boon

Deploying Amber: Edge-Based, Real-Time Anomaly Detection on Intel CPUs

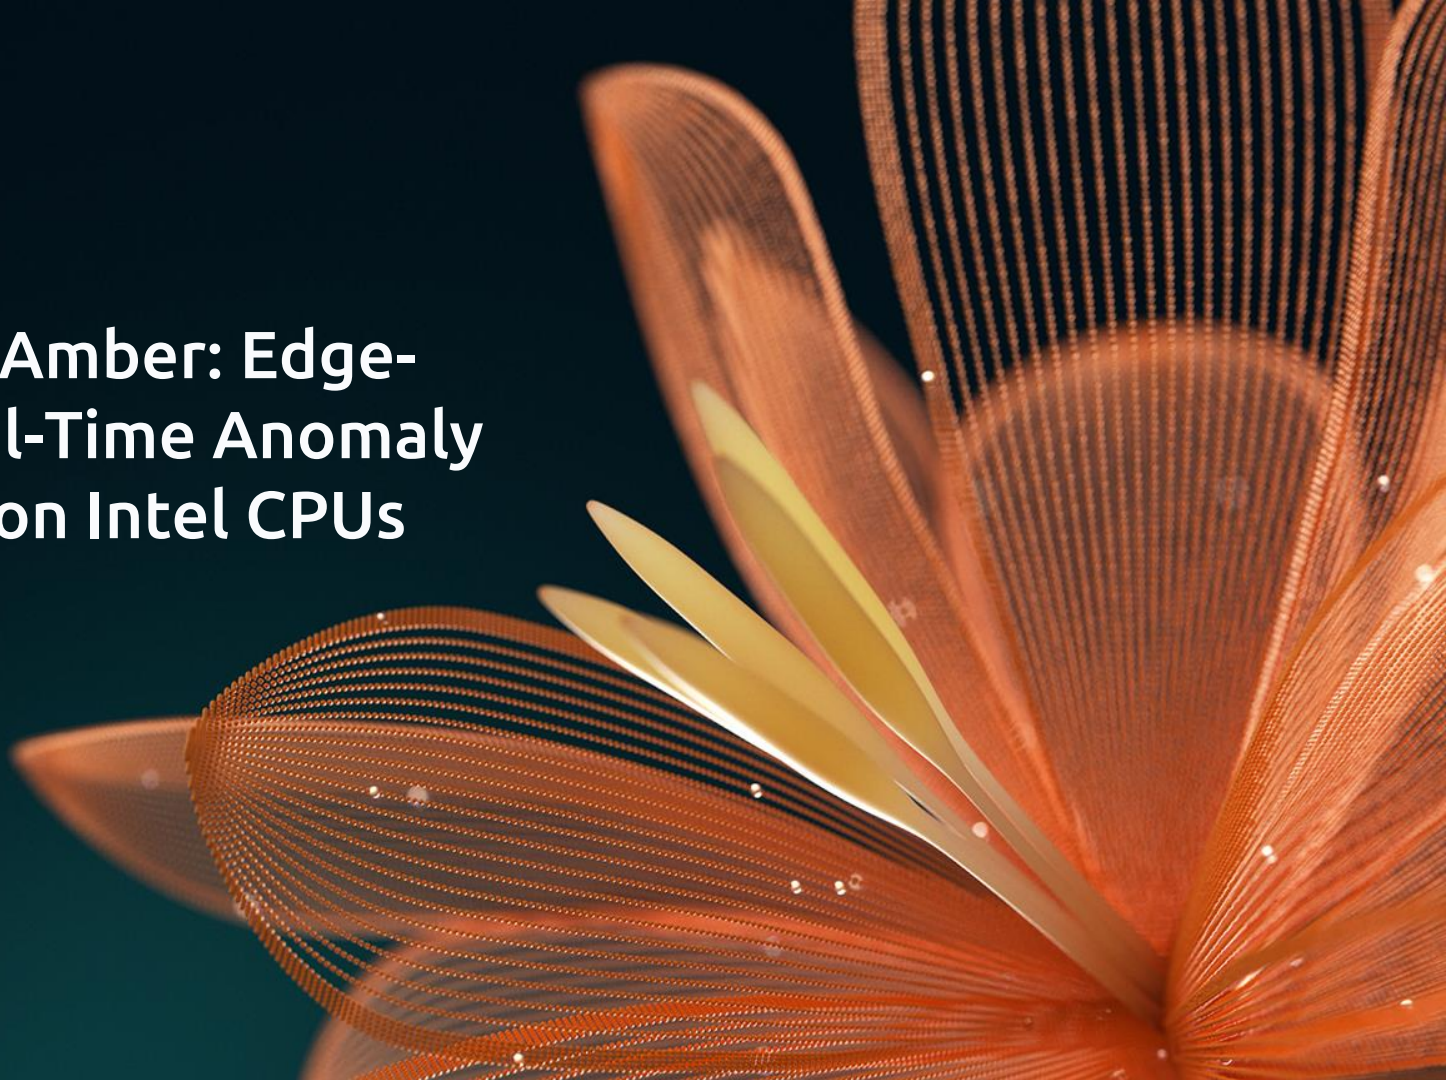

Dr. Brian Turnquist

Chief Technology Officer

Dr. Brian Turnquist has over 25 years of experience in applying traditional analytic techniques and machine learning in both academic and industrial settings. His academic research in neuroscience, focusing on biosignals, has led to 15 publications and collaborations with major universities globally. Since 2016, Dr. Turnquist has been with Boon Logic, where he applies advanced machine learning techniques to industrial anomaly detection in asset telemetry signals and video streams.

Dr. Rodney Dockter

Director of Engineering

Dr. Rodney Dockter has an extensive background in machine learning algorithms, medical robotics, and industrial applications. At Boon Logic, he leads deployment efforts and brings significant experience from both academia and industry. His research in surgical robotics and machine learning complements his work in developing algorithms, hardware integration, and the deployment of advanced industrial automation solutions.

Ronan Burns

Lead Solution Architect

Ronan Burns is a seasoned expert in data science and machine learning, with a strong focus on industrial applications. With a Master's degree in Data Science, Ronan brings a deep understanding of AI integration and solution architecture to his work. He has successfully driven the implementation of machine learning models in various industrial environments, helping enterprise customers solve complex data-driven problems.

Agenda

- **Boon Logic:**

  A New Normal In Anomaly Detection
- **Nano:**

  A CPU-Native Core Technology
- **Amber:**

  Applying the Boon Nano to Predictive Analytics
- **Deployments:**

  Solving Challenges at the Far Edge, Powered by Intel
- **Questions:**

  Let's Dive In


boon
A new normal in Anomaly detection

When to Choose Boon AI Technology…

Type of Artificial Intelligence	Description	Common Applications
Unsupervised ML	Train models from a “blank slate” typically based on clustering in *n*-dimensional space	Data segmentation (aka clustering), anomaly detection
Semi-supervised ML	Start with a trained UML model and then add meaningful labels to the generated clusters	Classifier
Supervised ML	Train from a large SME-pre-labeled data set and build a classifier to try to replicate the SME labeling	Classifier
Reinforcement Learning	Train using a specific goal (reward function)	Self-driving cars
Generative AI	Start with a very large UML model of a data corpus and use it to generate sequences	LLMs such as ChatGPT

When to Choose Boon AI Technology…

Customer/Partner Requirement	Boon Logic Technology Fit and Features
Air-Gapped Environment	Model training and inferencing entirely at the edge. No cloud connection is ever needed.
Mission-Critical Application or Regulated Industry	Predictability and explainability of outputs. Unbiased models. No hallucination.
High Throughput / Real Time	Up to 100s of millions of inferences per second on CPU
Scalability	100% self-training models. No SME labeling required.
Sustainability	**No GPUs.** Trains and runs entirely on lightweight and commodity CPUs.
Accuracy (lowest FP, Highest TP)	Individualized models, customized to each monitored endpoint, produce the greatest accuracy.
Live (Streaming) Model Training	Single-pass training algorithm.

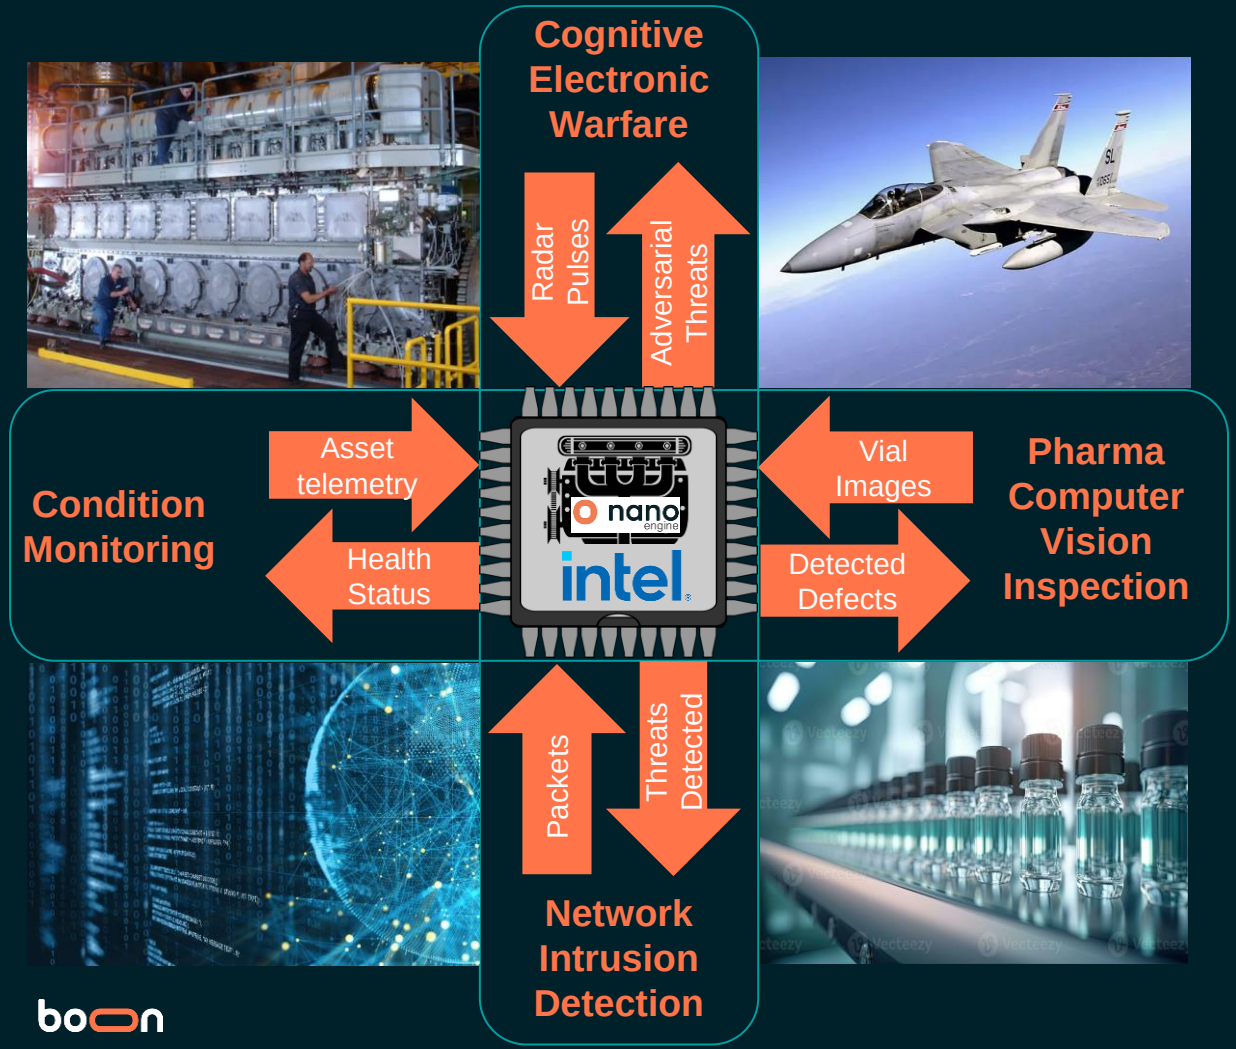

- ❑ Horizontally Applicable Tech
- ❑ Deployed through Partners
- ❑ Regulated Industries
- ❑ Mission-Critical Applications
- ❑ Transformational
- ❑ Annual License Model

boon


nano
A CPU-Native Core Technology

A New Algorithm nano

- Produces accurate clustering results.
- Incredibly fast and computationally lightweight.
- Segmentation 1000x faster than K-means, head-to-head, on the same COTS hardware and data sets.
- Supports thousands of clusters in a single model.
- Self-configures its hyperparameters in seconds.

A New Algorithm nano

Proprietary algorithm, not based on statistics, K-means, DBSCAN, graphs, forests, self-organizing maps, or autoencoders.

US Patent# US4541115A

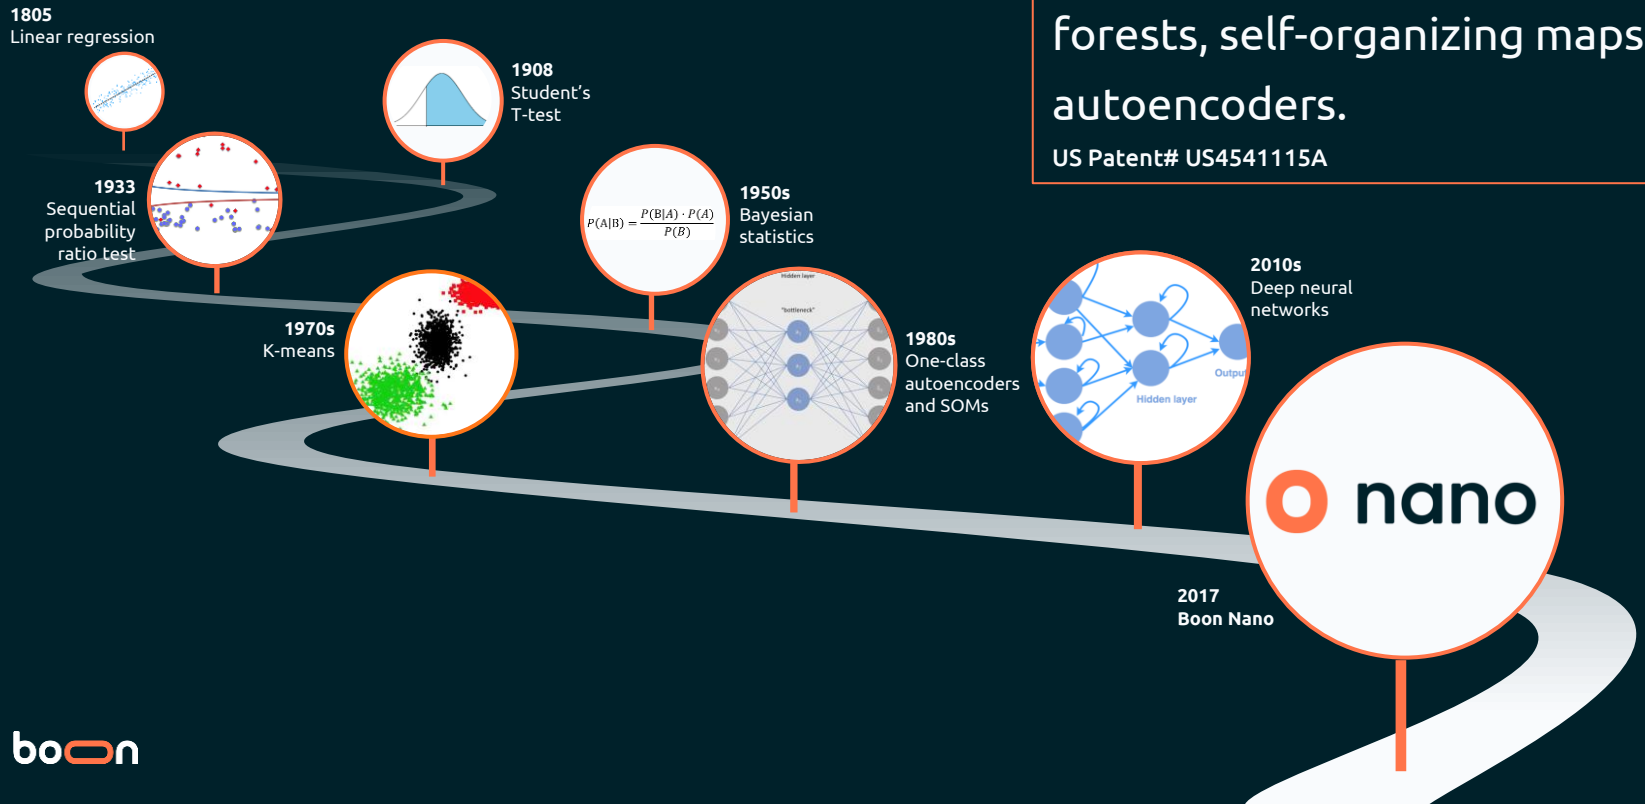

A New Algorithm nano

- Trains at the edge on CPU.
- Run at the edge on CPU.
- Does not use GPUs.
- Not based on neural networks.
- 1.4x faster on a single CPU core than K-means on 3,400-core GPU. (Cuomo, et al. "A GPU-Accelerated K-means algorithm" 2017)
- Outperforms FPGA-based segmentation algos.


amber
Applying the Boon Nano to Predictive Analytics

amber boon

Microservice for AI-based Predictive Analytics

amber

INCREASE QUALITY ASSURANCE

MERCK

PRIDE ENGINEERING

GSK

GE Healthcare

INCREASE ASSET RELIABILITY

TotalEnergies

GLENCORE

GREAT LAKES DREDGE & DOCK COMPANY, LLC

GE Healthcare

FAIRBANKS MORSE DEFENSE

MAXIMIZE PRODUCTION

PetroRio

FERREXPO

Atlas Copco

ENERFLEX

GE Healthcare

Enhance your existing platform

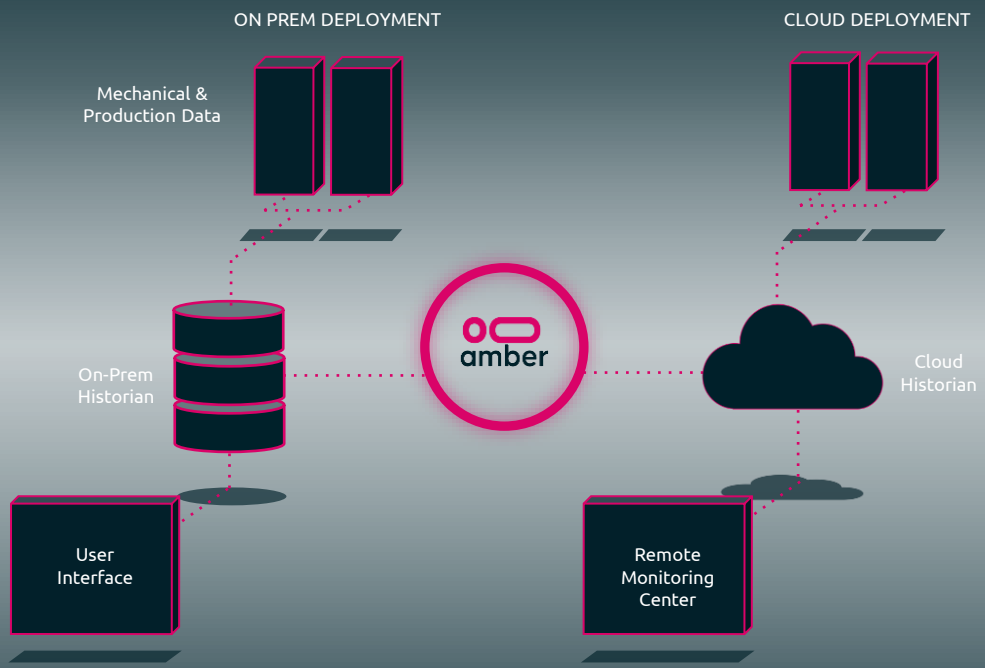

PROGRAMMING INTERFACES

- REST API
- Python SDK
- JavaScript SDK
- C++ SDK
- Go SDK
- R SDK
- C# SDK
- Java SDK

PROGRAMMING INTERFACES

- MQTT
- OPC UA
- MODBUS
- SWL

PROGRAMMING INTERFACES

Azure
Ignition
OSIsoft.
software AG
FactoryTalk

Steps to Use Amber

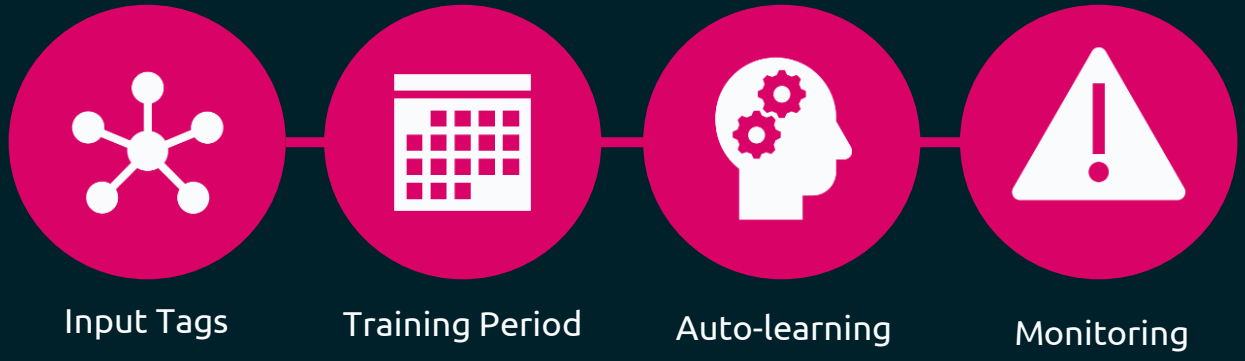

What is a Sensor Fusion Vector?

Values sampled at the same moment in time thereby creating a “snapshot” of the state of that asset at that moment. By assembling these measurements into a single vector (a sequence of comma-separated values), we create a **sensor fusion vector** for that asset at that moment in time

Sensor Fusion Vector →

Freq	RMS	Pos	Temp	Cur	RPM	PSI 1	PSI 2
52	119	134	74	161	33	88	50
58	116	139	74	152	34	95	52
62	100	152	75	170	26	100	66

Actionable Analytics

Compliance Score (CS) shows the health of the asset on a scale from 0% – 100%.

NORMAL	51-100%
CHANGING	25-50%
CRITICAL	0 – 25%

Feature Significance shows the top tags in the model that are causing CS to change.

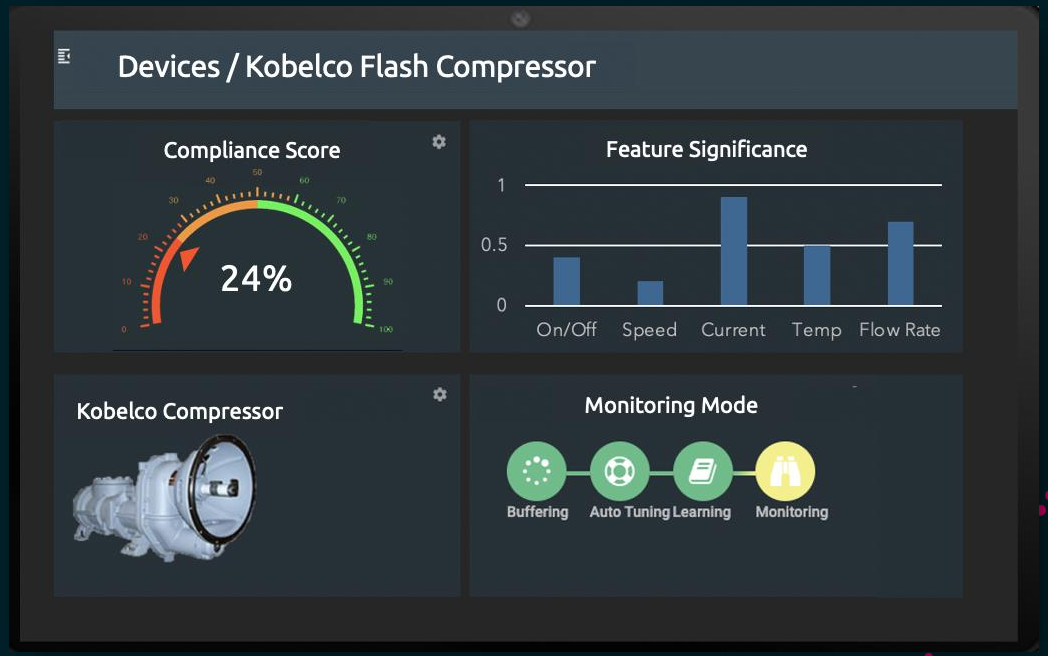

Detect what others can't

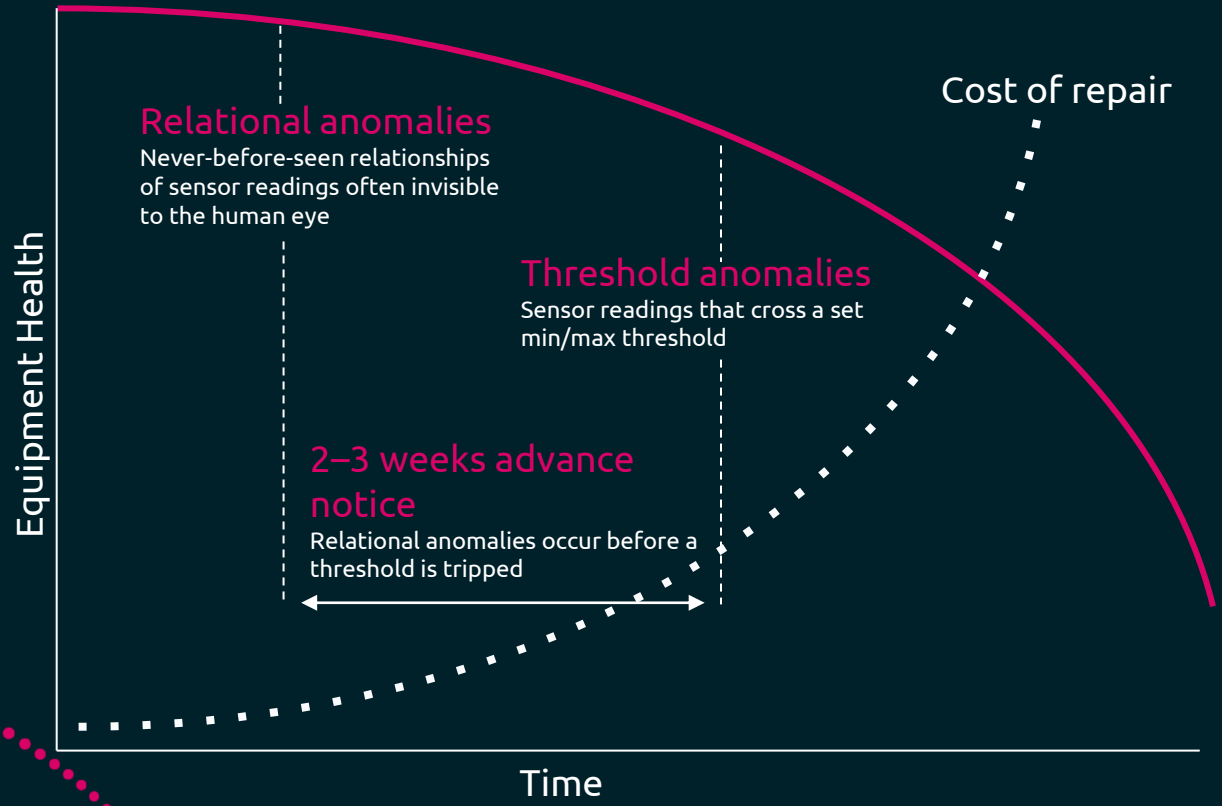

Centrifugal Compressor Use Case

CHALLENGE

The world's largest industrial gas company provides gas for a semiconductor manufacturer in Taiwan. Compressor failure is difficult to predict using PLC thresholding and user-built logic. The customer has a few backup compressors and requires advance warning for when the compressor should be switched to avoid production losses.

OUTCOME

- Amber identified an issue on Stage 3 of the compressor which was confirmed by maintenance teams as being a tear on the cooling fin.
- As a result of the warning, the team successfully prevented $740,000 in repair costs.

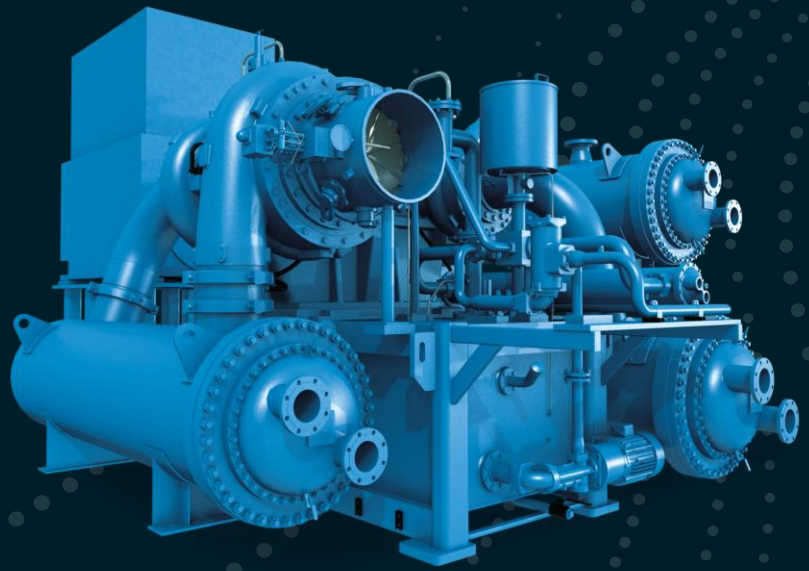

AMBER'S ALARMS
January 29 - "Asset changing"
February 11 - "Asset changing"
MARCH 9 - "Asset critical, issue in Stage 3 of Compressor"

MEANWHILE …
The customer's threshold-based vibration alarming system didn't notice anything abnormal with the asset.

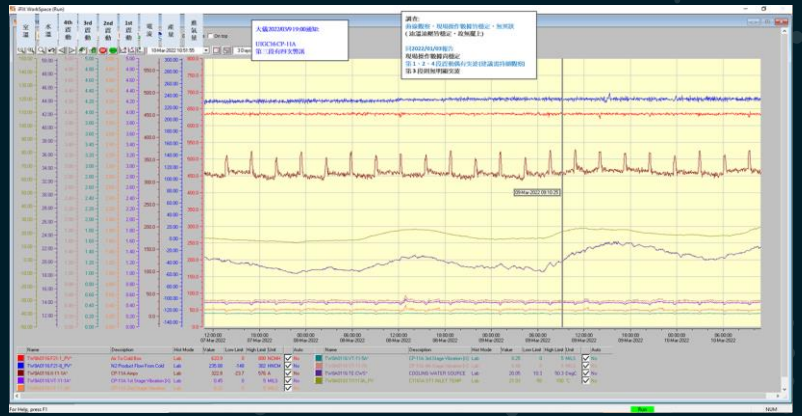

CONTINUOUS ALARMS: MARCH 9–13

MAINTENANCE TEAM CHECK-UP: MARCH 15

After numerous alarms from Amber indicating an issue in Stage 3 of the compressor, maintenance finally investigated the asset to find visible cracks on the Stage 3 cooling bundle. The damage from this event was able to be detected by Amber using the equipment's original sensors.

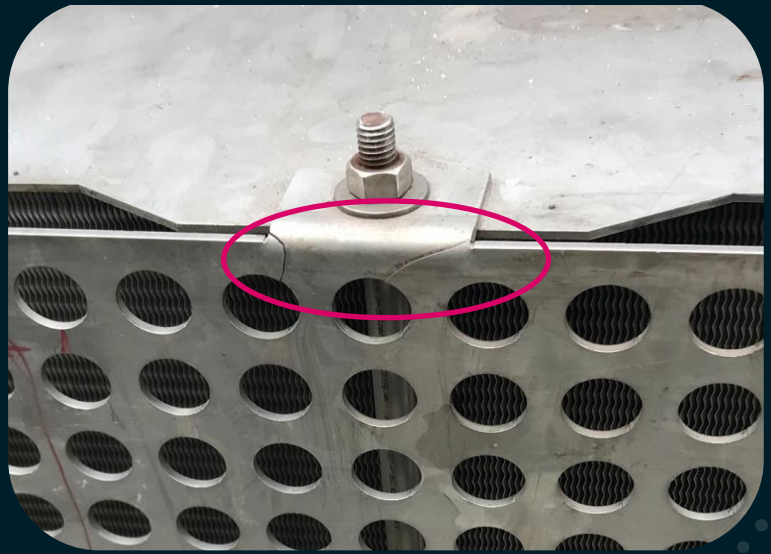

COMPRESSOR OVERHAUL: MARCH 15

After further investigation, several tears were found on the cooling fin. Amber was successfully used to reduce maintenance costs for this event. The cooling bundle cost $60,000 to replace. Had the breakage been detected later by maintenance teams, replacement of the core unit would have been required, costing over $800,000 in repairs.

Diesel Engine Monitoring Use Case

CHALLENGE

Large diesel engines used to power heavy mobile equipment (HME) for mining and military use normally suffer from numerous unplanned failure events, causing operational bottlenecks for mobile fleets and millions in lost production. A large German engineering company integrated sensors on a diesel engine and trialed Amber to determine if failure could be predicted, and if so how far in advance.

OUTCOME

- First alarm generated 6 months before failure
- Root Cause Analysis indicated failure was related to the fuel system

METHOD

Amber trained and built a high-dimensional model that represents all the complexity of the asset's normal operating state. This normal operating model was then used to detect anomalies, or infrequently occurring sensor fusion vectors.

- 5 months (250,000 samples) of data representing normal operation
- 10 features the customer thought were most relevant in monitoring
- <1 hour for a Boon Logic data scientist to engineer model

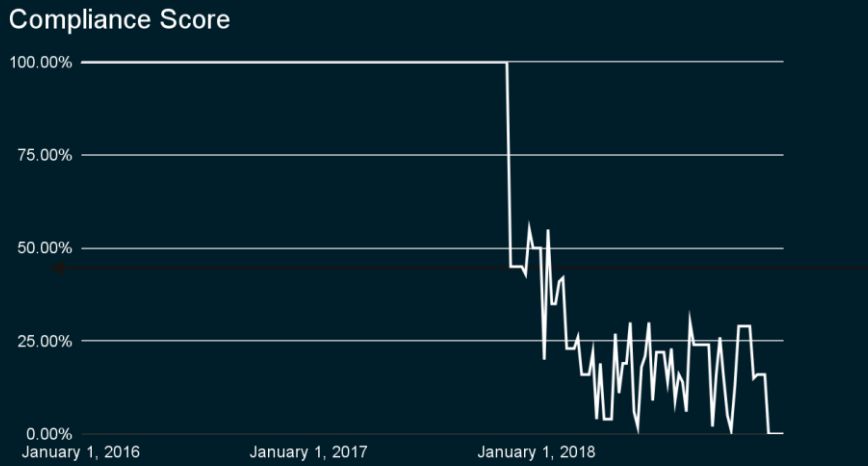

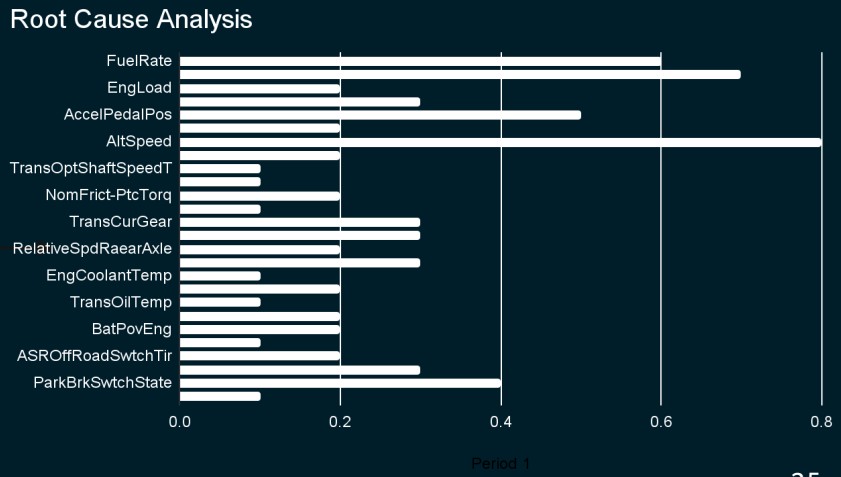

Root Cause Analysis

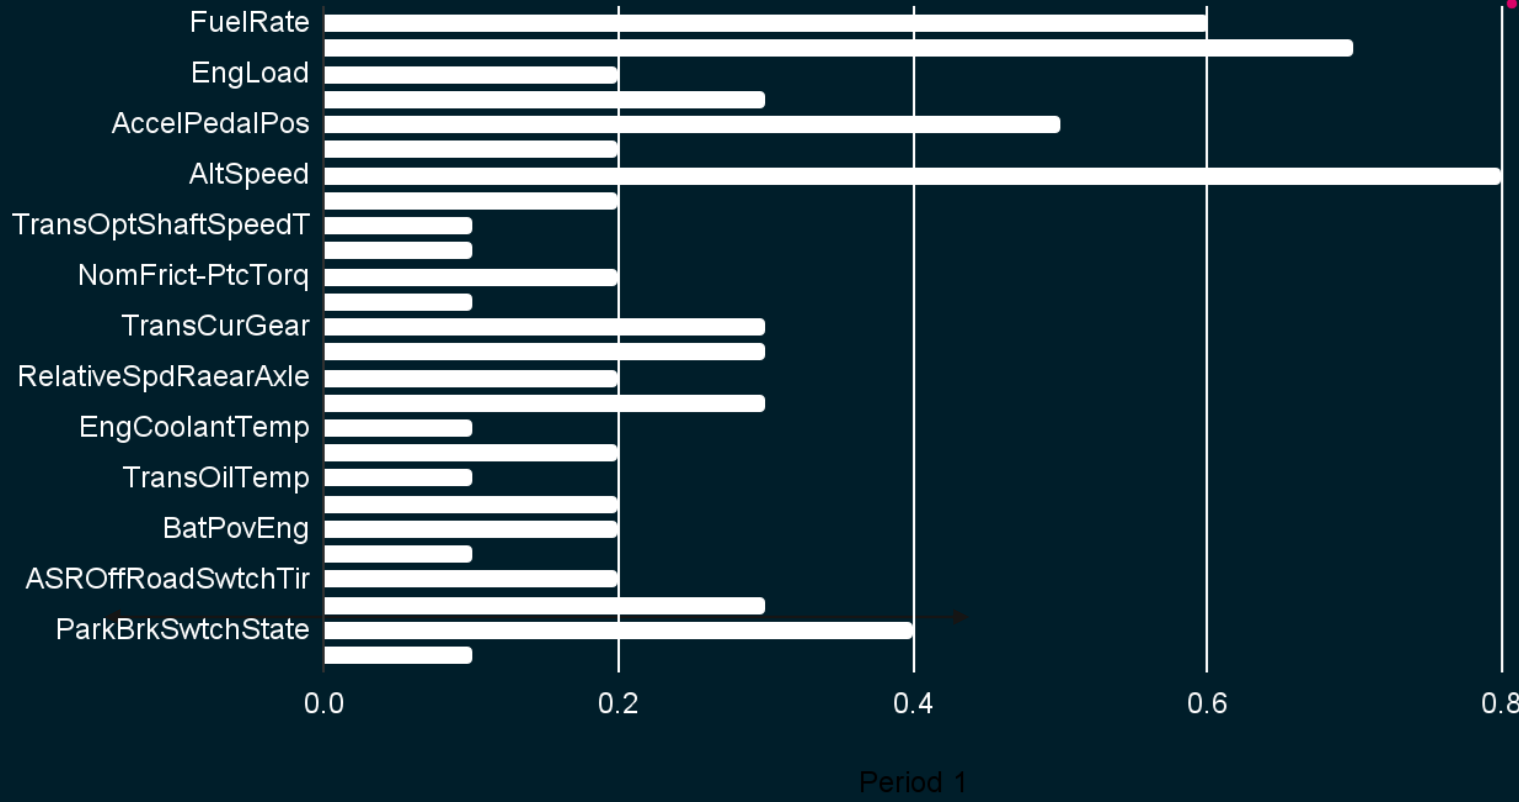

Deployments

Solving challenges at the far edge, powered by Intel.

Deploying Amber

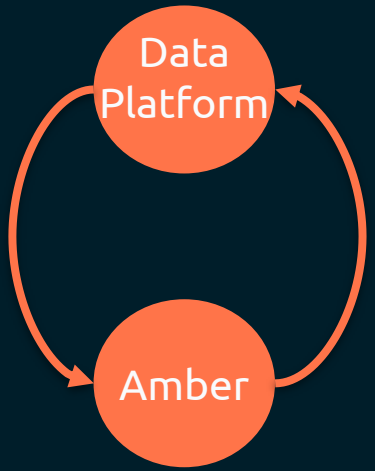

Connect Amber

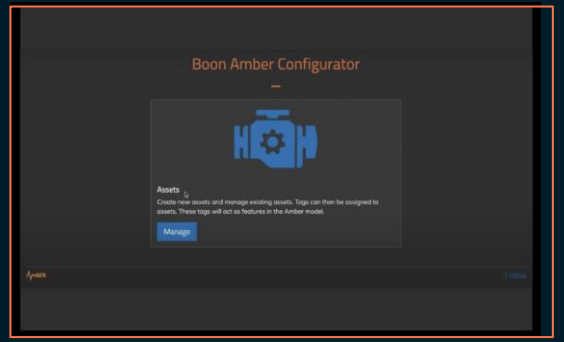

Configure Asset

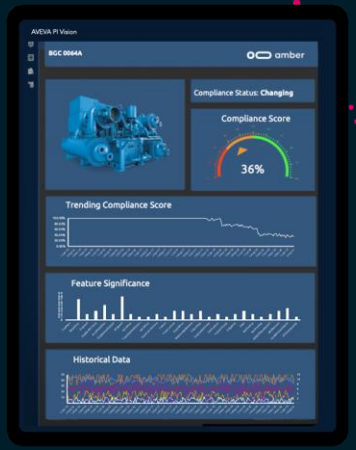

Visualize Asset Health

Enhance your existing platform

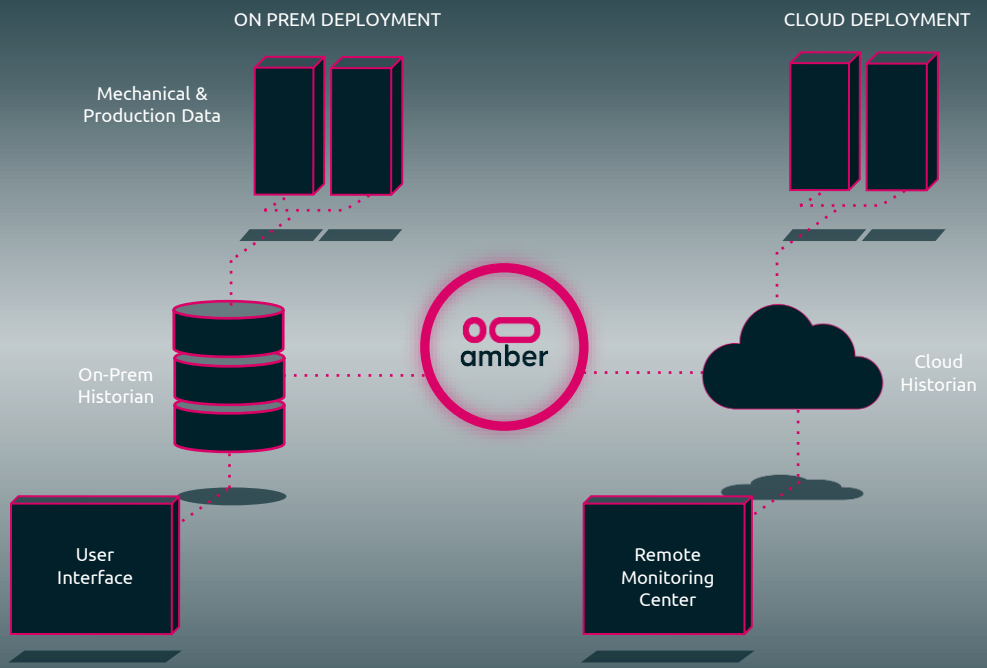

PROGRAMMING INTERFACES

- REST API
- Python SDK
- JavaScript SDK
- C++ SDK
- Go SDK
- R SDK
- C# SDK
- Java SDK

PROGRAMMING INTERFACES

- MQTT
- OPC UA
- MODBUS
- SWL

PROGRAMMING INTERFACES

Azure
Ignition
OSIsoft.
software AG
FactoryTalk

Partner Platforms

Pushing AI boundaries to solve critical issues at the far edge, minimizing latency and enhancing operational efficiency.

Intel-powered edge deployments

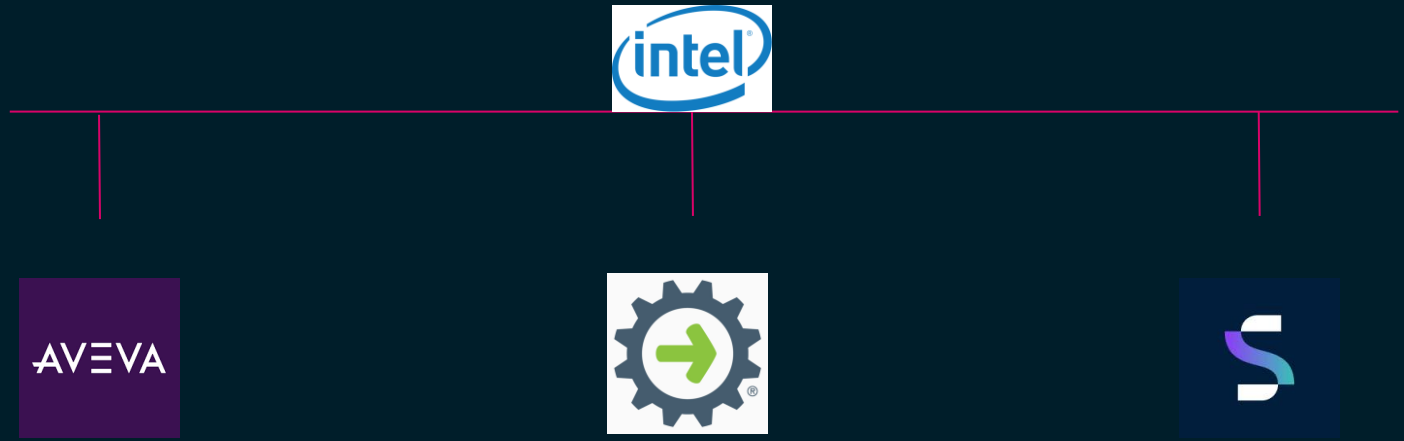

Aveva PI Server
+ PI Vision Dashboards

Inductive Automation
Ignition
+ Vision Client

Cumulocity IOT Server

Edge Deployment Showcase

Case Study: Largest Gold Mine in APAC Region

- **On-Premise Deployment:**
  Amber is deployed on-premise at the customer's remote gold mine in the APAC region. The system is integrated seamlessly into a highly secure, air-gapped environment, ensuring no external internet access.
- **Air-Gapped Security:**
  Operating in a fully air-gapped environment, Amber's server communicates solely through a closed ethernet connection to the customer’s PI Server, ensuring a secure and isolated setup.
- **Asset Monitoring and Anomaly Detection:**
  Amber monitors 15 unique high-value assets, with each asset having its own real-time AI model. These models are custom-built on the fly, offering precise anomaly detection and early warnings of potential failures.
- **Edge Computing with Intel Hardware:**
  Amber is running on a fanless industrial server powered by Intel® Core™ i7 processors. This setup ensures robust performance in the mine’s harsh environmental conditions while providing high computational power for real-time anomaly detection.

Architecture

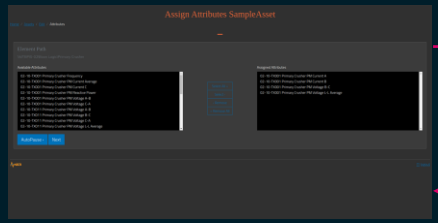

Asset Configuration

PI Amber Microservice

amber

AVEVA

PI Web API

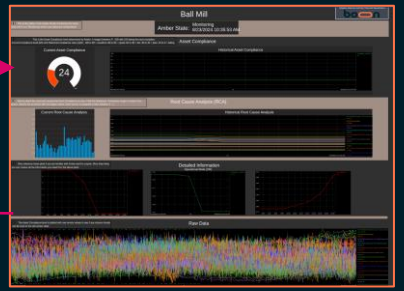

PI Vision Dashboards

Deployment Strategy

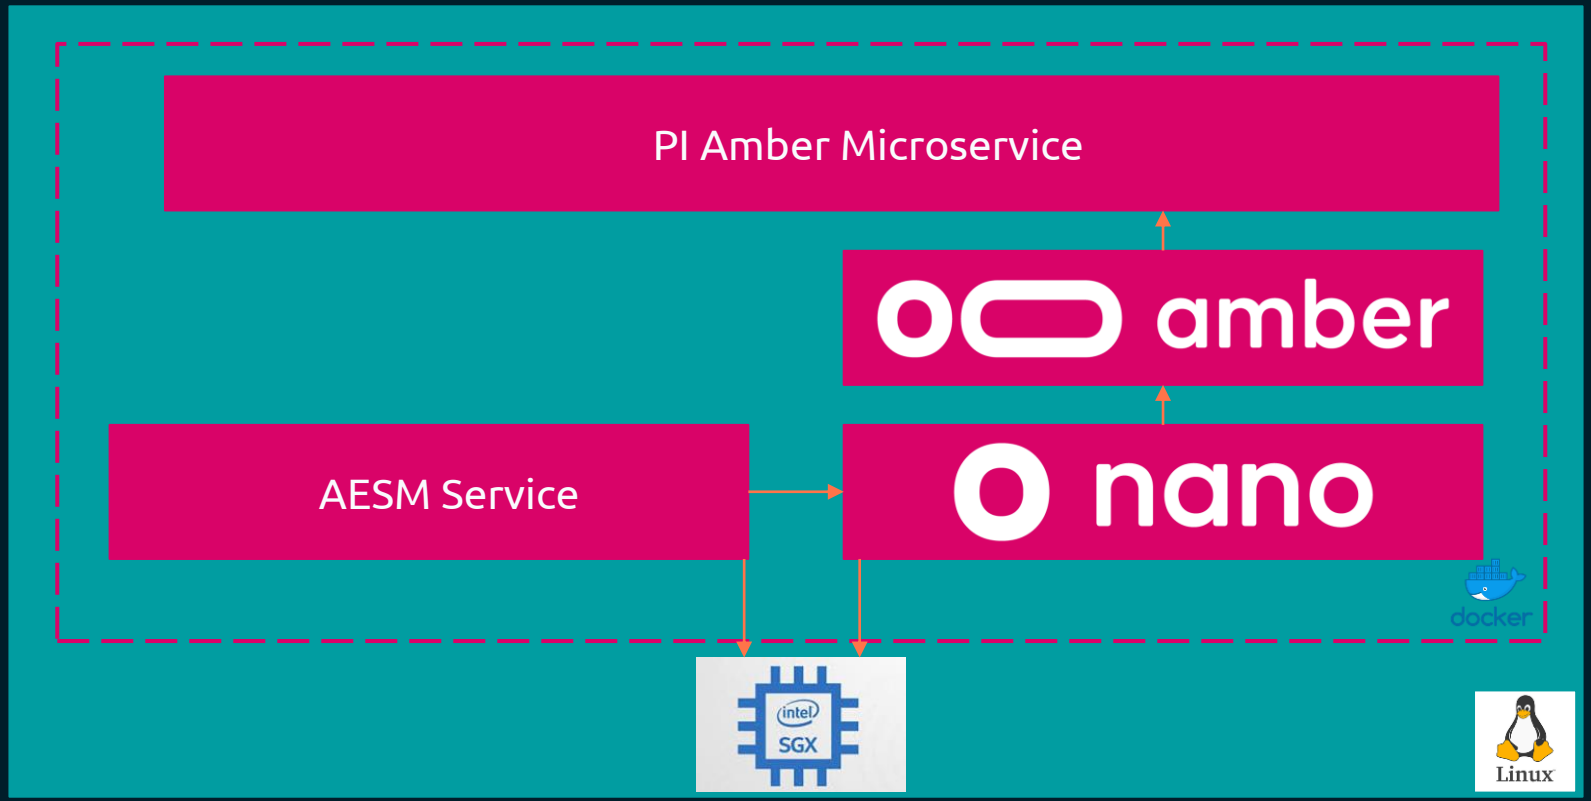

Demo

Contact Us for a Free Data Assessment

Unlock the power of AI for your business. Schedule a complimentary **AI Readiness Meeting** where we'll evaluate your company's potential to adopt **Amber**.

Ronan Burns

Lead Solution Architect

Ronan@boonlogic.com

BoonLogic.com

Questions?

Let’s dive in!

boon

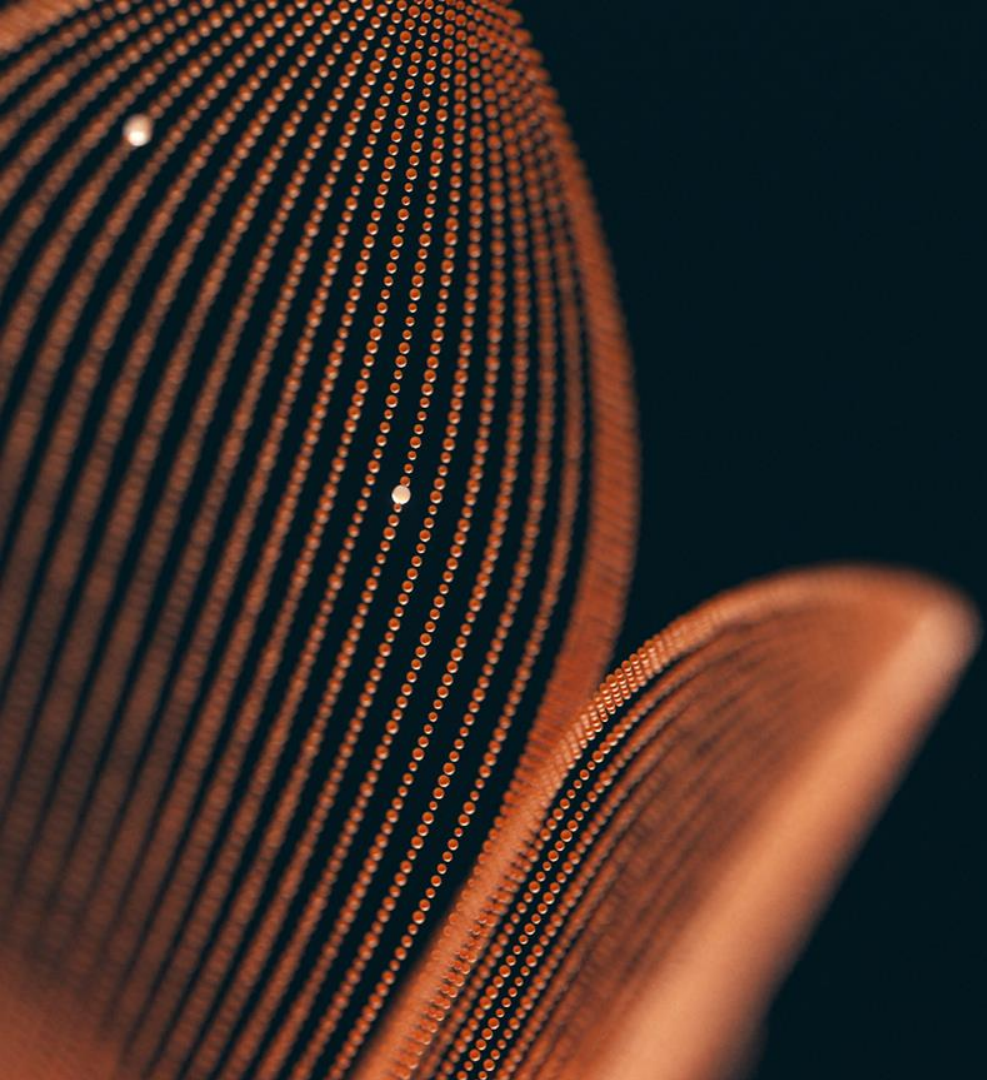

Thank You

boon