AWS re:Invent

IOT356-P

Machine learning at the edge for industrial applications

Richard Elberger

Principal Solutions Architect – IoT

Amazon Web Services

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

Agenda

Tenets

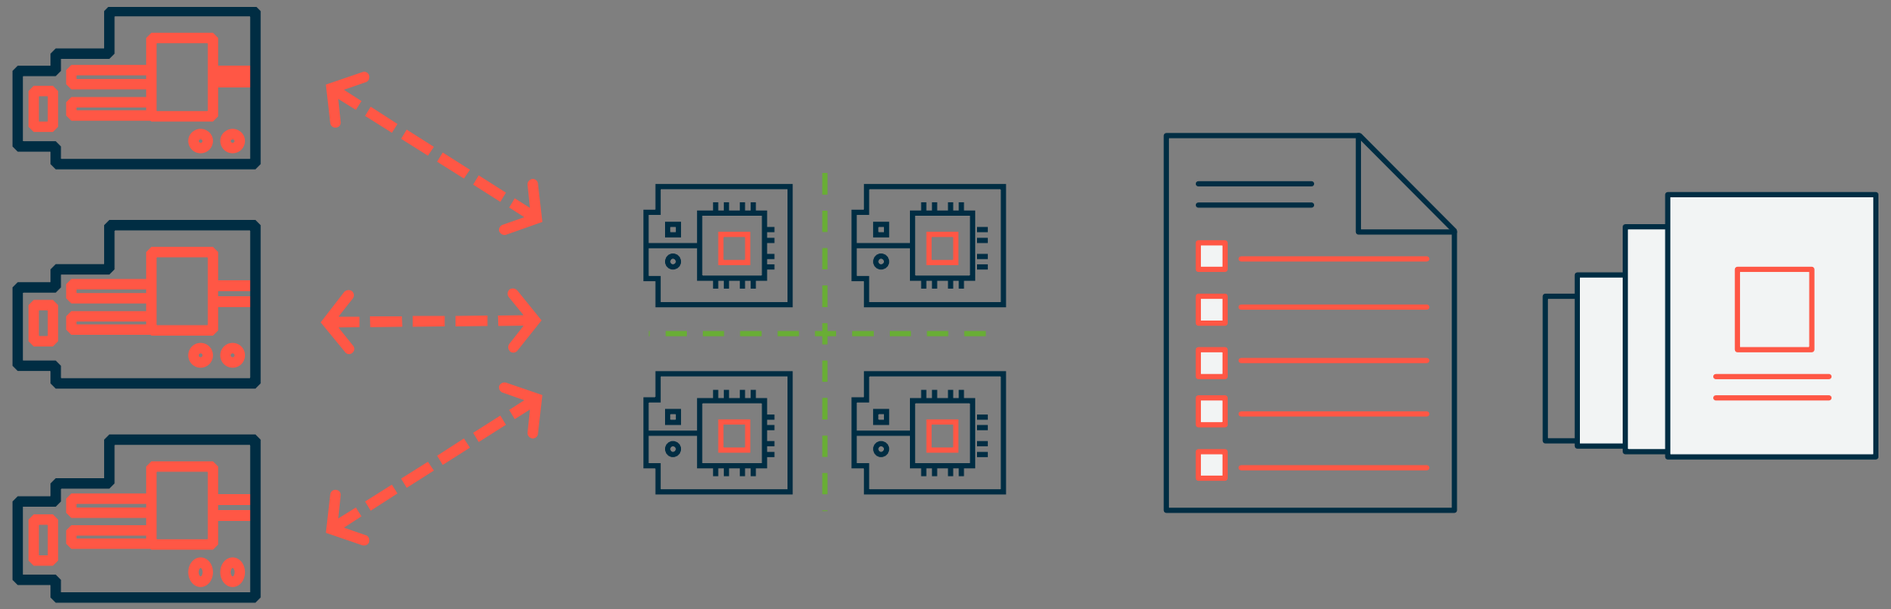
Industrial IoT architecture

AIoT lifecycle – a four-part story

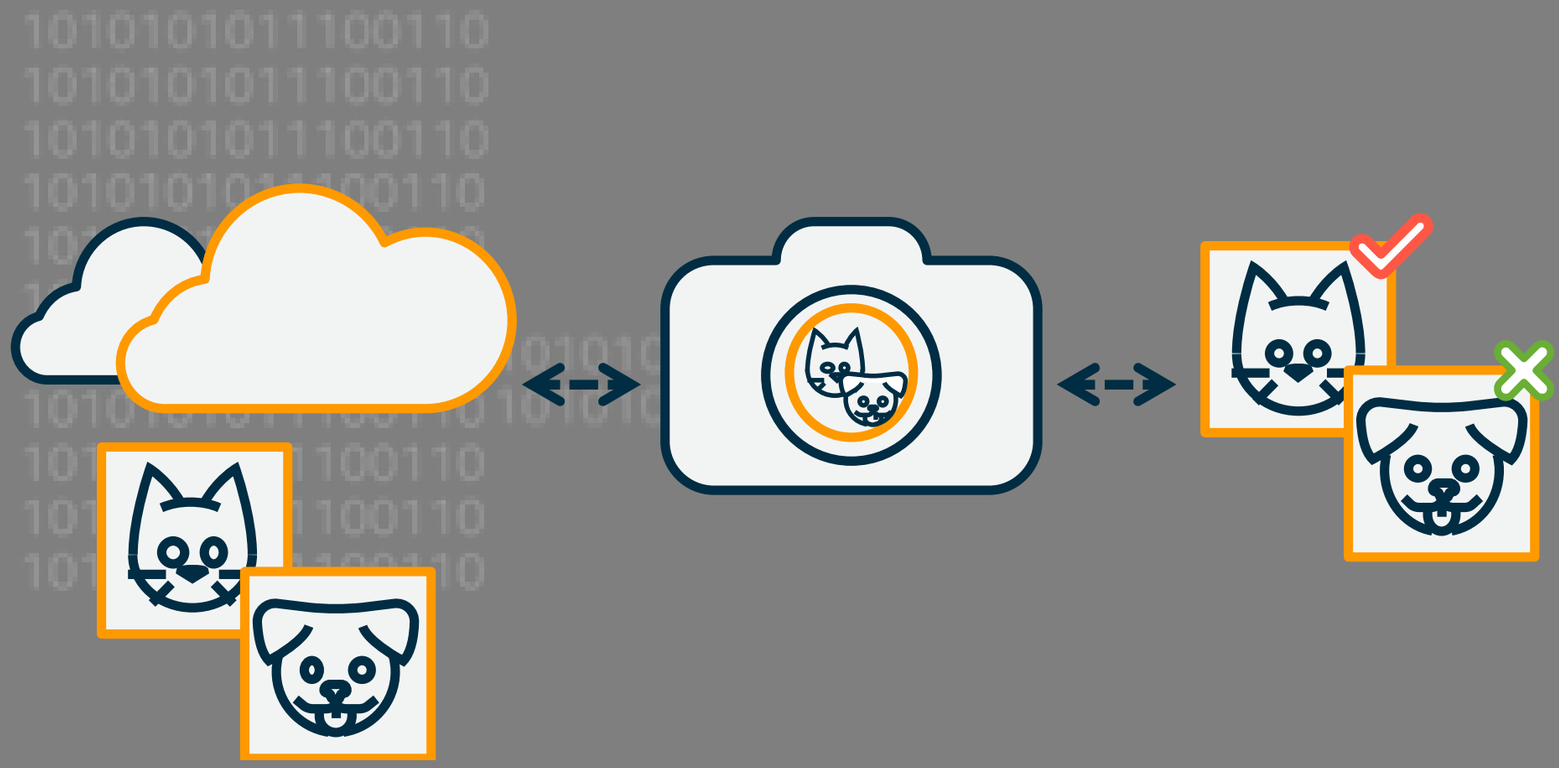
Tenets and reference case study

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

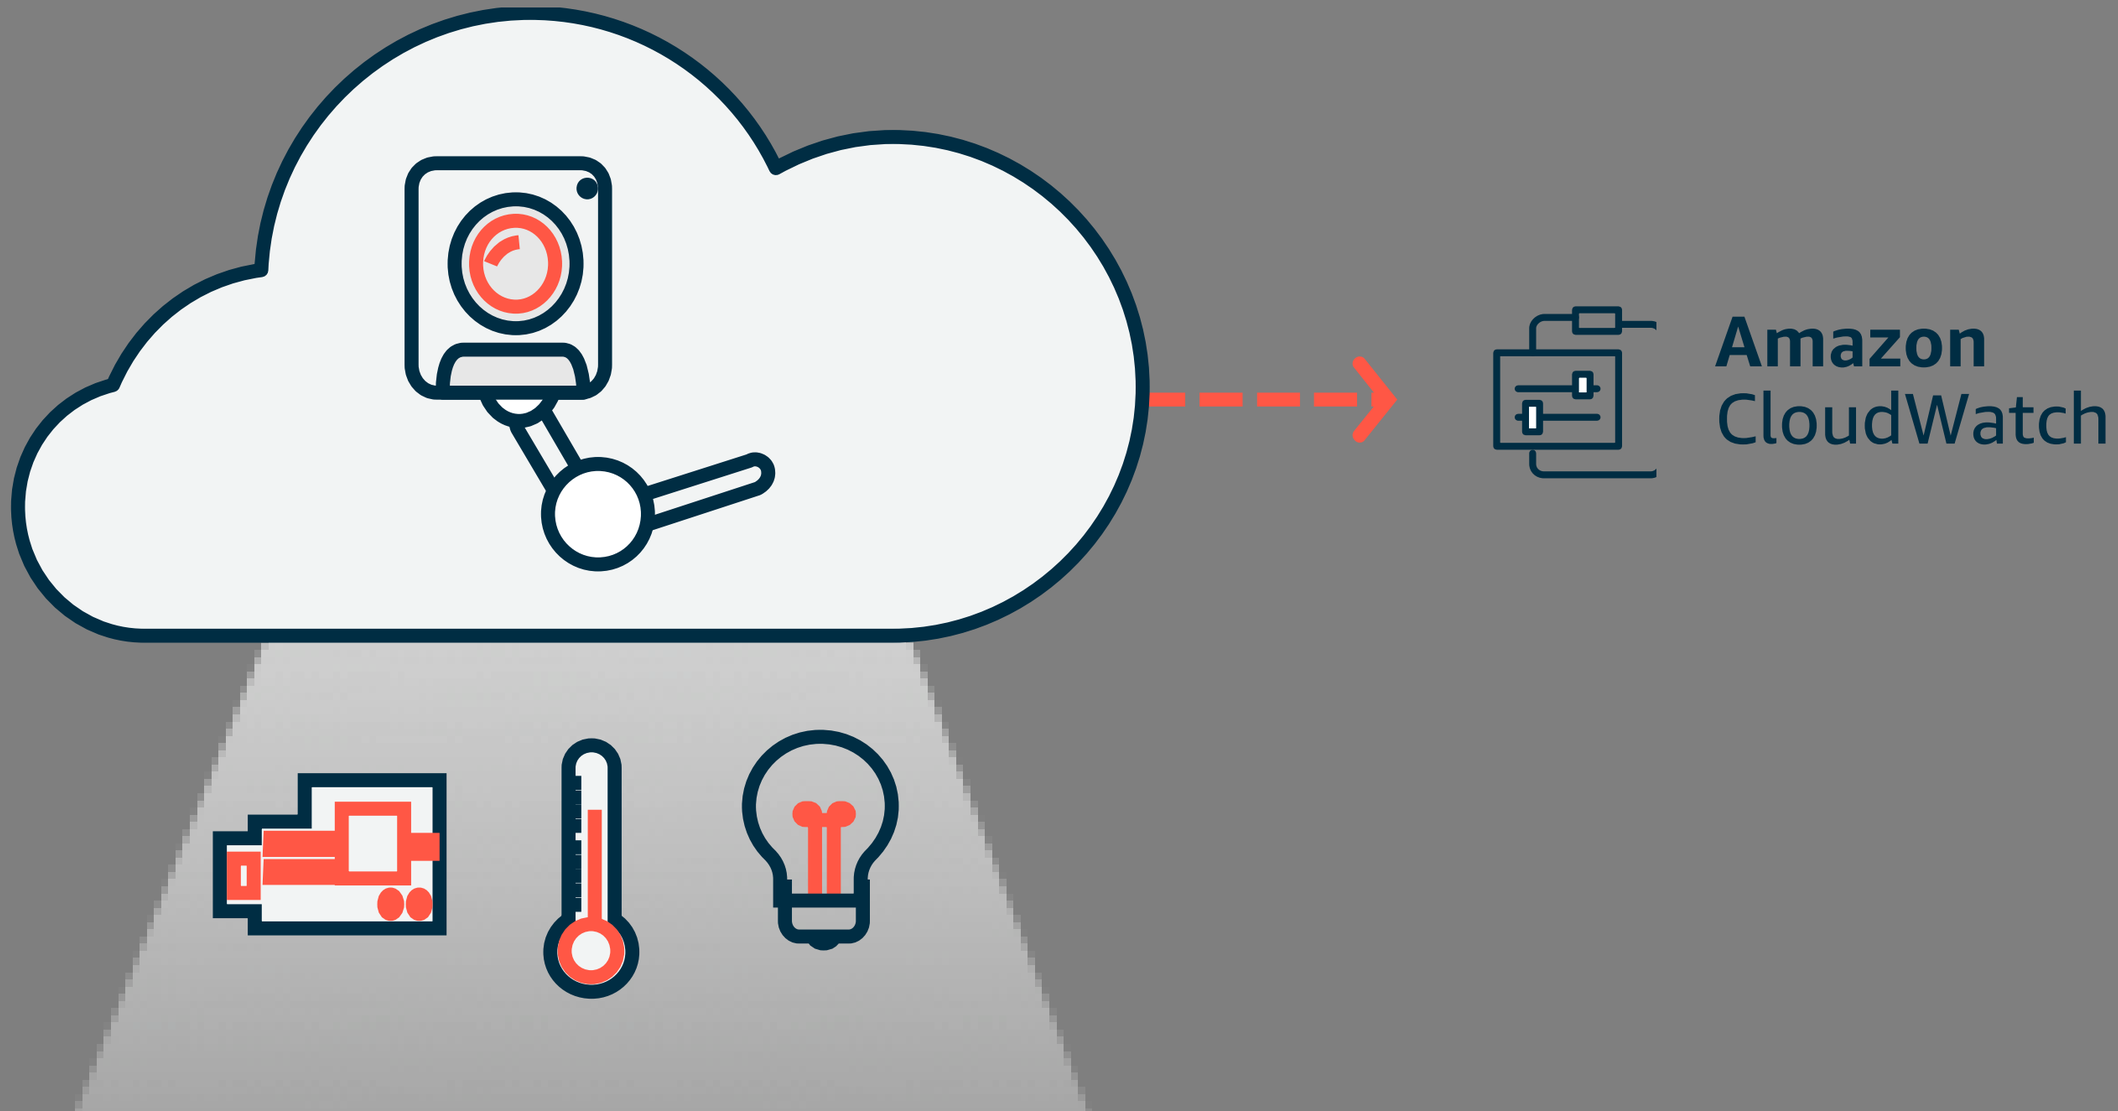
Effectively maintain the system over its lifecycle



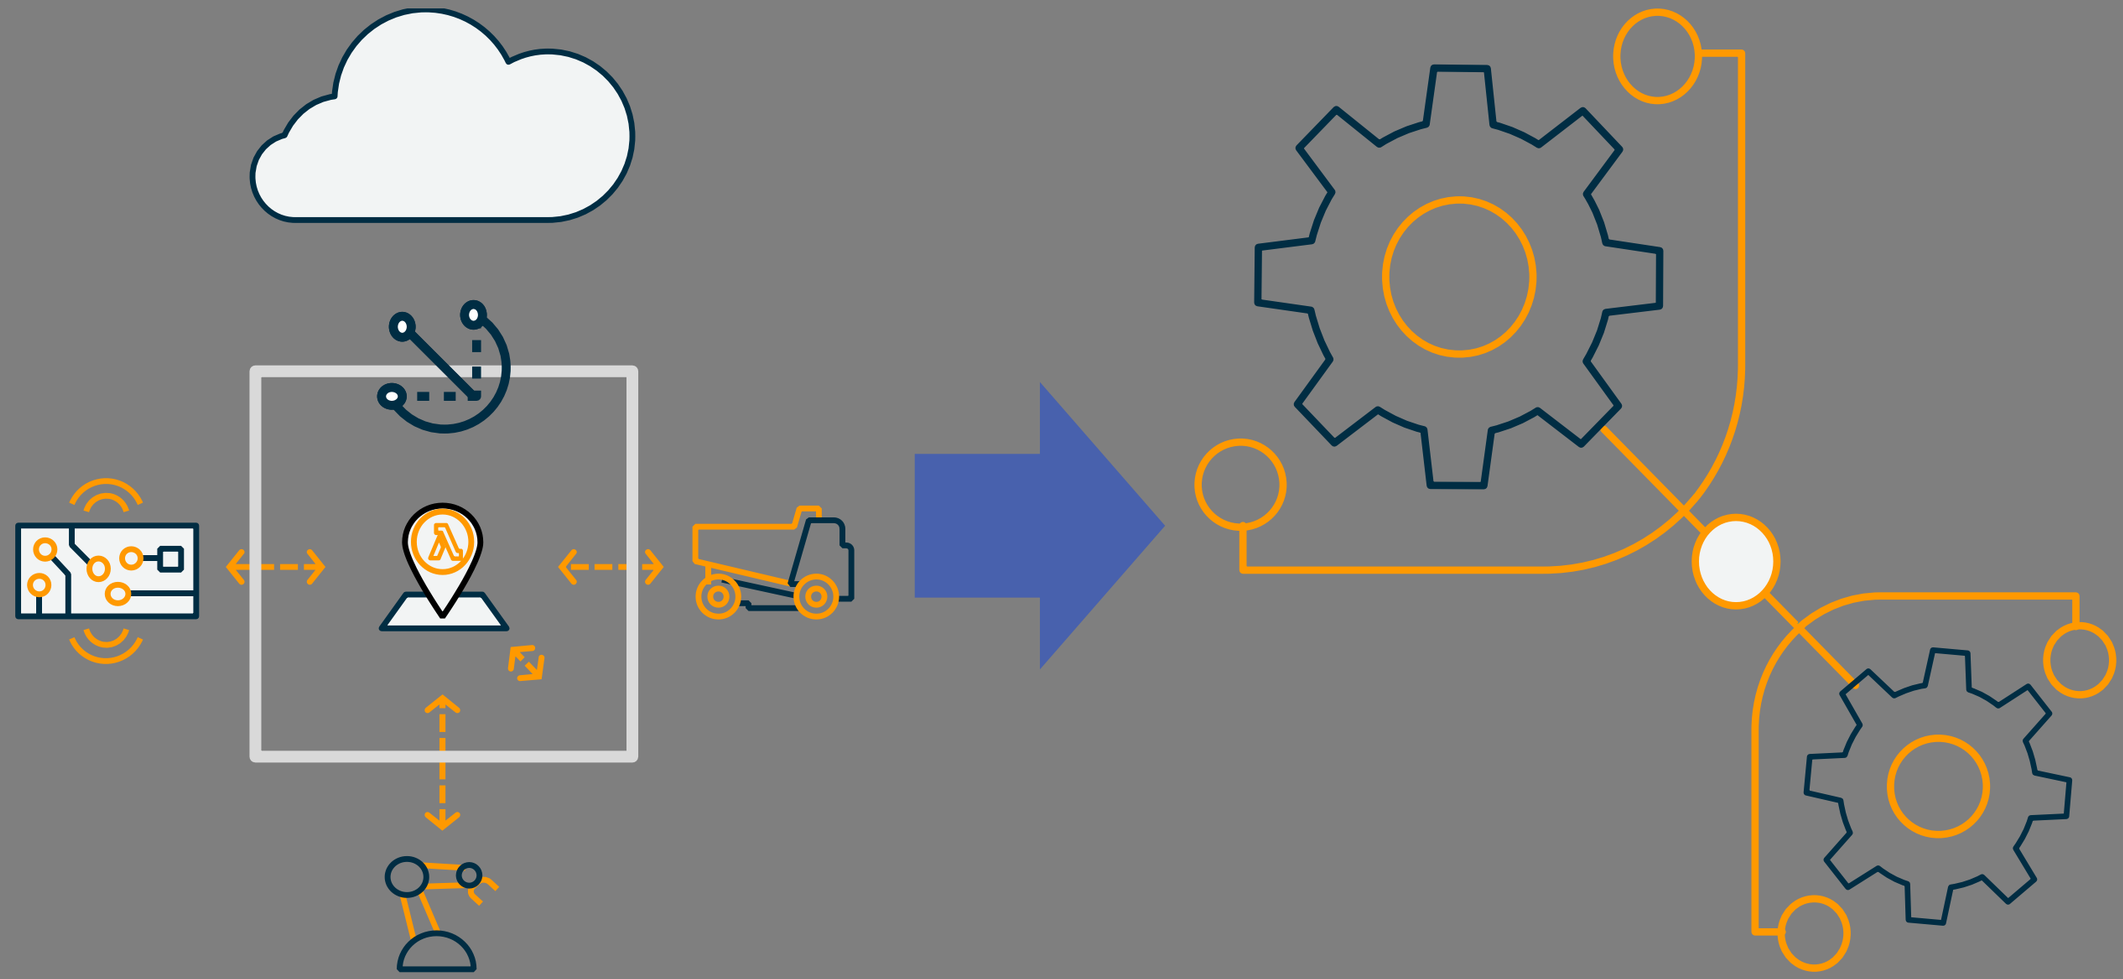
Exploit system cost effectiveness via new intelligence

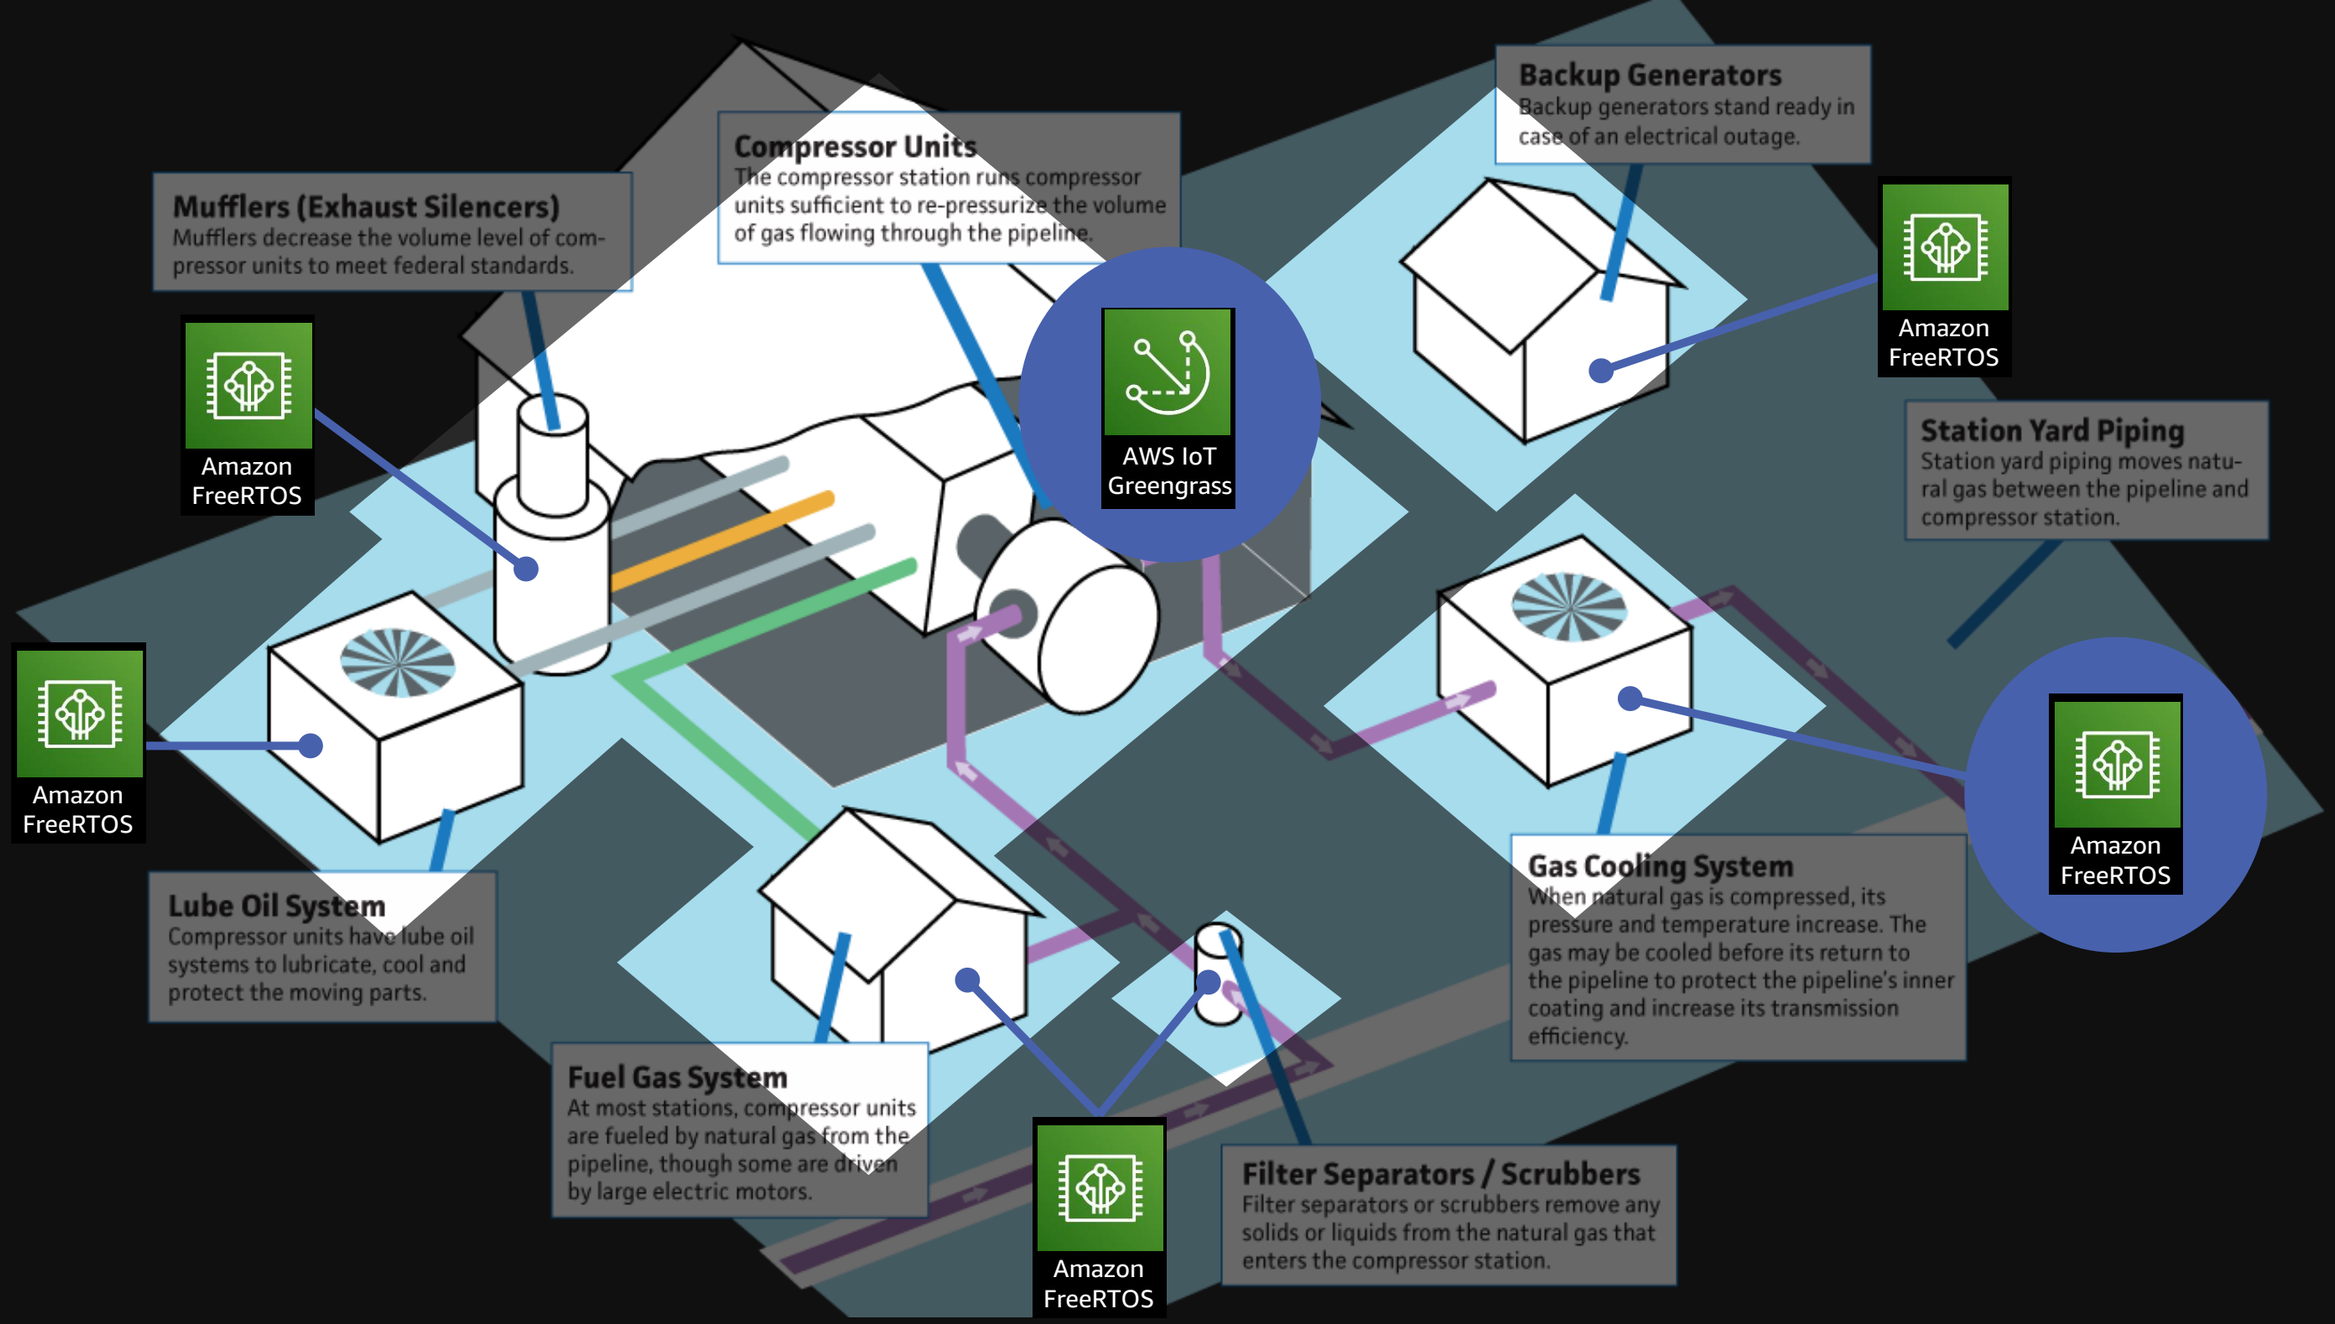


Better decisions through central dashboards and monitoring



Derive value through new capabilities



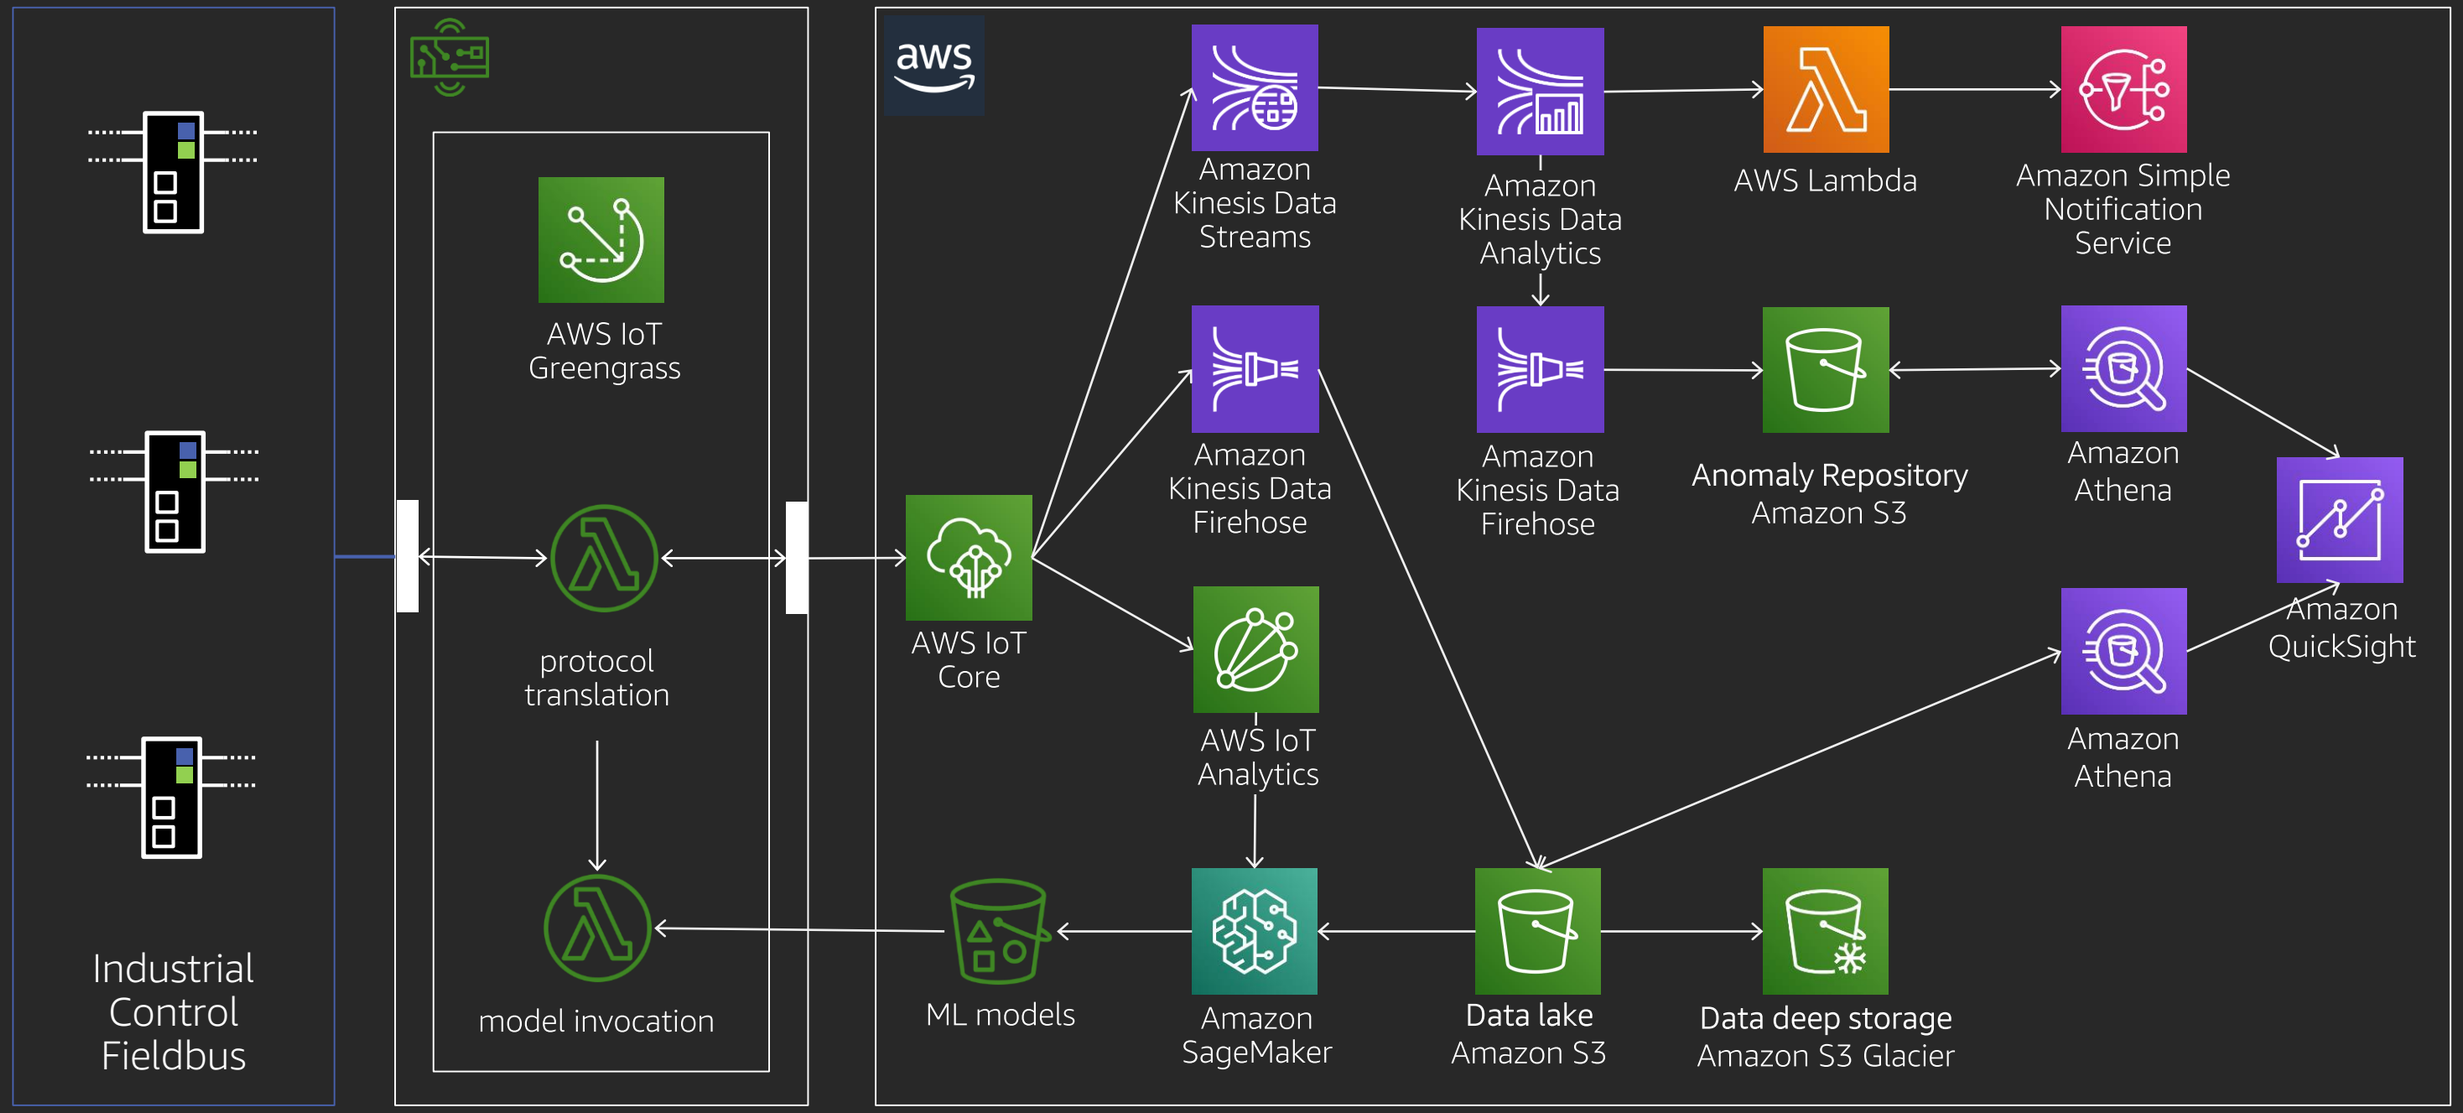

Mufflers (Exhaust Silencers)
Mufflers decrease the volume level of compressor units to meet federal standards.
Compressor Units
The compressor station runs compressor units sufficient to re-pressurize the volume of gas flowing through the pipeline.
Backup Generators
Backup generators stand ready in case of an electrical outage.
Amazon FreeRTOS
AWS IoT Greengrass
Amazon FreeRTOS
Station Yard Piping
Station yard piping moves natural gas between the pipeline and compressor station.
Amazon FreeRTOS
Amazon FreeRTOS
Lube Oil System
Compressor units have lube oil systems to lubricate, cool and protect the moving parts.
Gas Cooling System
When natural gas is compressed, its pressure and temperature increase. The gas may be cooled before its return to the pipeline to protect the pipeline's inner coating and increase its transmission efficiency.
Fuel Gas System
At most stations, compressor units are fueled by natural gas from the pipeline, though some are driven by large electric motors.
Amazon FreeRTOS
Filter Separators / Scrubbers
Filter separators or scrubbers remove any solids or liquids from the natural gas that enters the compressor station.

Industrial IoT architecture

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

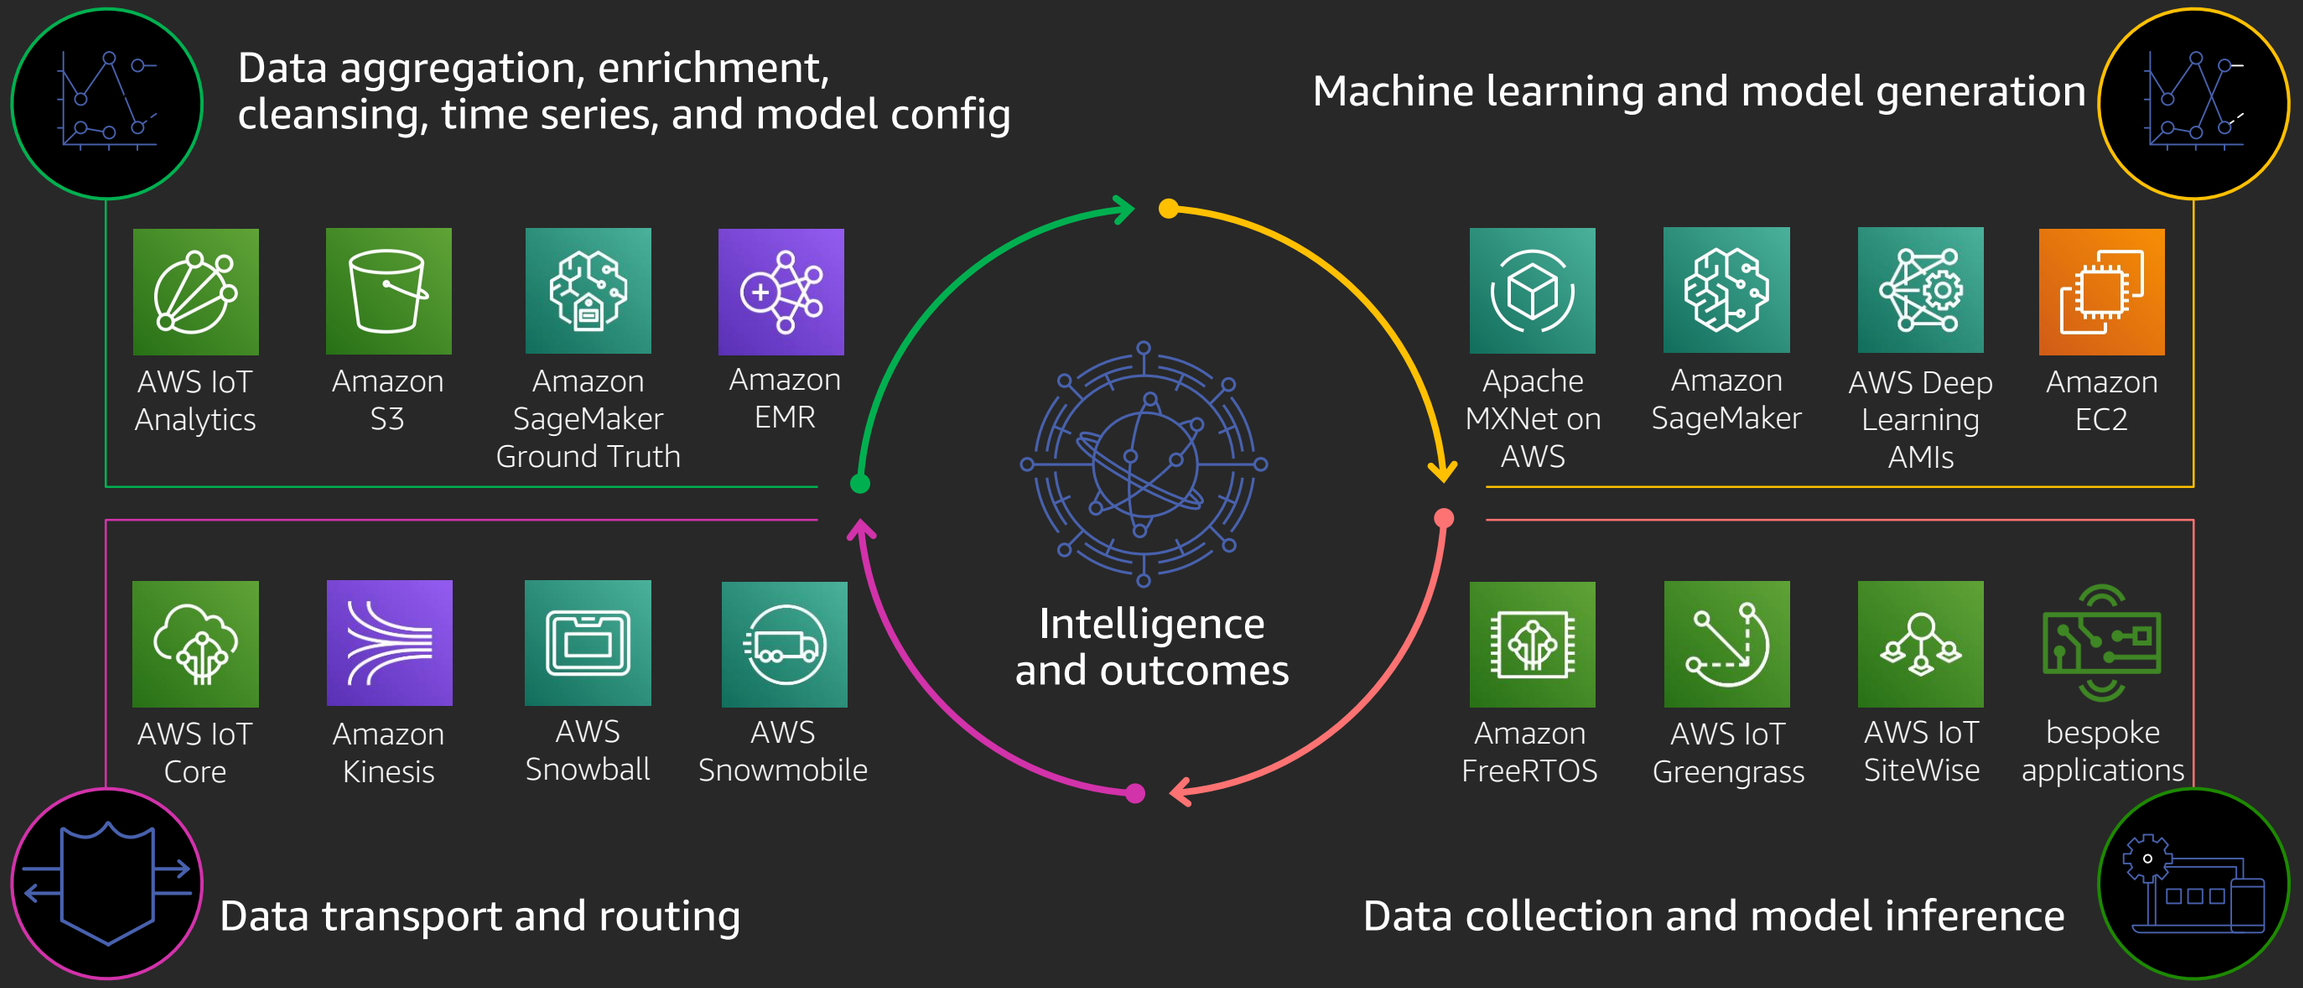

Industrial Control Fieldbus
AWS IoT Greengrass
protocol translation
model invocation
aws
AWS IoT Core
Amazon Kinesis Data Streams
Amazon Kinesis Data Analytics
AWS Lambda
Amazon Simple Notification Service
Amazon Kinesis Data Firehose
Amazon Kinesis Data Firehose
Anomaly Repository Amazon S3
Amazon Athena
Amazon QuickSight
AWS IoT Analytics
Amazon Athena
ML models
Amazon SageMaker
Data lake Amazon S3
Data deep storage Amazon S3 Glacier

Machine learning lifecycle

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AIoT – Machine Learning at the Edge

Train in the cloud and infer at the edge



AIoT – Machine Learning at the Edge

Train in the cloud and infer at the edge



Data aggregation, enrichment, cleansing, time series, and model config

AWS IoT Analytics

Amazon S3

Amazon SageMaker Ground Truth

Amazon EMR

Machine learning and model generation



Amazon SageMaker

XILINX DNNDK

Amazon EC2



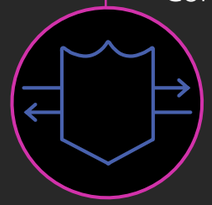
Intelligence and outcomes

AWS IoT Core

Amazon Kinesis

AWS Snowball

AWS Snowmobile



Data transport and routing

Amazon FreeRTOS Zynq7000 (DPU)

AWS IoT Greengrass UltraScale+ (DPU)



Data collection and model inference

Enabling machine learning at the edge: Feeding the data lake

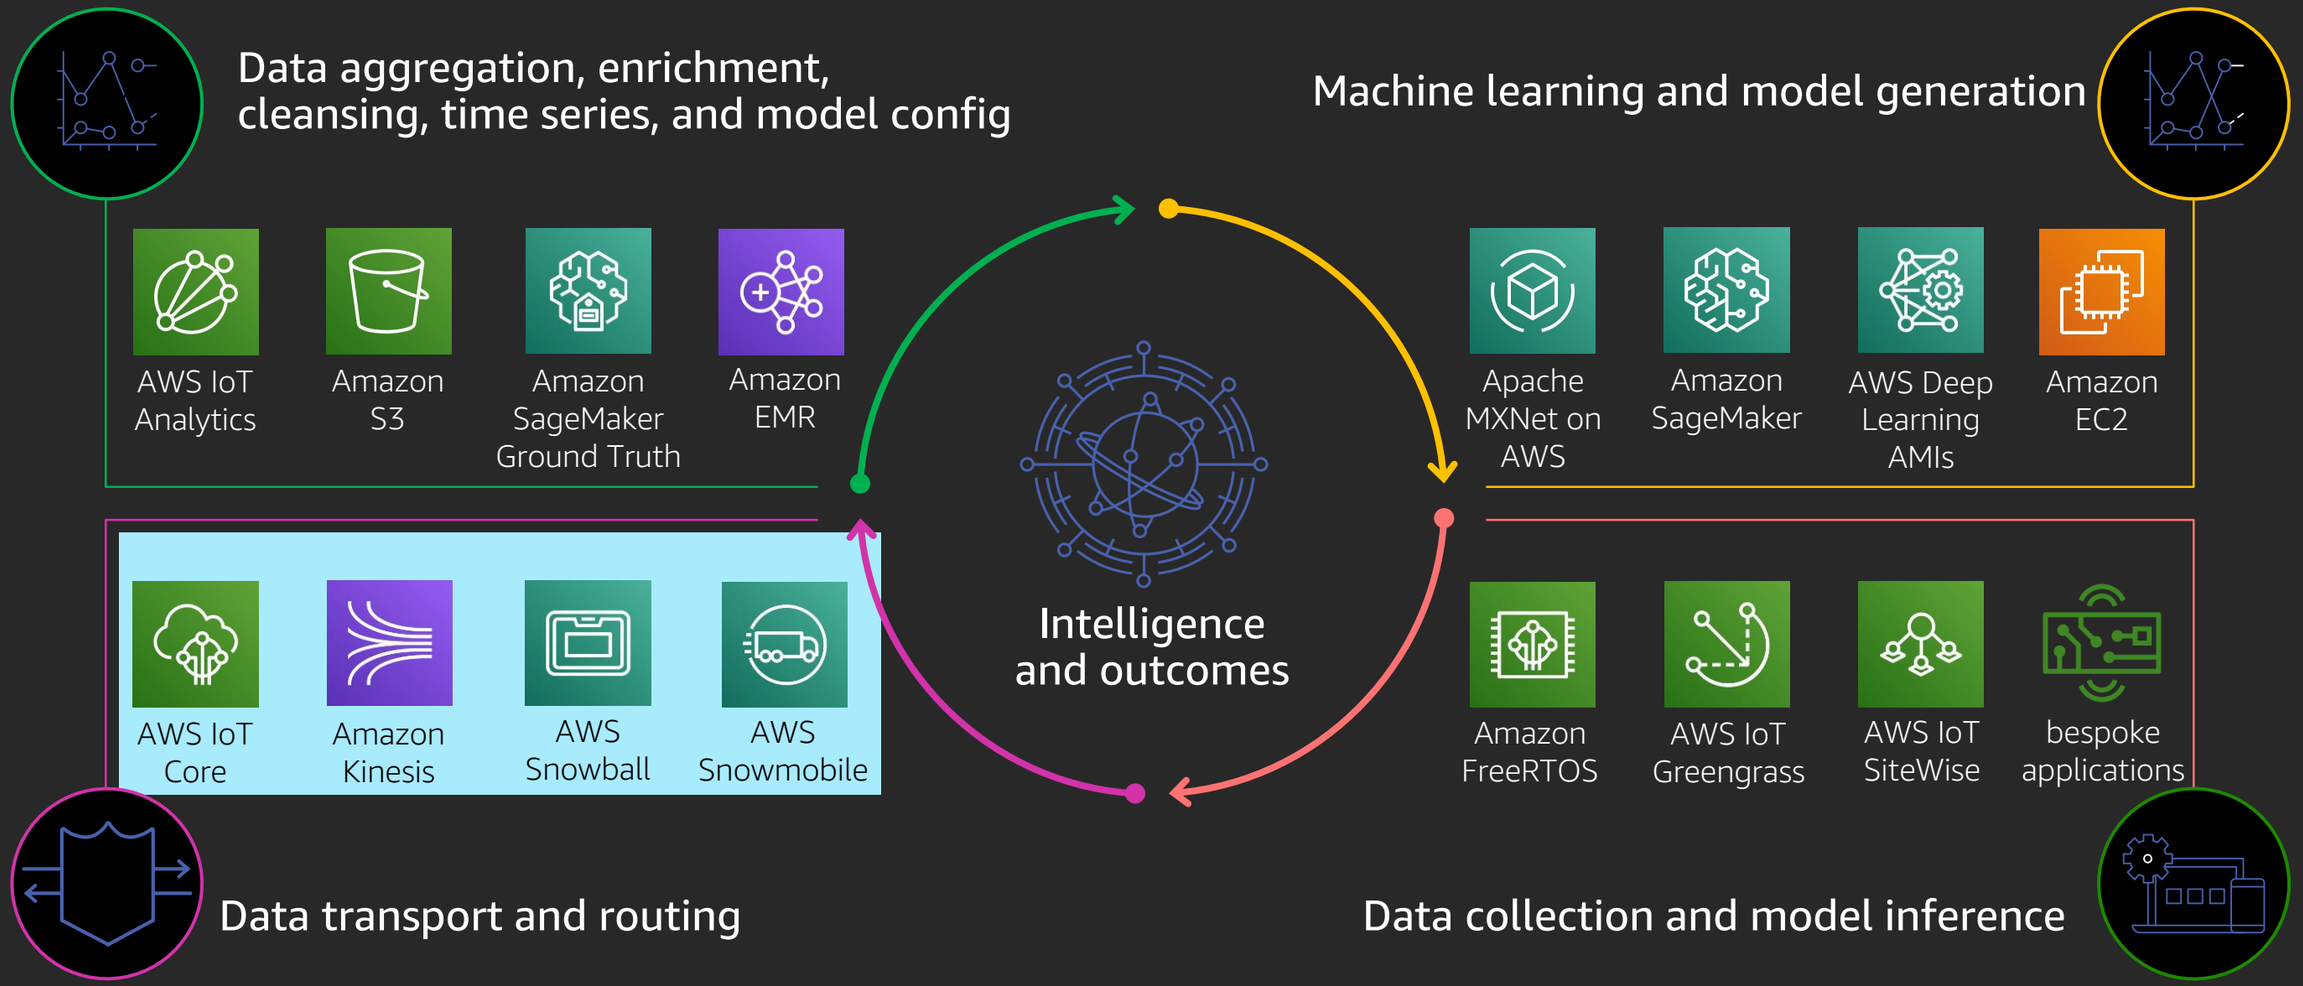
AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

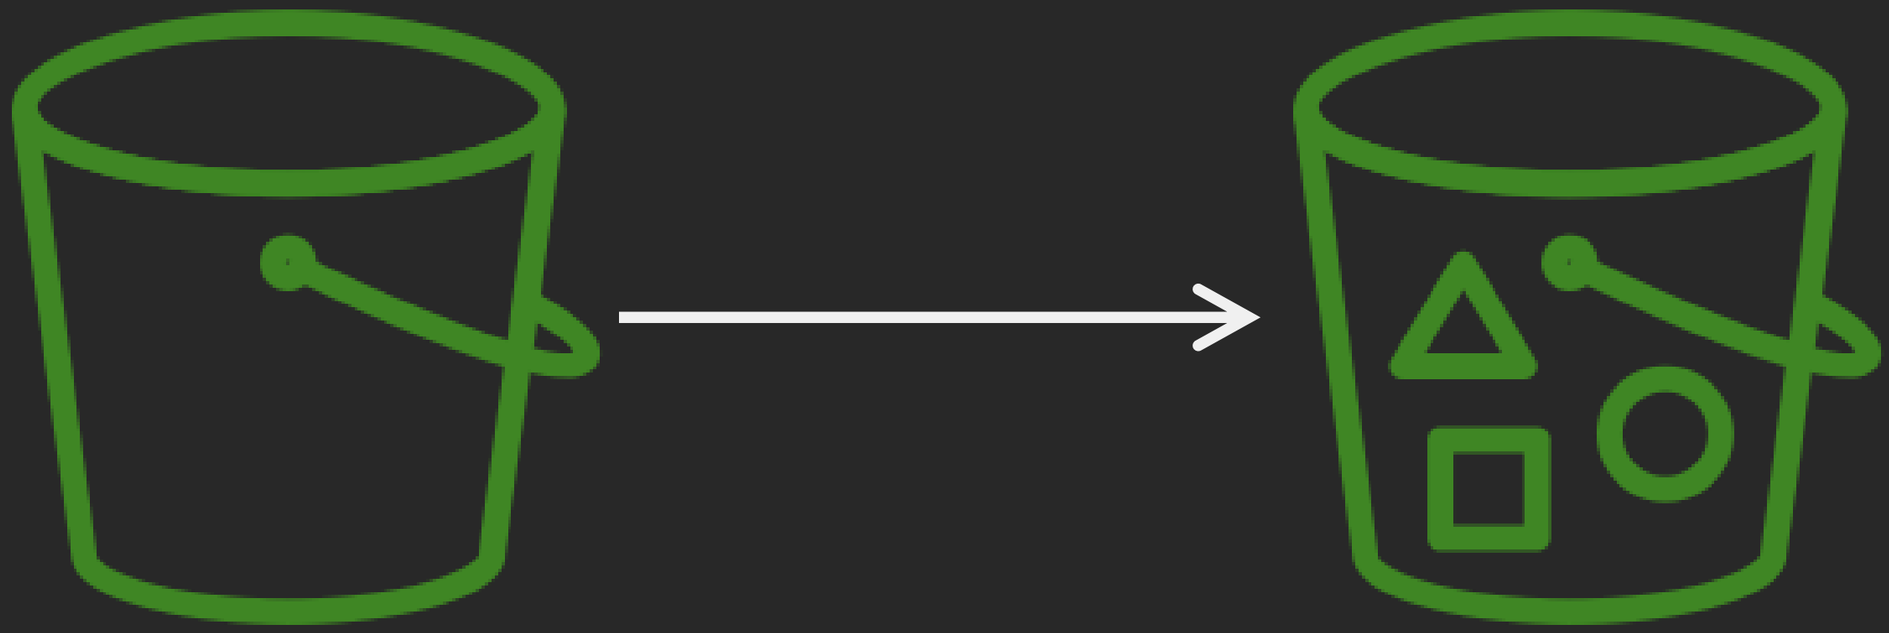
aws

AIoT – Machine Learning at the Edge

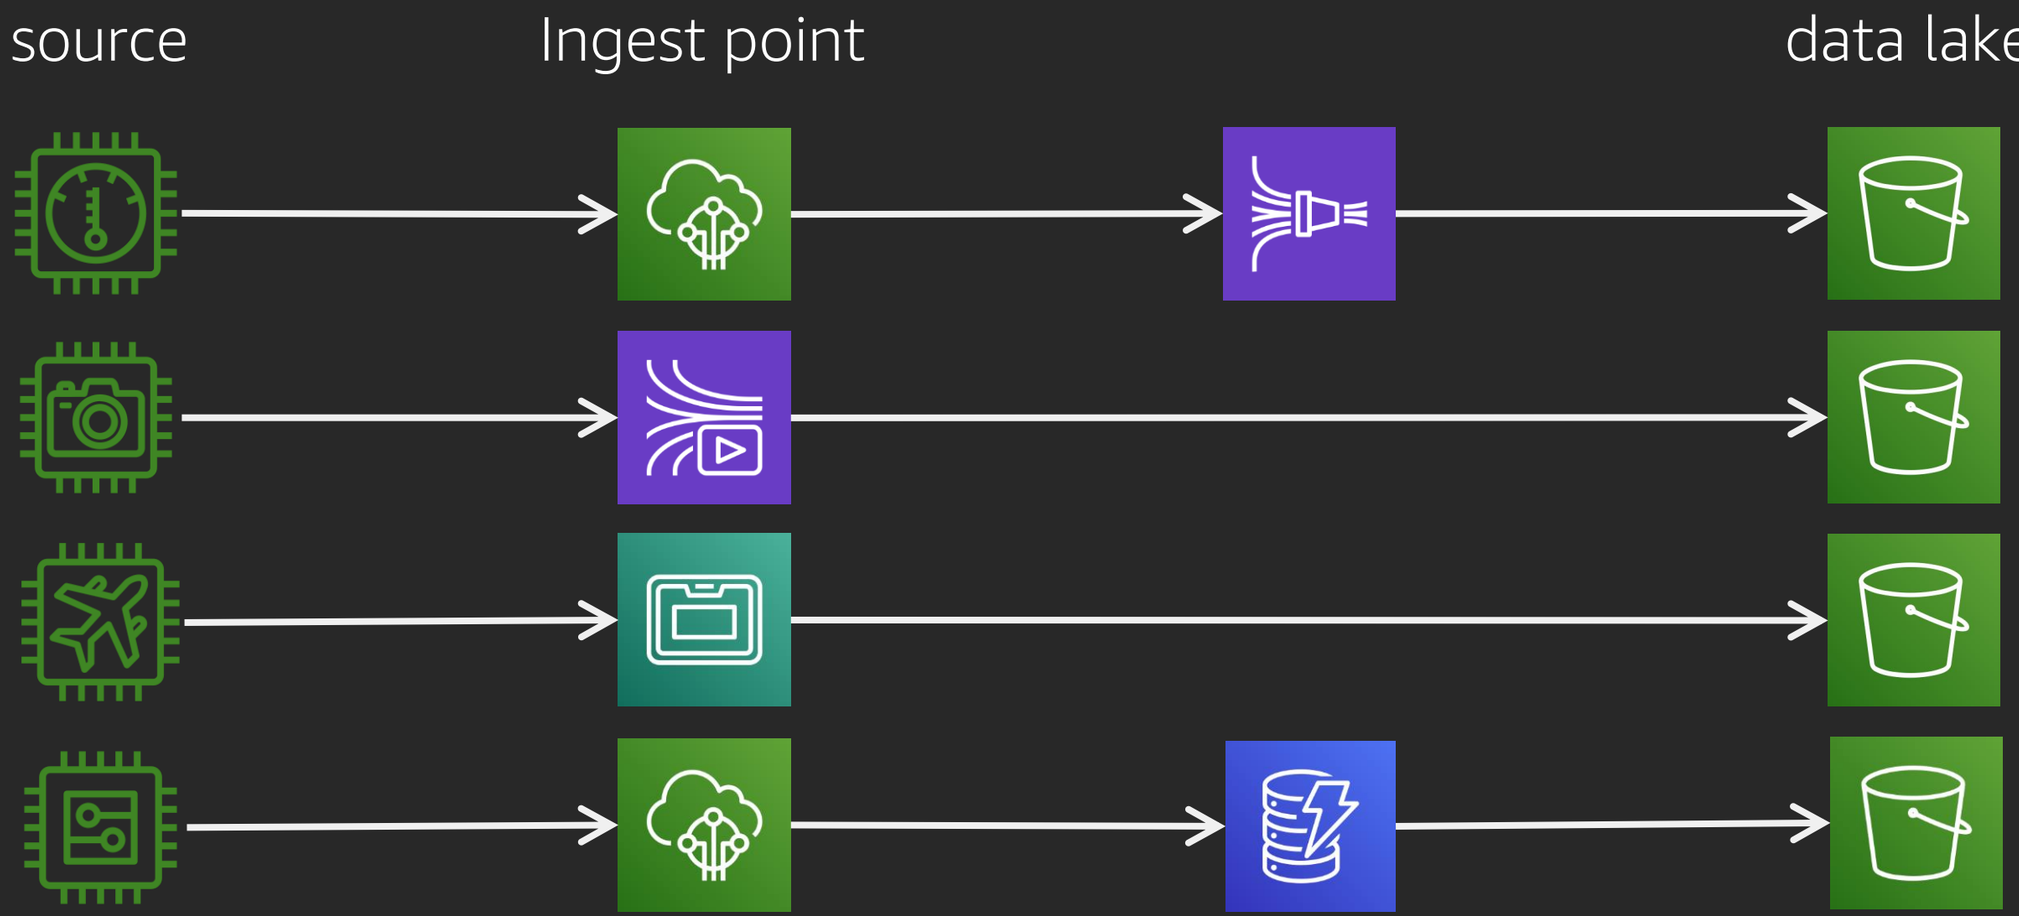
Train in the cloud and infer at the edge



In the beginning



Judgment on how to bring in data



Enabling machine learning at the edge: Curating the data lake

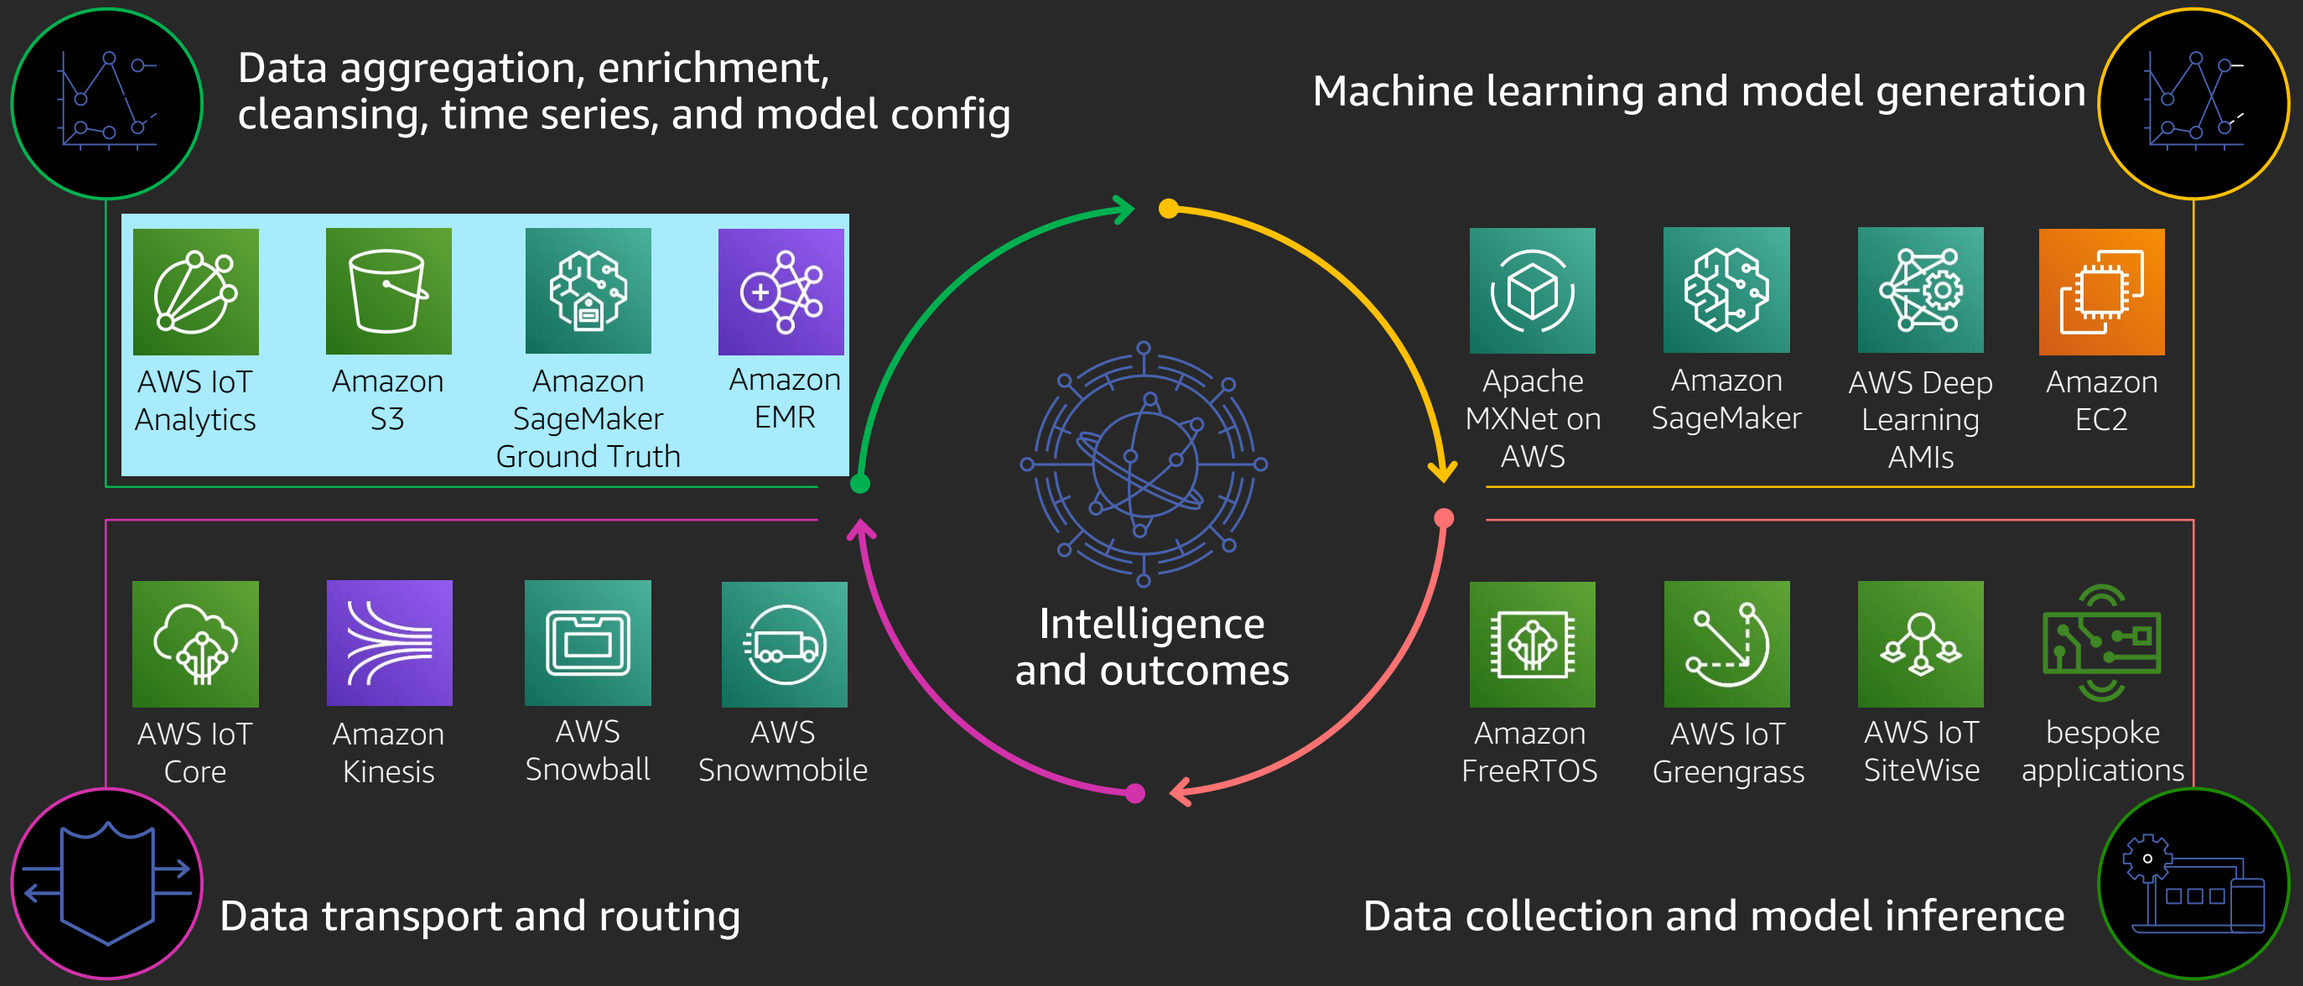
AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

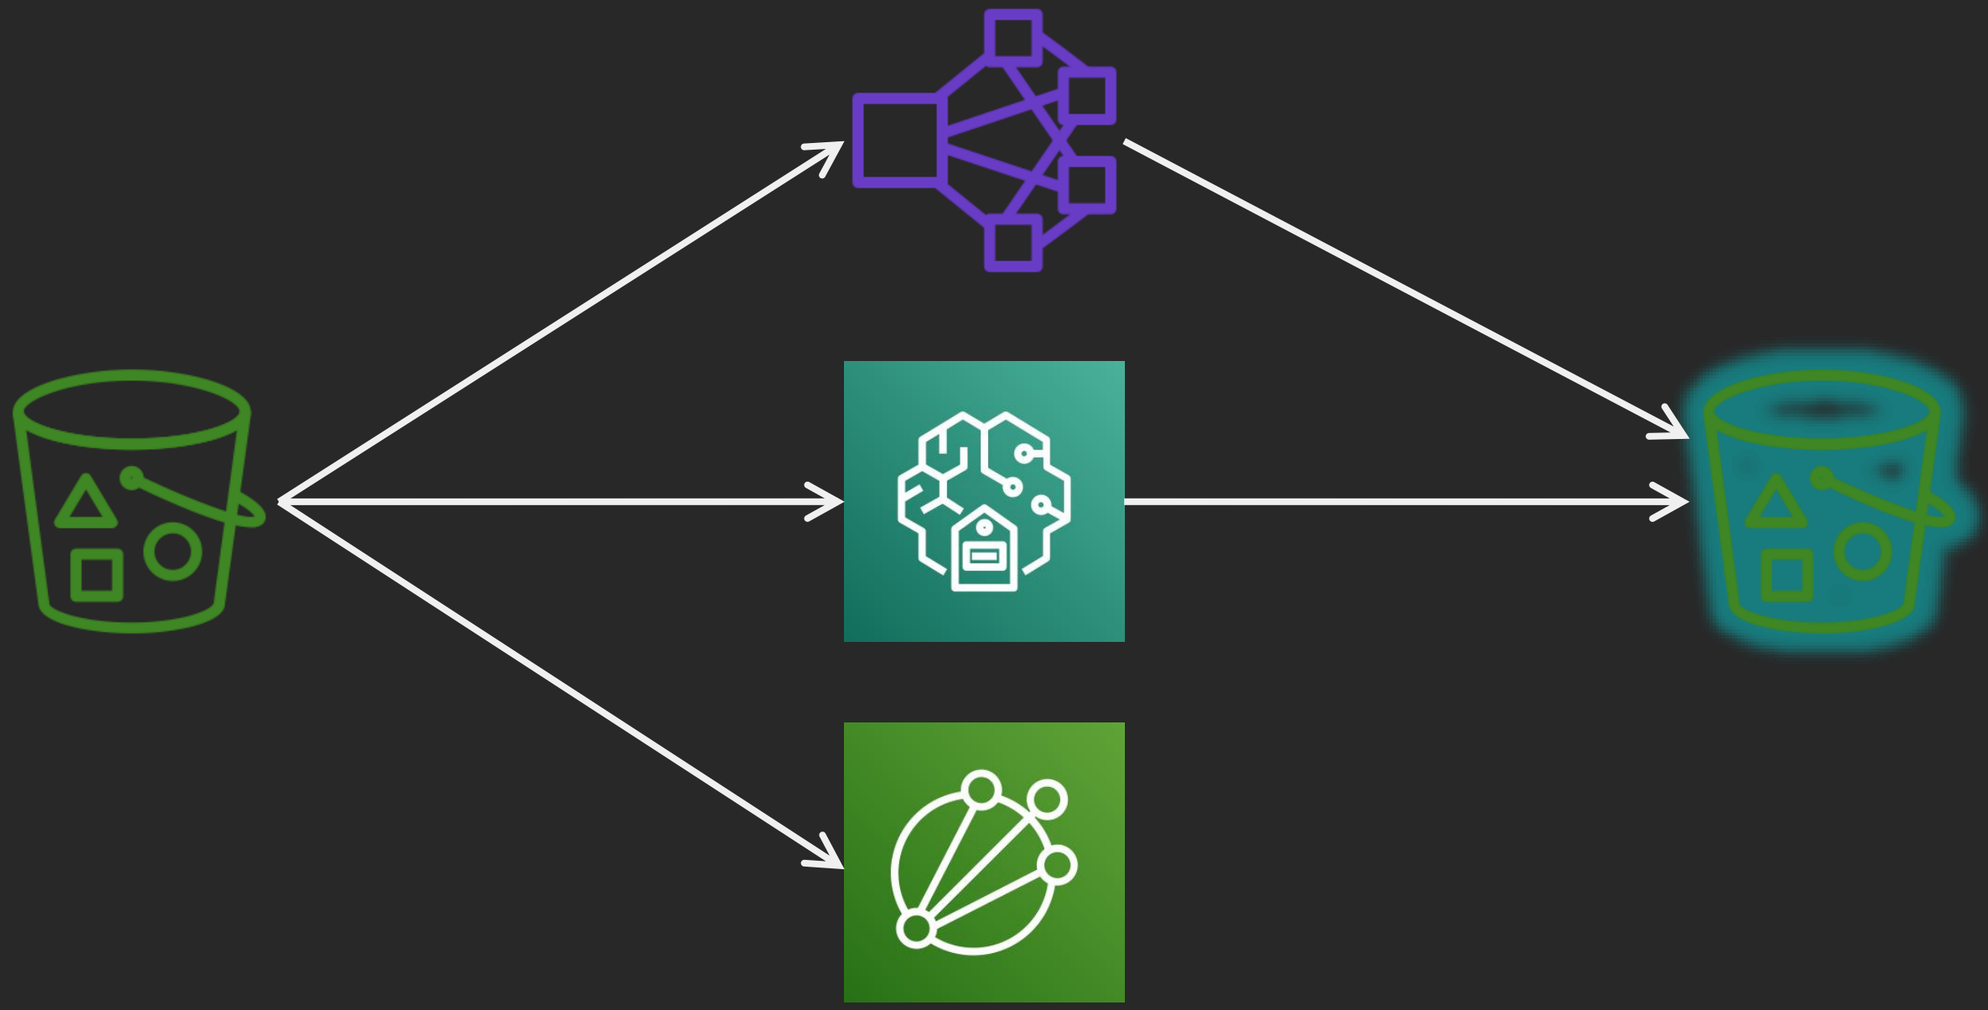
aws

AIoT – Machine Learning at the Edge

Train in the cloud and infer at the edge



Ingesting and methodically curating



Iteratively evaluating and refining

Where data science meets art



Enabling machine learning at the edge: Model training and staging

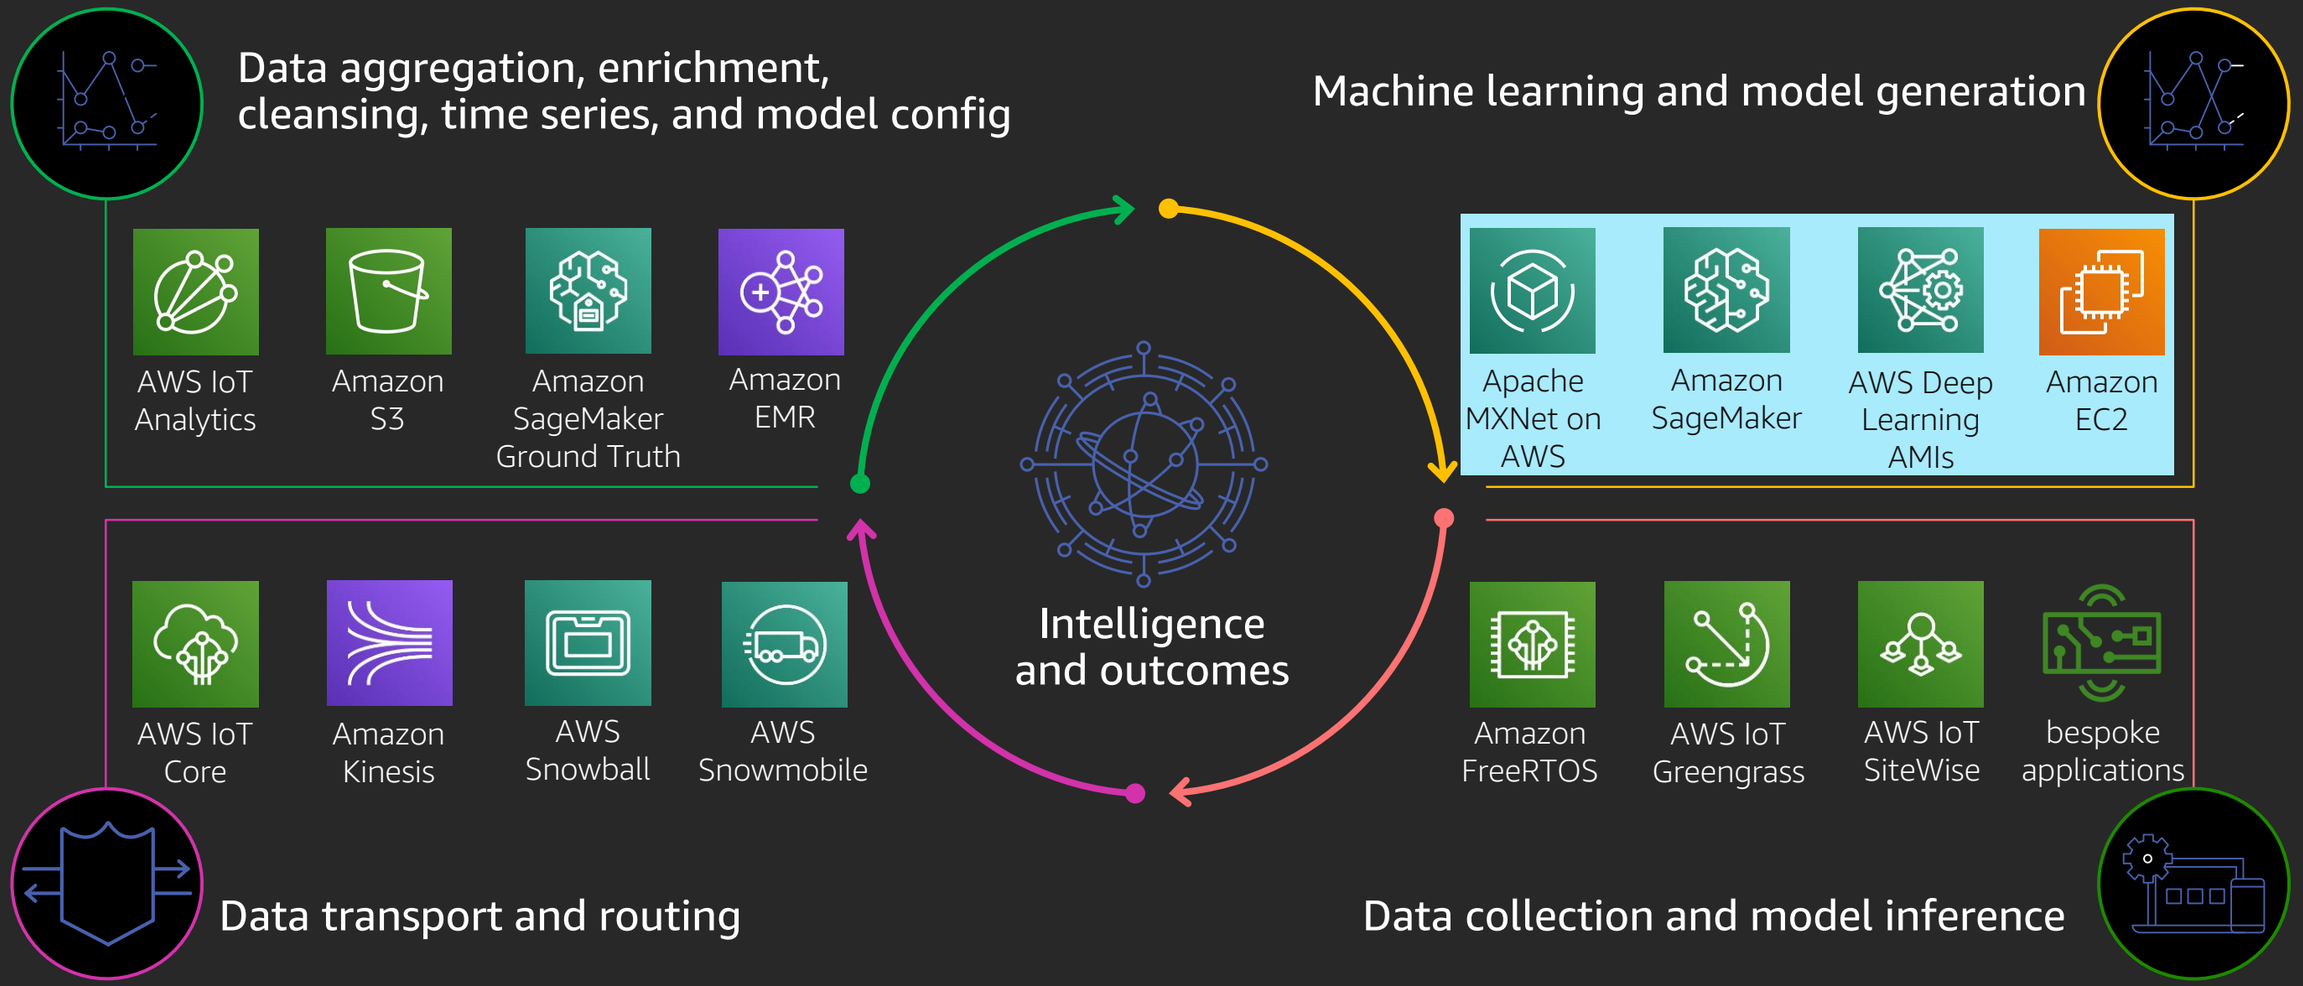
AWS re:Invent

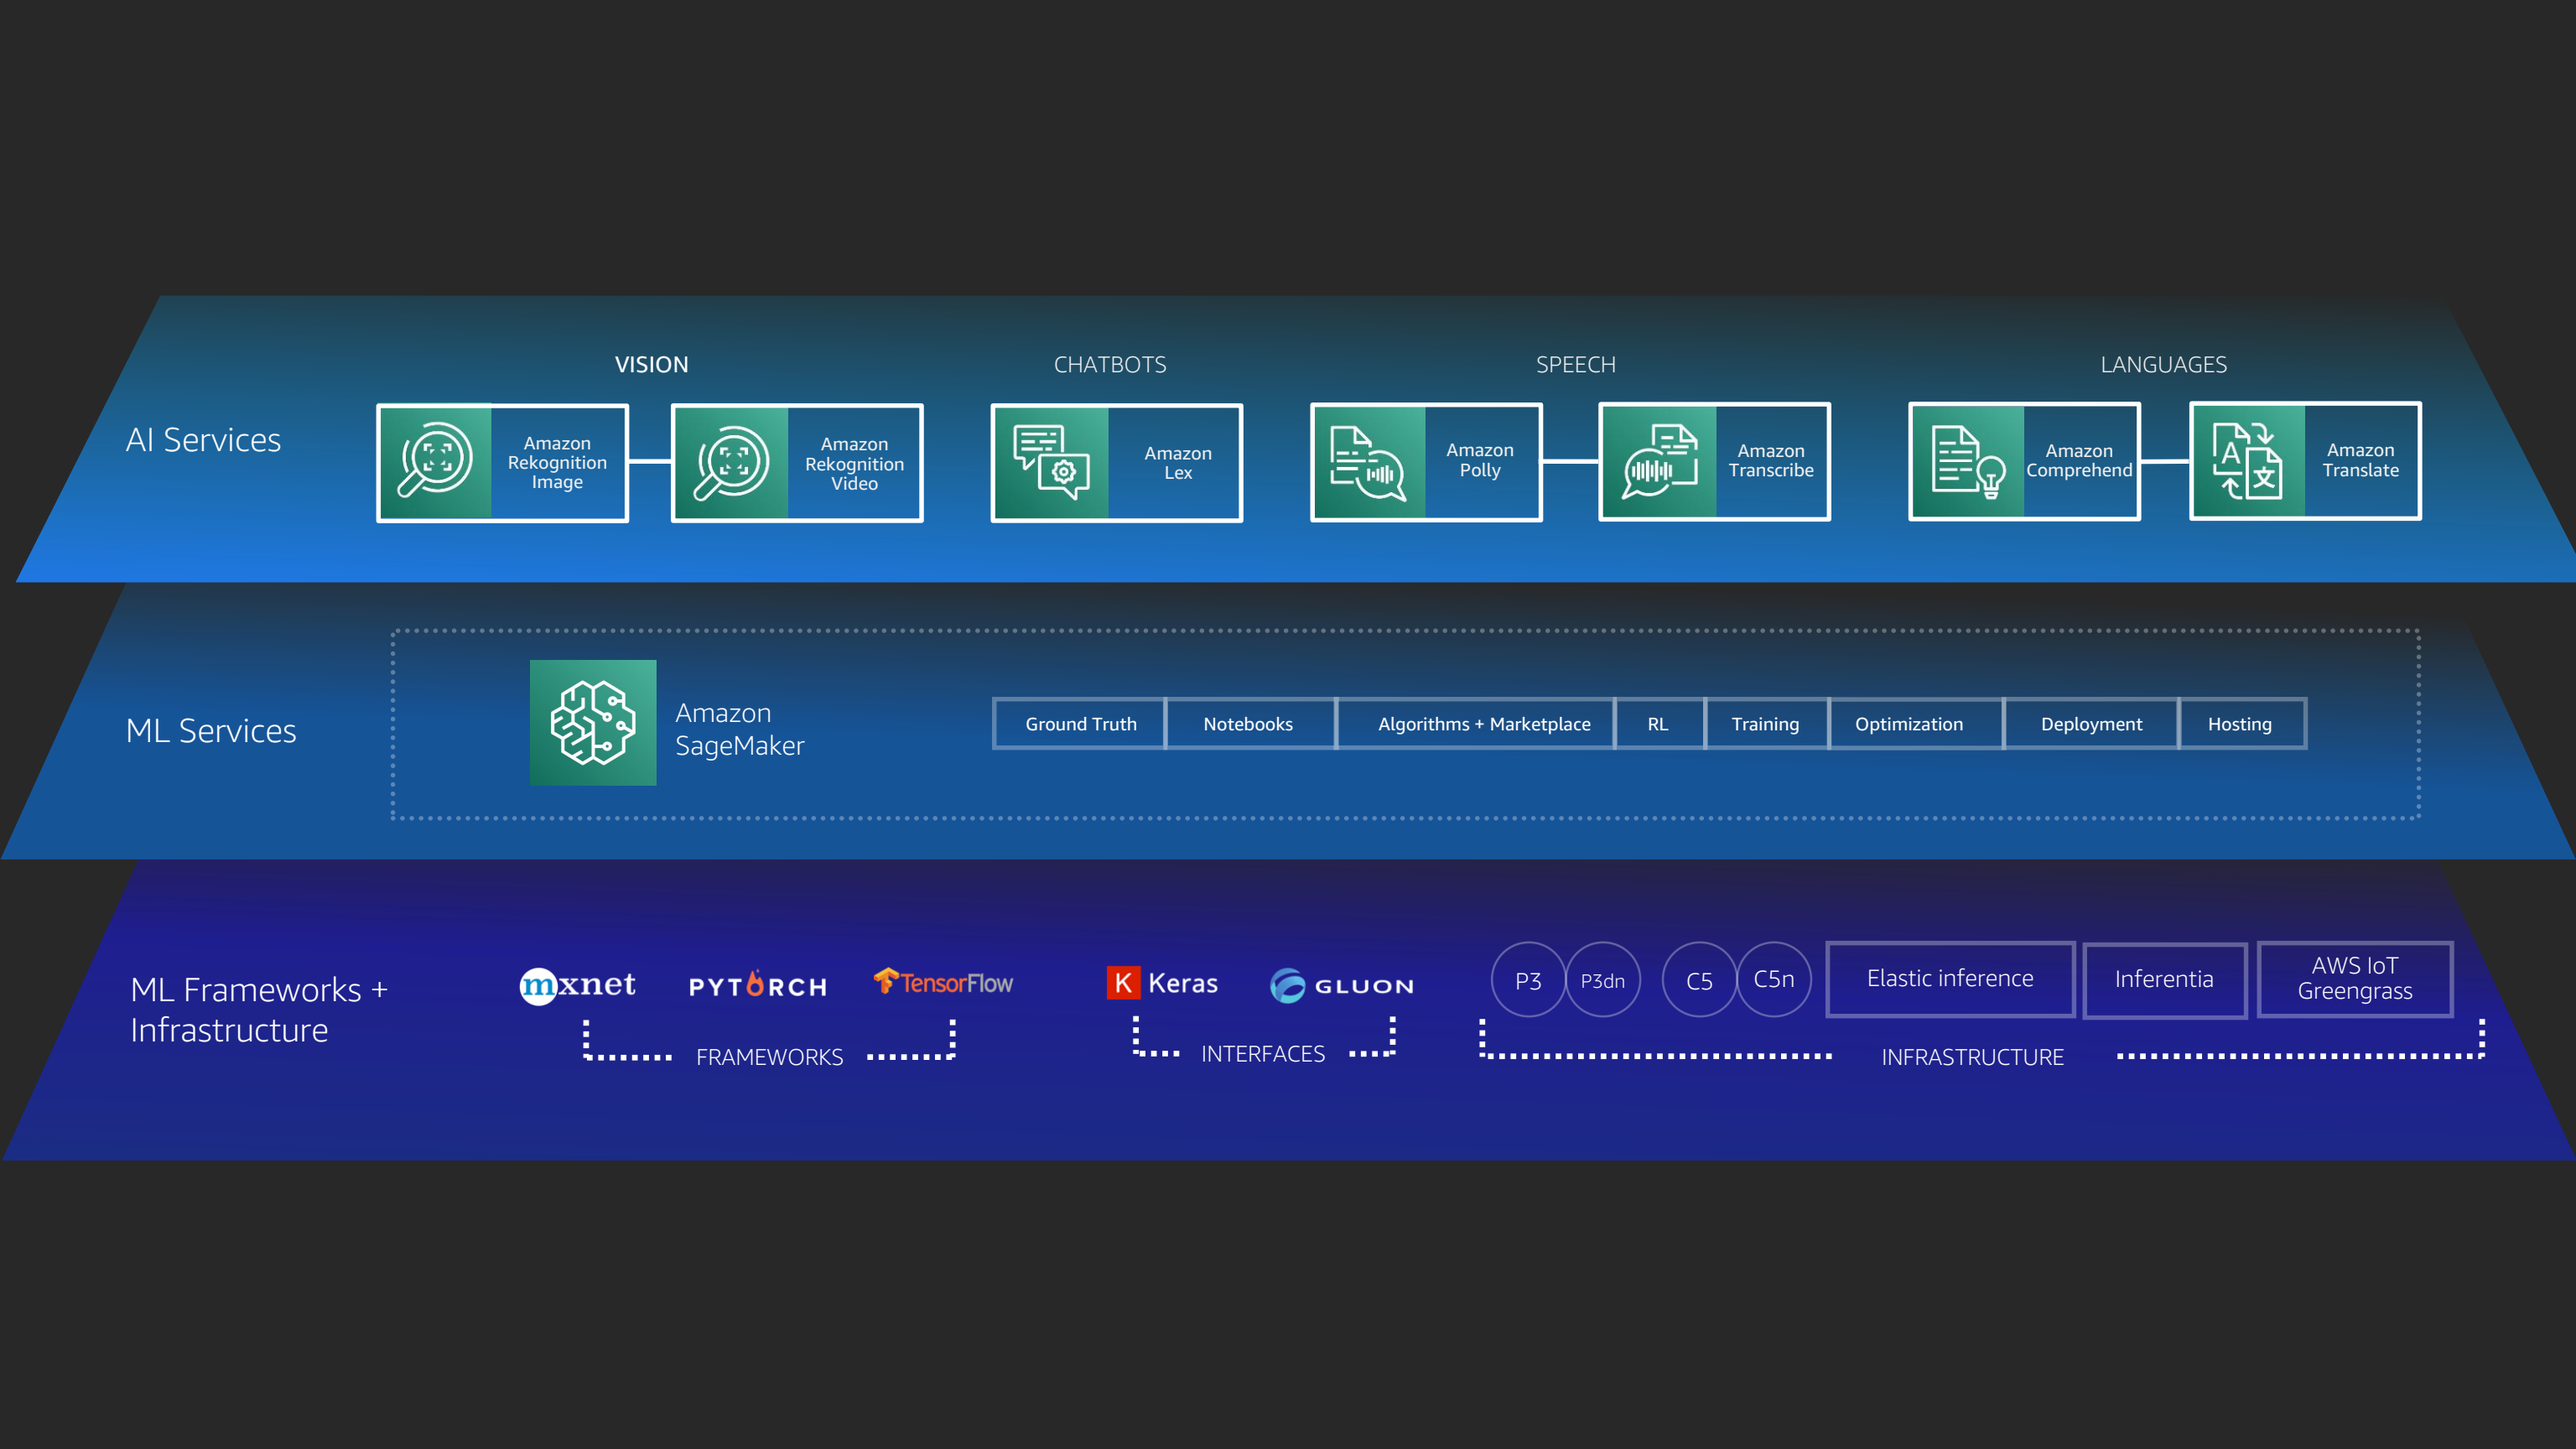
© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

AIoT – Machine Learning at the Edge

Train in the cloud and infer at the edge

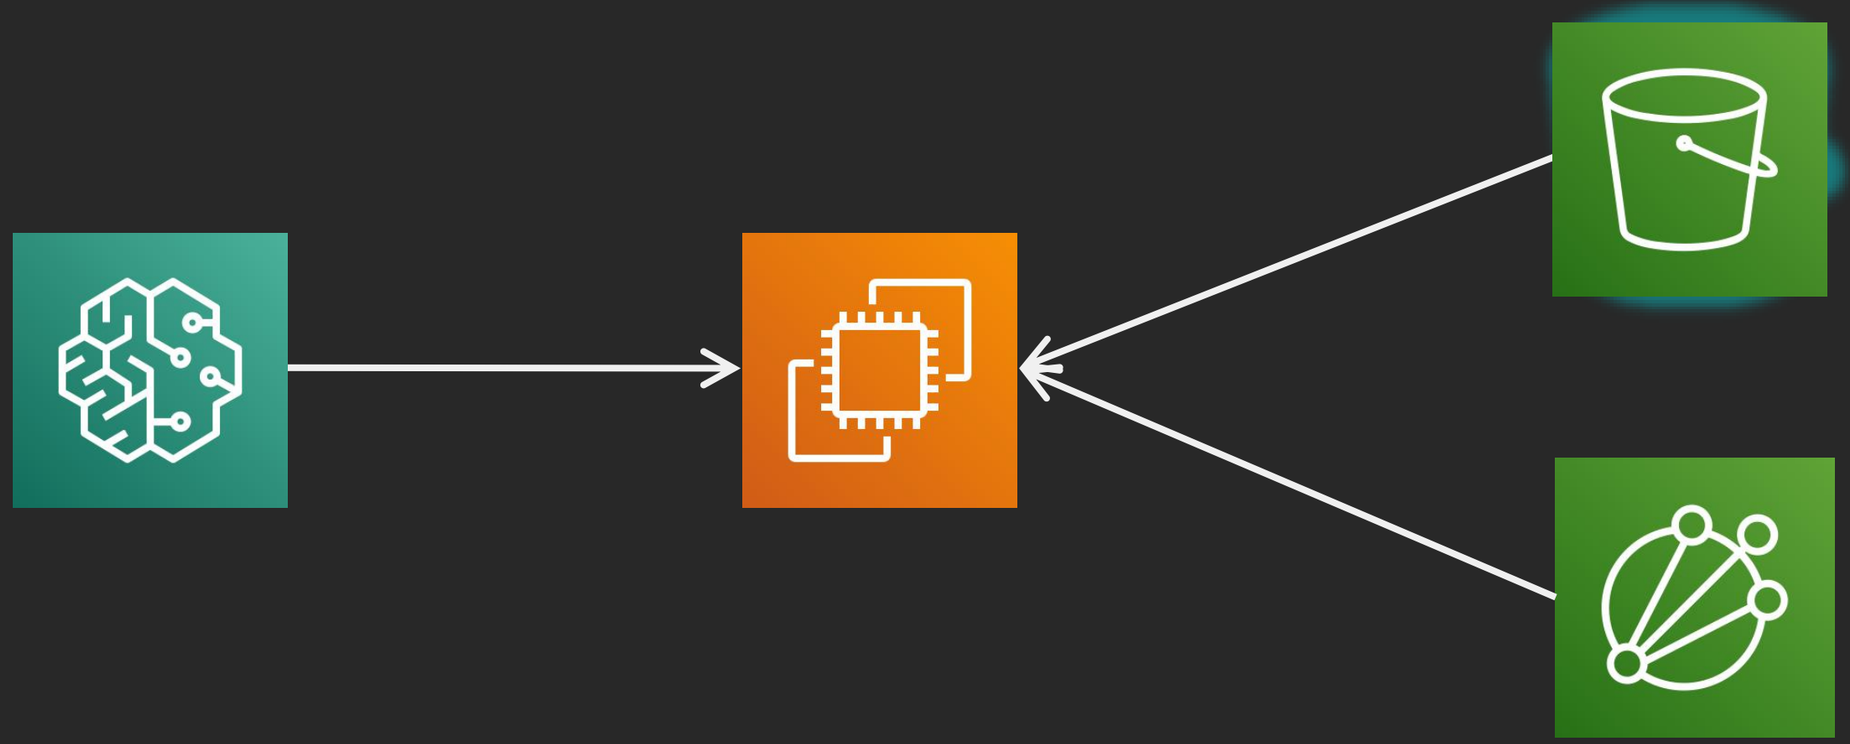



AI Services
VISION
Amazon Rekognition Image
Amazon Rekognition Video
CHATBOTS
Amazon Lex
SPEECH
Amazon Polly
Amazon Transcribe
LANGUAGES
Amazon Comprehend
Amazon Translate
ML Services
Amazon SageMaker
Ground Truth
Notebooks
Algorithms + Marketplace
RL
Training
Optimization
Deployment
Hosting
ML Frameworks + Infrastructure
mxnet
PYTORCH
TensorFlow
FRAMEWORKS
Keras
GLUON
INTERFACES
P3
P3dn
C5
C5n
Elastic inference
Inferentia
AWS IoT Greengrass
INFRASTRUCTURE

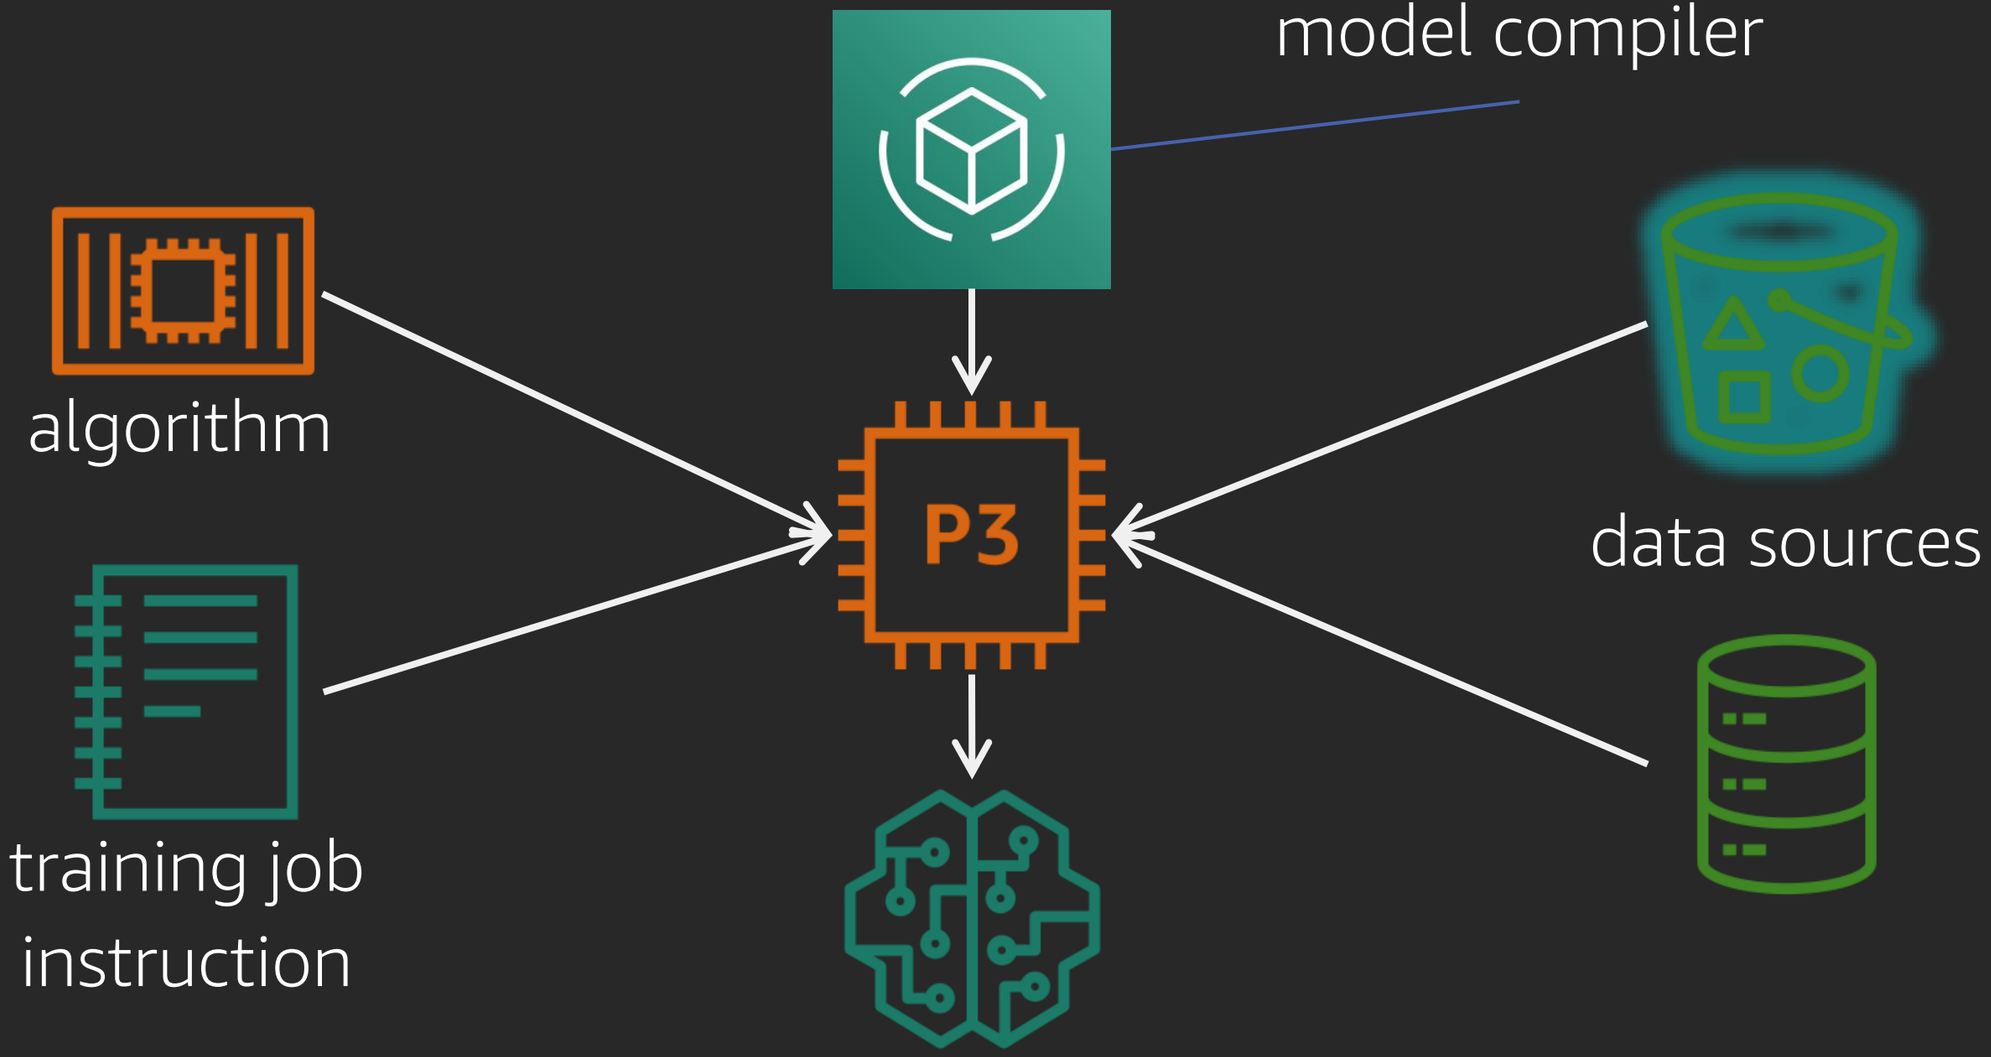
Deep learning frameworks and toolchains



Feeding the art: data sets



Training job: compilation



Staging



Enabling machine learning at the edge: Inferring at the edge

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

AIoT – Machine Learning at the Edge

Train in the cloud and infer at the edge



Data aggregation, enrichment, cleansing, time series, and model config

AWS IoT Analytics

Amazon S3

Amazon SageMaker Ground Truth

Amazon EMR

Machine learning and model generation



Apache MXNet on AWS

Amazon SageMaker

AWS Deep Learning AMIs

Amazon EC2



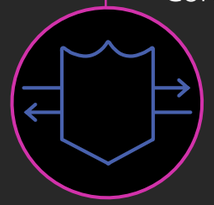
Intelligence and outcomes

AWS IoT Core

Amazon Kinesis

AWS Snowball

AWS Snowmobile



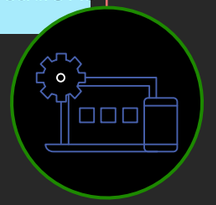
Data transport and routing

Amazon FreeRTOS

AWS IoT Greengrass

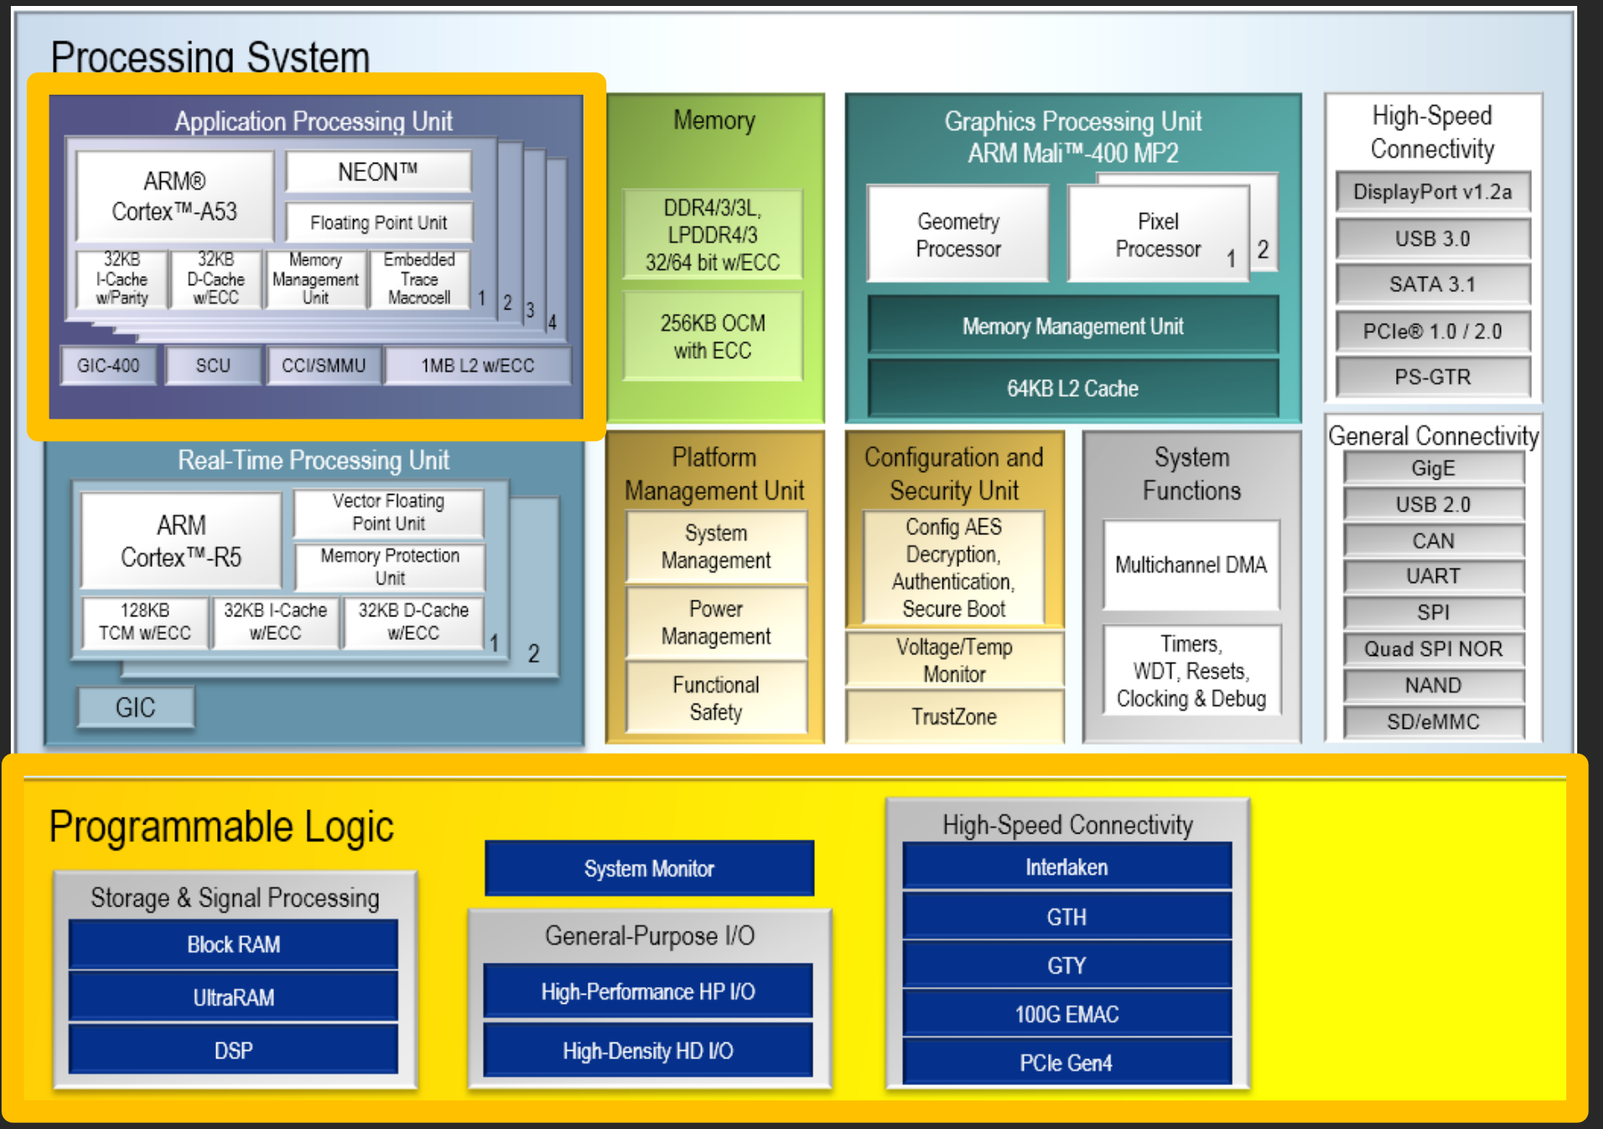
AWS IoT SiteWise

bespoke applications



Data collection and model inference

Mixed criticality system – block diagram



Mixed criticality system – software placement

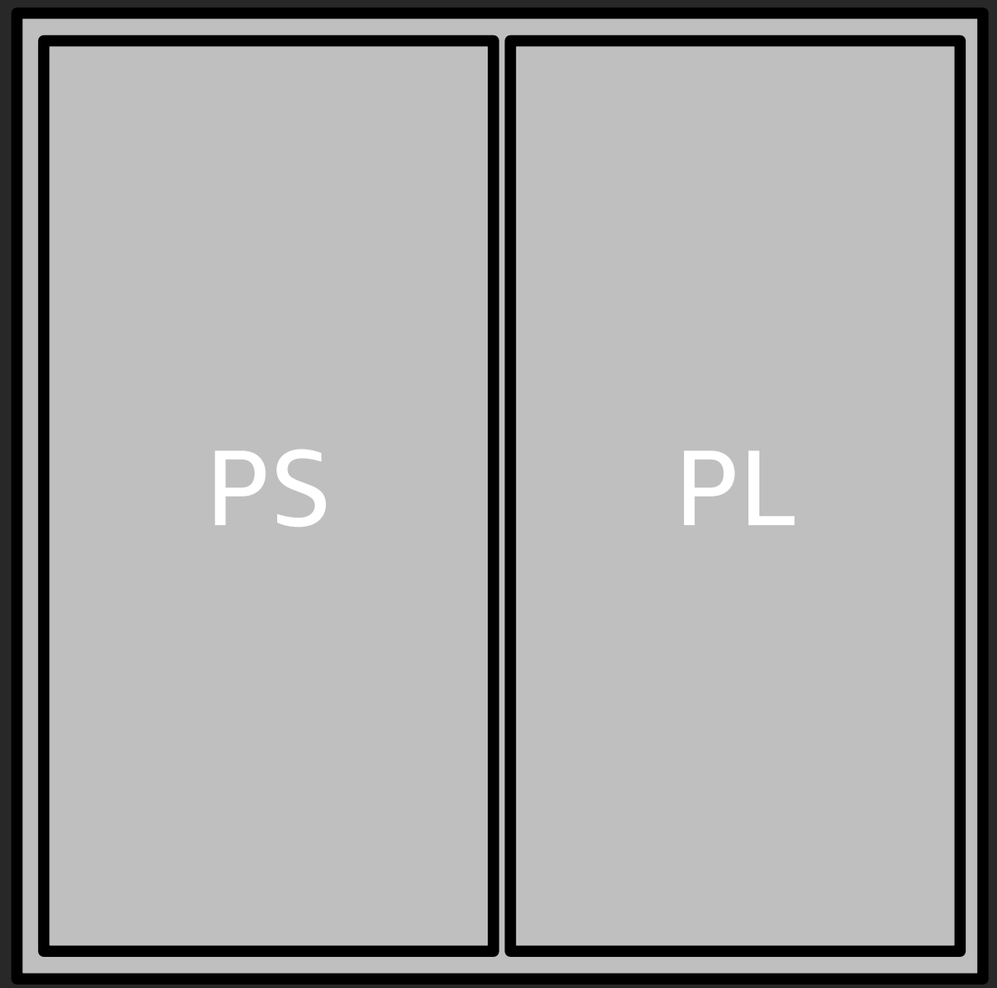


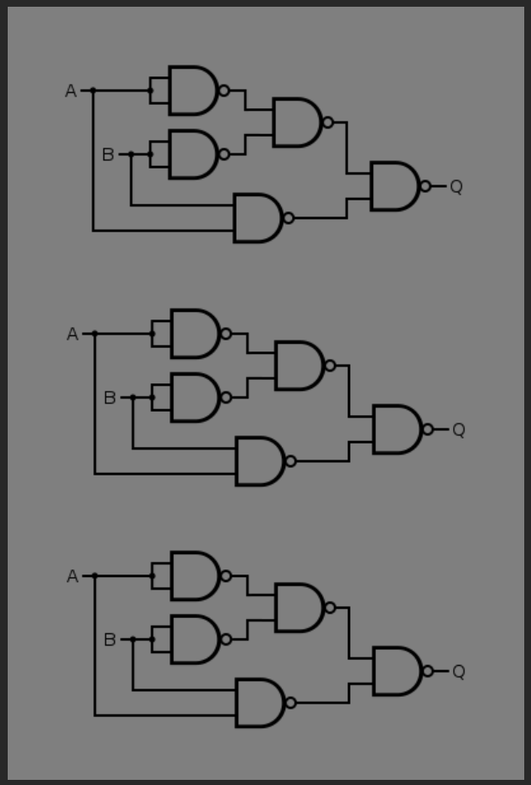
PetaLinux

Debian

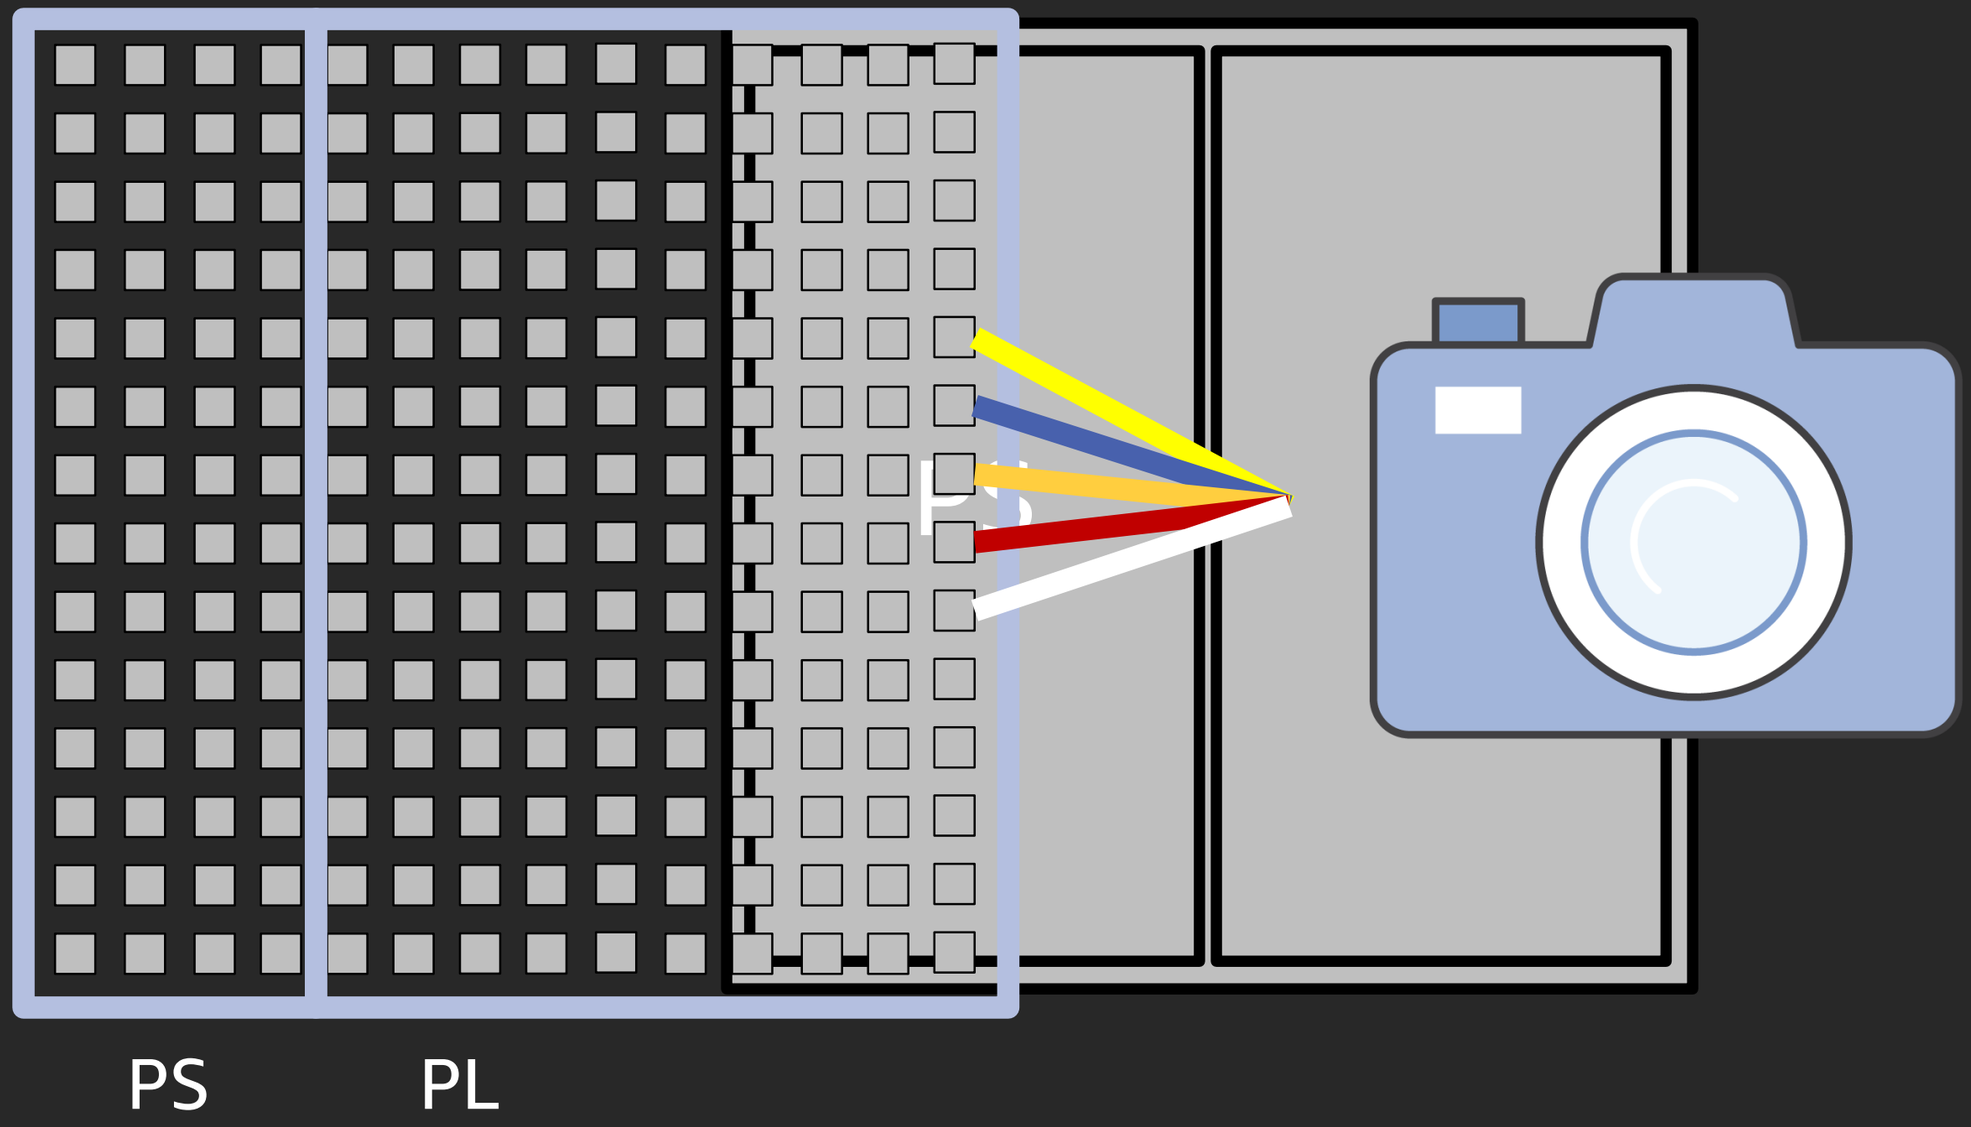
Ubuntu



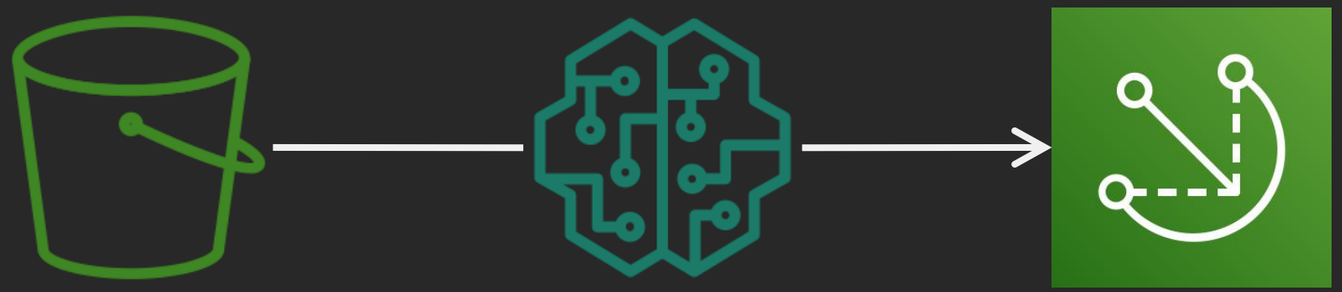




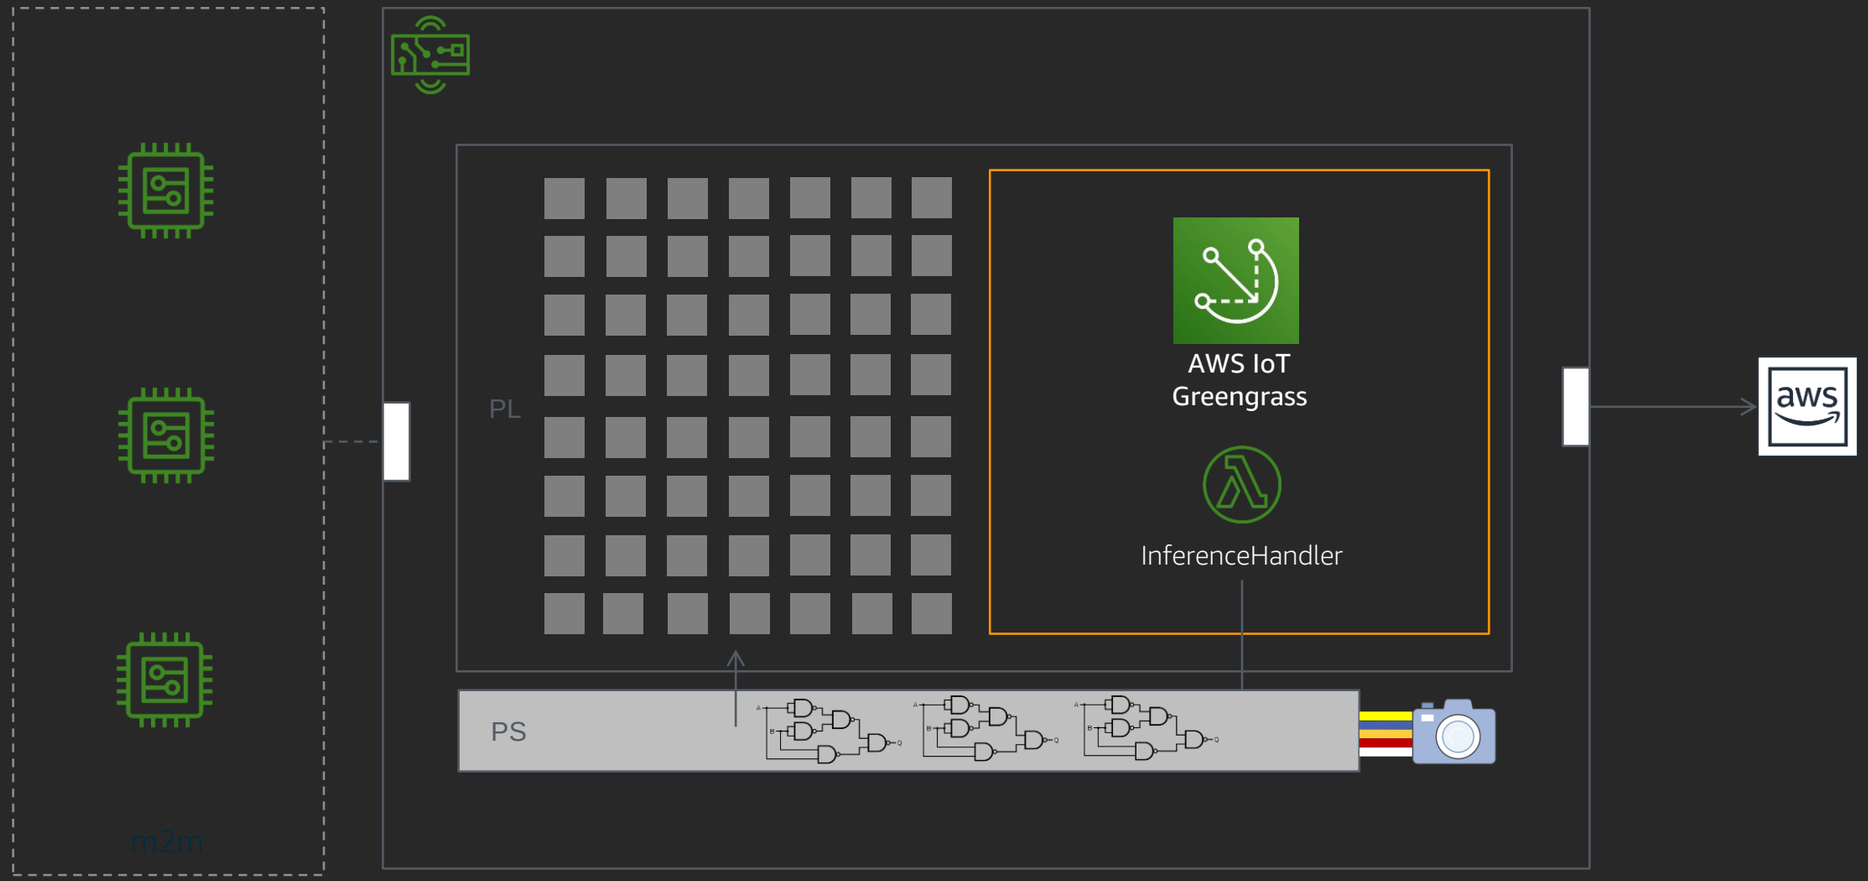
Mixed criticality system – pin wiring



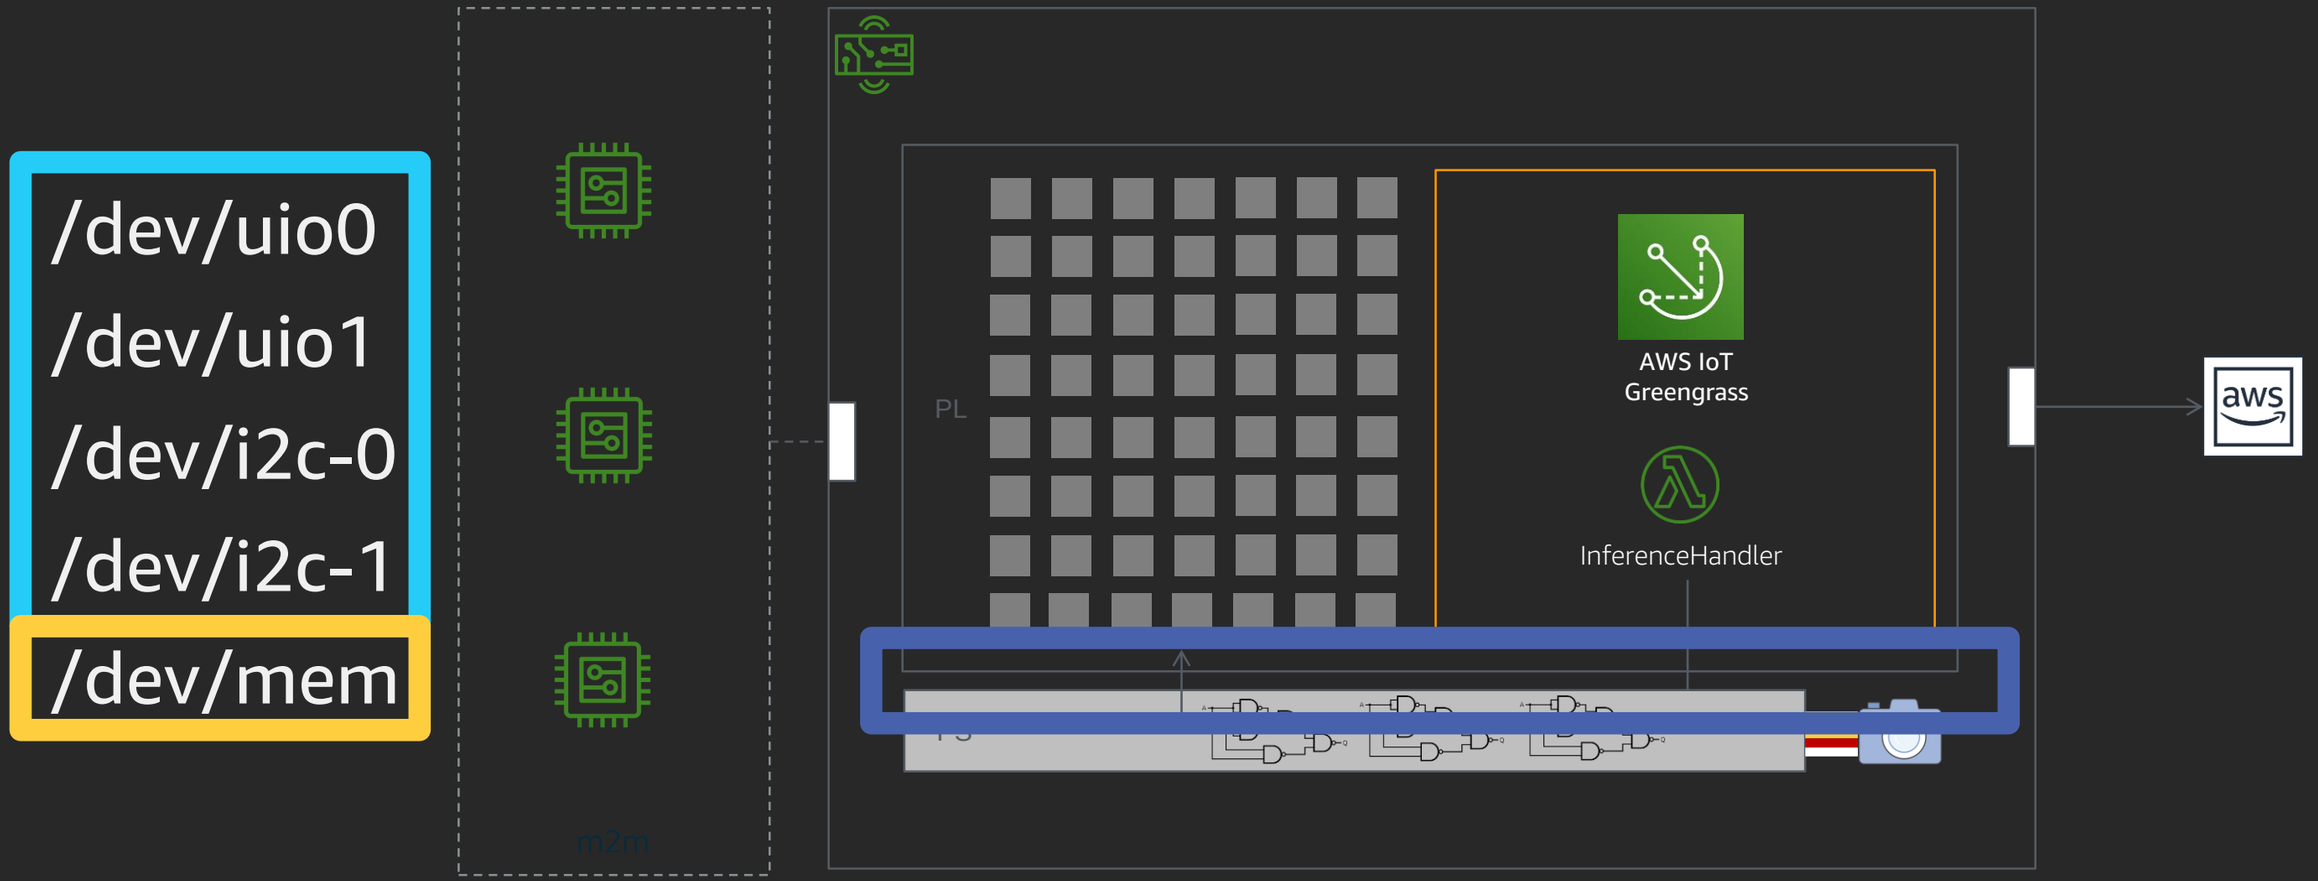
Updating the model



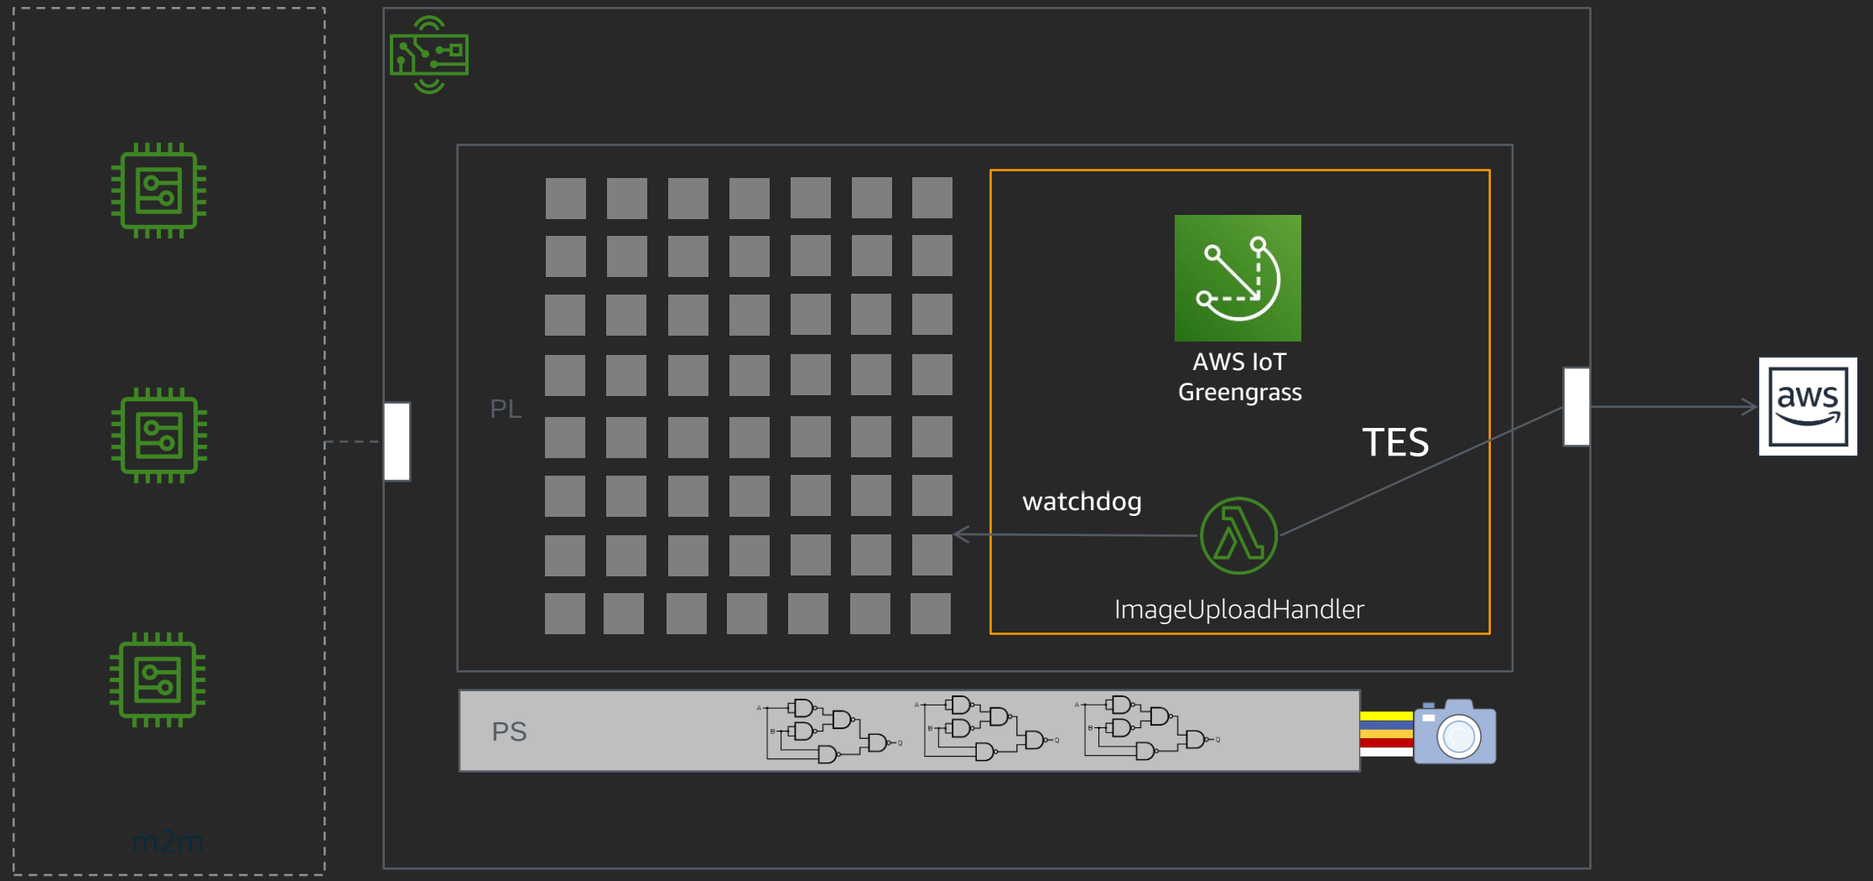
Performing inference and staging classified images



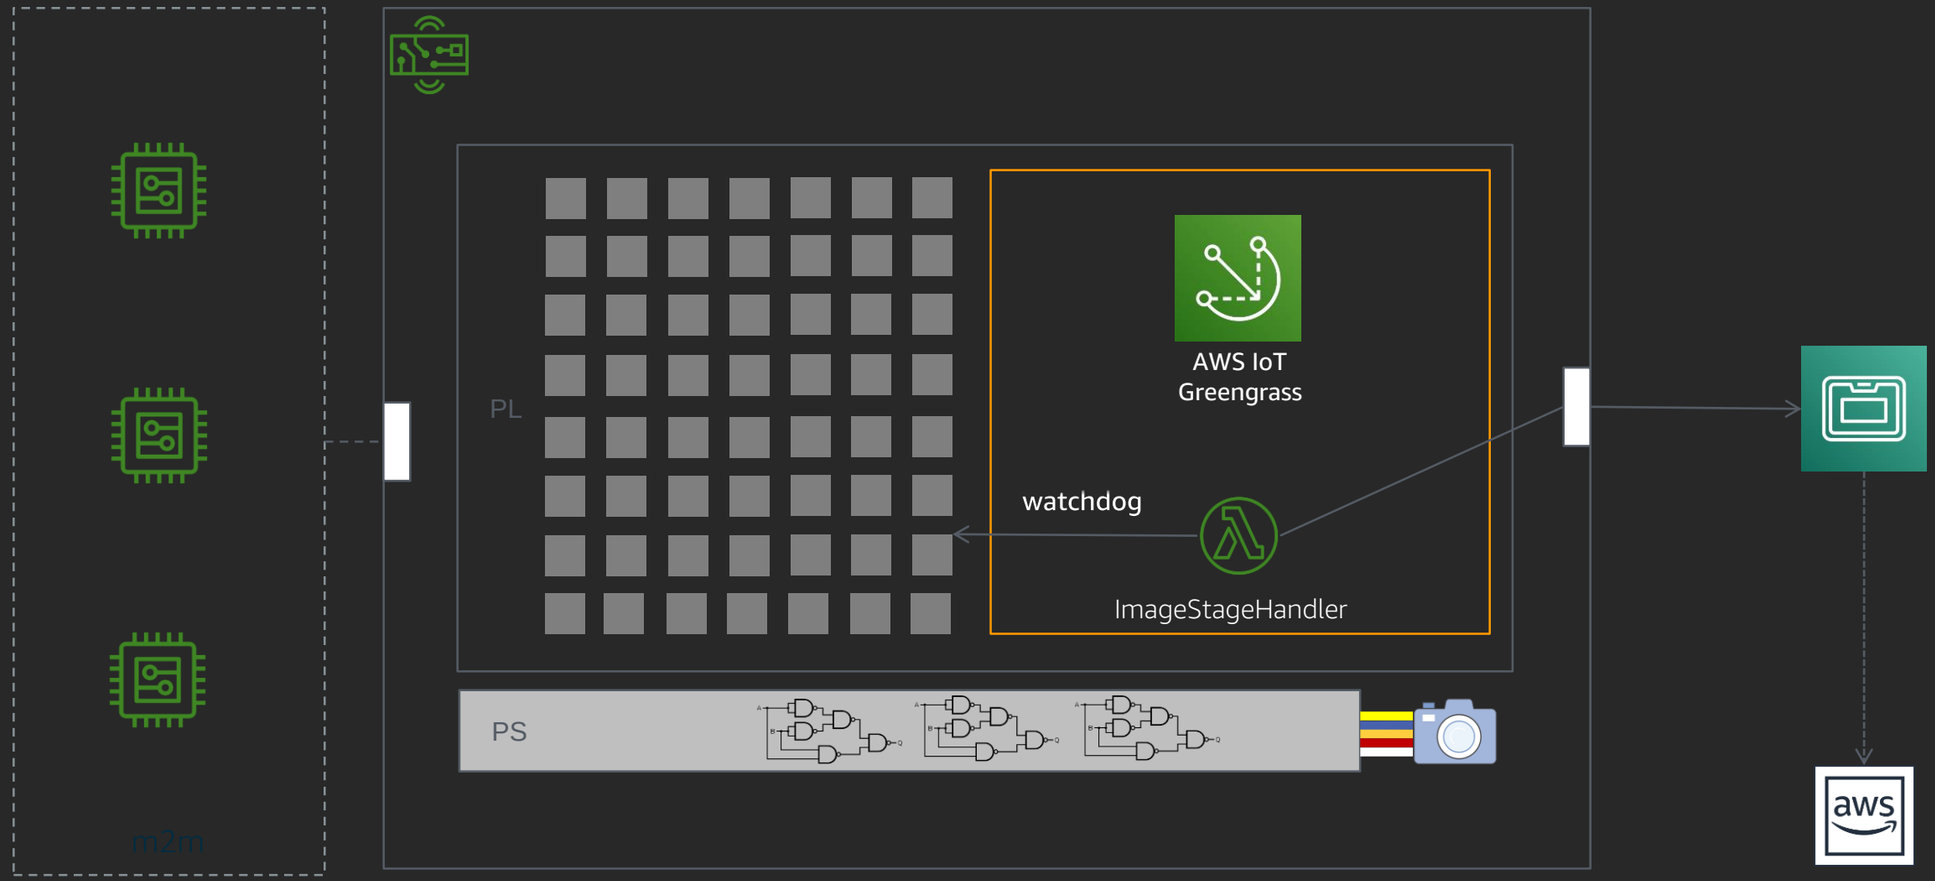
Device level dependencies



Reporting inference results



Staging data for building up the data lake



Recap

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Recap



Data aggregation, enrichment, cleansing, time series, and model config

AWS IoT Analytics

Amazon S3

Amazon SageMaker Ground Truth

Amazon EMR



Intelligence and outcomes

Machine learning and model generation



Apache MXNet on AWS

Amazon SageMaker

AWS Deep Learning AMIs

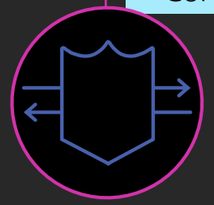
Amazon EC2

AWS IoT Core

Amazon Kinesis

AWS Snowball

AWS Snowmobile



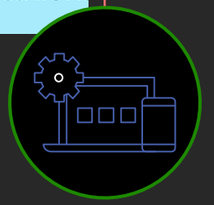
Data transport and routing

Amazon FreeRTOS

AWS IoT Greengrass

AWS IoT SiteWise

bespoke applications



Data collection and model inference

Thank you!

Richard Elberger

GitHub: rpcme
LinkedIn: https://www.linkedin.com/in/richardelberger/

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws



Please complete the session survey in the mobile app.

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.