapbox
DEEPMAP
TOMTOM
Mapillary
Leica Geosystems

Simulation

rFpro
aiMotive
OPAL-RT TECHNOLOGIES
FIVE
parallel domain
AVSIMULATION
cognata

In-vehicle software

TTTech
QNX
rti

DevOps, MLOps orchestration

Akridata
tracetronic
GUIDING YOUR VISION

Data ingestion

dSPACE
b-plus
Elektrobit
edgeconnex
EQUINIX
Velodyne Lidar

Systems integrators

Luxoft
A DXC Technology Company
KPIT
tcs TATA CONSULTANCY SERVICES
bertrandt
accenture
Capgemini

NVIDIA
Applied Intuition
IPG
SVL SIMULATOR

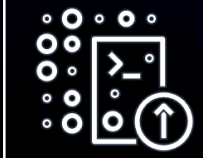
aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

How Aurora leverages AWS

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora scaling challenges



HYPERSCALE

Thousands of cores of compute for model development and validation. TBs of data to collect, ingest, and store every mile driven by our fleets translates into PB-scale data processing, storage, and management problems.



AGILITY & SPEED

We don’t want every ML engineer and roboticist to also be an expert in HPC, distributed systems, cloud, and Kubernetes



COST

PB-scale data storage and large-scale compute costs, significant capex of on-premises compute, requires significant human investment



COMPLEXITY

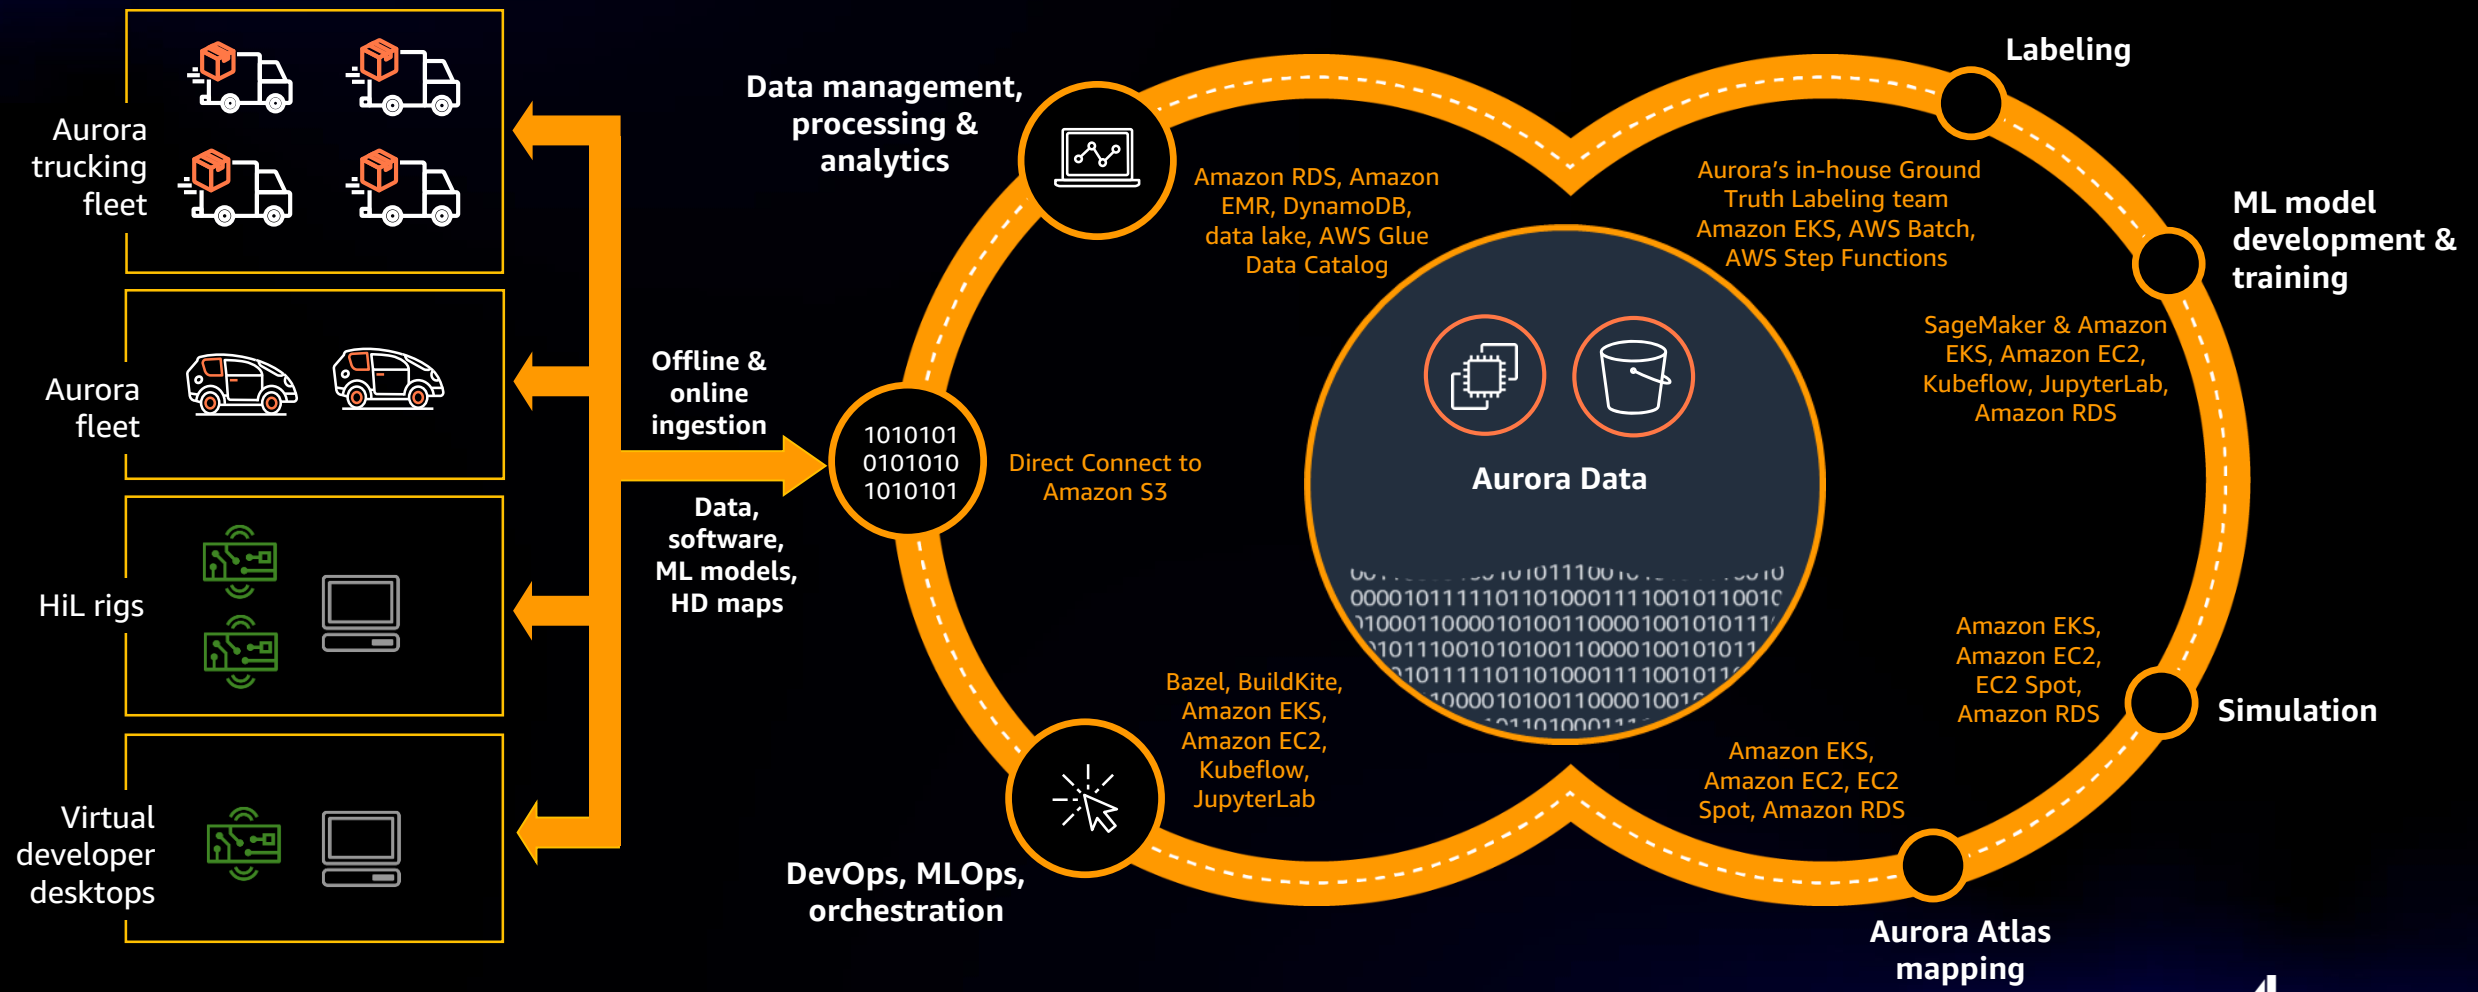
Different workloads across Aurora: simulation, perception scenarios, HD mapping, Ground Truth

Aurora

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

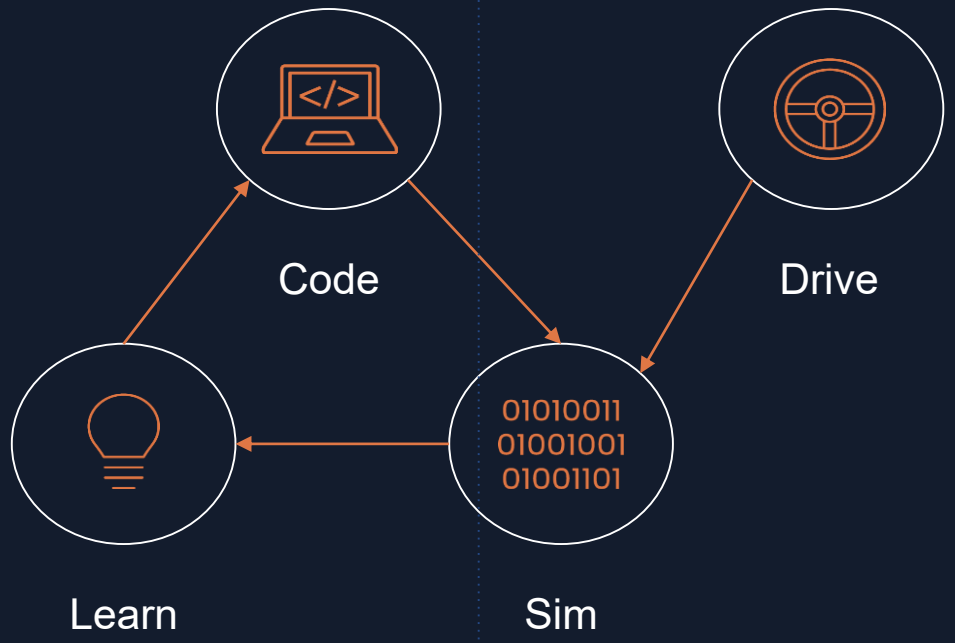
Aurora AWS development platform



aws

Aurora

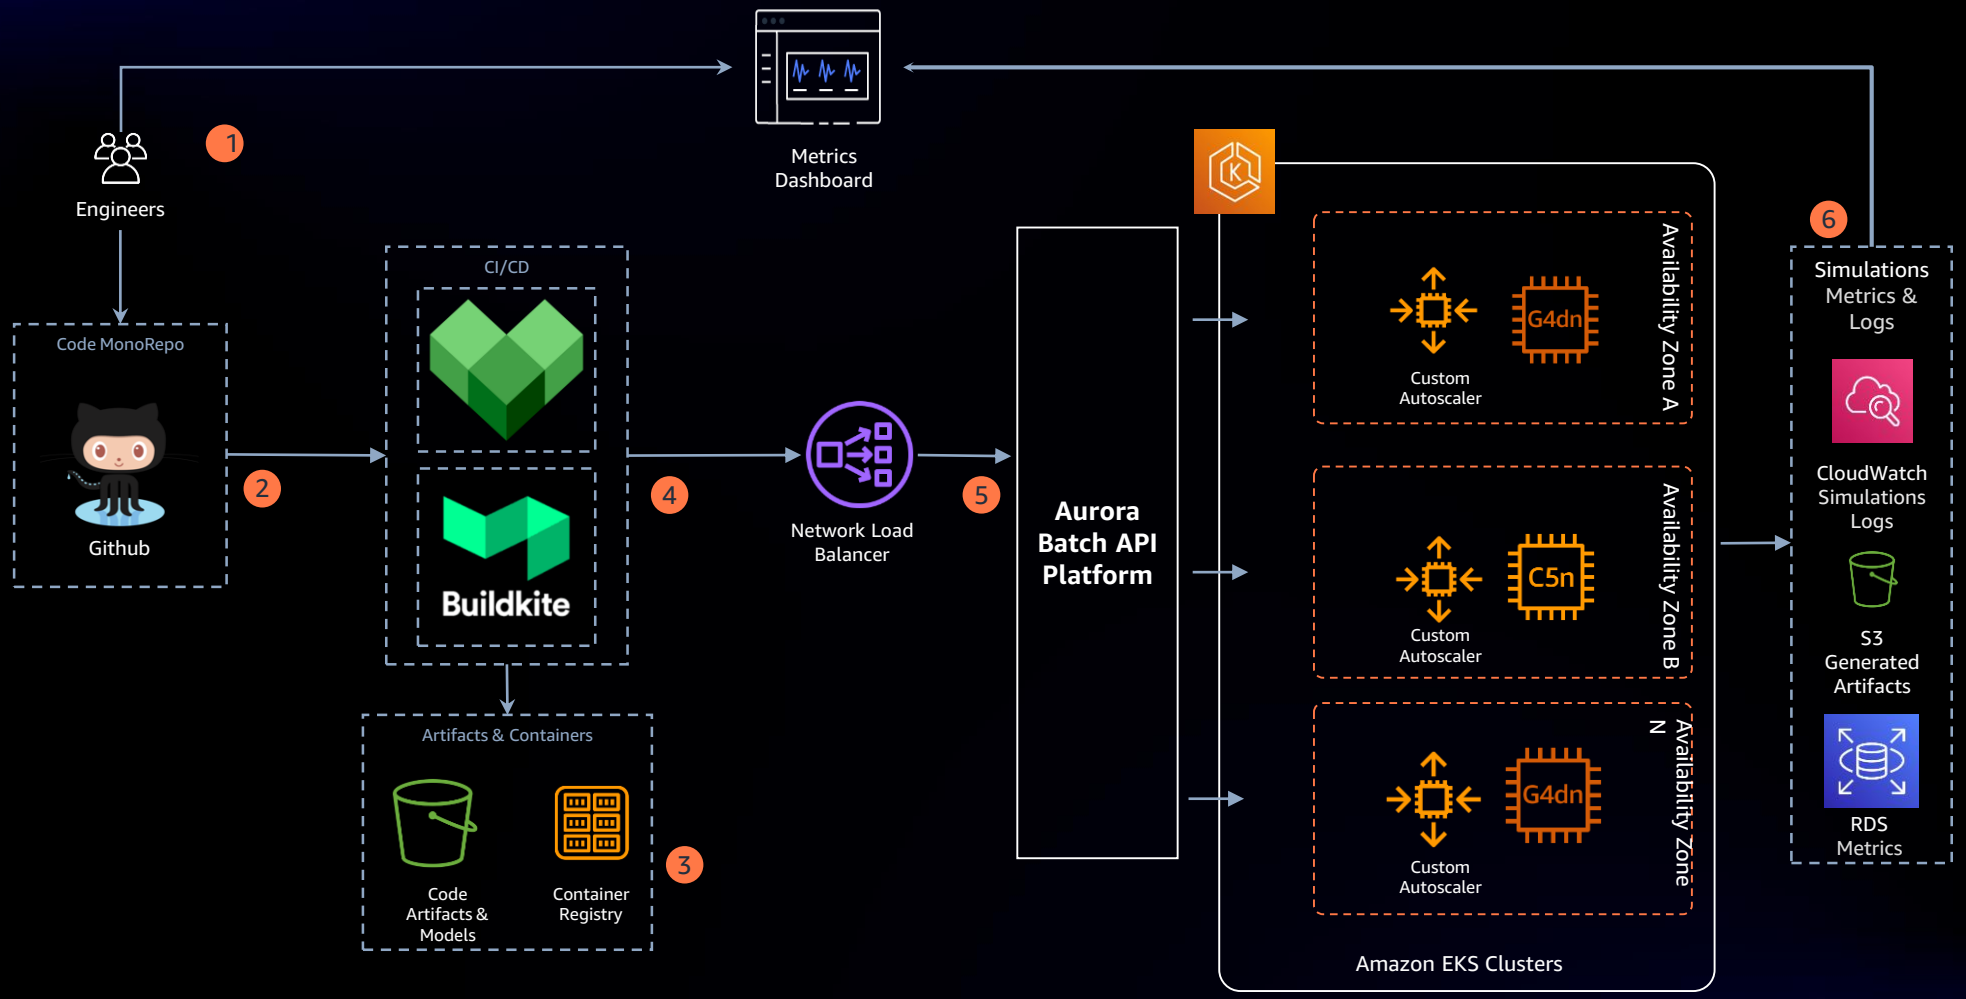
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.



FAST!

©2021 | Aurora Proprietary

Aurora's supercomputer



1. Workflows are triggered via user interaction on the dashboard, time intervals, or automatically through the CI/CD pipeline
2. Any code change will trigger automated CI/CD pipelines and unit tests as part of a fully hermetic build process
3. Artifacts such as Docker containers are built and pushed to Container Registry and Amazon S3
4. BuildKite CI/CD job triggers closed-loop job and defines DAG steps for BatchAPI
5. BatchAPI uses Aurora's custom auto scaling to provision capacity across regions, instance types, and EKS clusters. CPU or GPU jobs are scheduled based on job size and user defined input
6. Results and artifacts are gathered in object storage and database (RDS Postgres) then published to a metrics dashboard

Aurora

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora lessons learned

- Use all Availability Zones in a Region
- Multiple instance types key for larger-scale workloads
- Bursty workloads and can be challenging on Kubernetes
- Larger nodes and lower pod counts key to scaling Kubernetes
  - We even use metal instances!
- Utilize EC2 Spot Instances
- Prep your AWS accounts on limits, VPC CIDR, and scaling up Amazon EKS control plane to support thousands of worker nodes

Aurora

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Conclusions

- Autonomous systems development is a hyperscale problem, requires agility and speed, attention on cost management, and enforcement of hard safety requirements
- AWS provides a complete set of services to support the main autonomous driving workloads: data storage and management, ML training, simulation
- It’s not just cloud services, but it’s necessary to integrate the vehicles in a “Big Loop” perspective (shadow mode, snapshots – “intelligent edge,” in general)
- Autonomous driving is an ecosystem play, and we want to be the most ISV-friendly cloud provider, with APN and AWS Marketplace opportunities

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Work hard
Have fun
Make history

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Autonomous data lake: Resources

Blog posts

Automating Data Ingestion and Labeling for Autonomous Vehicle Development
https://go.aws/3ojgjcr

Deploying Autonomous Driving and ADAS Workloads at Scale with Amazon Managed Workflows for Apache Airflow
https://go.aws/30aGWYO

Building an Automated Image Processing and Model Training Pipeline for Autonomous Driving
https://go.aws/30nXHzU

Building an Automated Scene Detection Pipeline for Autonomous Driving – ADAS Workflow
https://go.aws/3HaqicK

Deploy and Visualize ROS Bag Data on AWS Using rviz and Webviz for Autonomous Driving
https://go.aws/3F9P9vv

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.