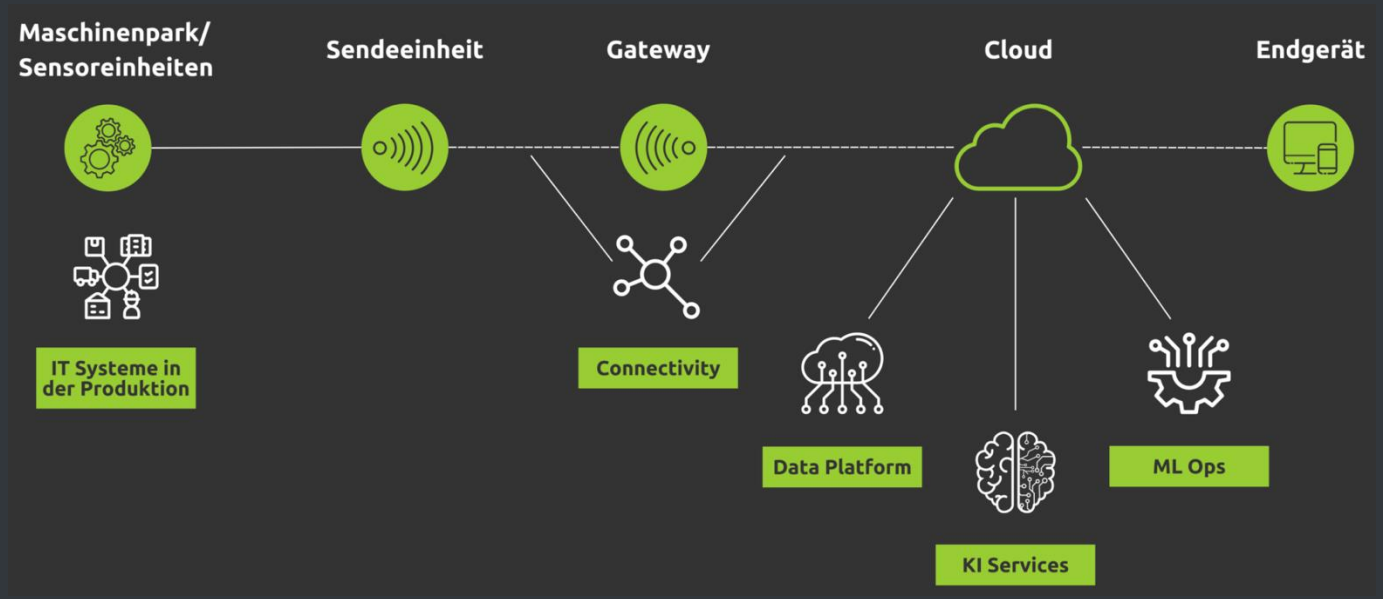

CoffAI – Technical Deep Dive

How to get from Sensor data to an AI-based realtime application

embeddedworld
Exhibition&Conference

esentri

PRESENTATION
AGENDA

- **Who are we and what are we doing?**
- Background information: The coffAI Project
- Technical Deep Dive
  - From the sensors to the Cloud
  - Cloud Architecture
- Let's talk about costs
  - Required knowledge and implementation time
  - Infrastructure costs
- Conclusion

WHO ARE WE AND WHAT ARE WE DOING?

esentri – Data & AI - Industrial Analytics & IoT

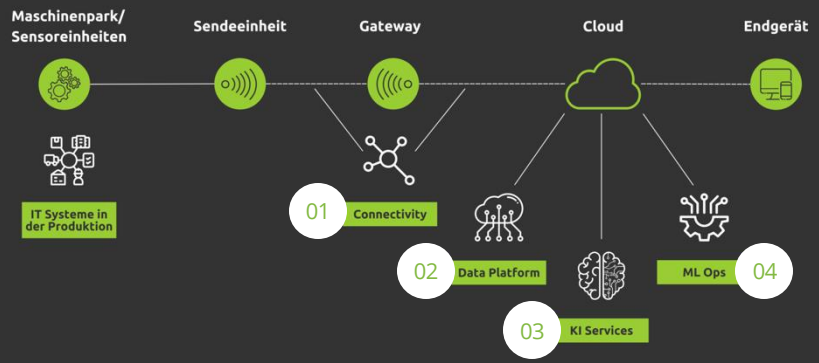

IN-DEPTH SOURCES

Conceptual view: Der Weg vom Sensor zum Digitalen Produkt

Practical example: coffAI - Klassifizierung von Vibrationsdaten

01

CONNECTIVITY & EDGE

Cloud connection of IoT devices & edge deployments and their control from the cloud

02

DATA PLATTFORM

Building data platforms from raw sensor values to highly aggregated KPIs

03

KI SERVICES

Implementation of machine learning services with a focus on vibration/time series

04

ML OPS

Secure operation of ML-based applications and ensuring maintainability, extensibility and security

PRESENTATION
AGENDA

- Who are we and what are we doing?
- **Background information: The coffAI Project**
- Technical Deep Dive
  - From the sensors to the Cloud
  - Cloud Architecture
- Let's talk about costs
  - Required knowledge and implementation time
  - Infrastructure costs
- Conclusion

BACKGROUND INFORMATION: THE COFFAI PROJECT

Real-world monitoring solution at KSB

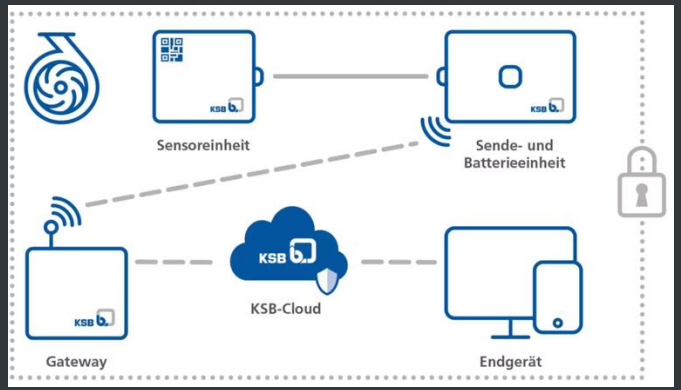

- Plug & play monitoring solution for pumps (https://www.ksb.com/de-ch/guard)

Customer Services:

- (Near) Real-time dashboards
- digital asset management
- machine learning-based services such as condition monitoring and anomaly detection

BACKGROUND INFORMATION: THE COFFAI PROJECT

CoffAI project at esentri AG

BACKGROUND INFORMATION: THE COFFAI PROJECT

CoffAI project at esentri AG

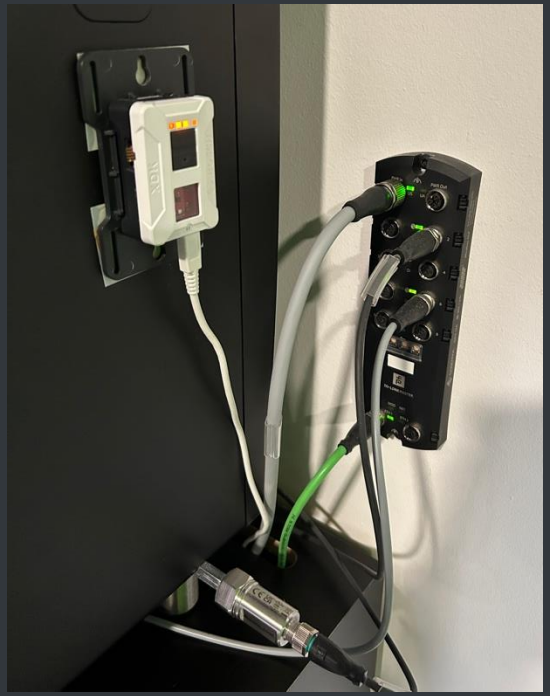

Focused Parts:

- Real-time dashboard
- machine learning-based service to predict coffee type based on vibration and magnetic field
- IaC & CI/CD

Ignored Parts:

- Encoding / decoding
- Proper Data Platform
- ML Ops
- Staging

PRESENTATION
AGENDA

- Who are we and what are we doing?
- Background information: The coffAI Project
- **Technical Deep Dive**
  - From the sensors to the Cloud
  - Cloud Architecture
- Let's talk about costs
  - Required knowledge and implementation time
  - Infrastructure costs
- Conclusion

TECHNICAL DEEP DIVE

From the sensors to the Cloud

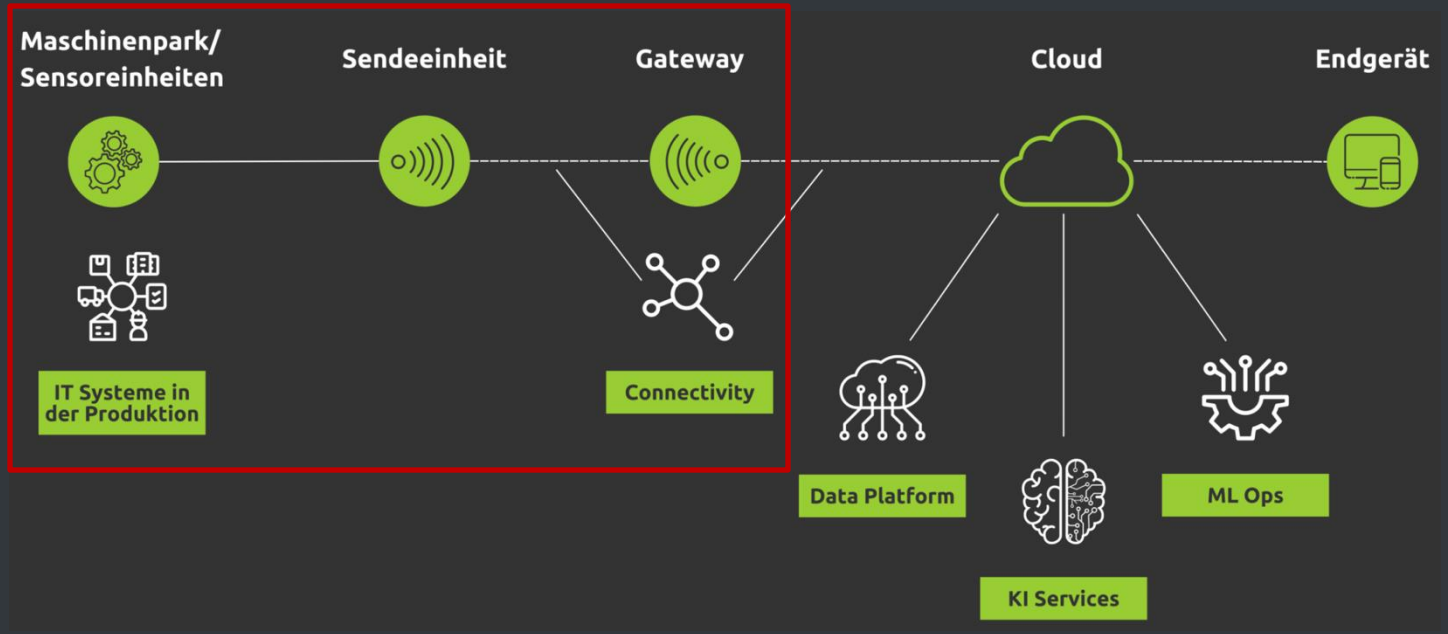

TECHNICAL DEEP DIVE

From the sensors to the Cloud

MQTT

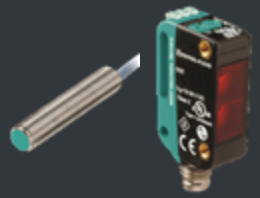

IO-Link

Netzwerk-Modul

CloudRail.box Gateway Industrie-PC

Rechenleistung

Edge Device

OPC UA

MQTT

TECHNICAL DEEP DIVE

From the sensors to the Cloud

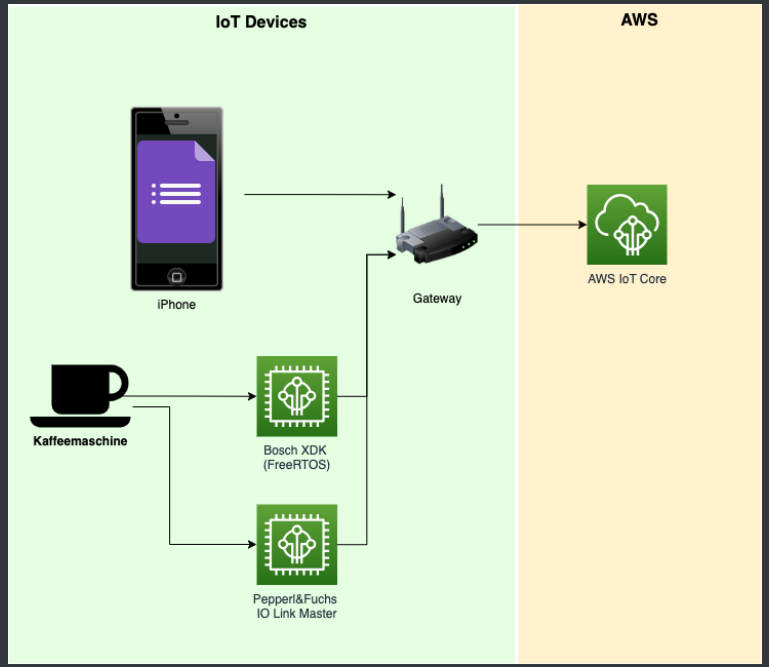

First Idea:

- Run a MQTT Broker on the Gateway and do all the real-time tasks there. Just take the long-term statistics from the cloud and safe historic data in the cloud for model training.
- But: In AWS IoT Core you have a MQTT broker running out of the box (serverless). That means, with the right certificates you can also stream directly to the cloud!

Crucial question: What data aggregation level do we really need for AI based classification?

- After some tests we realized that for this case we would not need the raw signals but we could just work with RMS and Energy values (10Hz, sampling rate 2kHz)
- The idea was born to do everything besides RMS and Energy calculation in the cloud.

TECHNICAL DEEP DIVE

Cloud Architecture

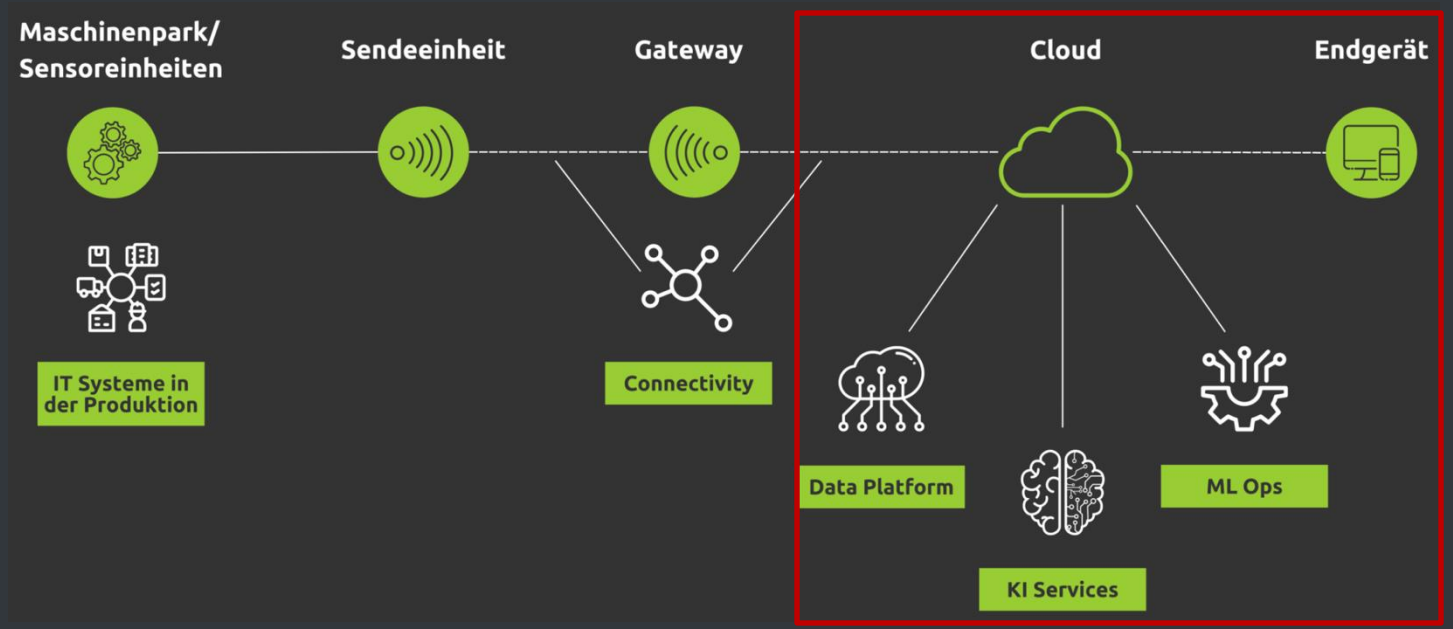

TECHNICAL DEEP DIVE

Cloud Architecture

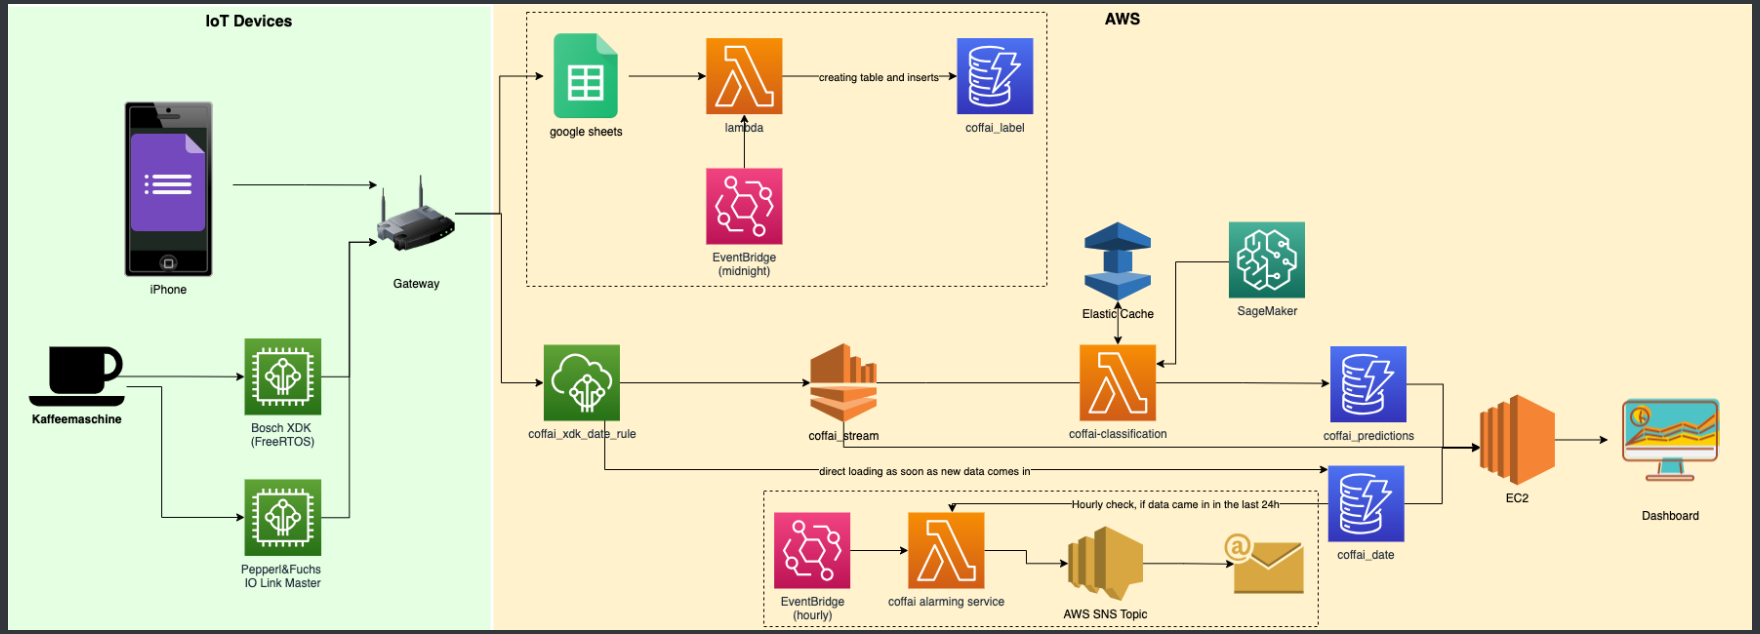

- Highly scalable Architecture
- Whole inferencing part could be brought on edge. But: with this architecture we have latency of below 1 second!

PRESENTATION
AGENDA

- Who are we and what are we doing?
- Background information: The coffAI Project
- Technical Deep Dive
  - From the sensors to the Cloud
  - Cloud Architecture
- **Let's talk about costs**
  - Required knowledge and implementation time
  - Infrastructure costs
- Conclusion

LET'S TALK ABOUT COSTS

Required Knowledge and implementation time

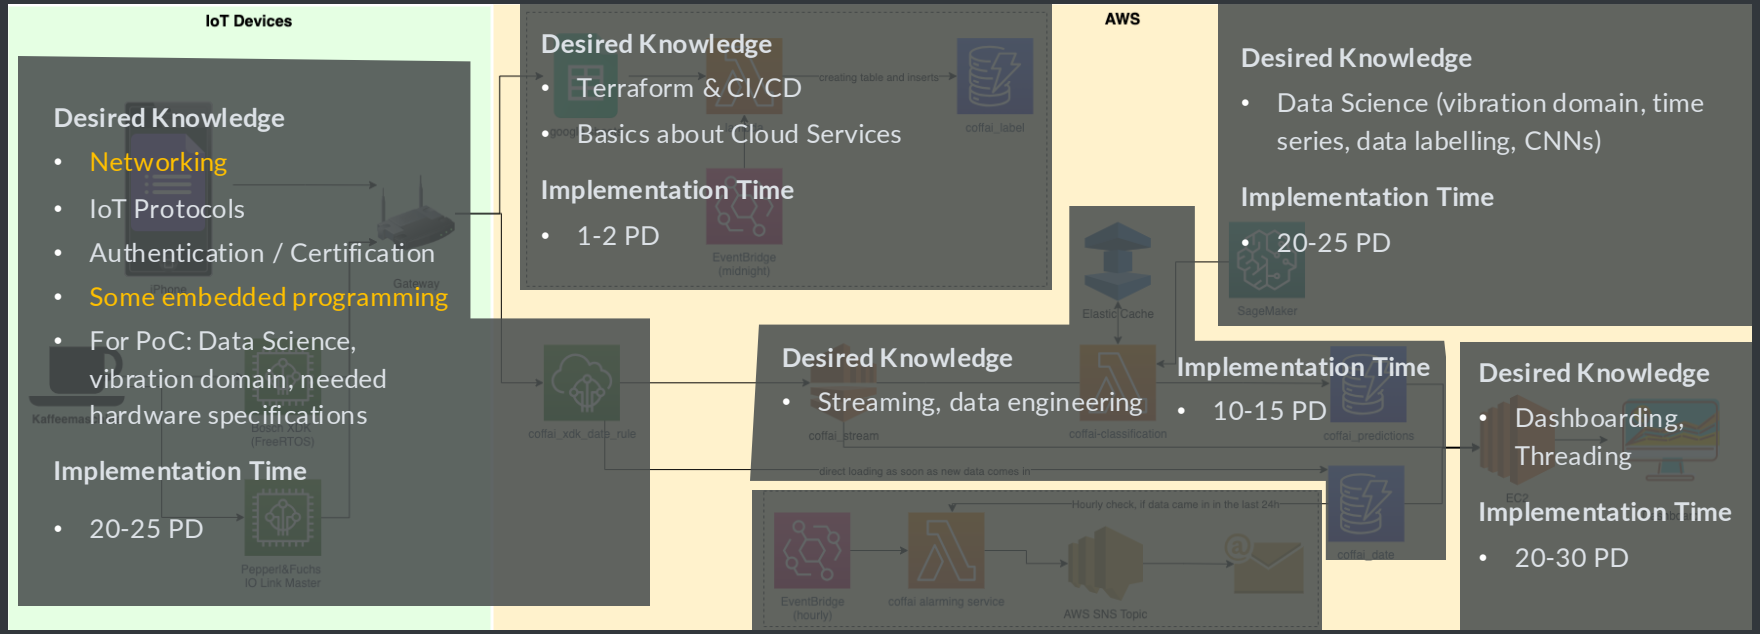

-> 71 – 97 PD

LET'S TALK ABOUT COSTS

Infrastructure Cost

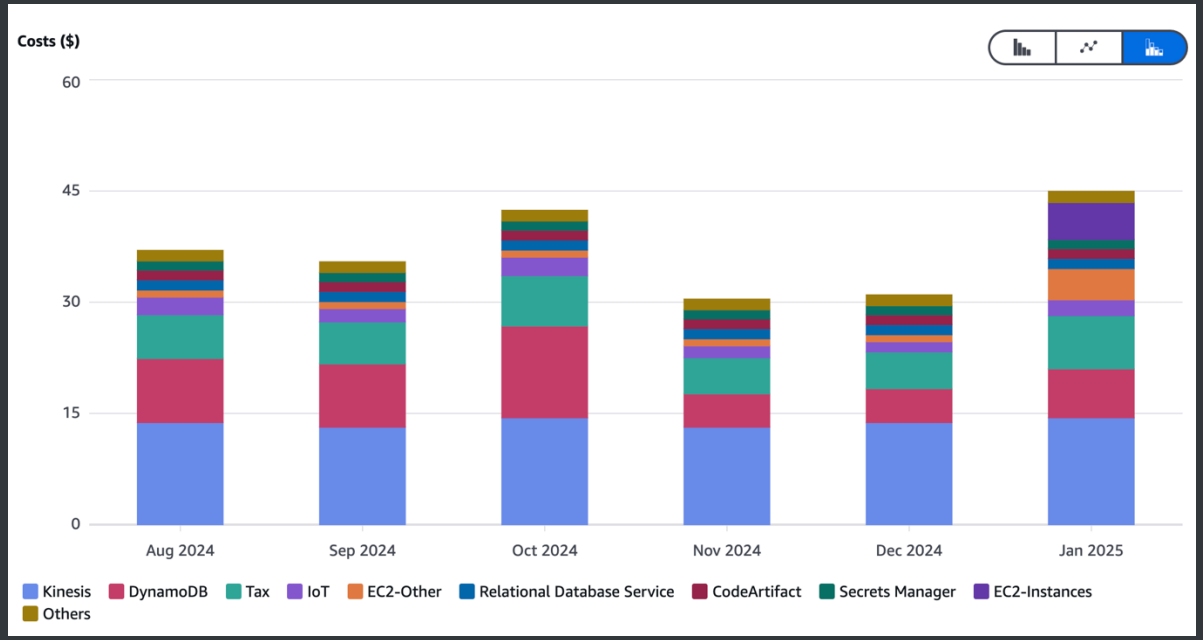

Nov / Dez 2024 as representative months of pure infrastructure cost

Cost Optimizations:

- Usage of Serverless Services whenever it is possible (Lambda functions, dynamoDB)
- Shot down of the streaming parts, whenever there is no office Time
- Run the dashboard on a local machine in the office instead of an EC2 Instance
- 10 datapoints per second (only if there is some vibration)

LET‘S TALK ABOUT COSTS

Infrastructure Cost

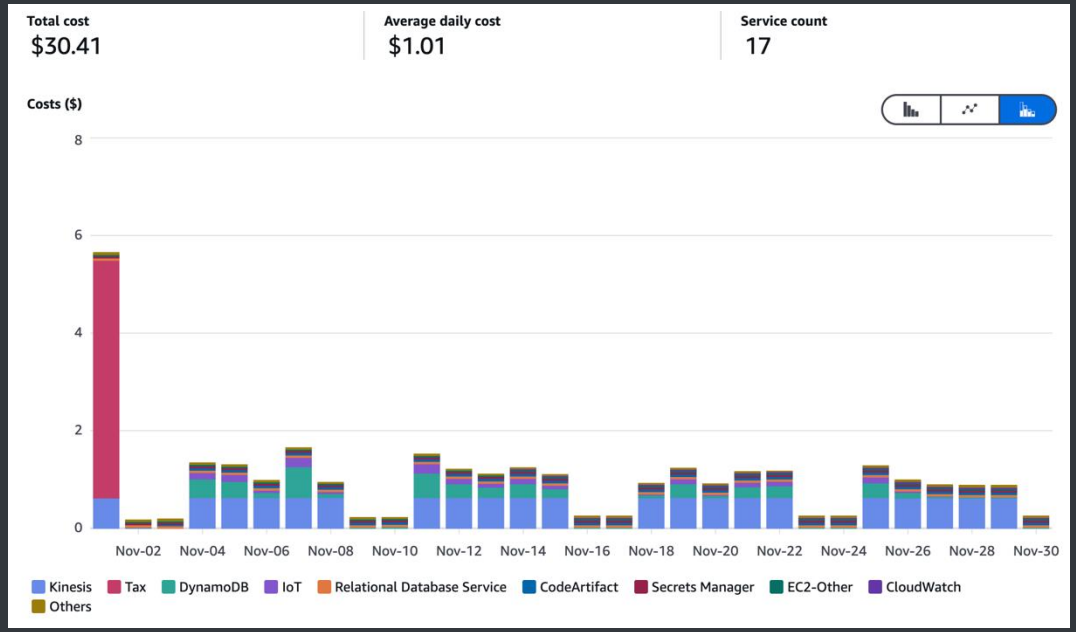

PRESENTATION
AGENDA

- Who are we and what are we doing?
- Background information: The coffAI Project
- Technical Deep Dive
  - From the sensors to the Cloud
  - Cloud Architecture
- Let's talk about costs
  - Required knowledge and implementation time
  - Infrastructure costs
- **Conclusion**

CONCLUSION

What should you take away from this lecture?

01 Make sure that Domain and Data Science knowledge is there from the very beginning. No „let's first collect data"

02 Realtime has its price. How much real-time does your service really need?

- CoffAI: Realtime dashboard & Coffee classification necessary. But for anomaly detection near real-time would be completely sufficient
- KSB: Near real-time dashboard for the customer is sufficient. Response topics have to be handled in real-time

03 When it comes to vibration data, invest in data compression and message packaging. Especially for predictive Maintenance applications, raw data is very valuable

04 For the PoC phase: Take some not production ready sensor that is highly configurable in order to find out what data resolution you really need

Have I aroused your interest?
Let‘s connect!

Simon Kneller
esentri AG
Head of Industrial Analytics & IoT

+49 160 967 648 04

simon.kneller@esentri.com

www.linkedin.com/in/simon-kneller/