influxdata®

Quix

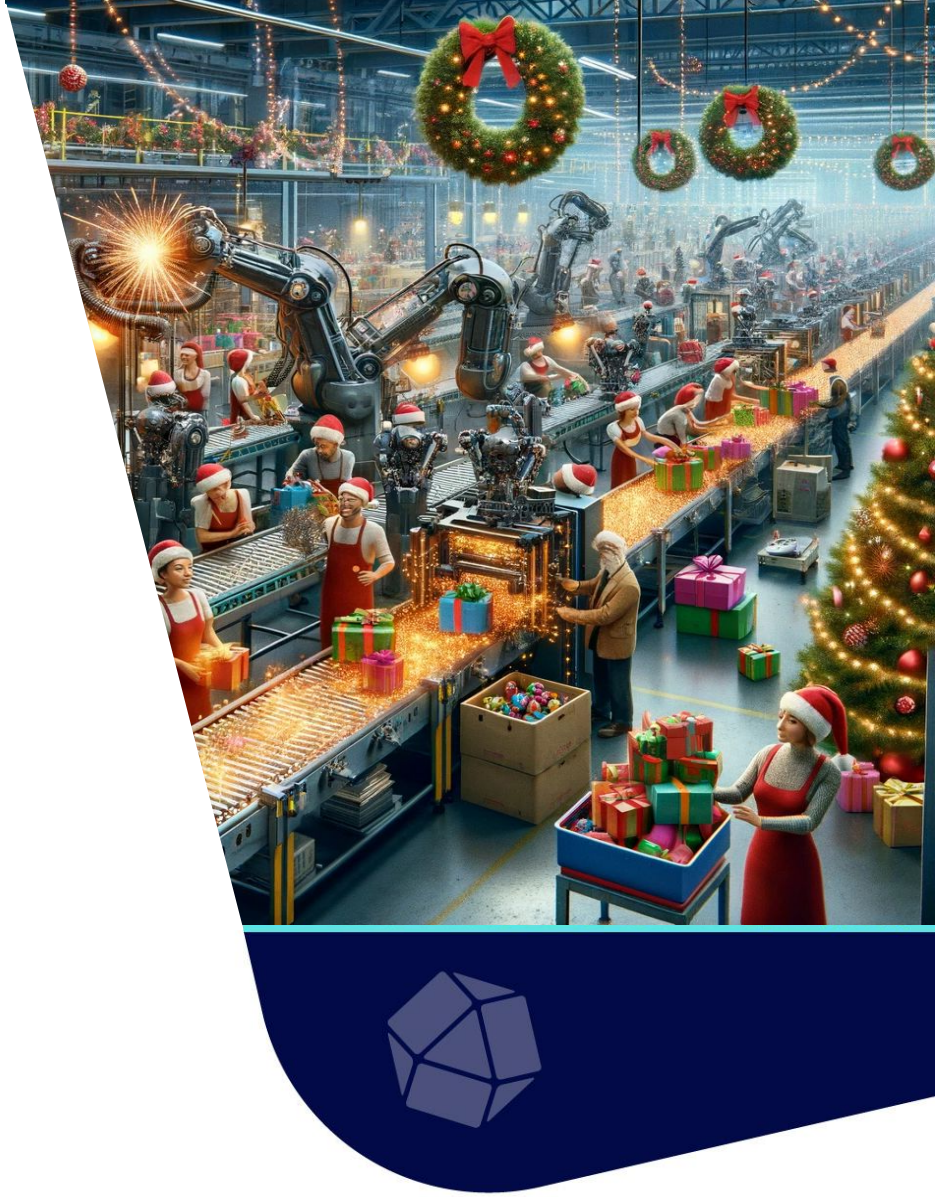

Saving the Holidays with Quix and InfluxDB: The OpenTelemetry Anomaly Detection Story

December 2023

Tun Shwe
VP of Data, Quix

Past life: Head of Data and Data Engineer.

Passion: Event-driven/real-time/streaming technologies and all things audio/music.

Driven: To create a new normal where data is processed as soon as it is generated.

Belief: Less is more. Get started sooner.

Quix | influxdata®

Jay Clifford
Developer Advocate, InfluxData

Past life: Sales Engineer for IIoT Solutions.

Passion: Apache Ecosystem, Big Data and Demo tinkering.

Driven: To make observability and IoT solutions accessible to all.

Belief: An industry's success comes from the domain experts.

Agenda

An intro to OTEL

As our data pipeline grows what can be do to make sure that we know what's happening at each stage/

What really is data plumbing?

How does this relate to event streaming and TSDBs

1

Let's look at a problem

Problems drive learning.
Let's create a scenario with a problem to solve.

2

Let's solve that problem

We will deploy Quix (streaming platform) and InfluxDB (TSDB) to solve the problem

3

Next steps

Get your hands on the source code and get involved with our communities.

4

Quix | influxdata®

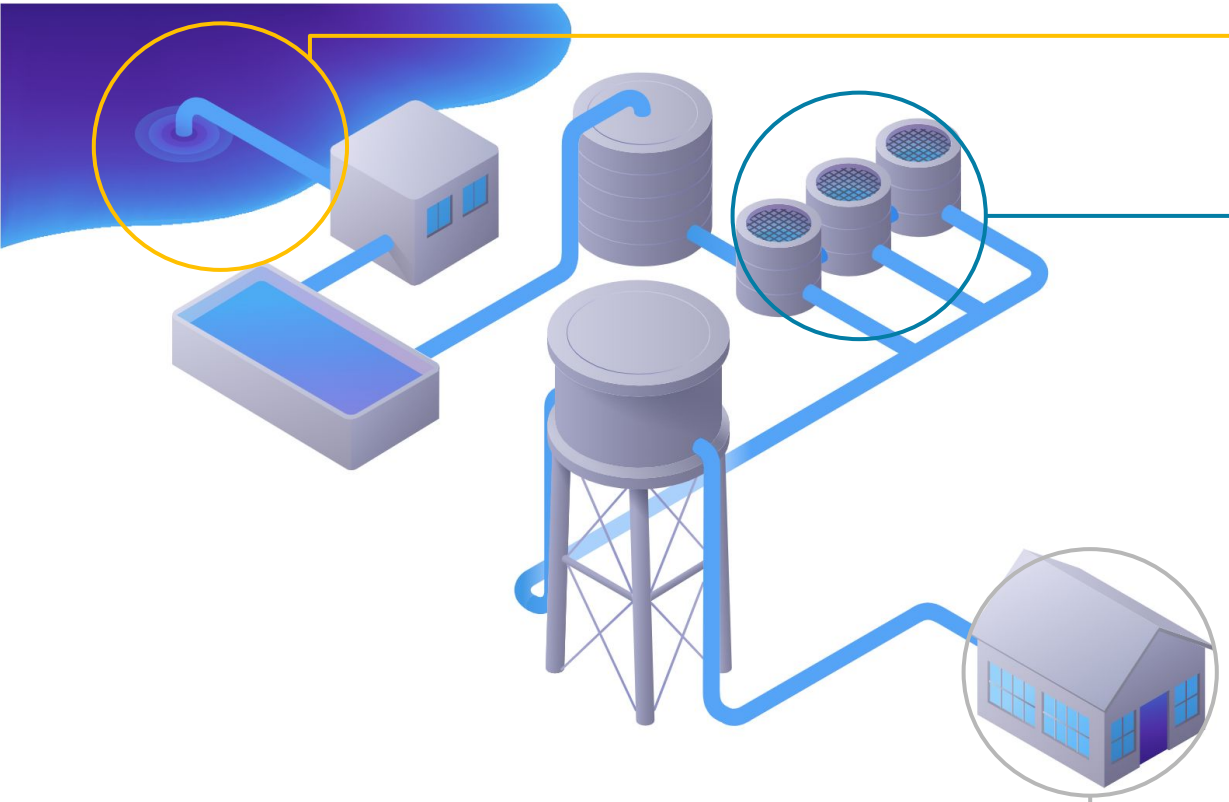

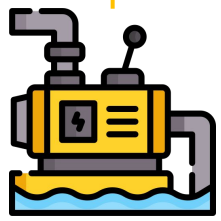

Source

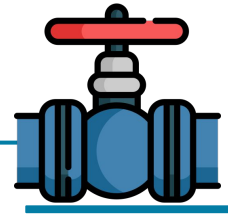

Transformation

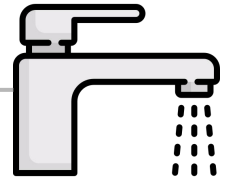

Destination

What are data pipelines?

Quix | influxdata®

Introduction to event streaming and TSDBs

Streaming technologies

Batching

Data is collected in a database.

Bounded data is periodically scheduled to be loaded from the database into the processing system.

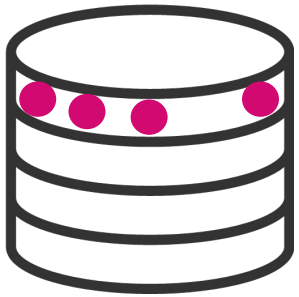

</>

t	Gx	Gy	Gz	GT	ΔG
t_1	0.1	1.0	0.2	1.3	
t_2	0.2	1.1	0.1	1.4	0.1
t_3	0.1	0.9	0.1	1.1	-0.3
t_n	0.3	1.0	0	1.3	1.3 - GT_{n-1}

- Computation on range of historical data (stateless).
- Process data at rest.
- Results are not in real time.

Quix | influxdata®

Streaming

Data is collected in a broker or transport, e.g. a Kafka topic.

Unbounded data is continuously consumed and processed as soon as it is published to the topic.

- Computation on each event (stateful where necessary).
- Process data in motion.
- Real-time results.

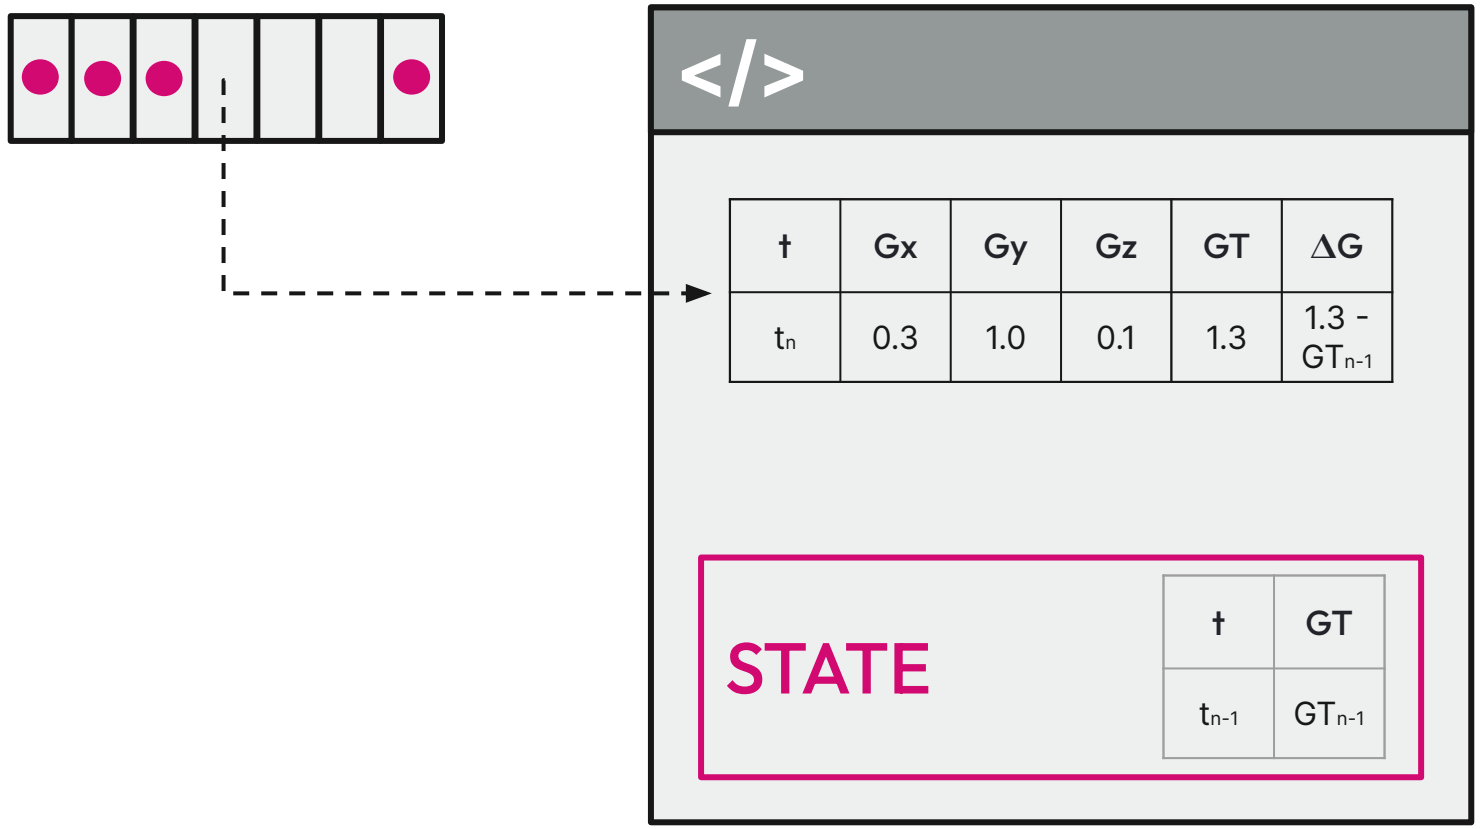

Quix | influxdata®

Time series data

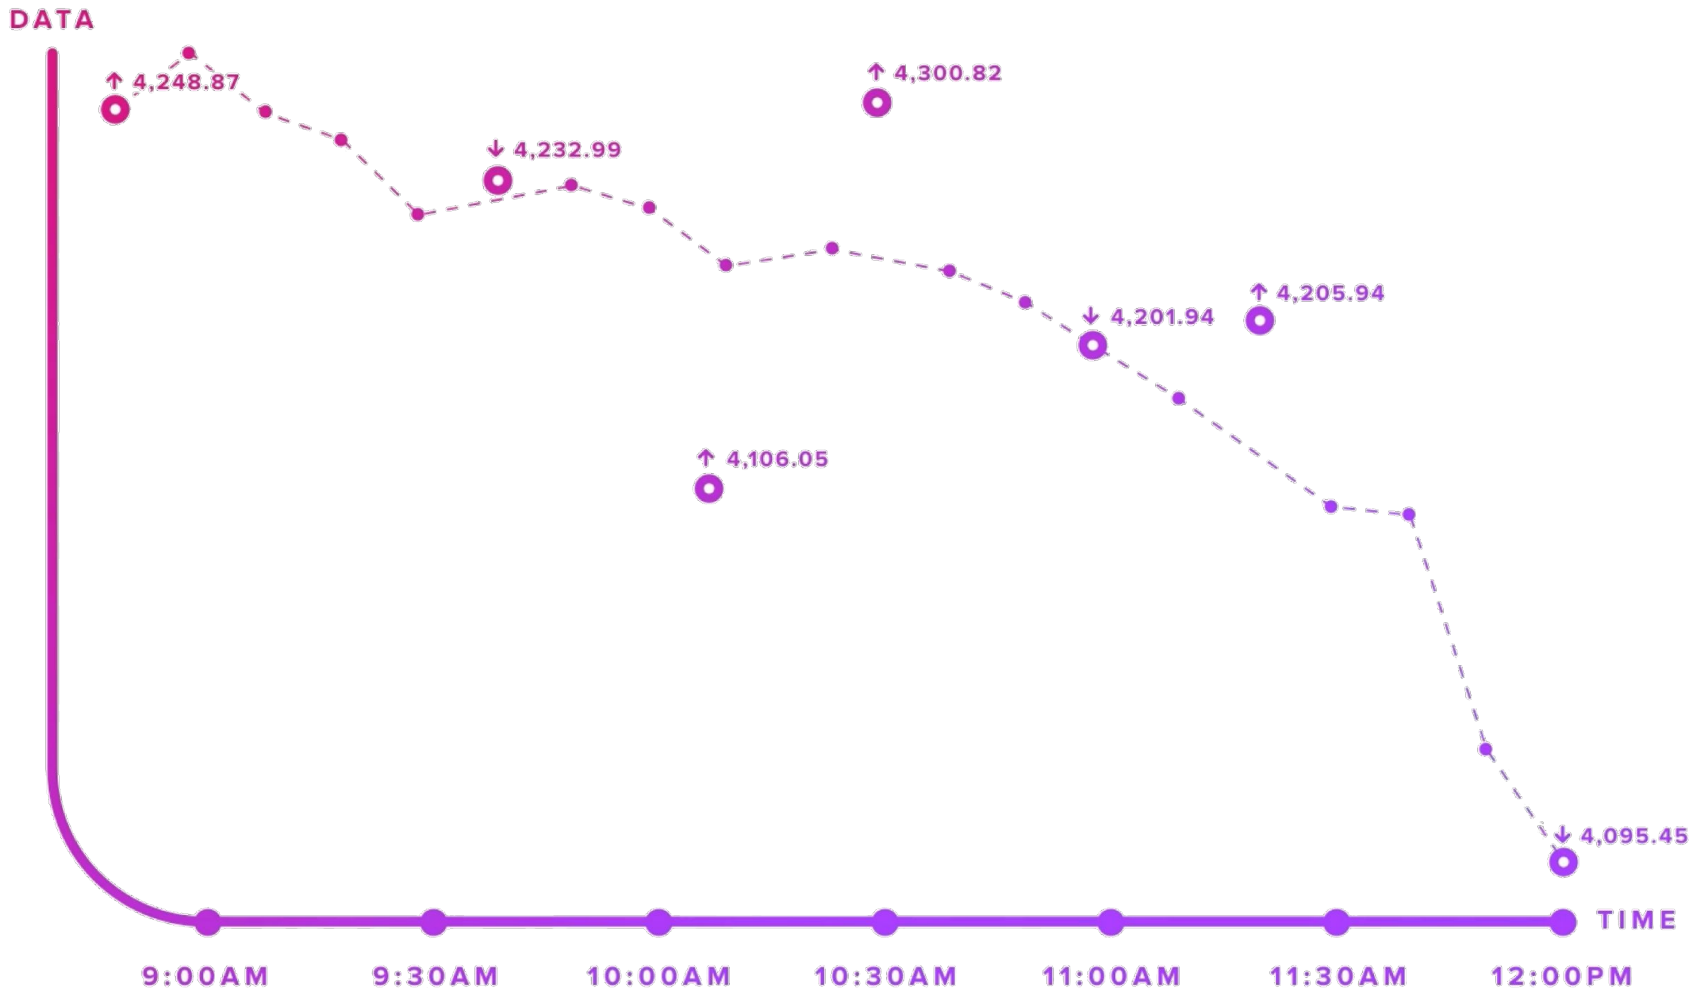

Quix | influxdata®

Types of time series

Metrics

Measurements at **regular** time intervals

Events

Measurements at **irregular** time intervals

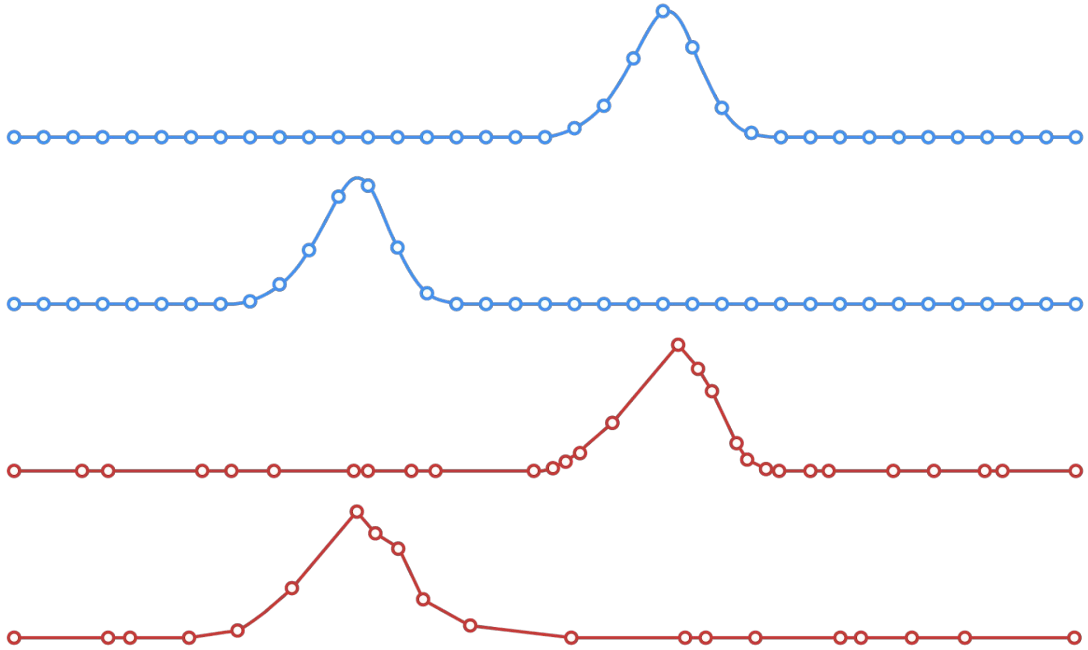

Quix | influxdata®

What is a time series database?

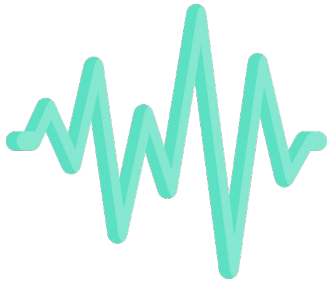

Time series data

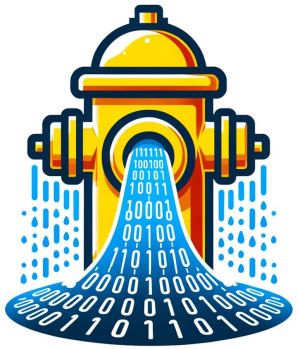

High write throughput

Efficient queries over time ranges

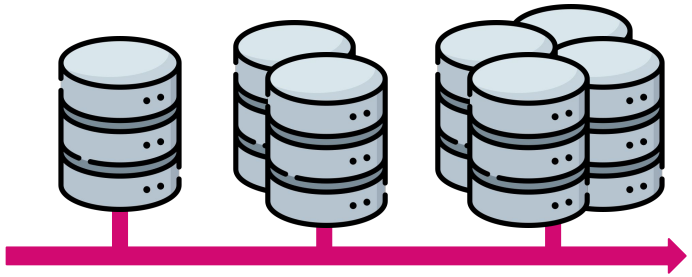

Scalability and performance

Quix | influxdata®

InfluxDB 3.0

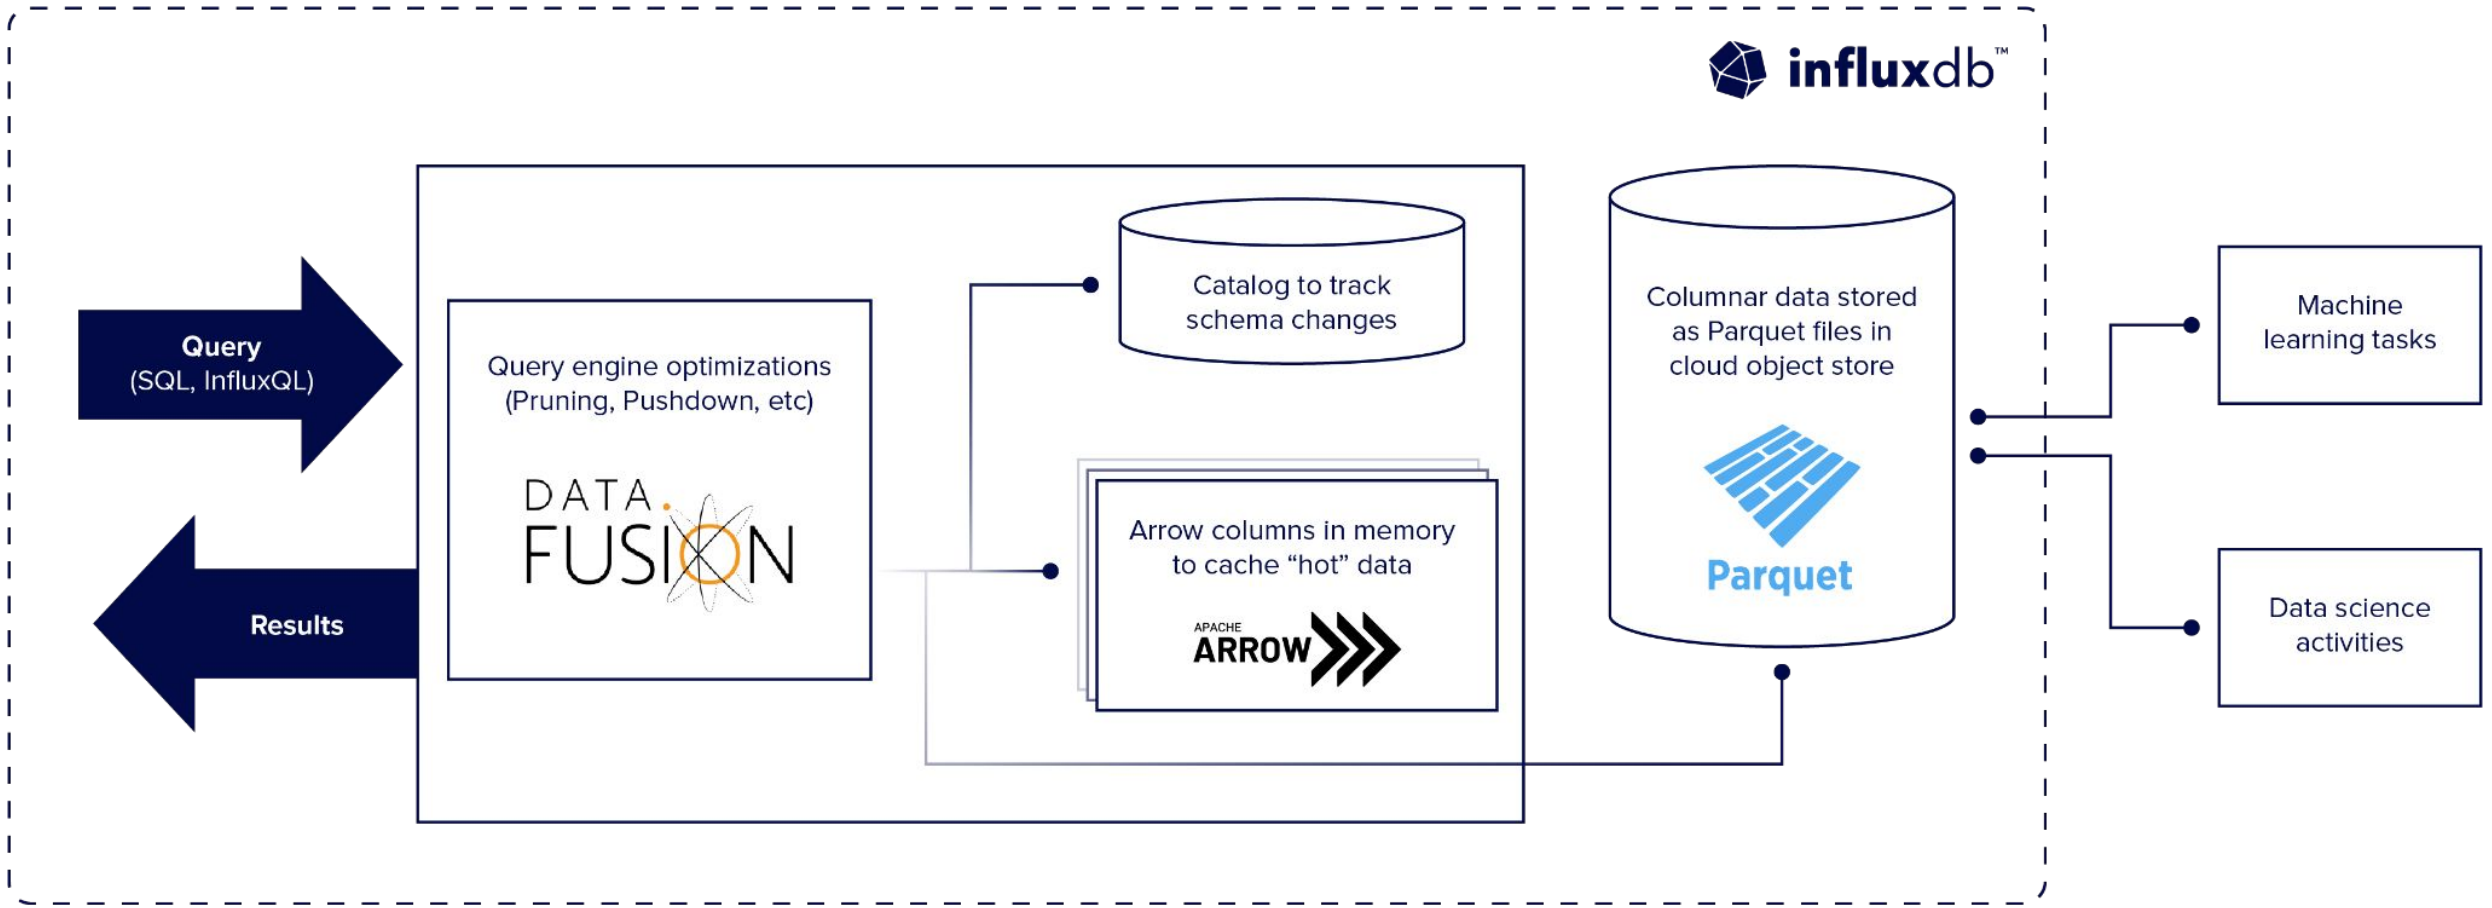

Quix | influxdata®

Quix ✌️

Quix Streams

Open source library for processing data in Kafka using pure Python. Inspired by FaaS and pandas

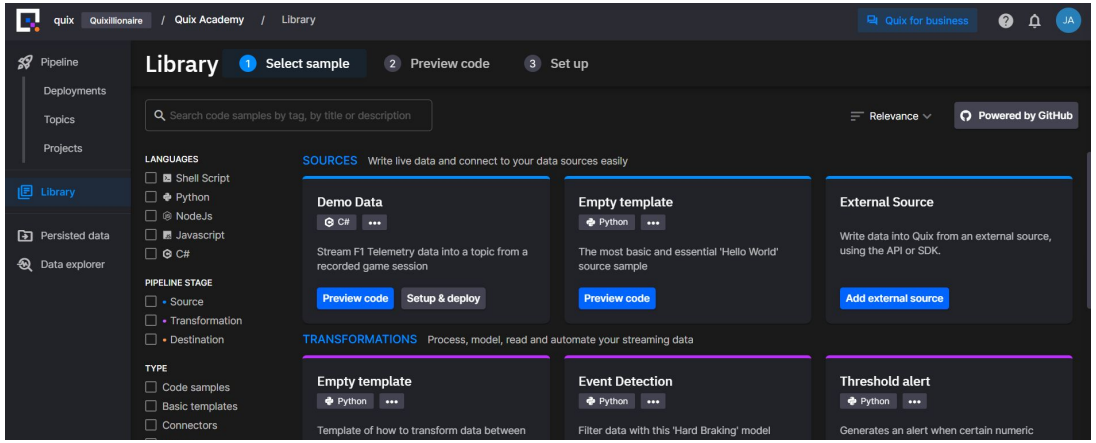

Quix Cloud

Platform to quickly build, test and deploy streaming data pipelines and applications without having to manage infrastructure

Quix | influxdata®

Quix platform architecture

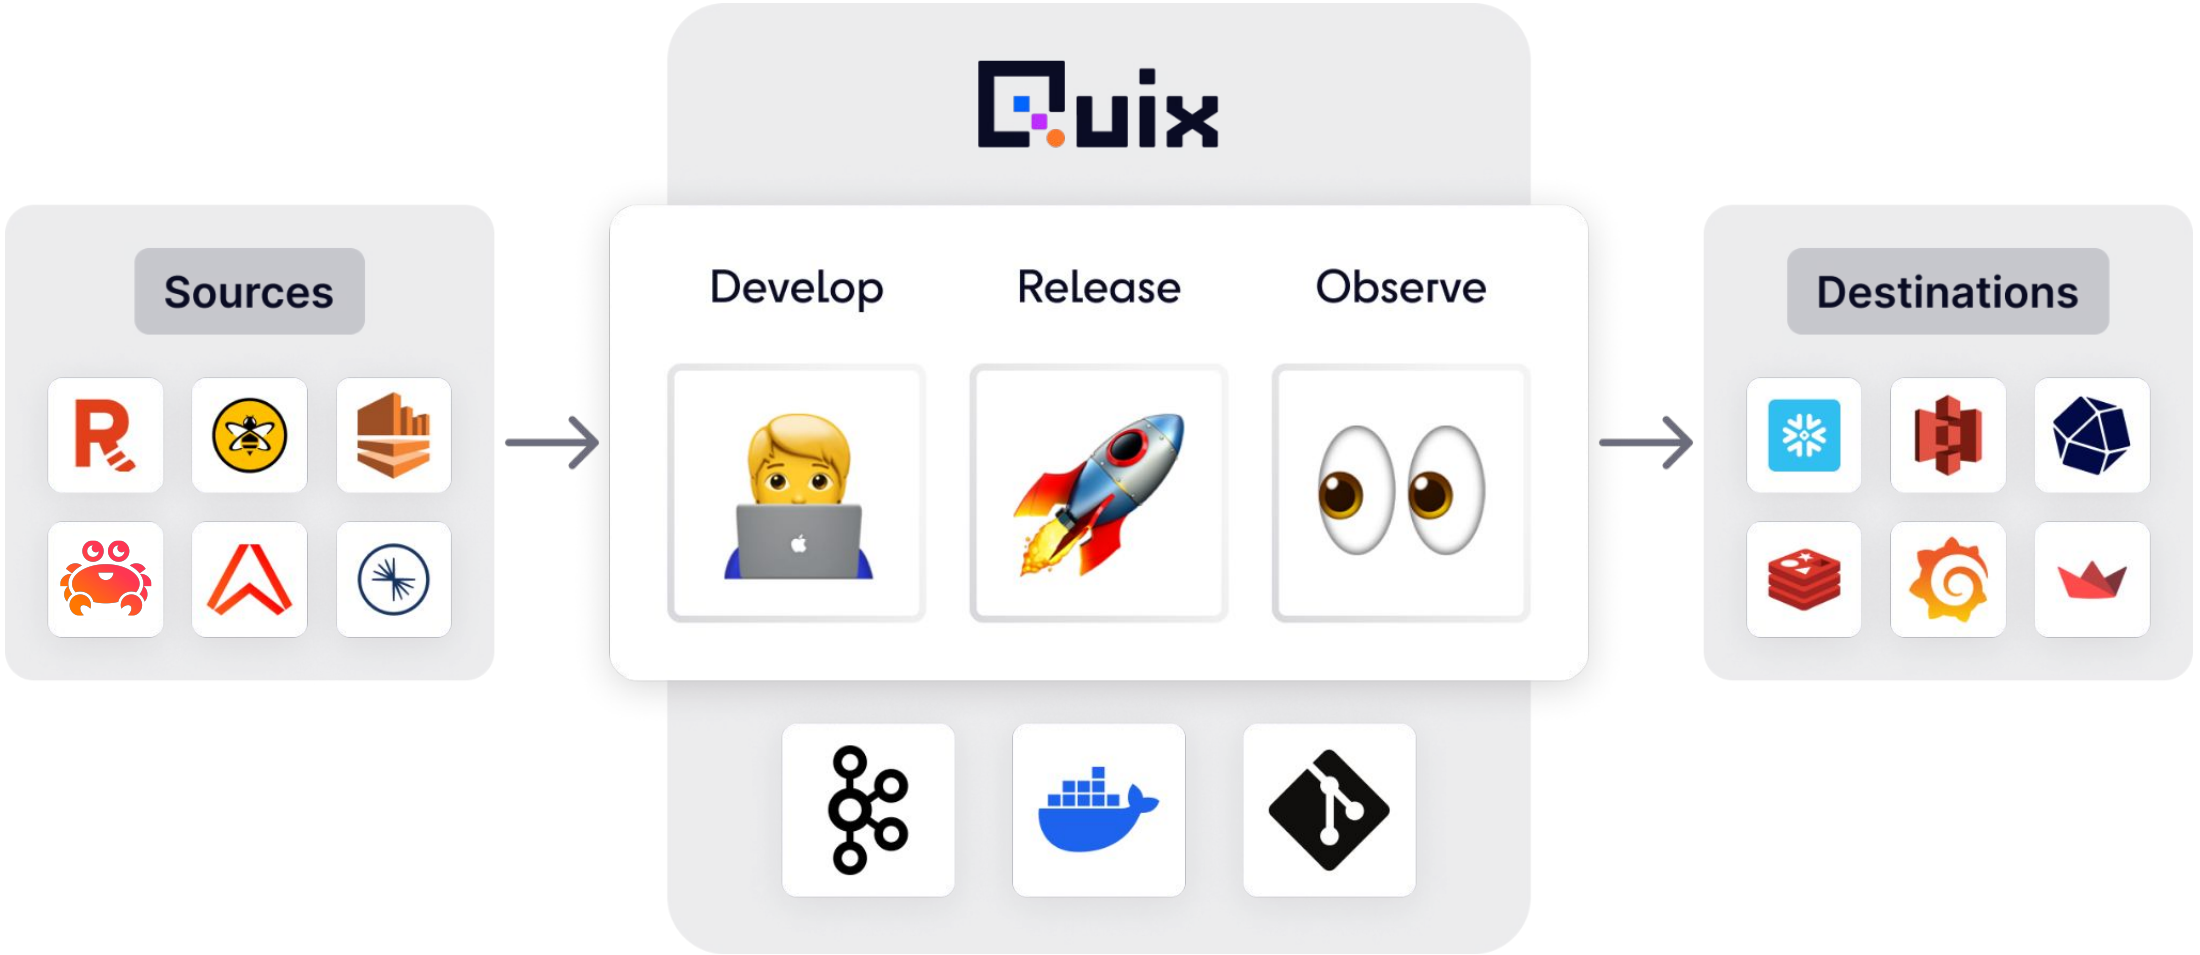

Quix | influxdata®

Building your own architecture is costly

8 months

3 Months ---> 3 Weeks ---> 7 Days --->

Build infrastructure

Technical complexities.

Design complexities.

Develop

Data consistency and synchronisation.

Release

Orchestration and management.

Observe

Effectively monitoring and debugging.

Platform team: 11 FTE

Engineering: 2 FTE + Data team: 2 FTE

Quix | influxdata®

Accelerated application development

Weeks → Hours → Minutes

Develop

Use free open source connectors & code samples to get started. Enable ML and GenAI faster.

Release

IaC: code, test and deploy event streaming applications. Powered by Kafka, Docker and Git.

Observe

A suite of observability tools for in-depth insights into your event-driven architecture.

Engineering: 2 FTE + Data team: 2 FTE

Predictable Cost & Risk

Quix | influxdata®

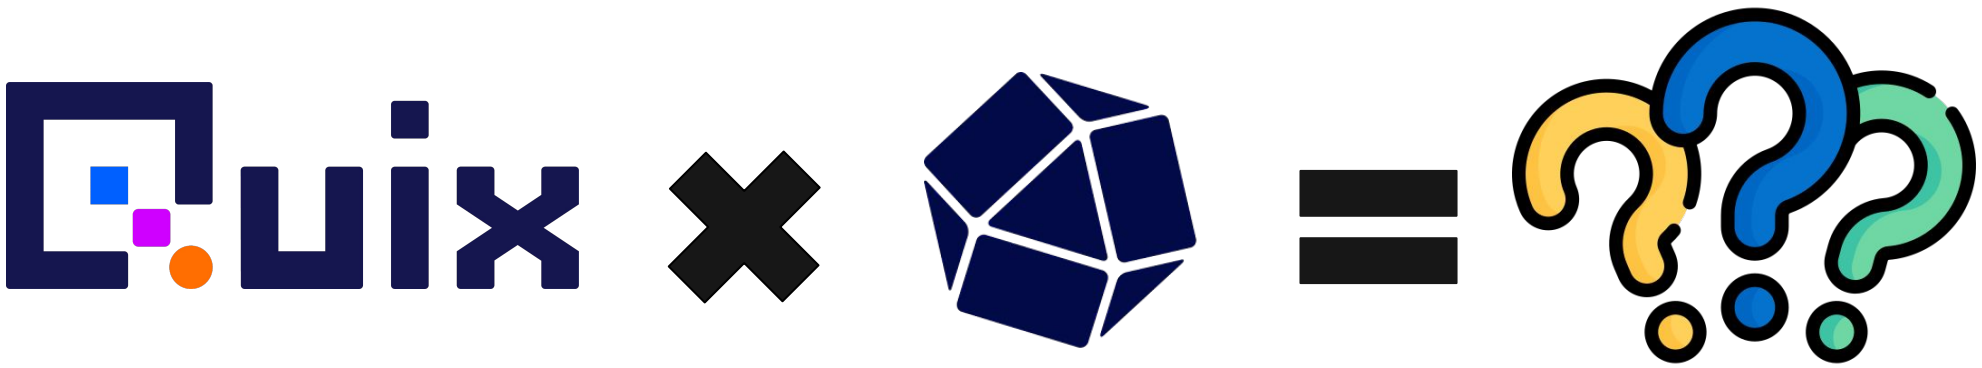

Let's look at a problem

📦 Packing Co — Anomaly Detection

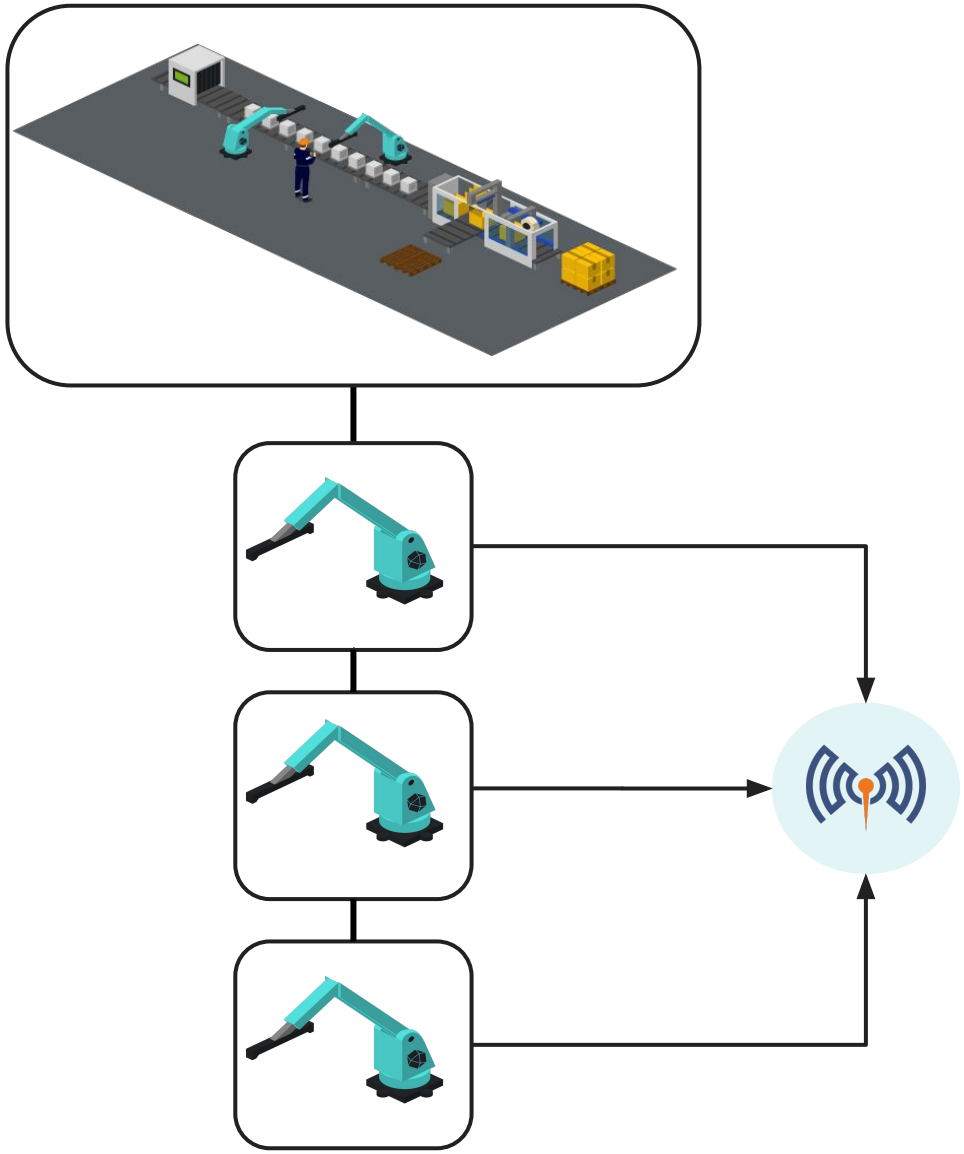

😢 Packing Co is having **recurring issues** with their packaging machines.

🤖 Unexpectedly, any of the machines will enter a **failing state**, which requires a manual reset by an engineer.

📊 The Plant Manager has advised, **when running normally** all machine sensors will follow **similar output patterns**. If a machine is at **fault** these will **fluctuate abnormally.**

🤔 How can we use **Quix** and **InfluxDB** to solve this?

Quix | influxdata®

Let’s solve that problem

Solution Architecture

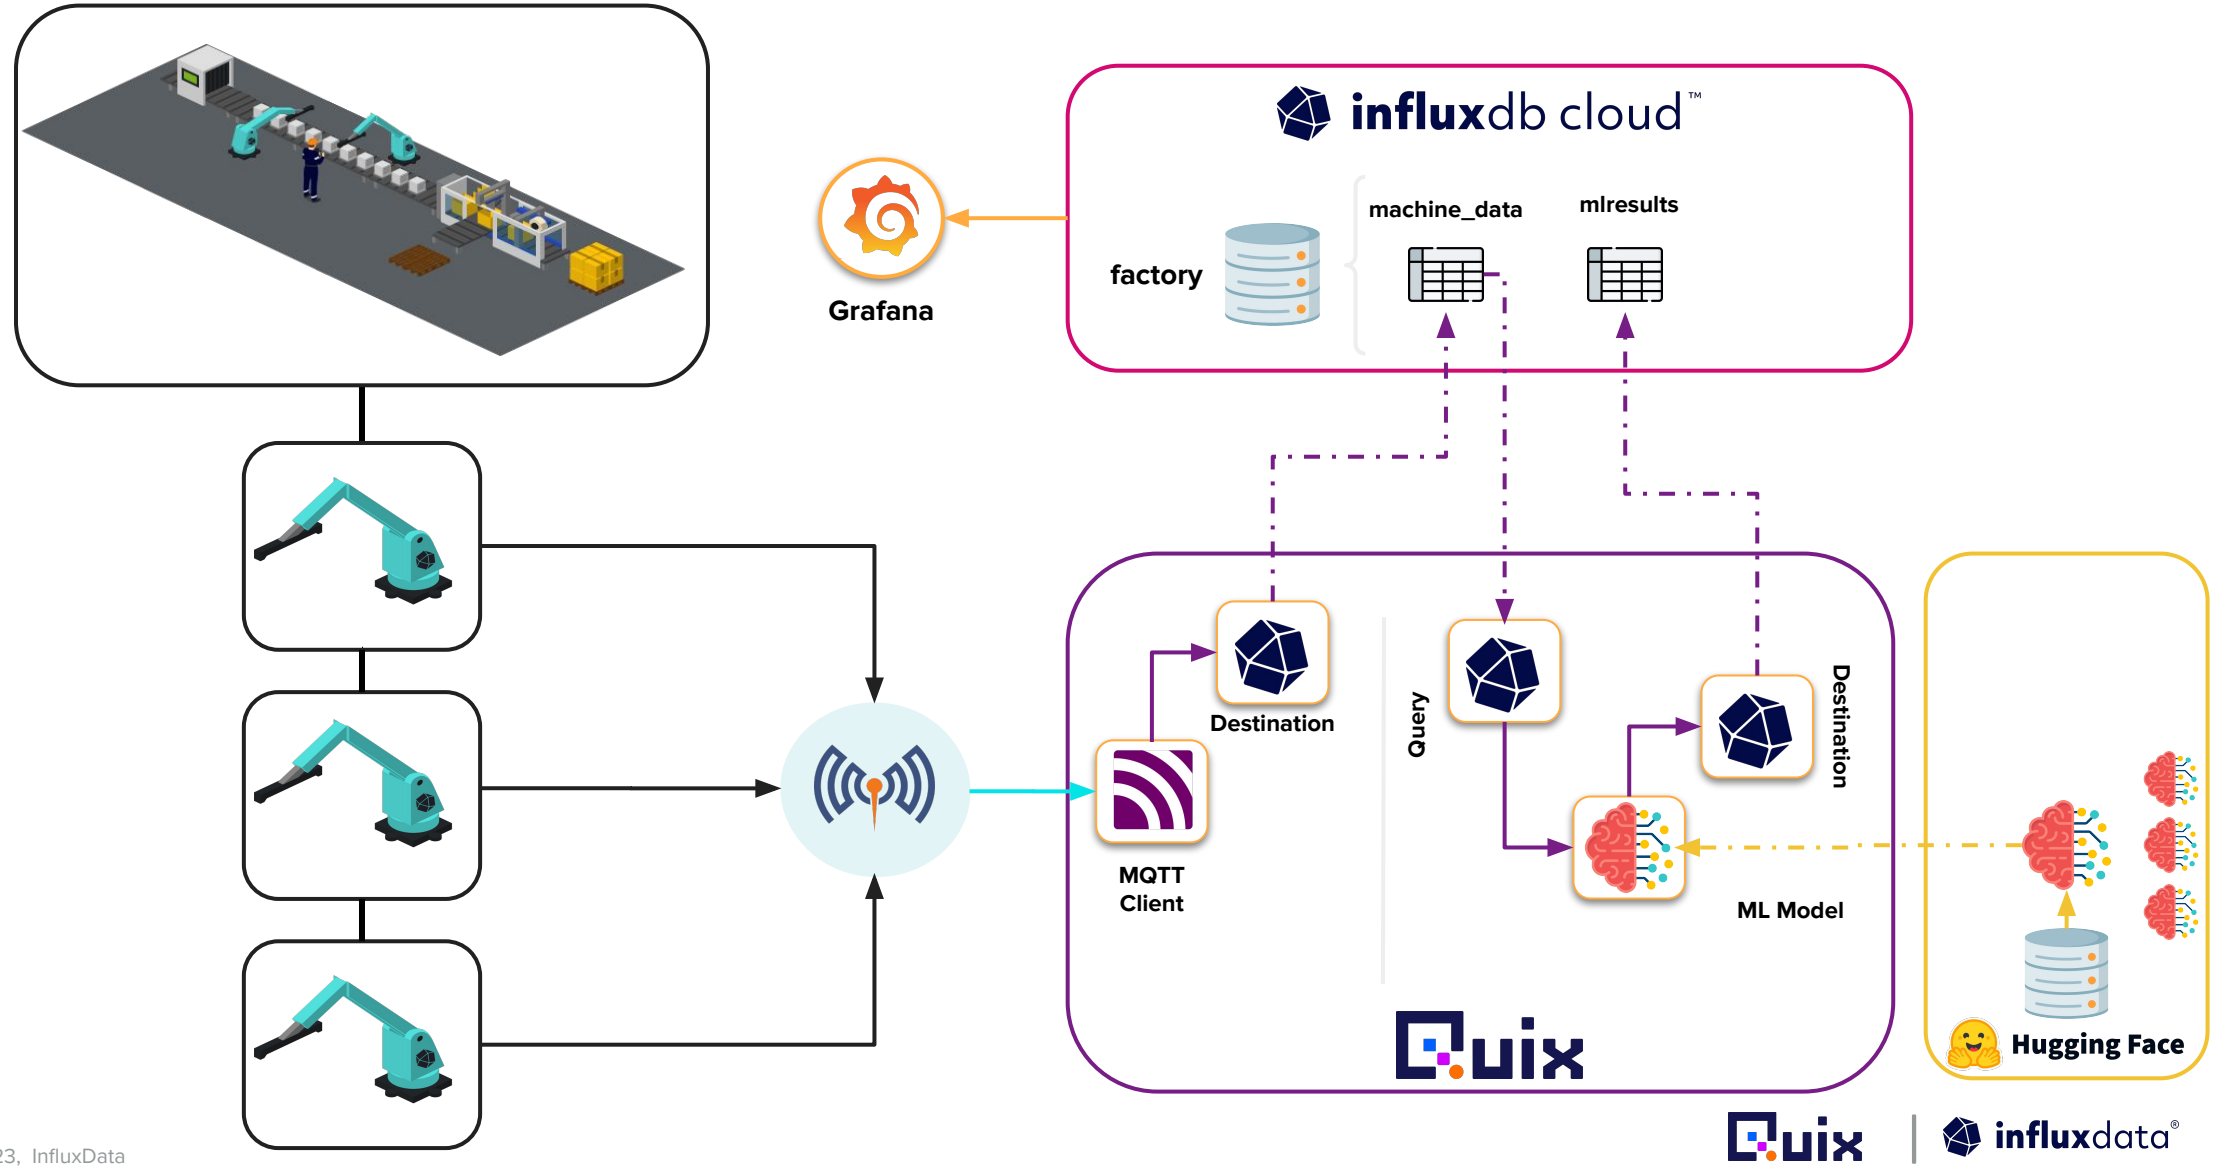

Quix | influxdata®

Data Ingest

Solution Architecture (Ingest)

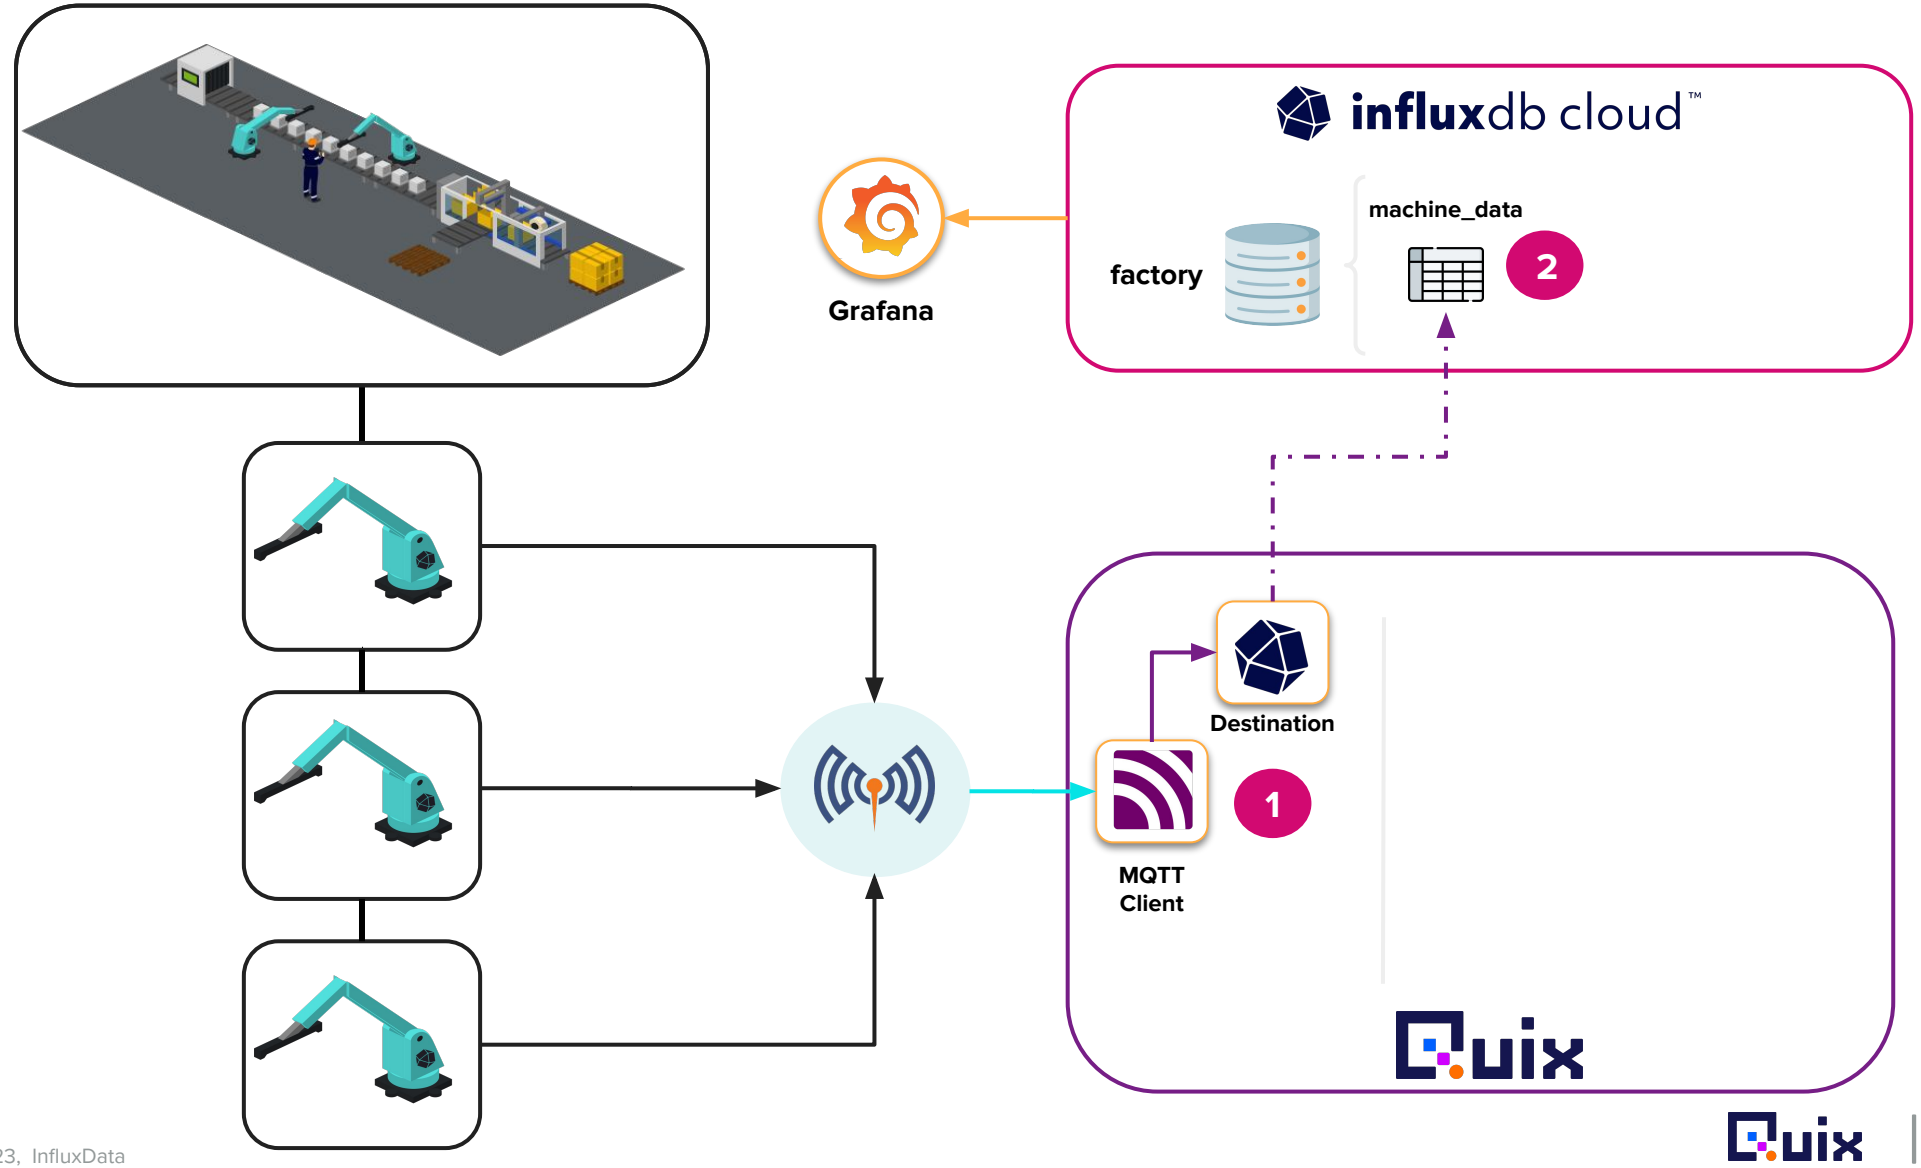

Quix | influxdata®

Choosing a Model

“Less is more” - Tun Shwe

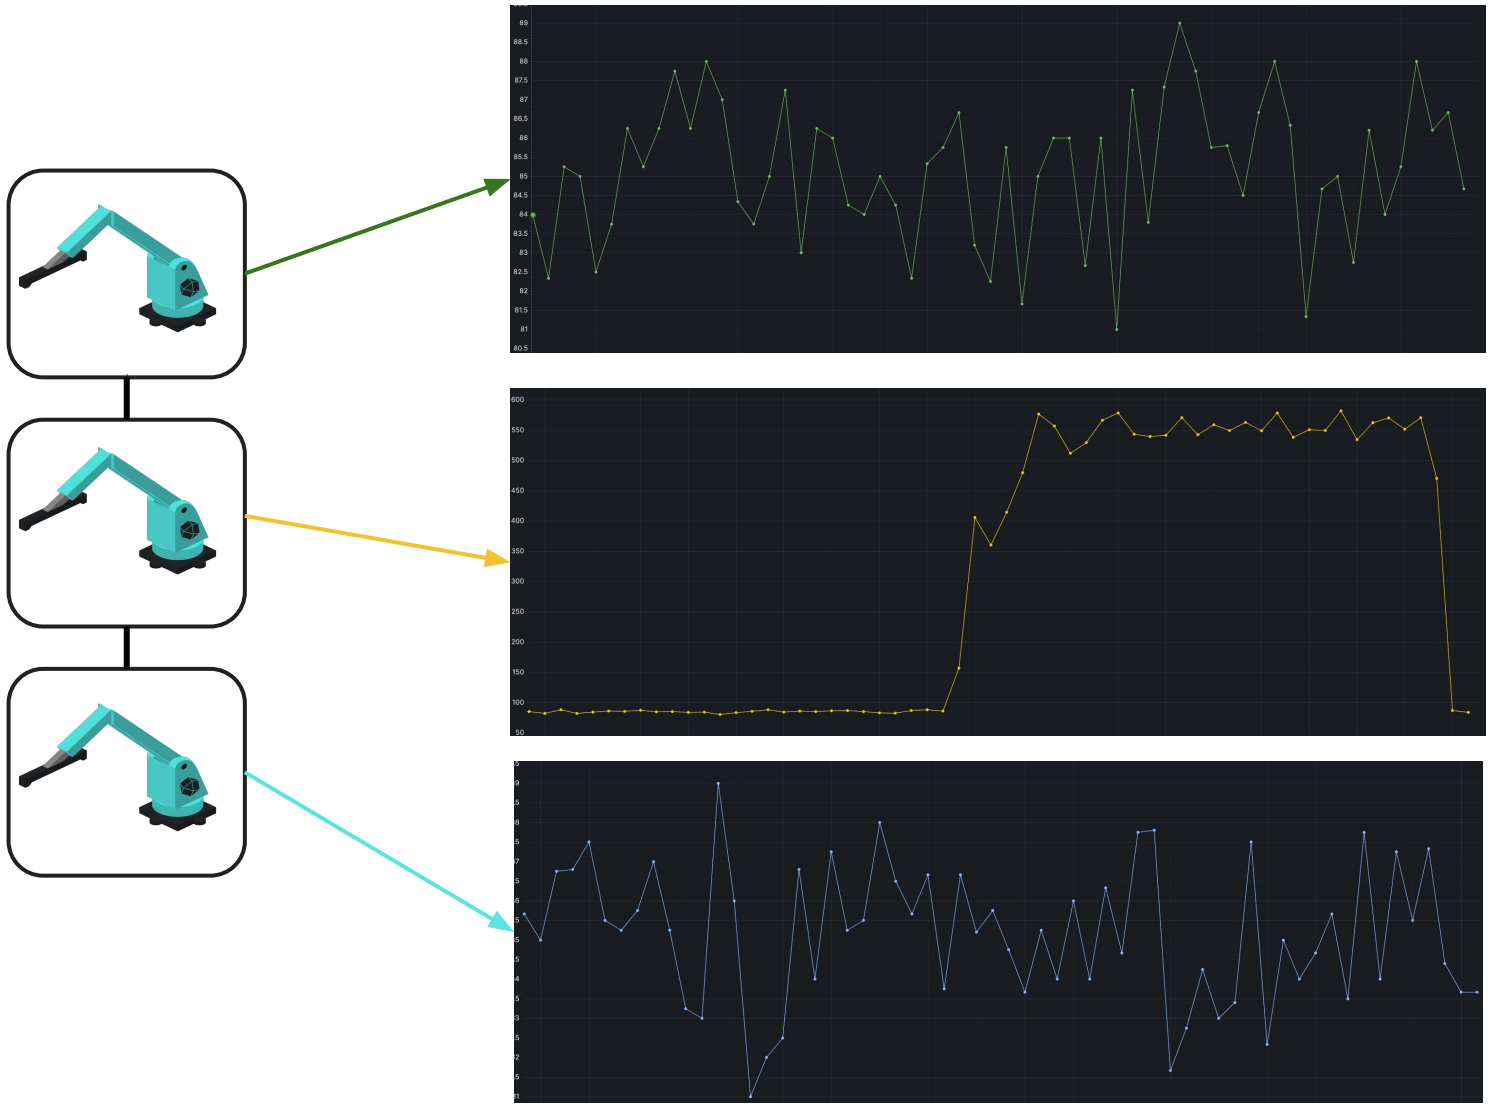

This could easily be solved with thresholding

Quix | influxdata®

“Success comes from the domain experts” - Jay Clifford

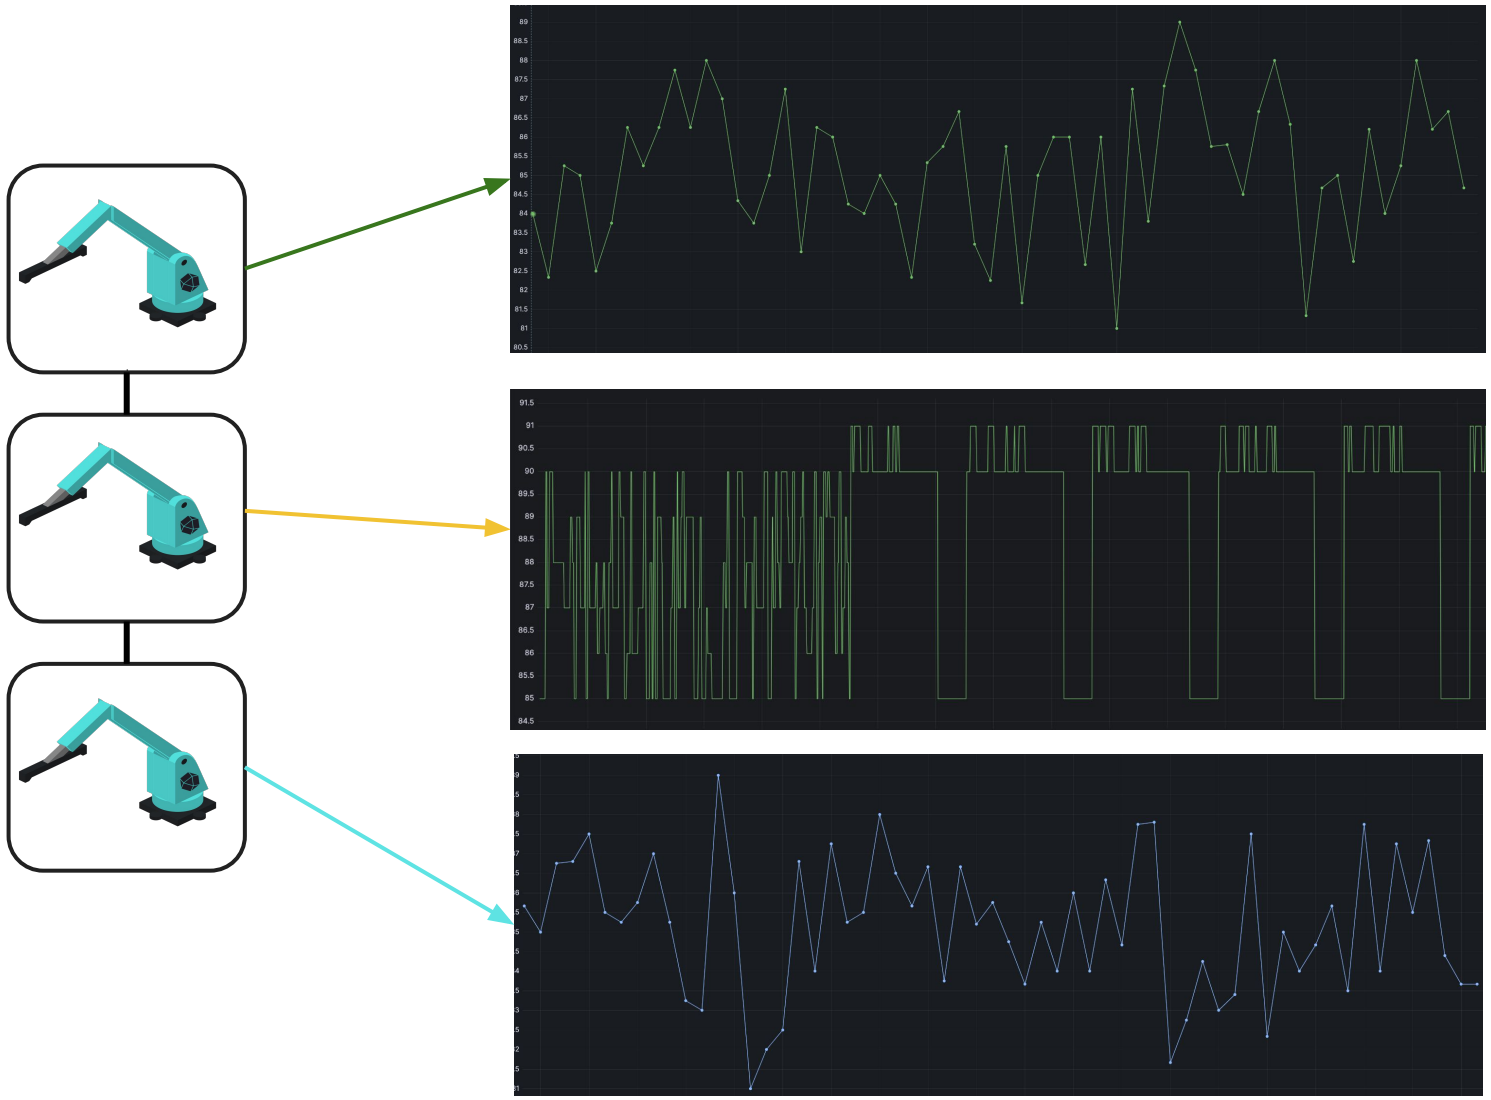

What do we do when our result becomes unpredictable by conventional means?

Quix | influxdata®

Machine learning – Autoencoder

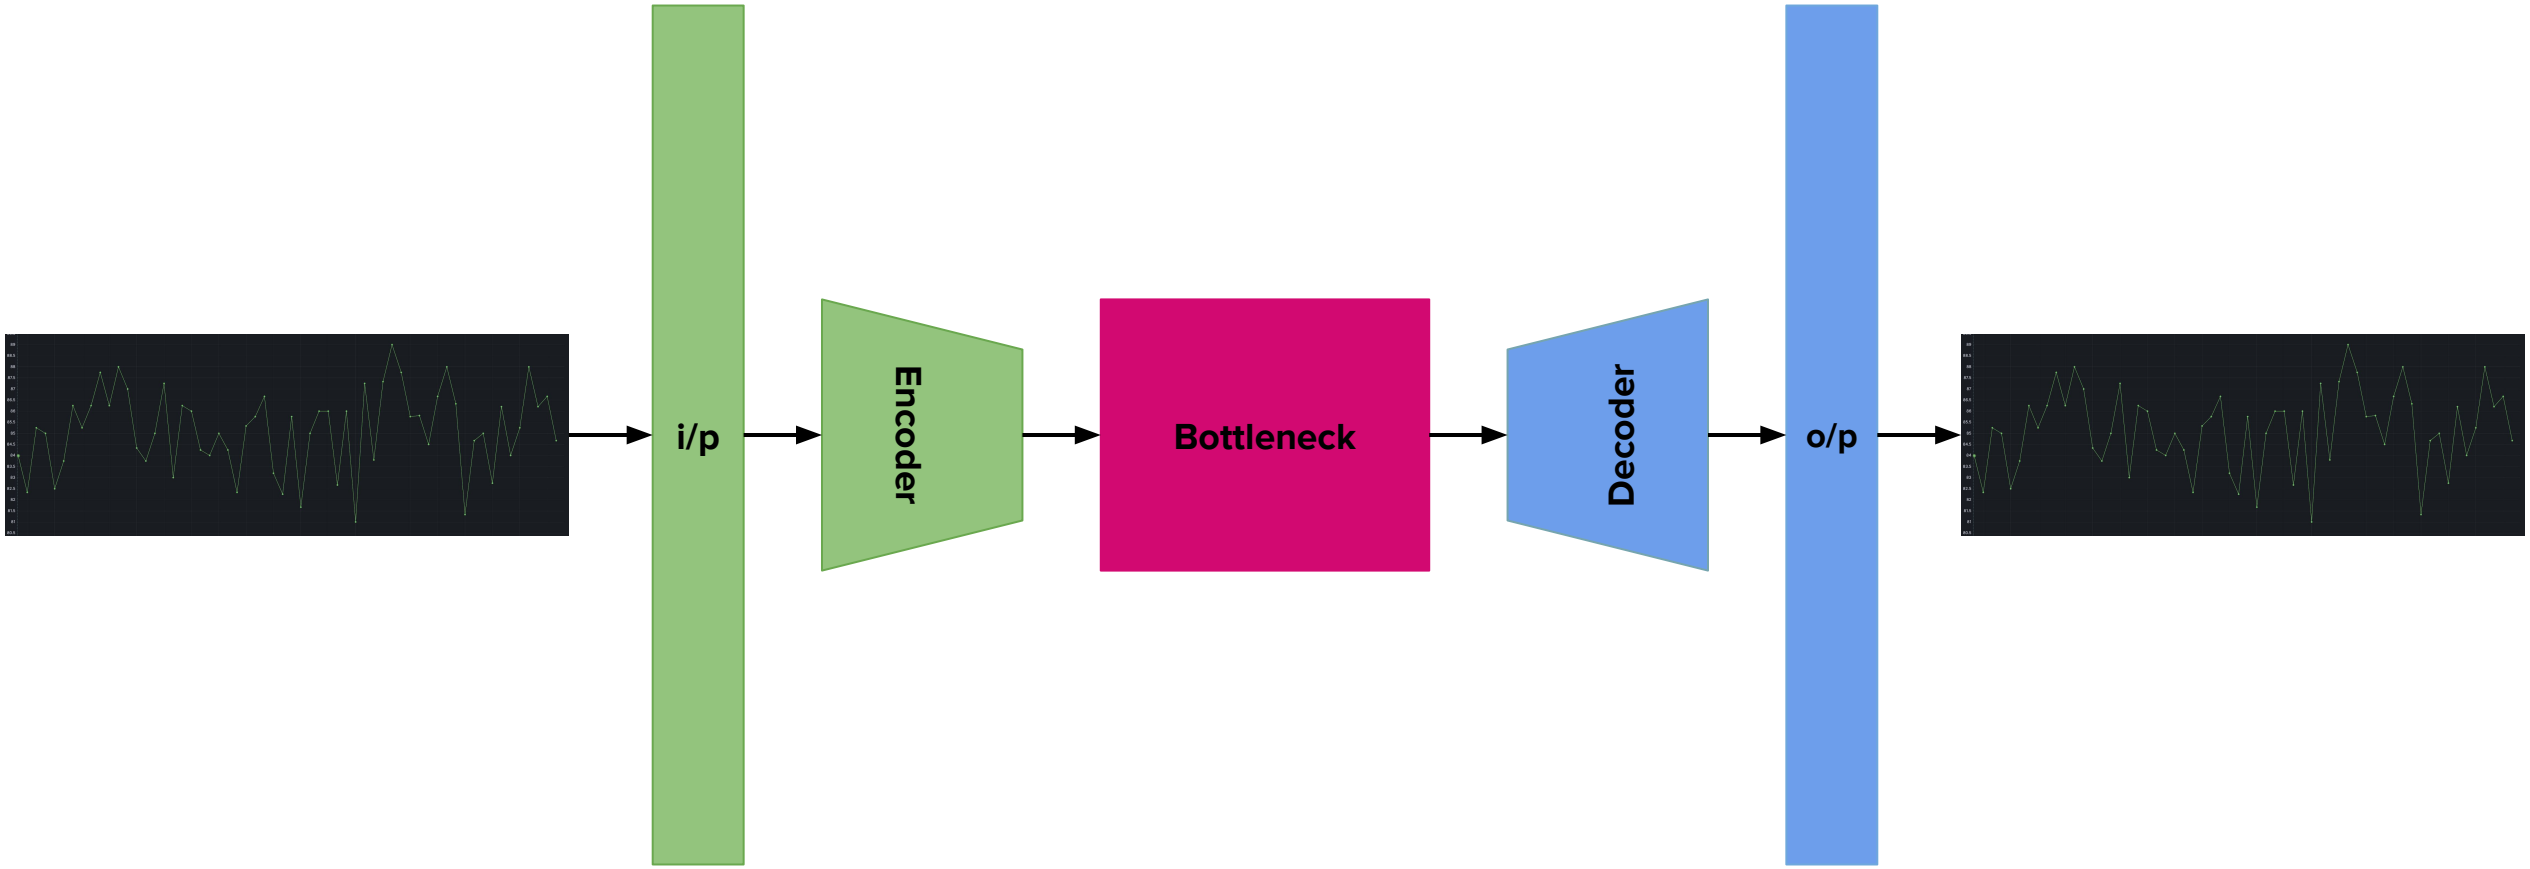

ML Deployment

Hugging Face

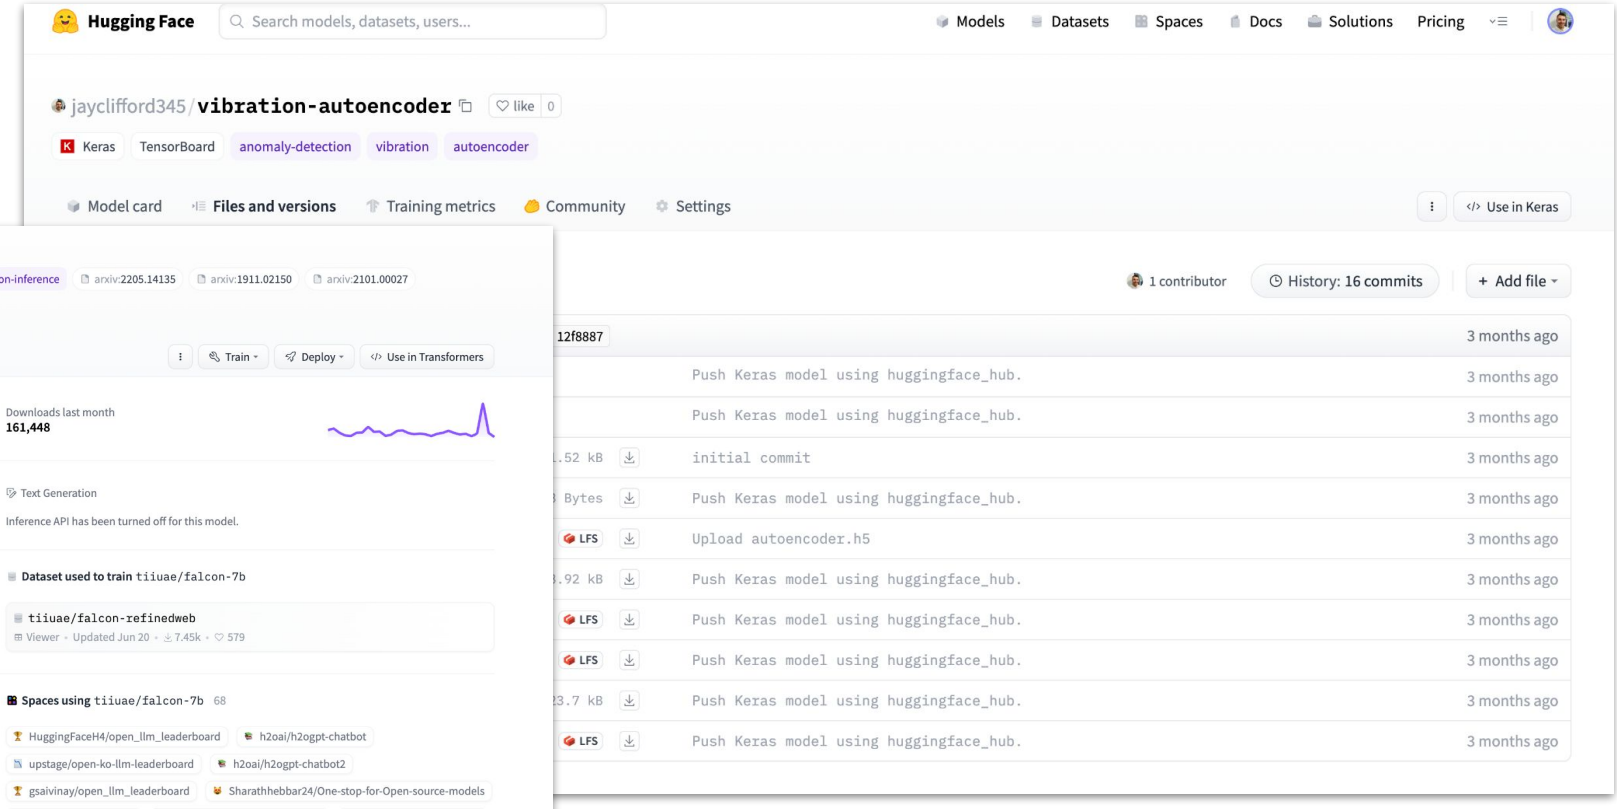

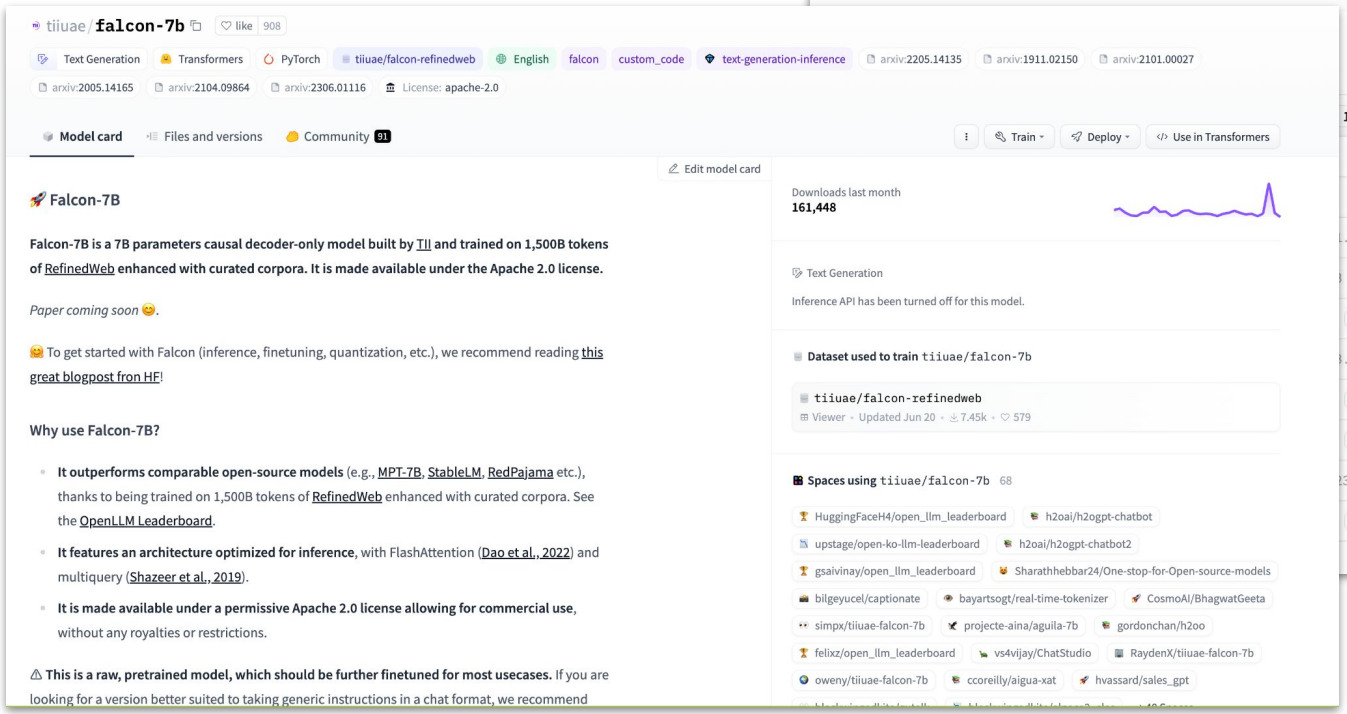

Solution Architecture (ML Deployment)

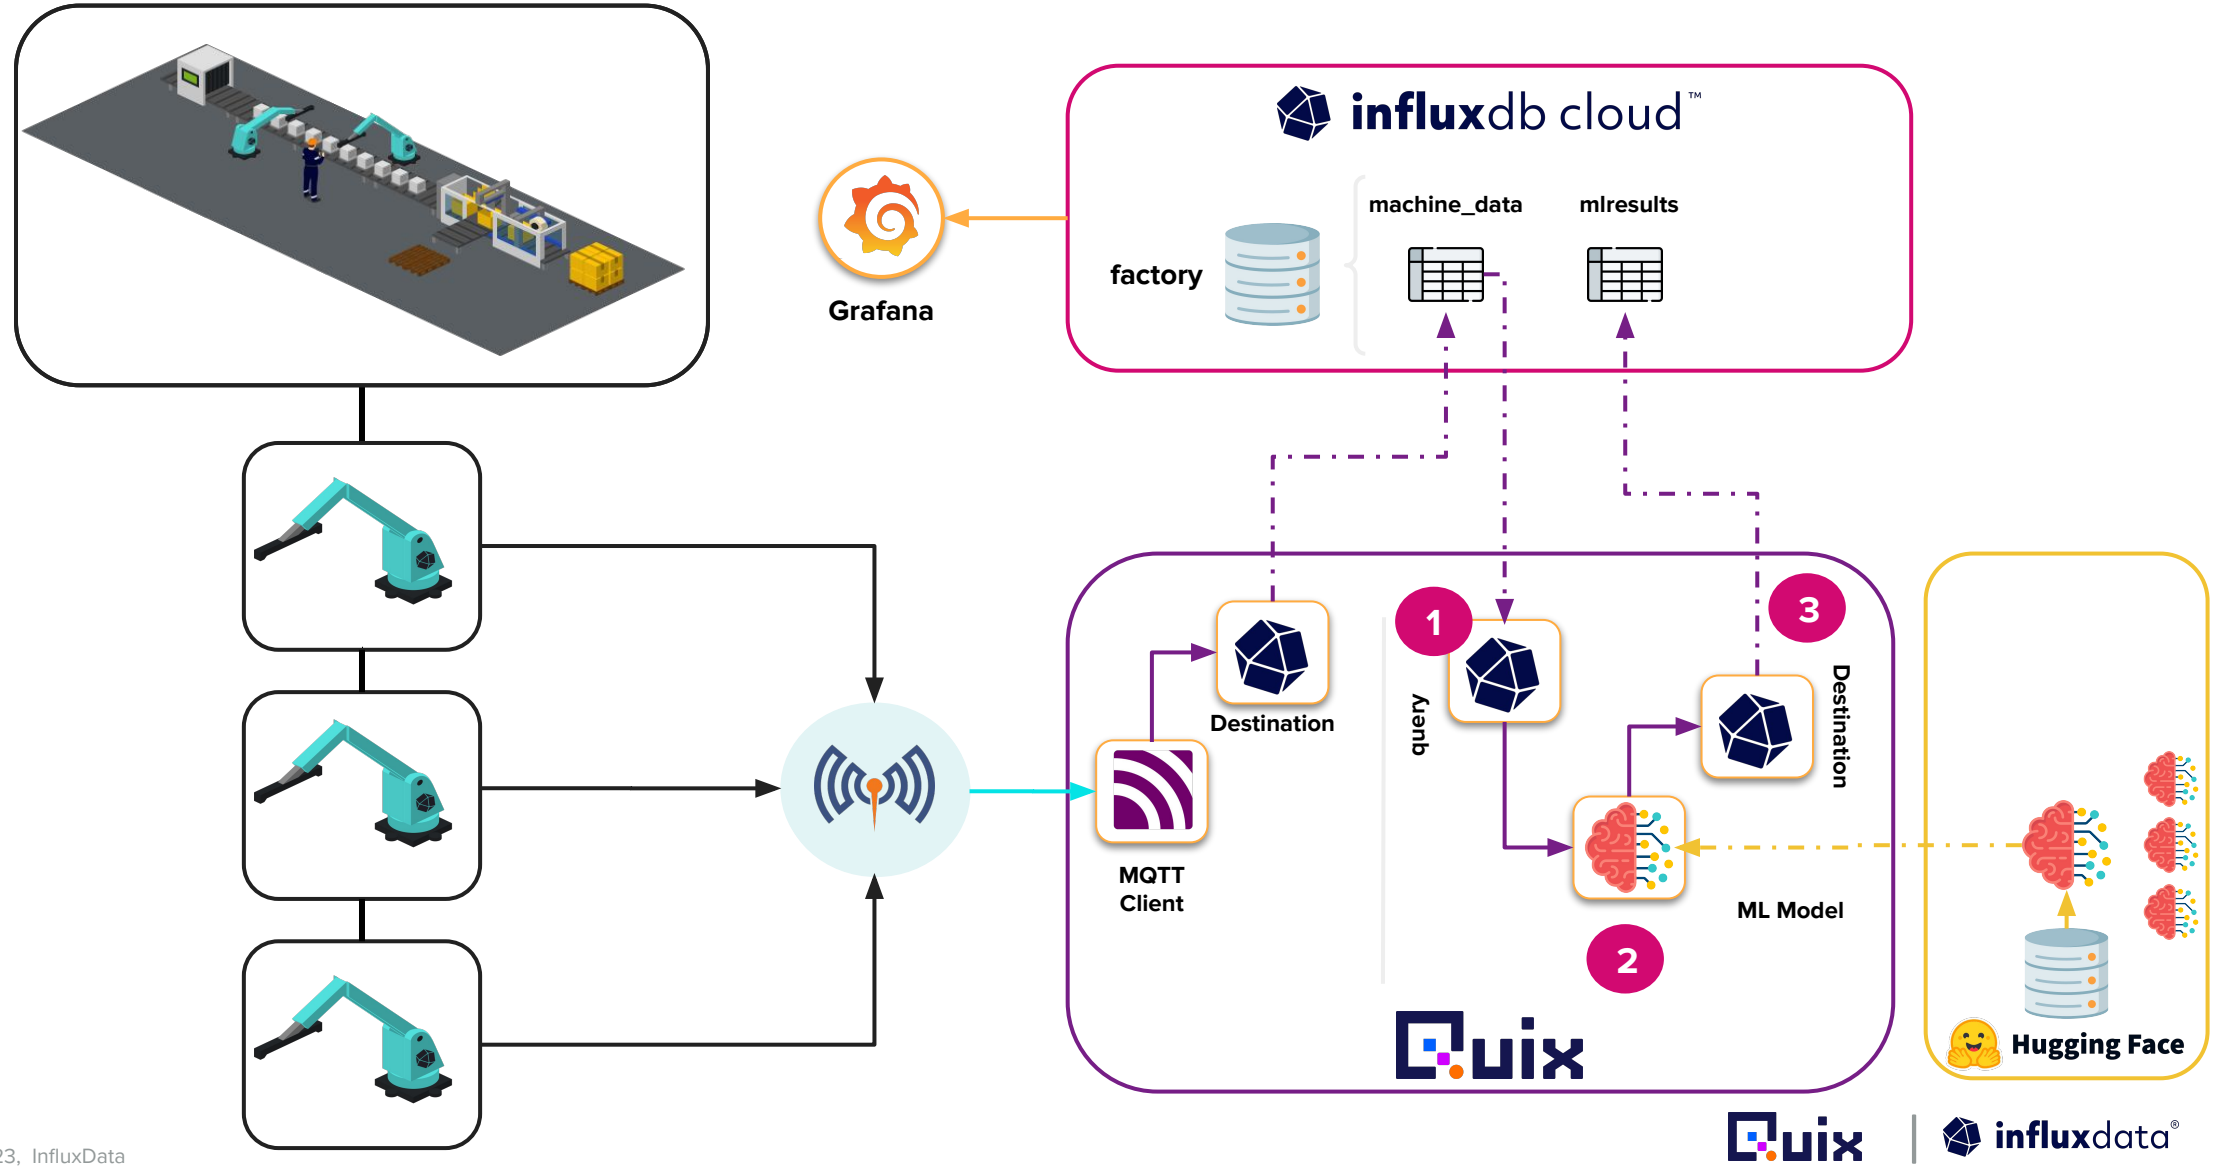

So what have solved?

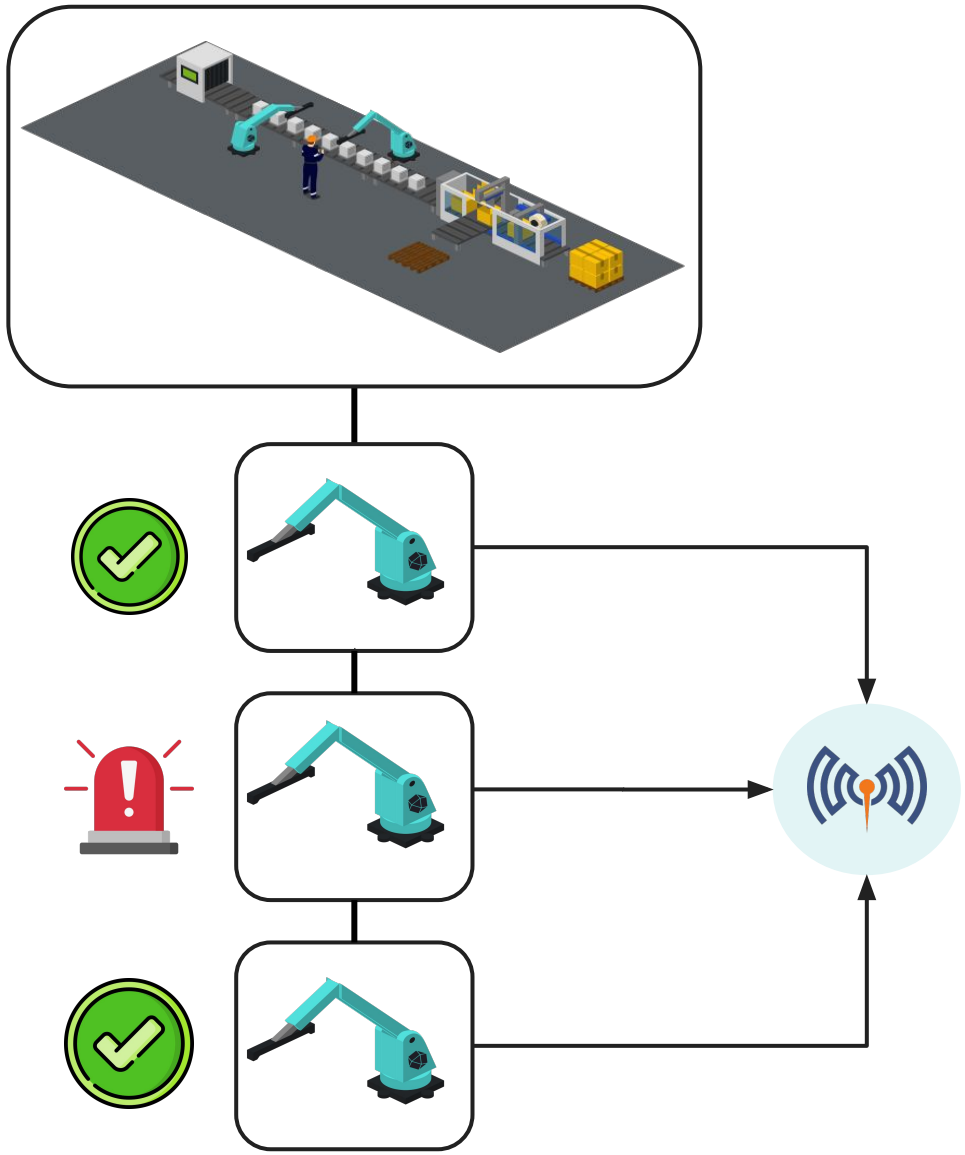

😃 Packing Co — Happy!

✅ Enabled the **ingest, transformation and storage** of their machine data.

✅ Deployed an initial **machine learning model** to detect potential malfunctions using vibration data from the machines.

✅ Provided the foundations of a **scalable data pipeline**.

✅ **Saved the holidays**.

The history of observability

What is observability?

🤓 In control theory, observability is a measure of how well internal states of a system can be inferred from knowledge of its external outputs.

😶🌫️ In distributed systems, observability is the ability to collect data about programs' execution, modules' internal states and the communication among components.

🤩 Observability is a full understanding of our systems.

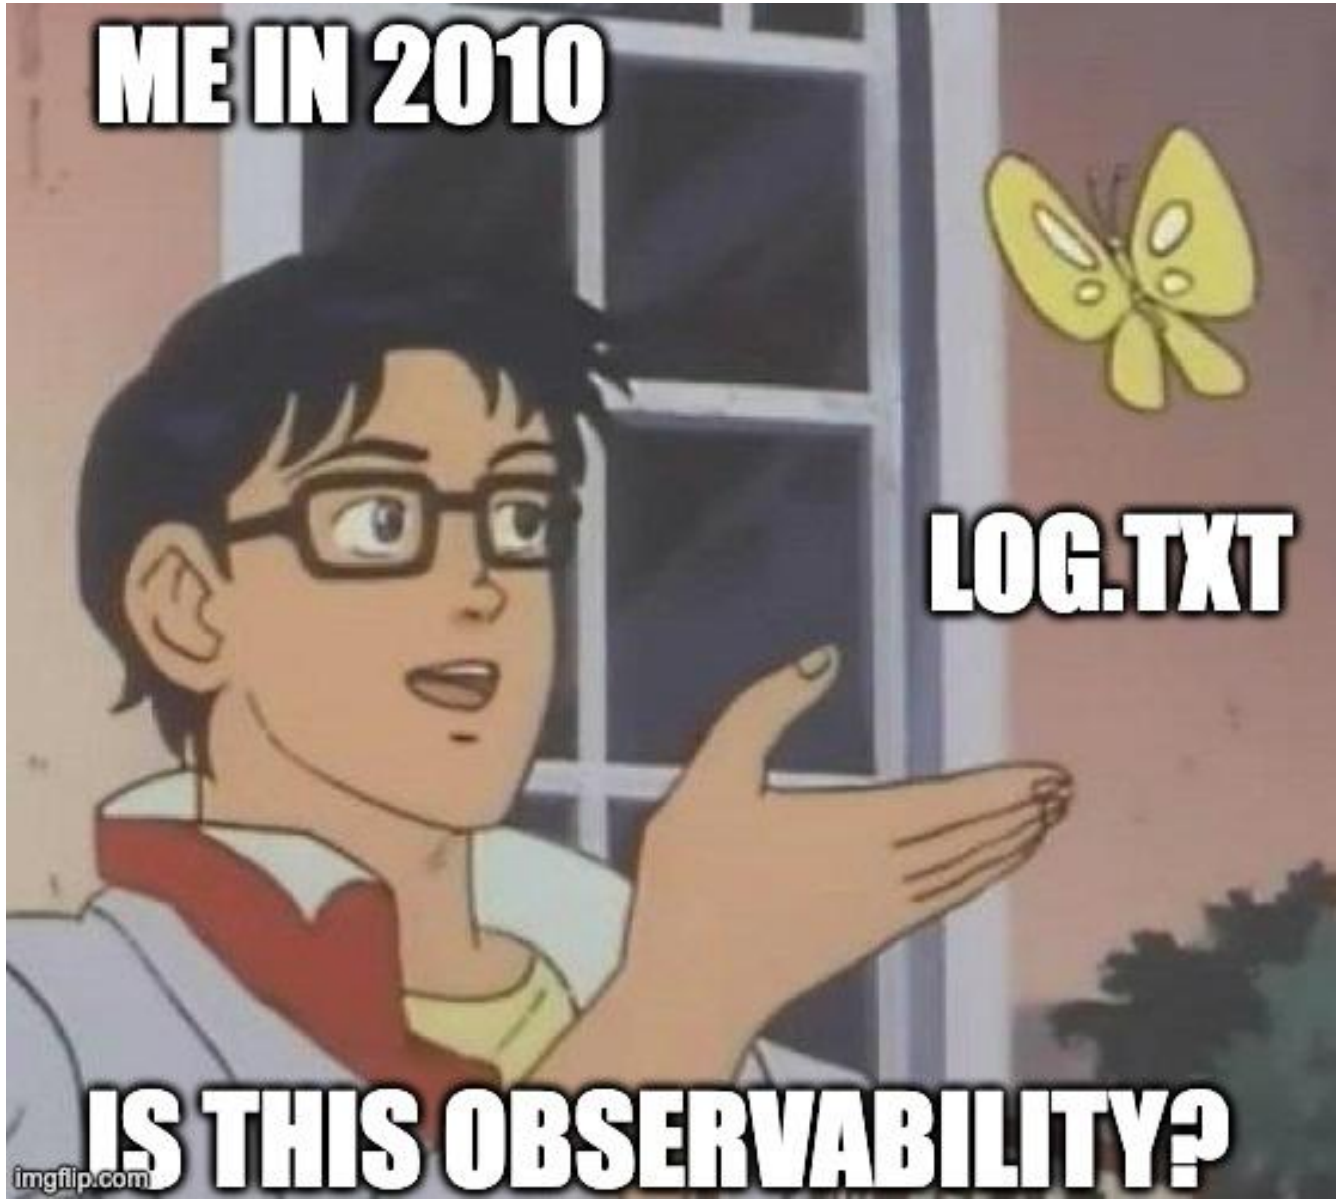

Quix | influxdata®

>_SSH

>_SSH

>_SSH

>_SSH

Quix | influxdata®

2015

Quix | influxdata®

2015-2017

JAEGER

OPENTRACING

OpenCensus

2019

JAEGER

OPENTRACING

OpenCensus

Quix | influxdata®

2019

OpenTelemetry
OTel
OTEL

Quix | influxdata®

Unification

2022:

- Logs push stdout logfmt or JSON Elastic, filesystem
- Metrics pull HTTP Prometheus exposition InfluxDB, Prometheus
- Traces push UDP Jaeger thrift or gRPC Elastic, Cassandra

2023:

- Logs push gRPC
- Metrics push gRPC
- Traces push gRPC

Quix | influxdata®

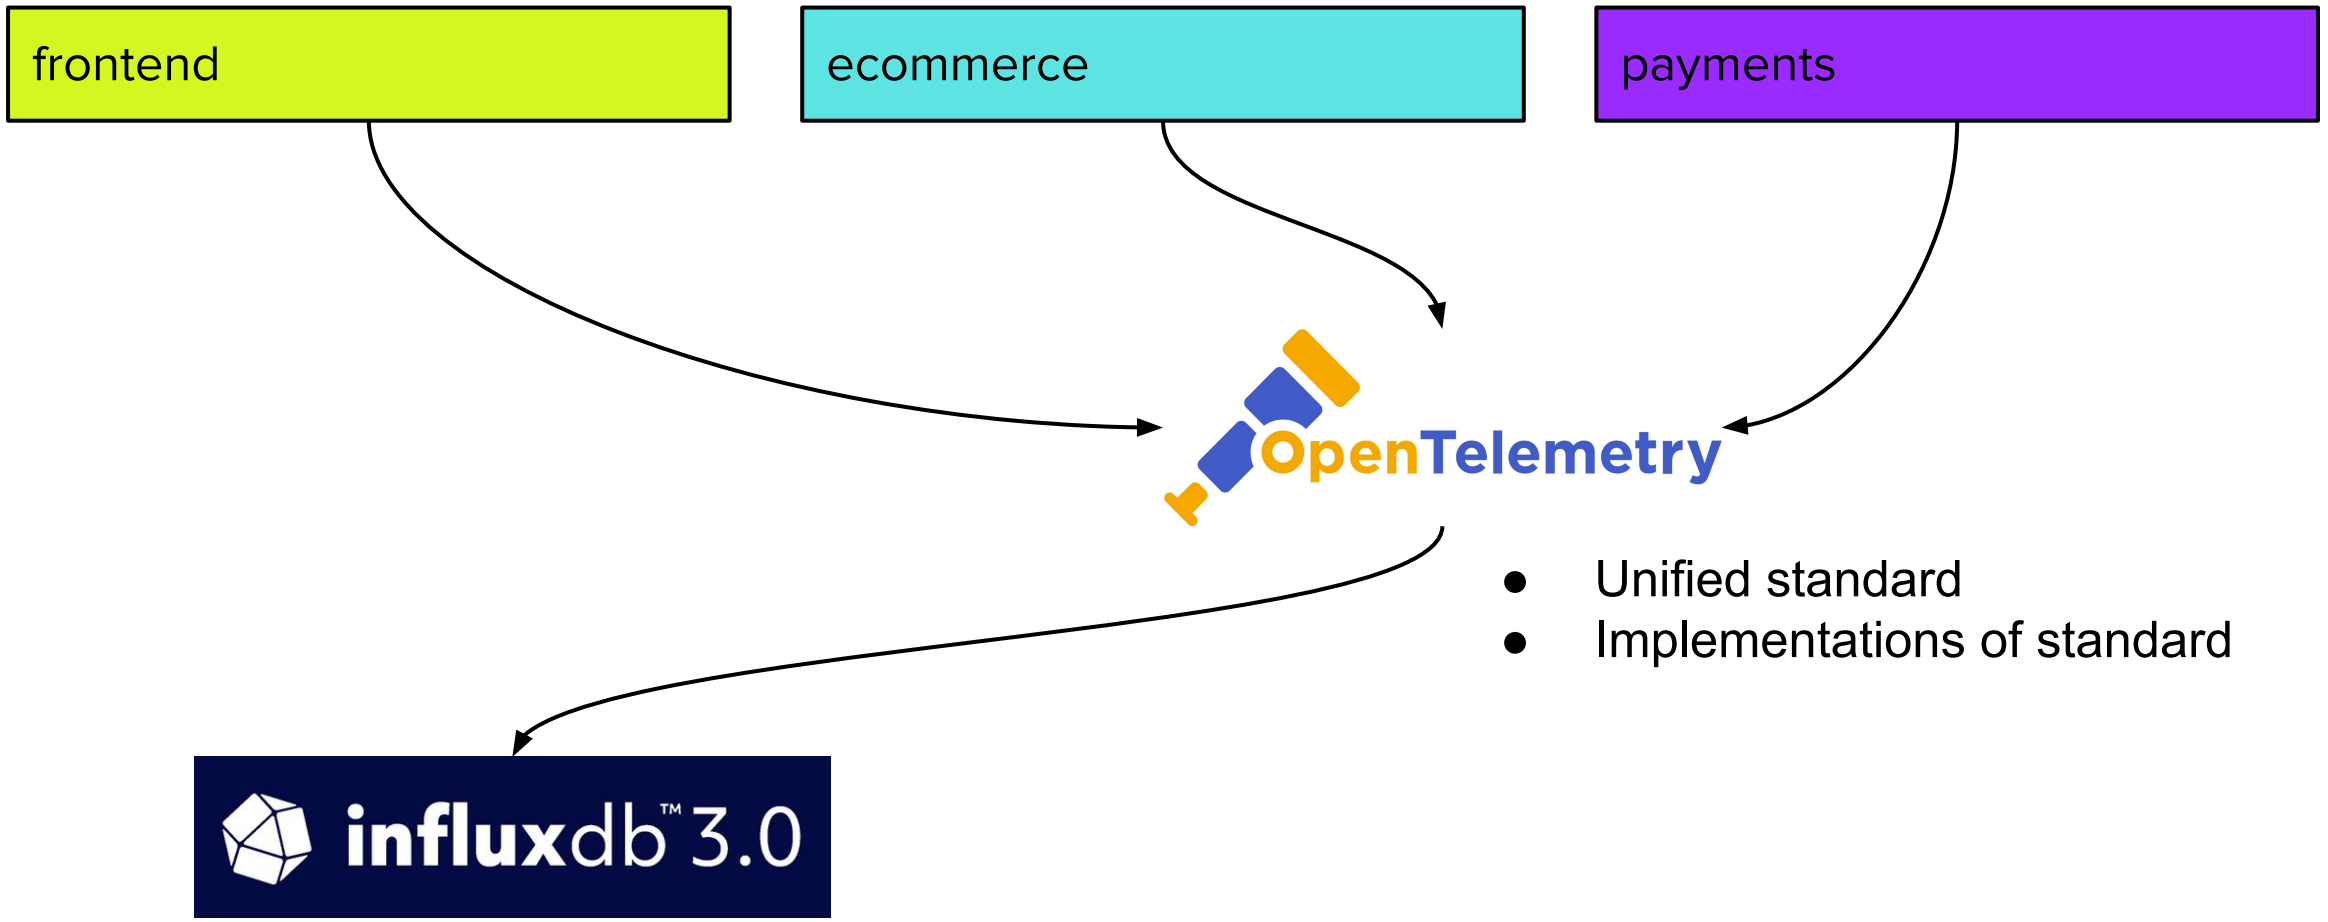

- Unified standard
- Implementations of standard

OpenTelemetry in practice

Birds Eye view

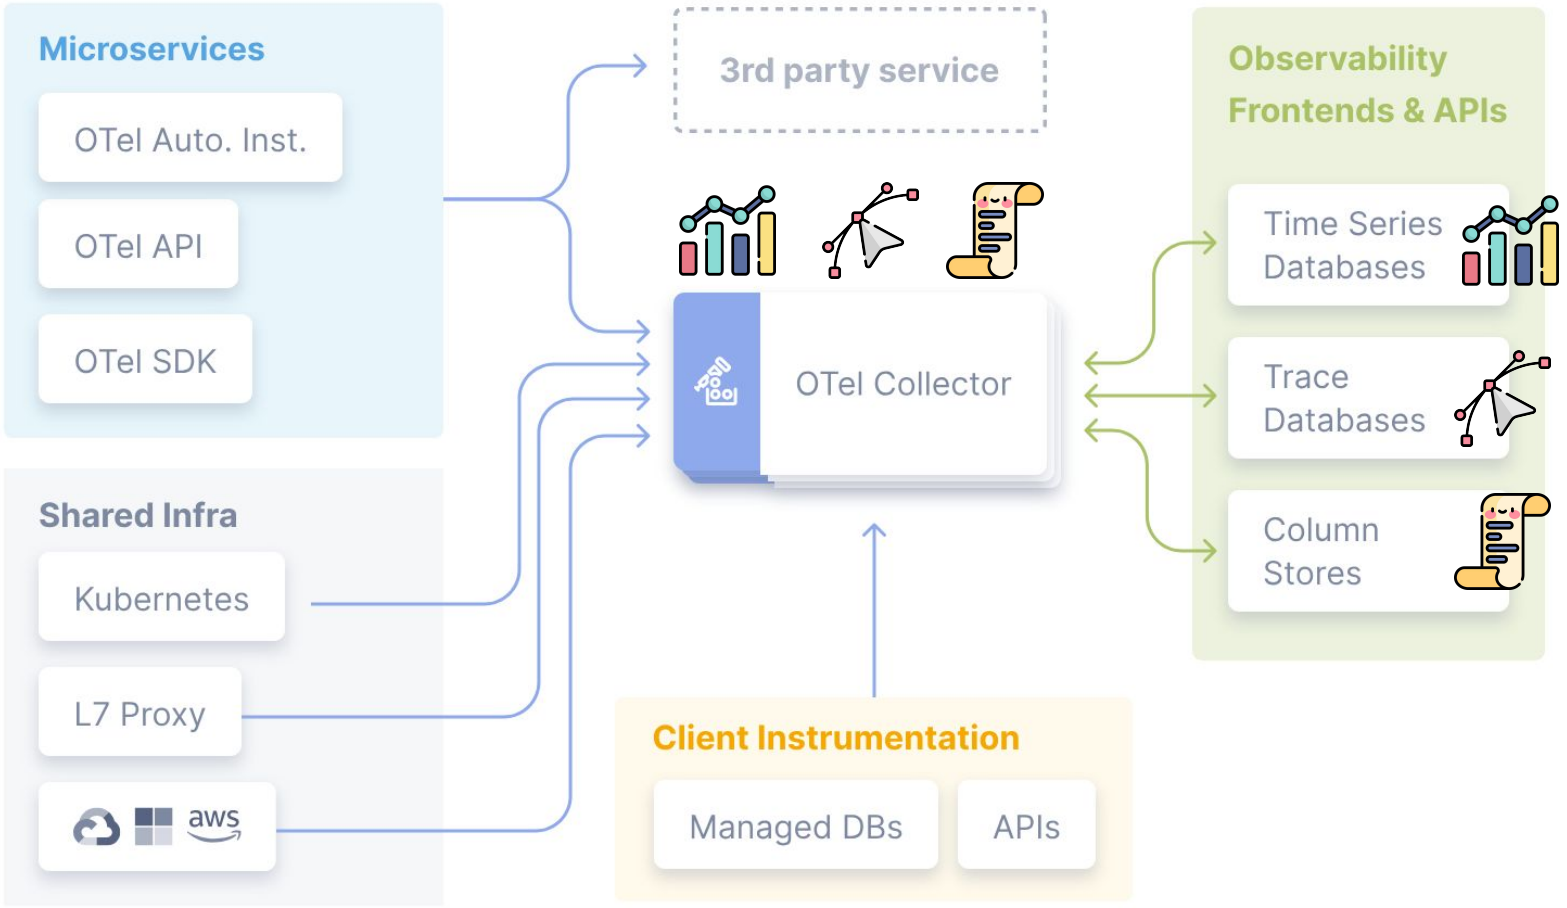

Quix | influxdata®

Our focus

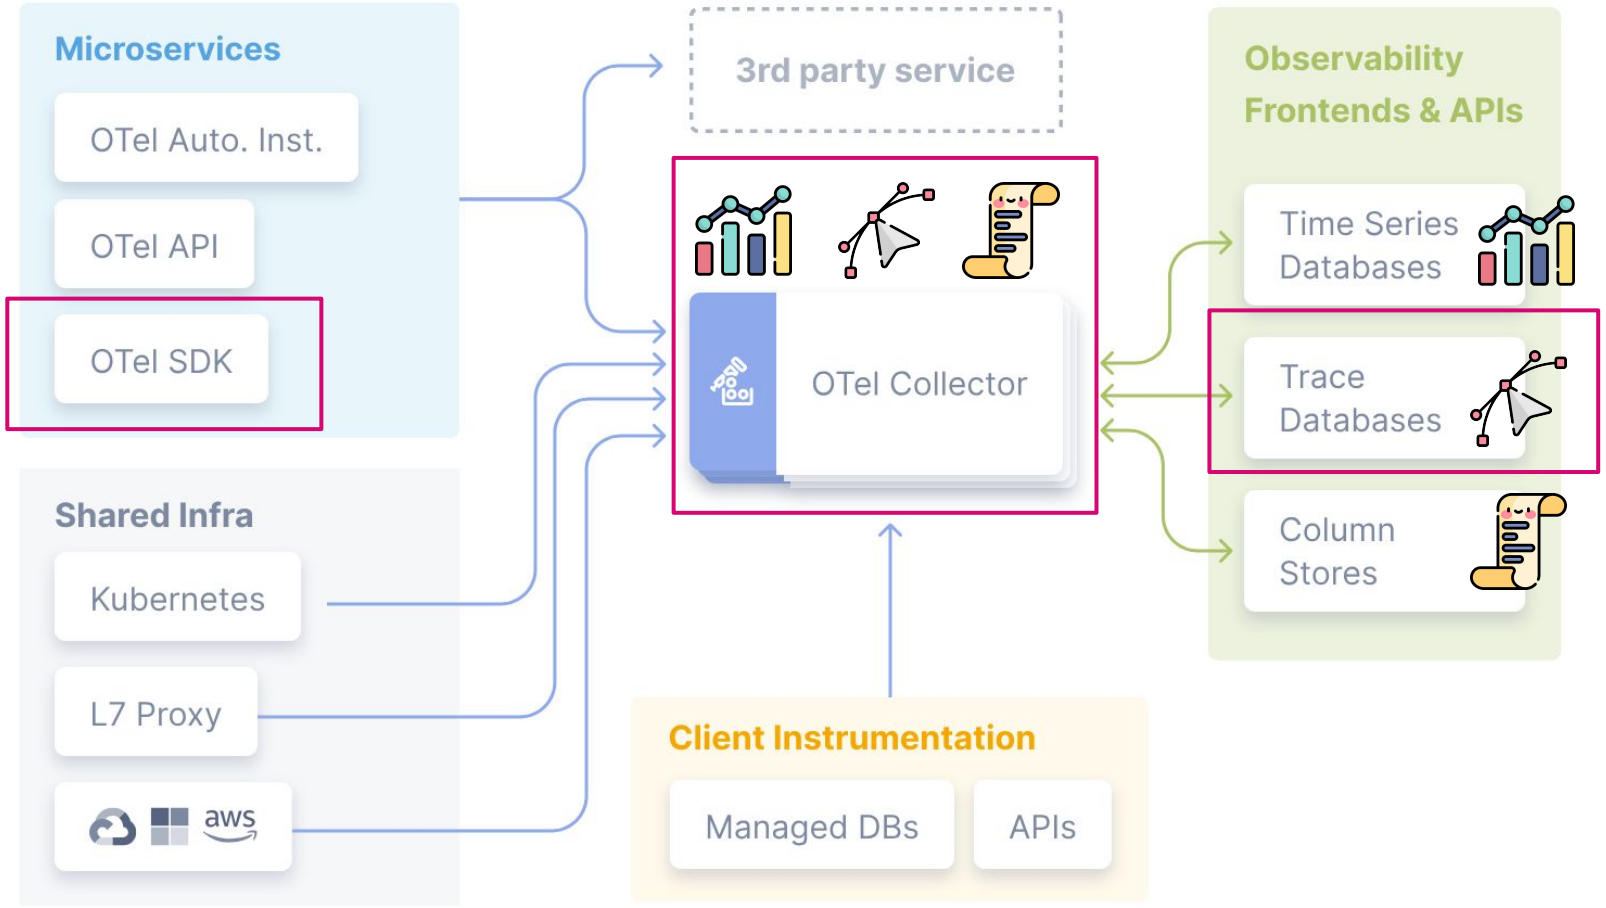

Quix | influxdata®

What exactly is a trace?

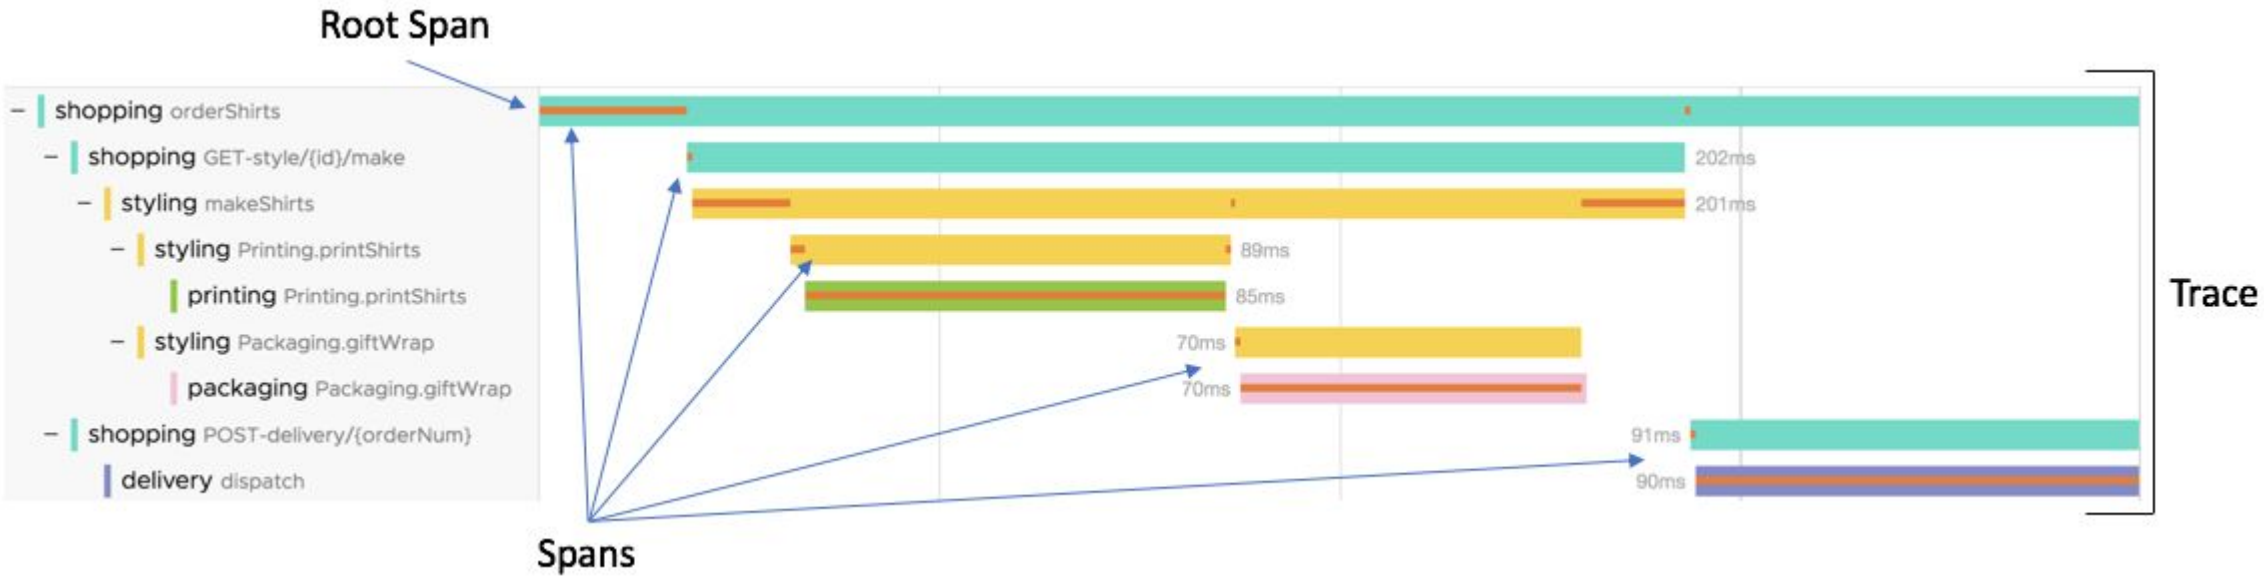

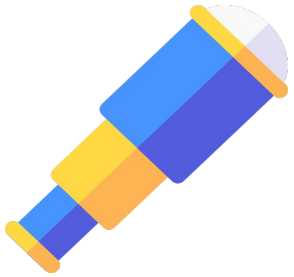

OpenTelemetry Collector

An open source agent that facilitates the collection, processing and export of telemetry data.

OpenTelemetry Collector			
Receivers	Processors	Connectors	Exporters
		 	influxdb™

Quix | influxdata®

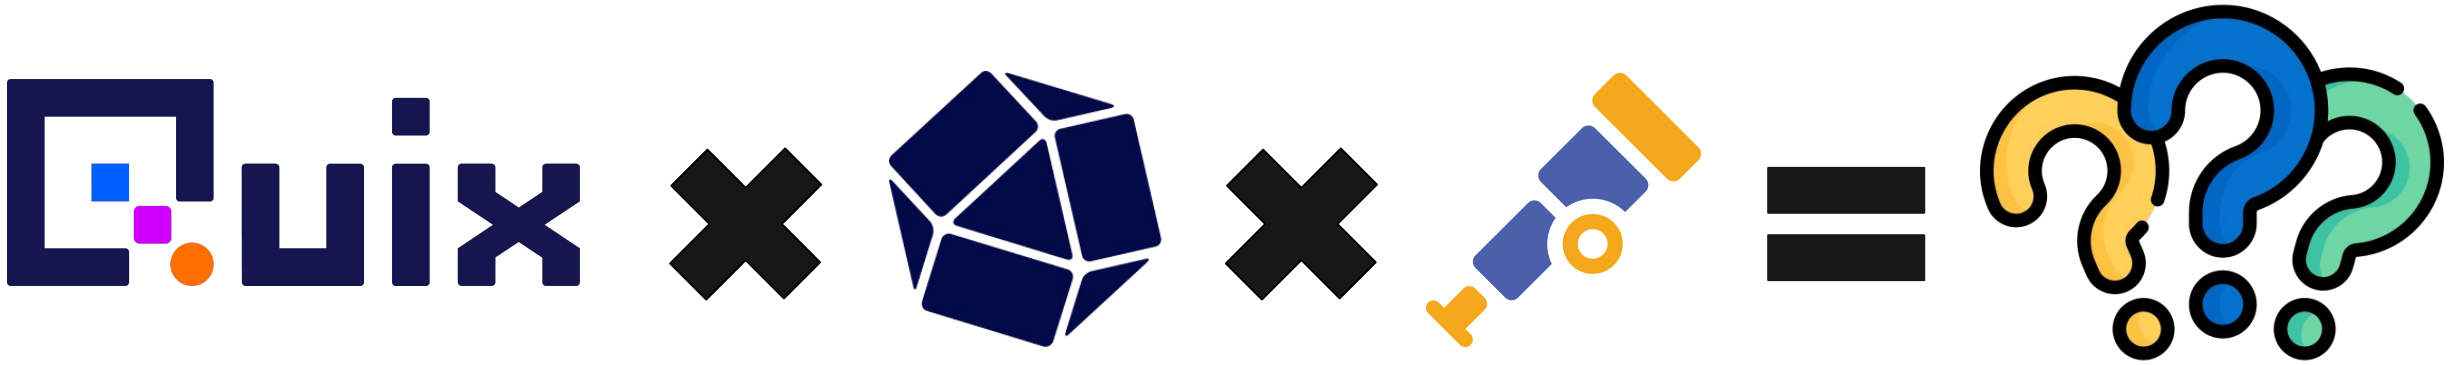

Solution Architecture (Otel)

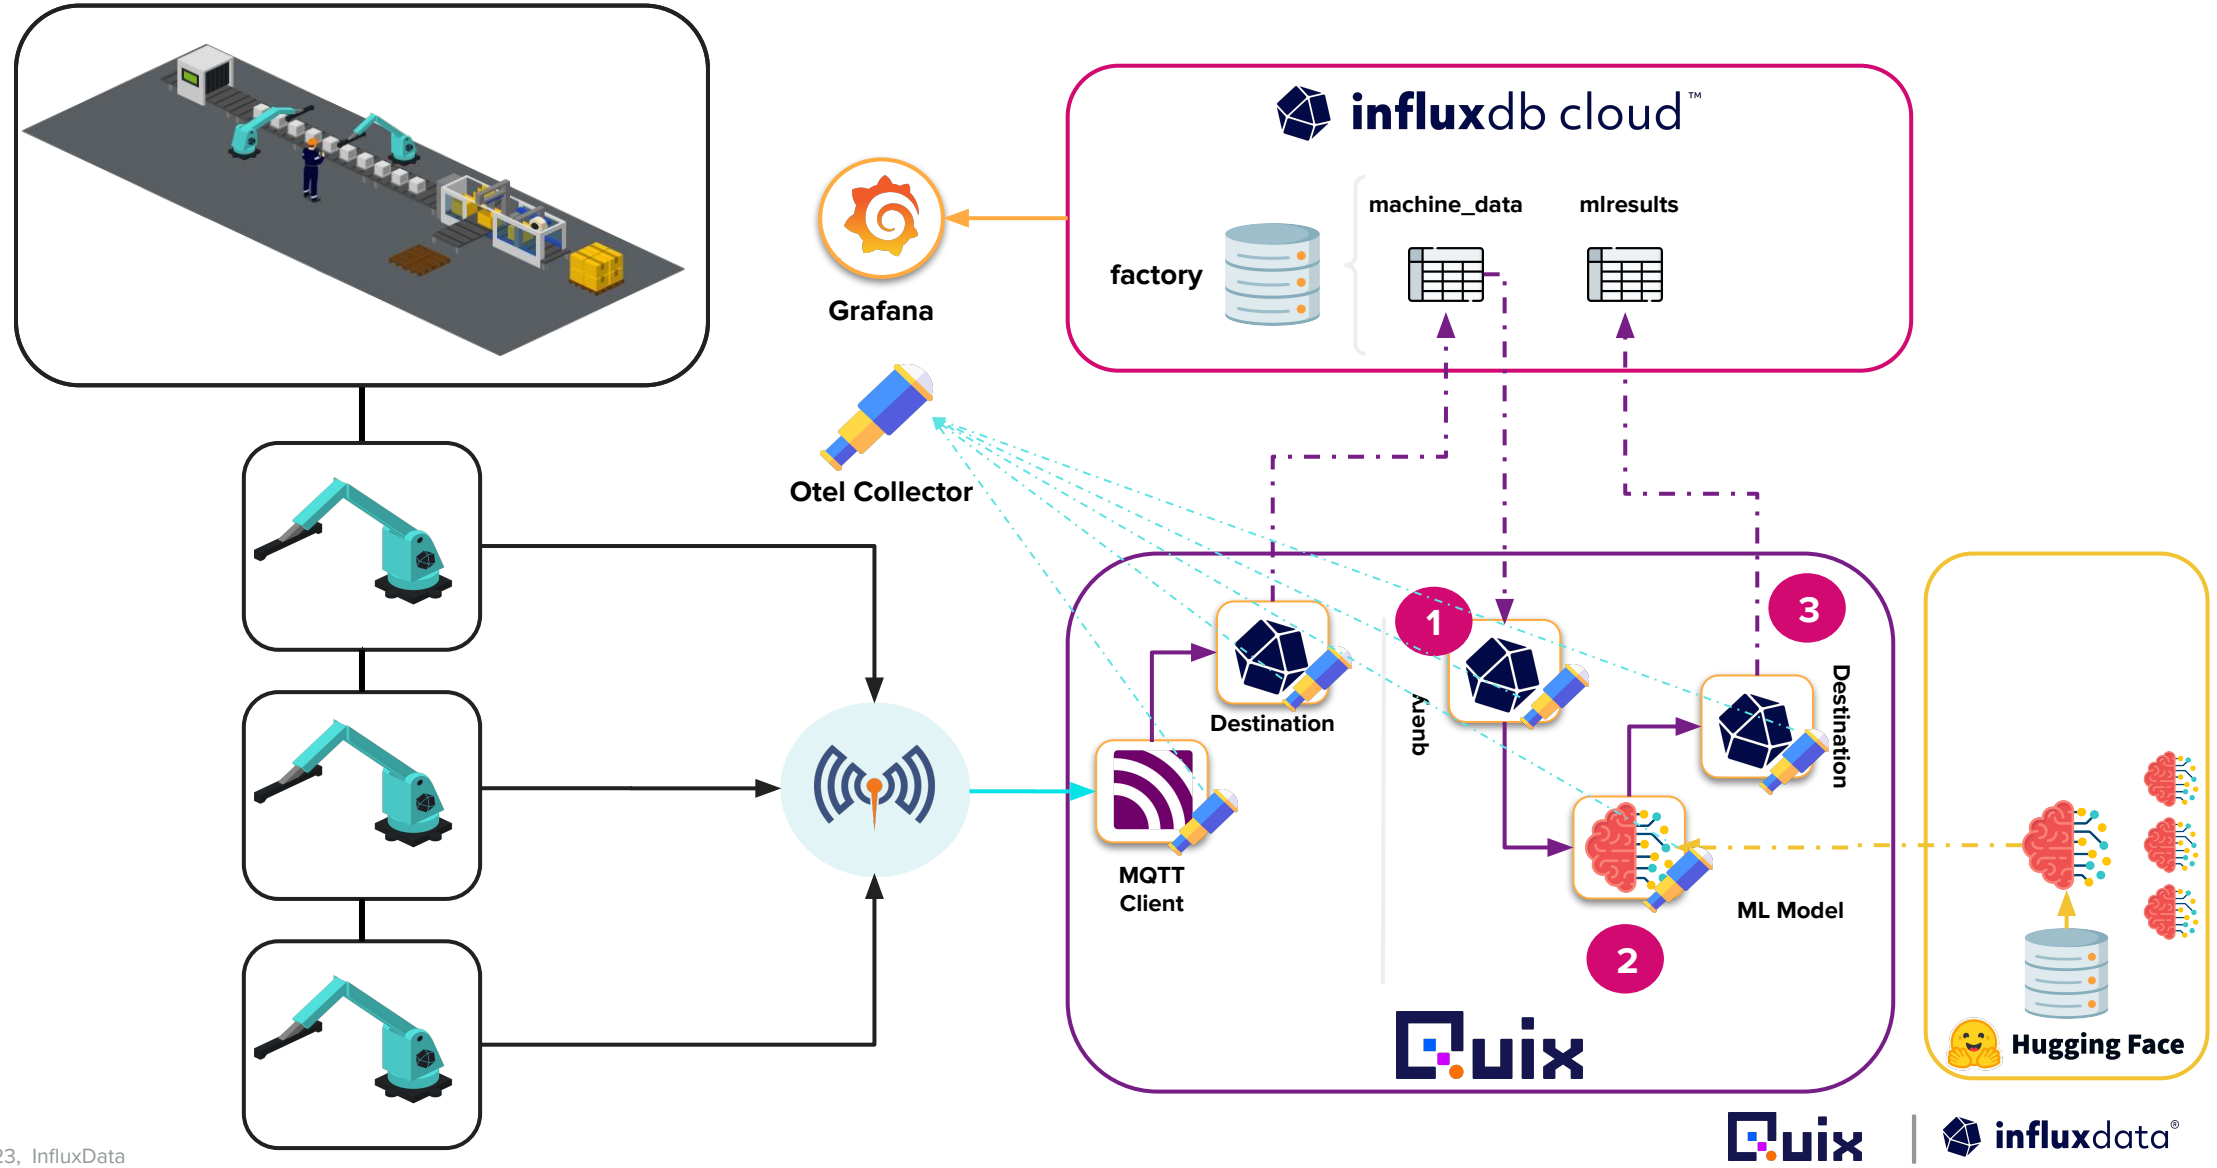

Quix | influxdata®

Next Steps

Where could we go next?

Parallel Model Deployment

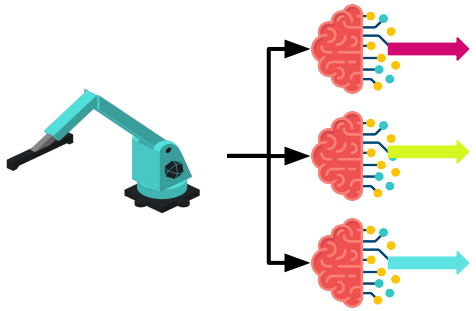

Closing the loop (Eventually)

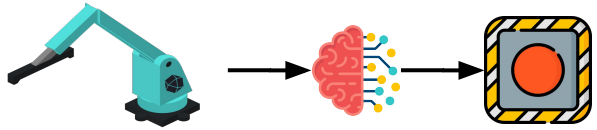

Quix | influxdata®

Try the demo yourself

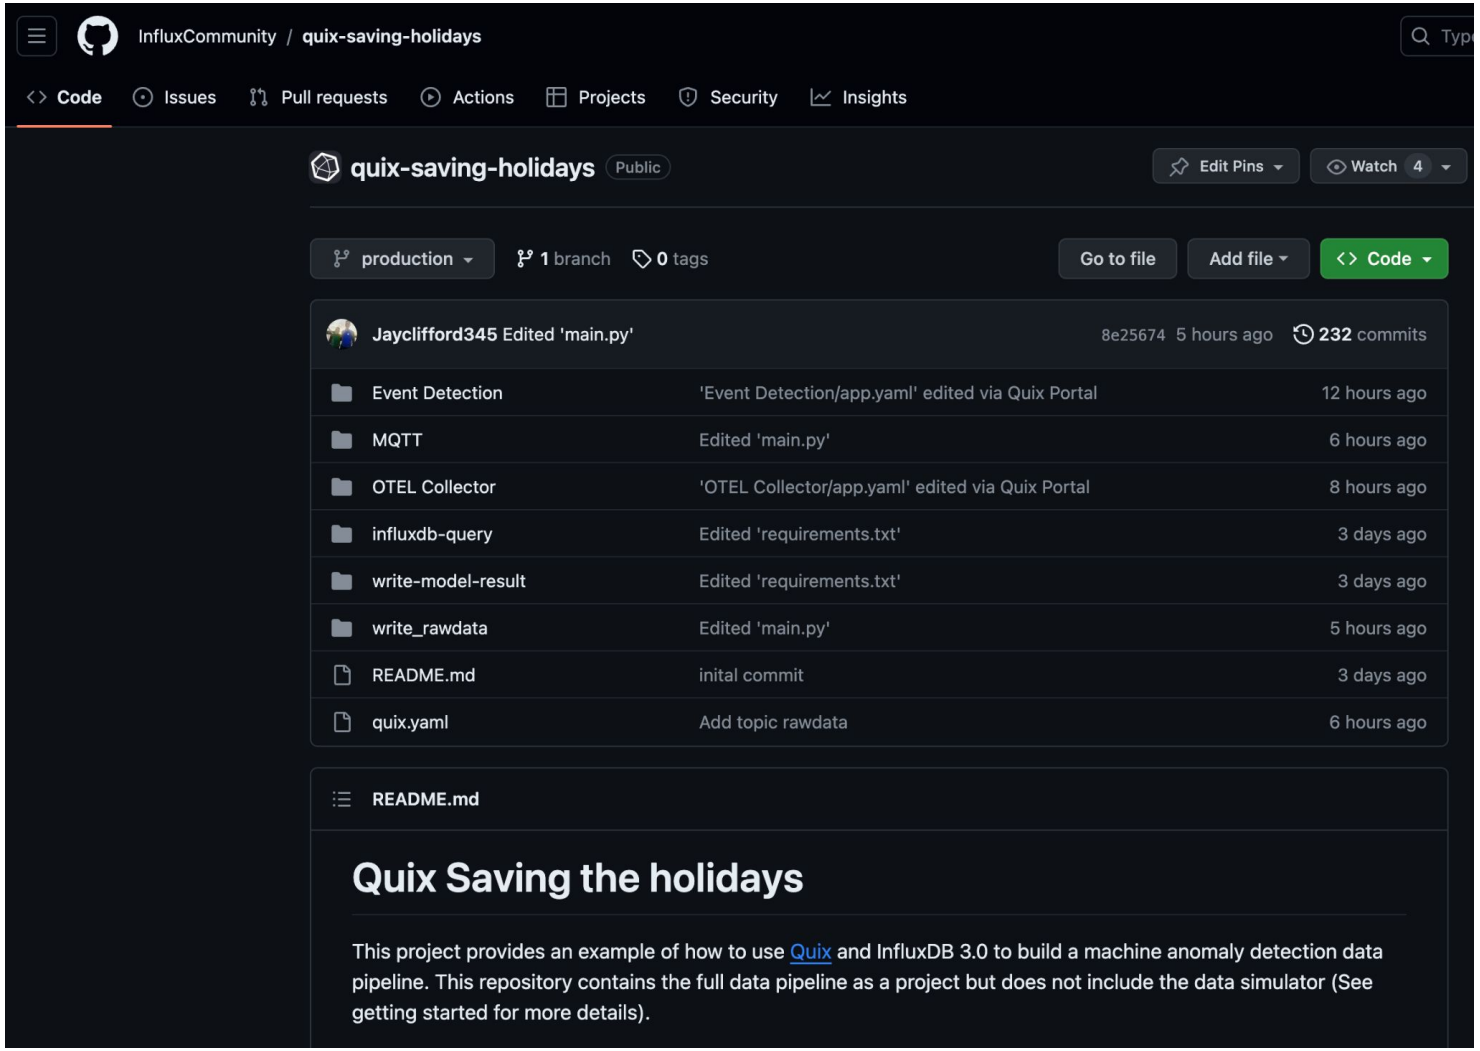

https://github.com/InfluxCommunity/quix-saving-holidays

Quix | influxdata®

Getting started

Sign up

Influxdata.com ---------- Get InfluxDB

Via cloud marketplace ---------- aws Google Cloud Microsoft Azure

Learn

- [x] Self-service content
- [x] Documentation
- [x] InfluxDB University https://influxdbu.com/
- [x] Community https://influxcommunity.slack.com/

Quix | influxdata®

Getting started

Sign up

quix.io ---------- Get Quix

Bring your own cloud ---------- aws Google Cloud Microsoft Azure

Learn

- [x] Templates https://quix.io/templates
- [x] Documentation https://quix.io/docs
- [x] Quix Streams https://github.com/quixio/quix-streams
- [x] Community https://quix.io/slack-invite

Quix | influxdata®

THANK YOU

Enjoy the Open Source Data Summit!

Quix | influxdata®

influxdata®

www.influxdata.com

Demo Screenshots

New Project

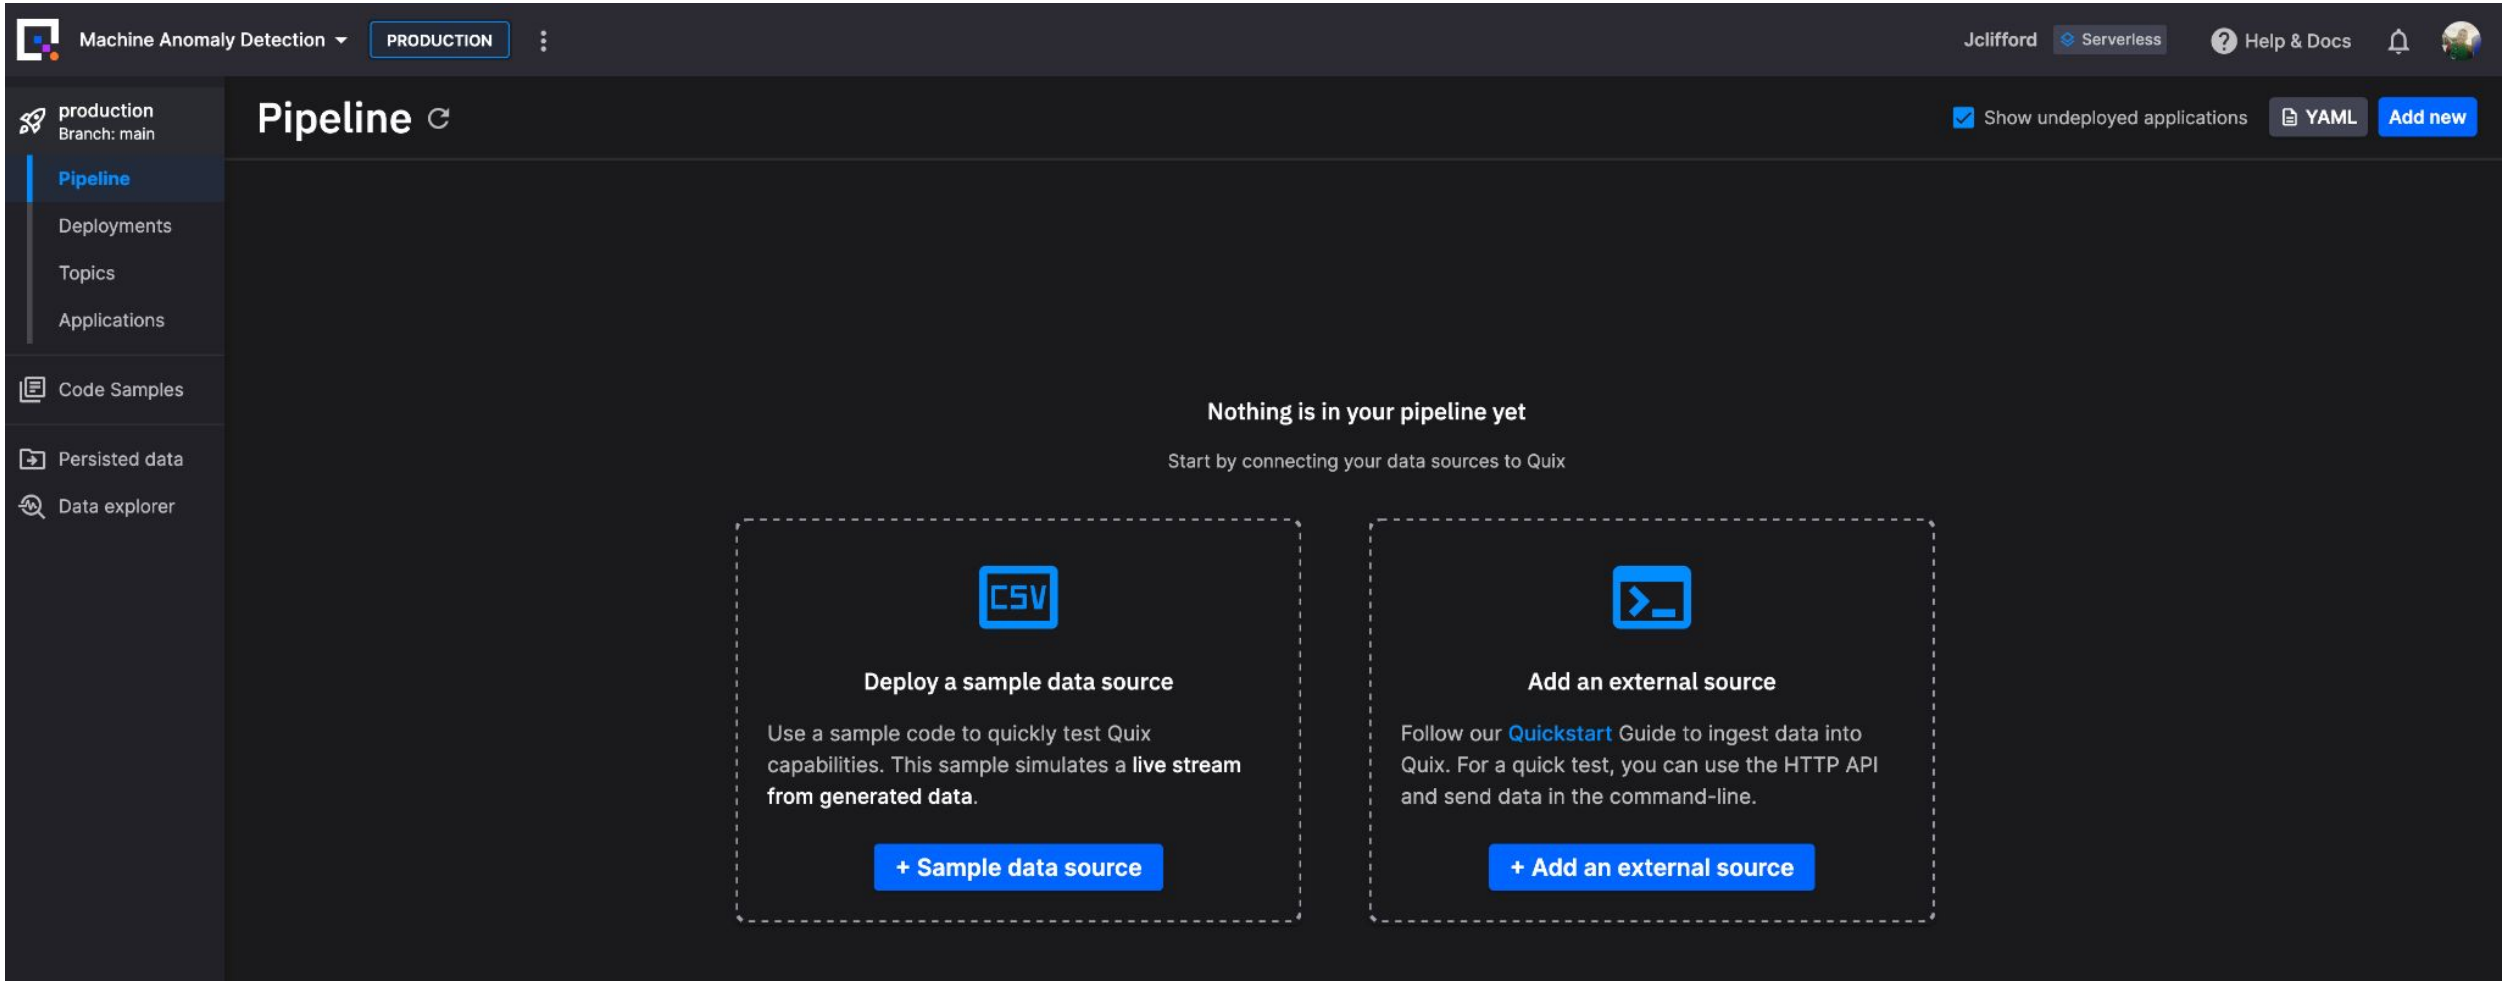

Quix | influxdata®

MQTT → InfluxDB

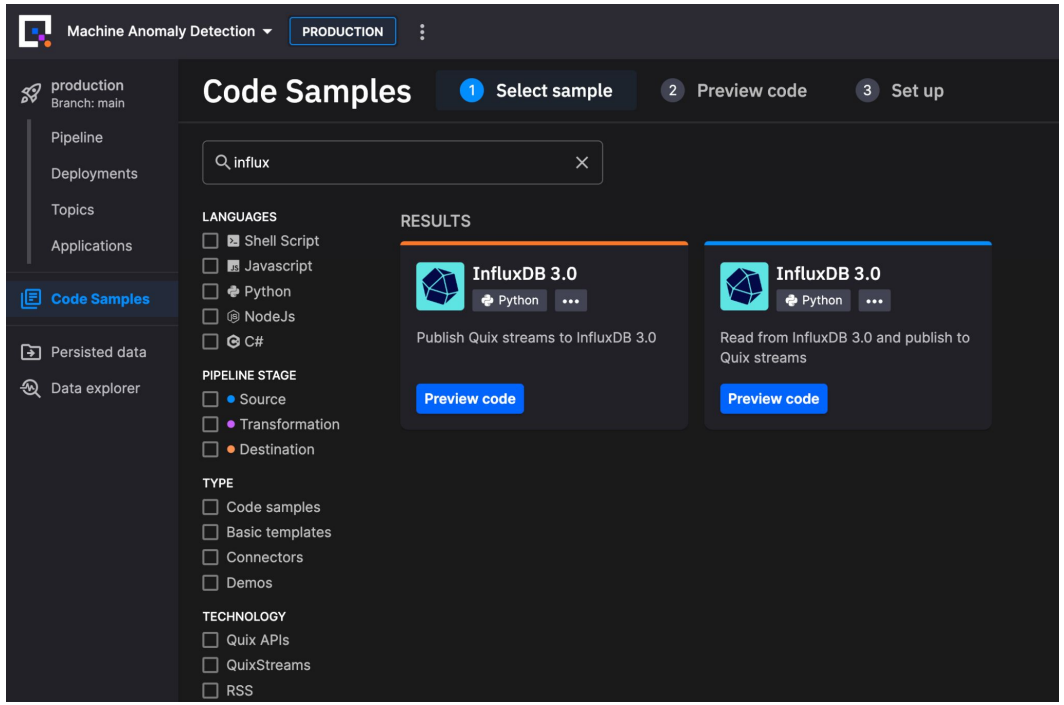

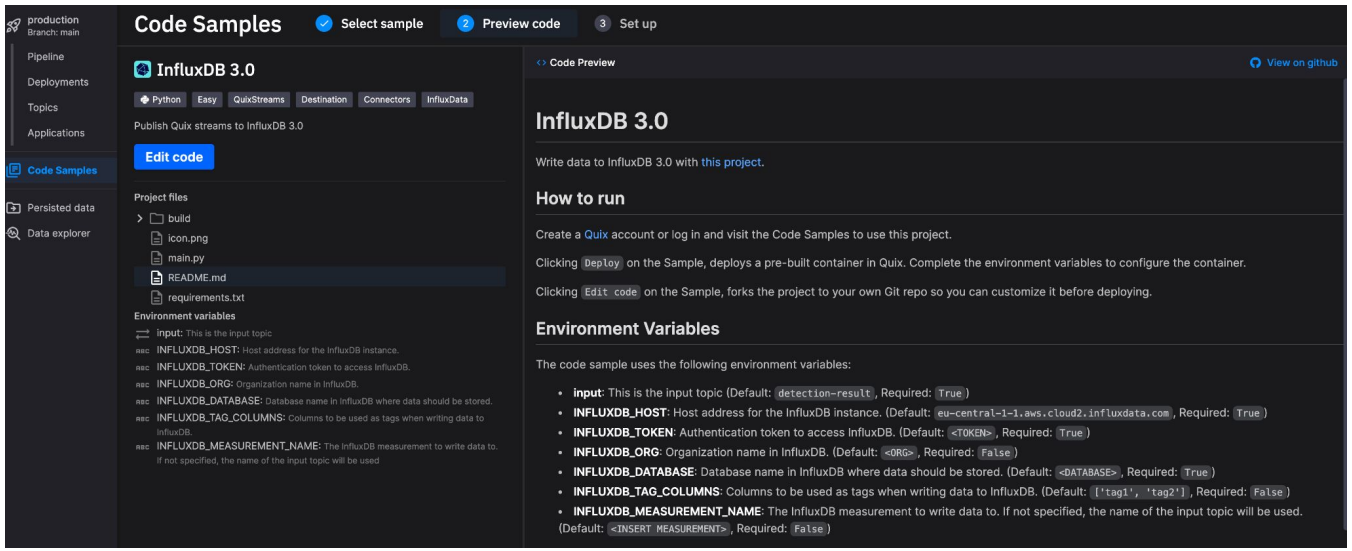

MQTT → InfluxDB

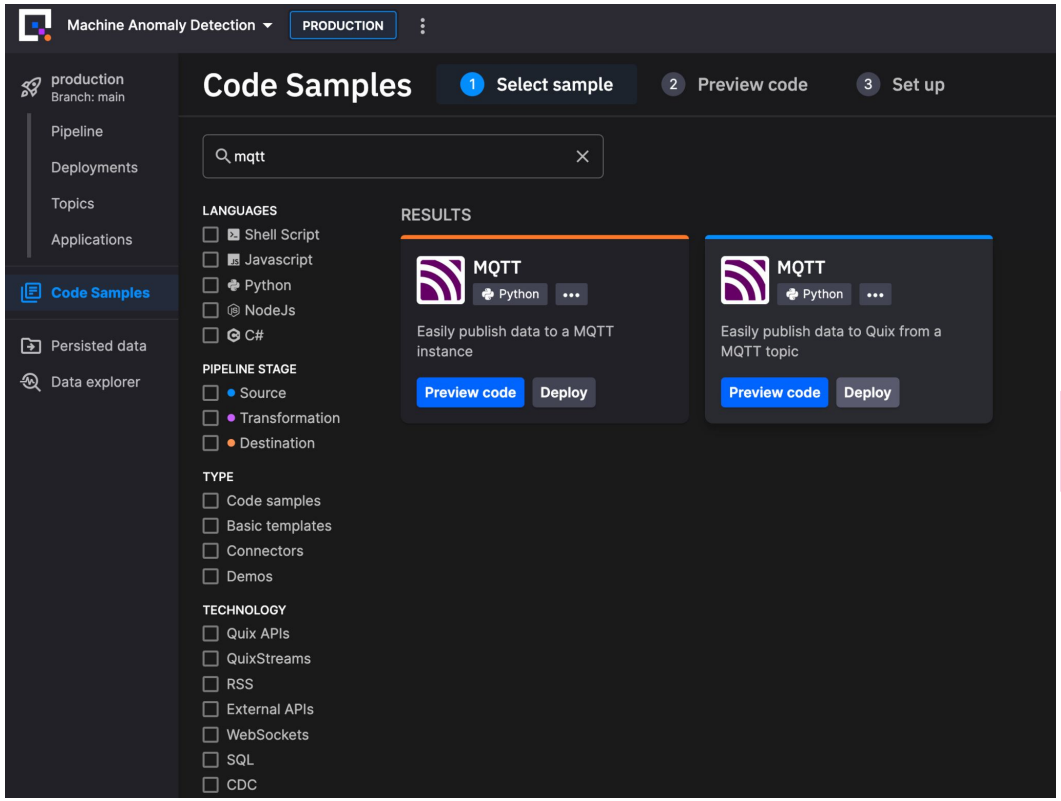

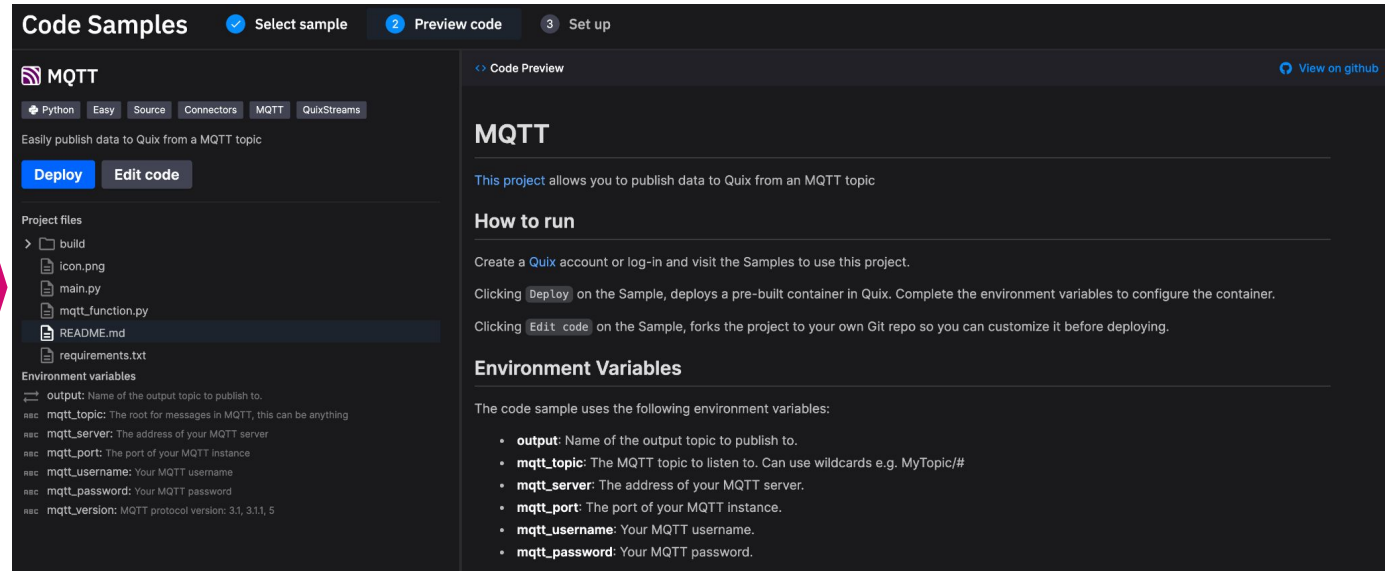

MQTT → InfluxDB

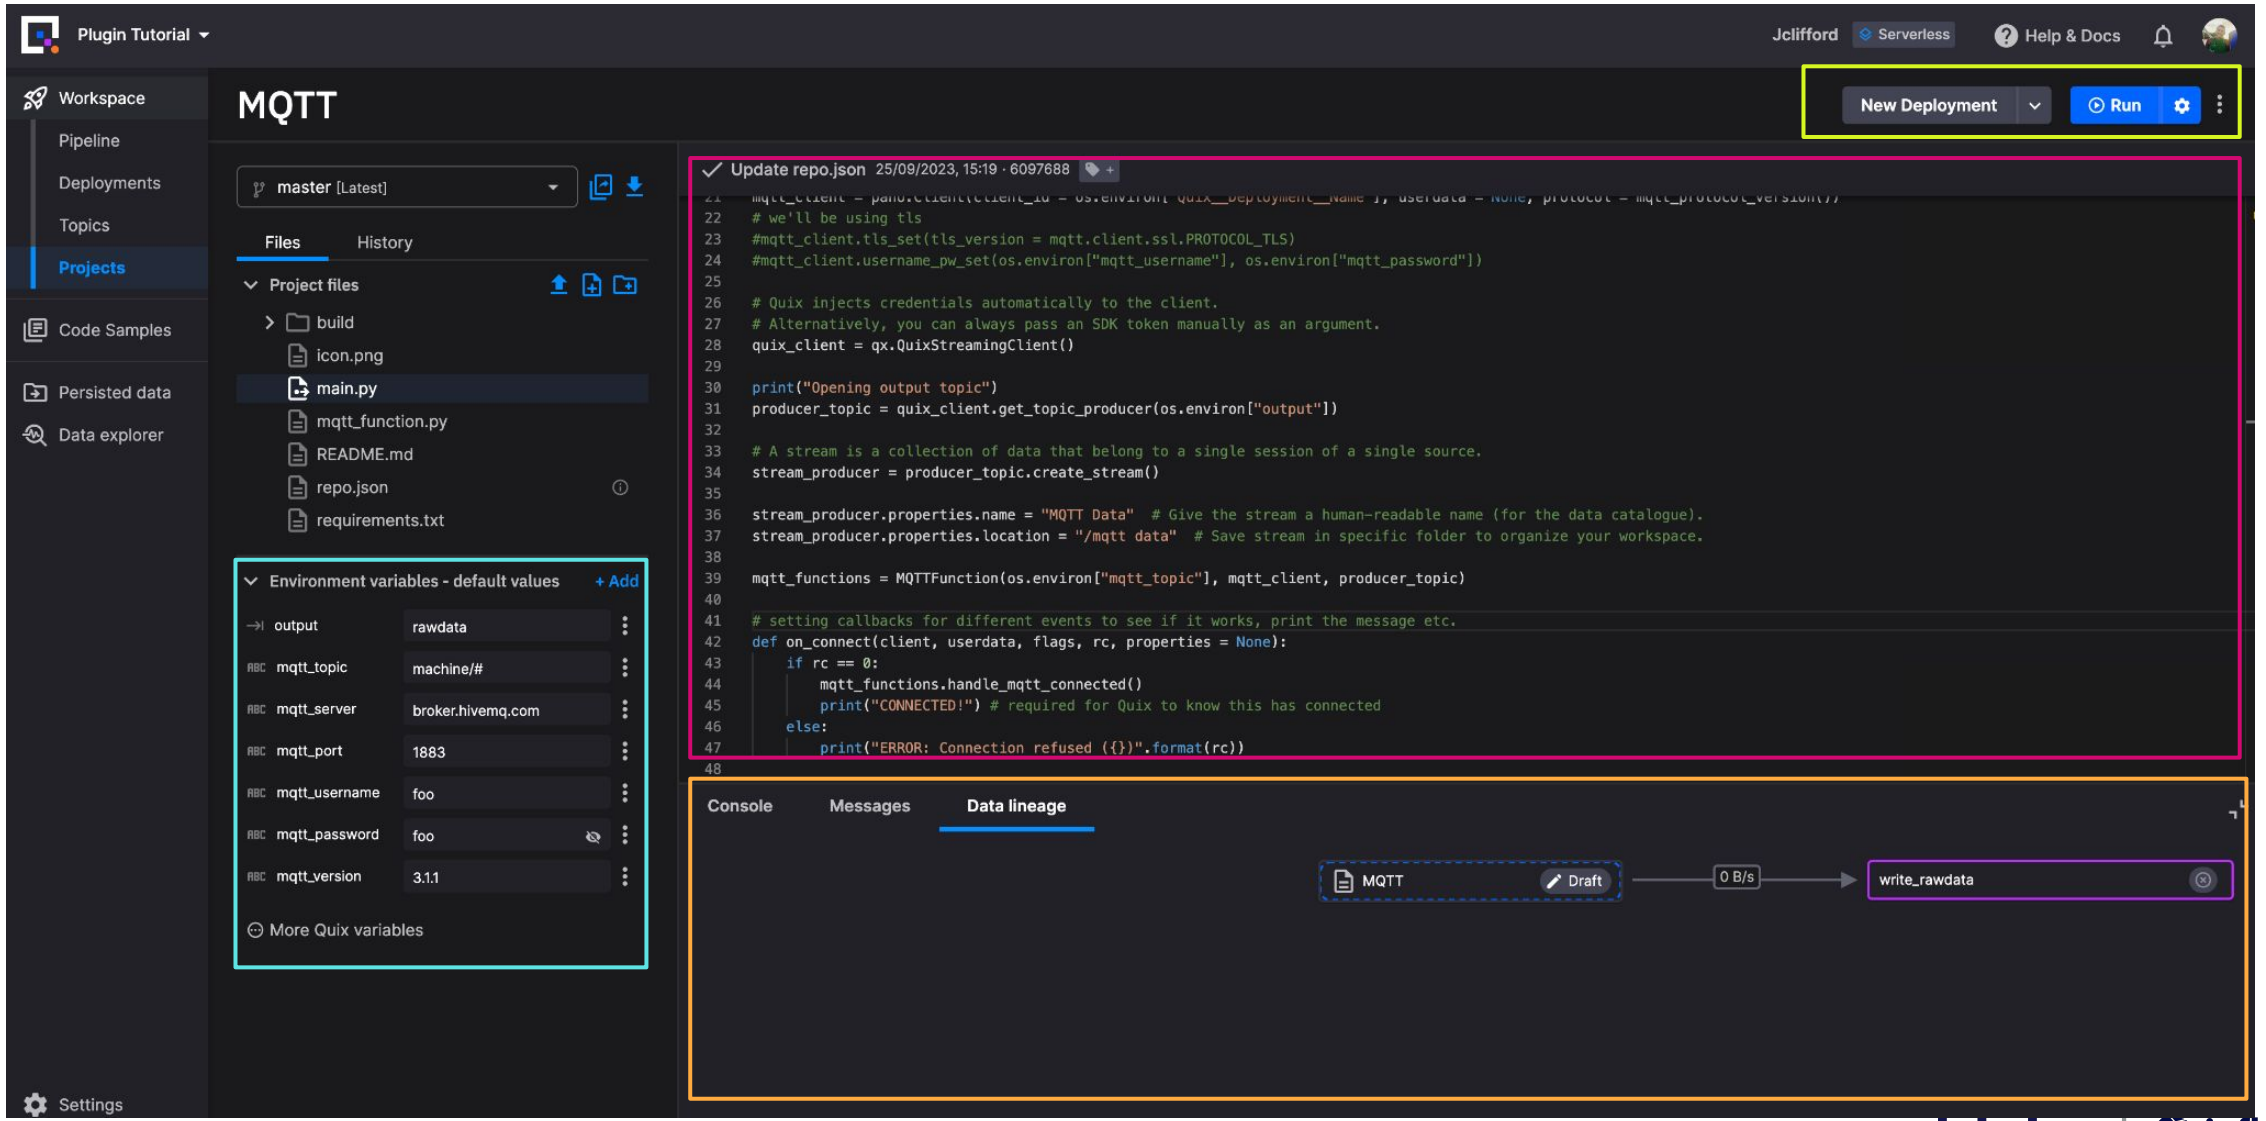

MQTT → InfluxDB

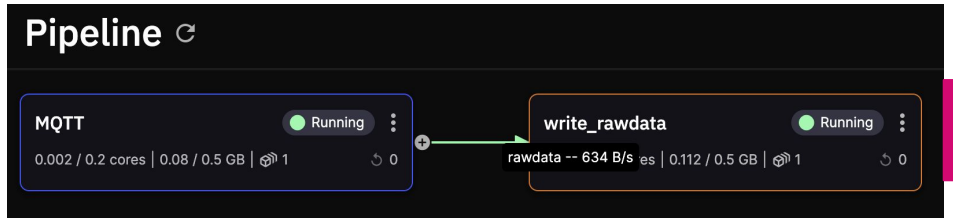

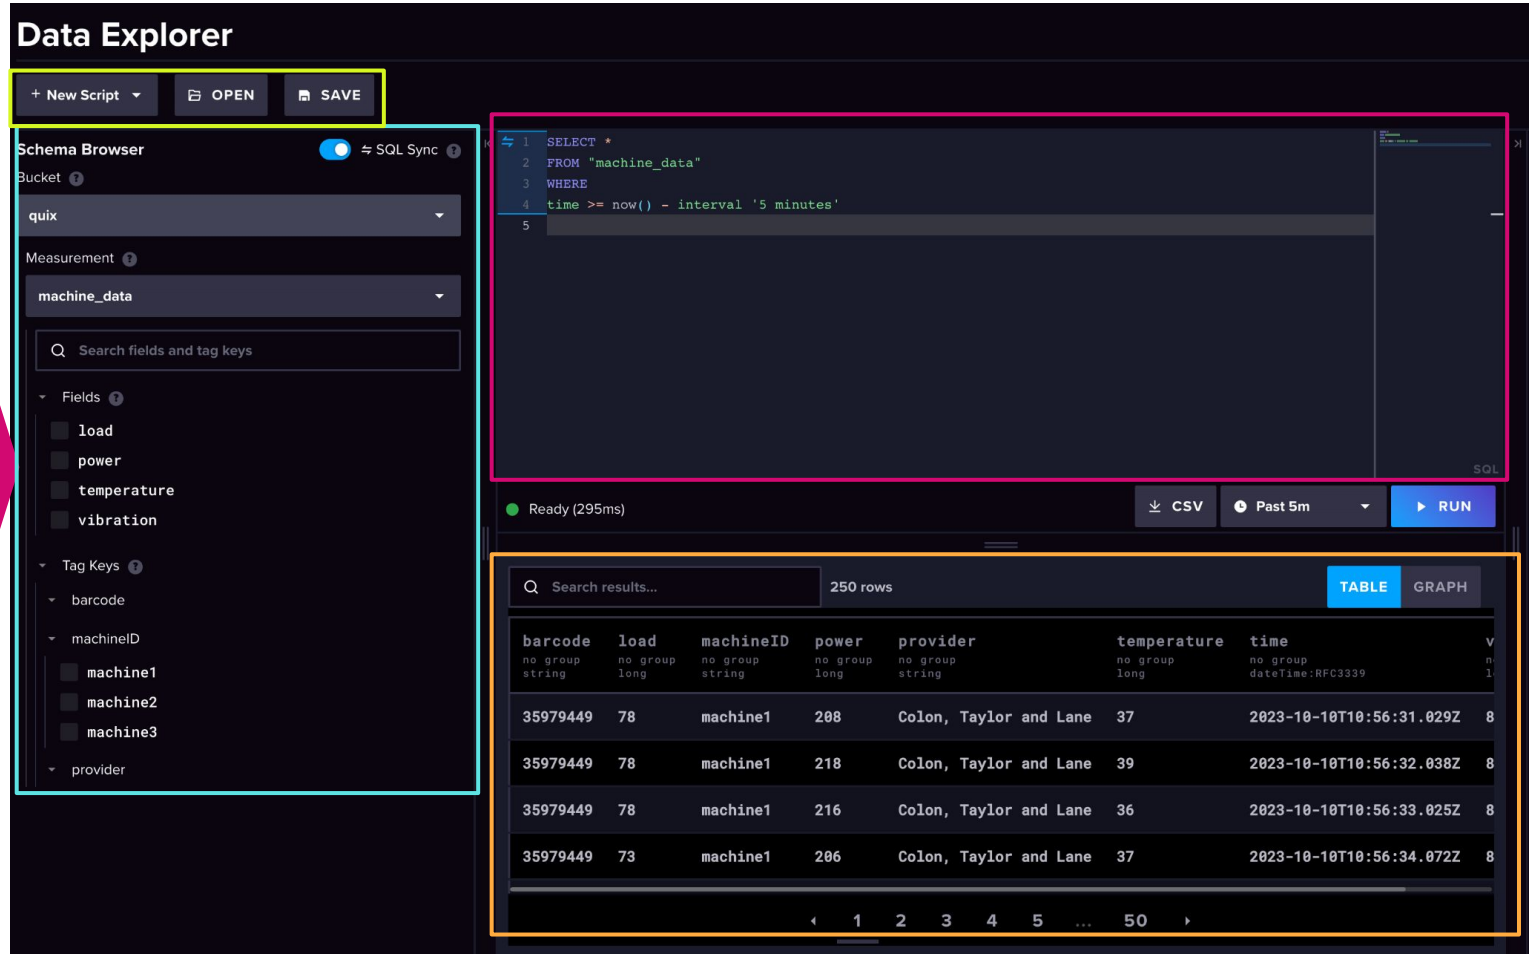

Quix | influxdata®

Training

```python
# Deeper Autoencoder architecture

# Create and fit a Normalization layer with your training data
normalization_layer = Normalization()
normalization_layer.adapt(normal_data)

# Deeper Autoencoder architecture with Normalization layer
input_layer = Input(shape=(1,))
normalized_input = normalization_layer(input_layer)
encoded = Dense(8, activation='relu', activity_regularizer=L1L2(l1=0.0, l2=0.1))(normalized_input)
encoded = BatchNormalization()(encoded)
encoded = Dense(4, activation='relu')(encoded)
decoded = Dense(4, activation='relu')(encoded)
decoded = BatchNormalization()(decoded)
decoded = Dense(1, activation='sigmoid')(decoded)

autoencoder = Model(input_layer, decoded)
autoencoder.compile(optimizer='adam', loss='mean_squared_error')
# Directory to store logs
log_dir = os.path.join(
    "logs",
    "fit",
    datetime.datetime.now().strftime("%Y%m%d-%H%M%S"),
)

# Creating the TensorBoard callback
tensorboard_callback = TensorBoard(log_dir=log_dir, histogram_freq=1)


history = autoencoder.fit(
    normal_data,
    normal_data,
    epochs=100, # specify the number of epochs
    batch_size=32, # specify the batch size
    callbacks=[tensorboard_callback] # Pass the TensorBoard callback
)
```

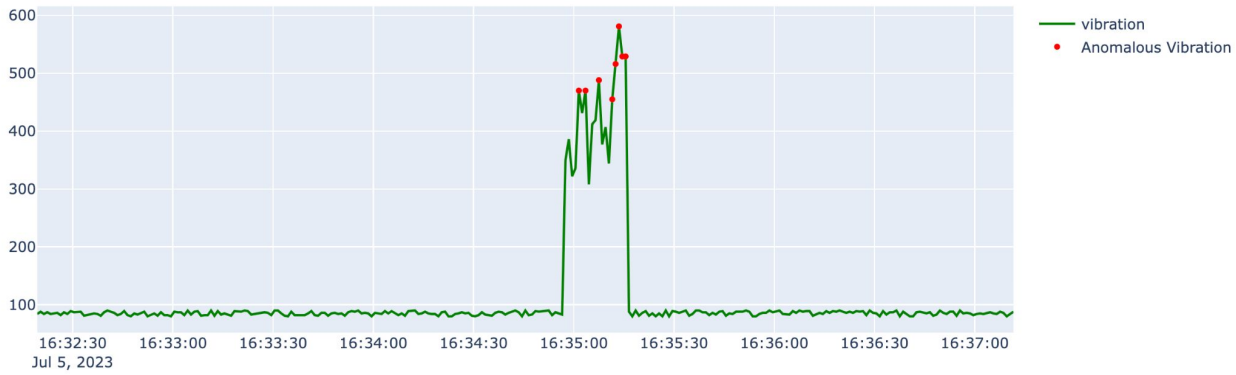

InfuxDB ➔ Quix

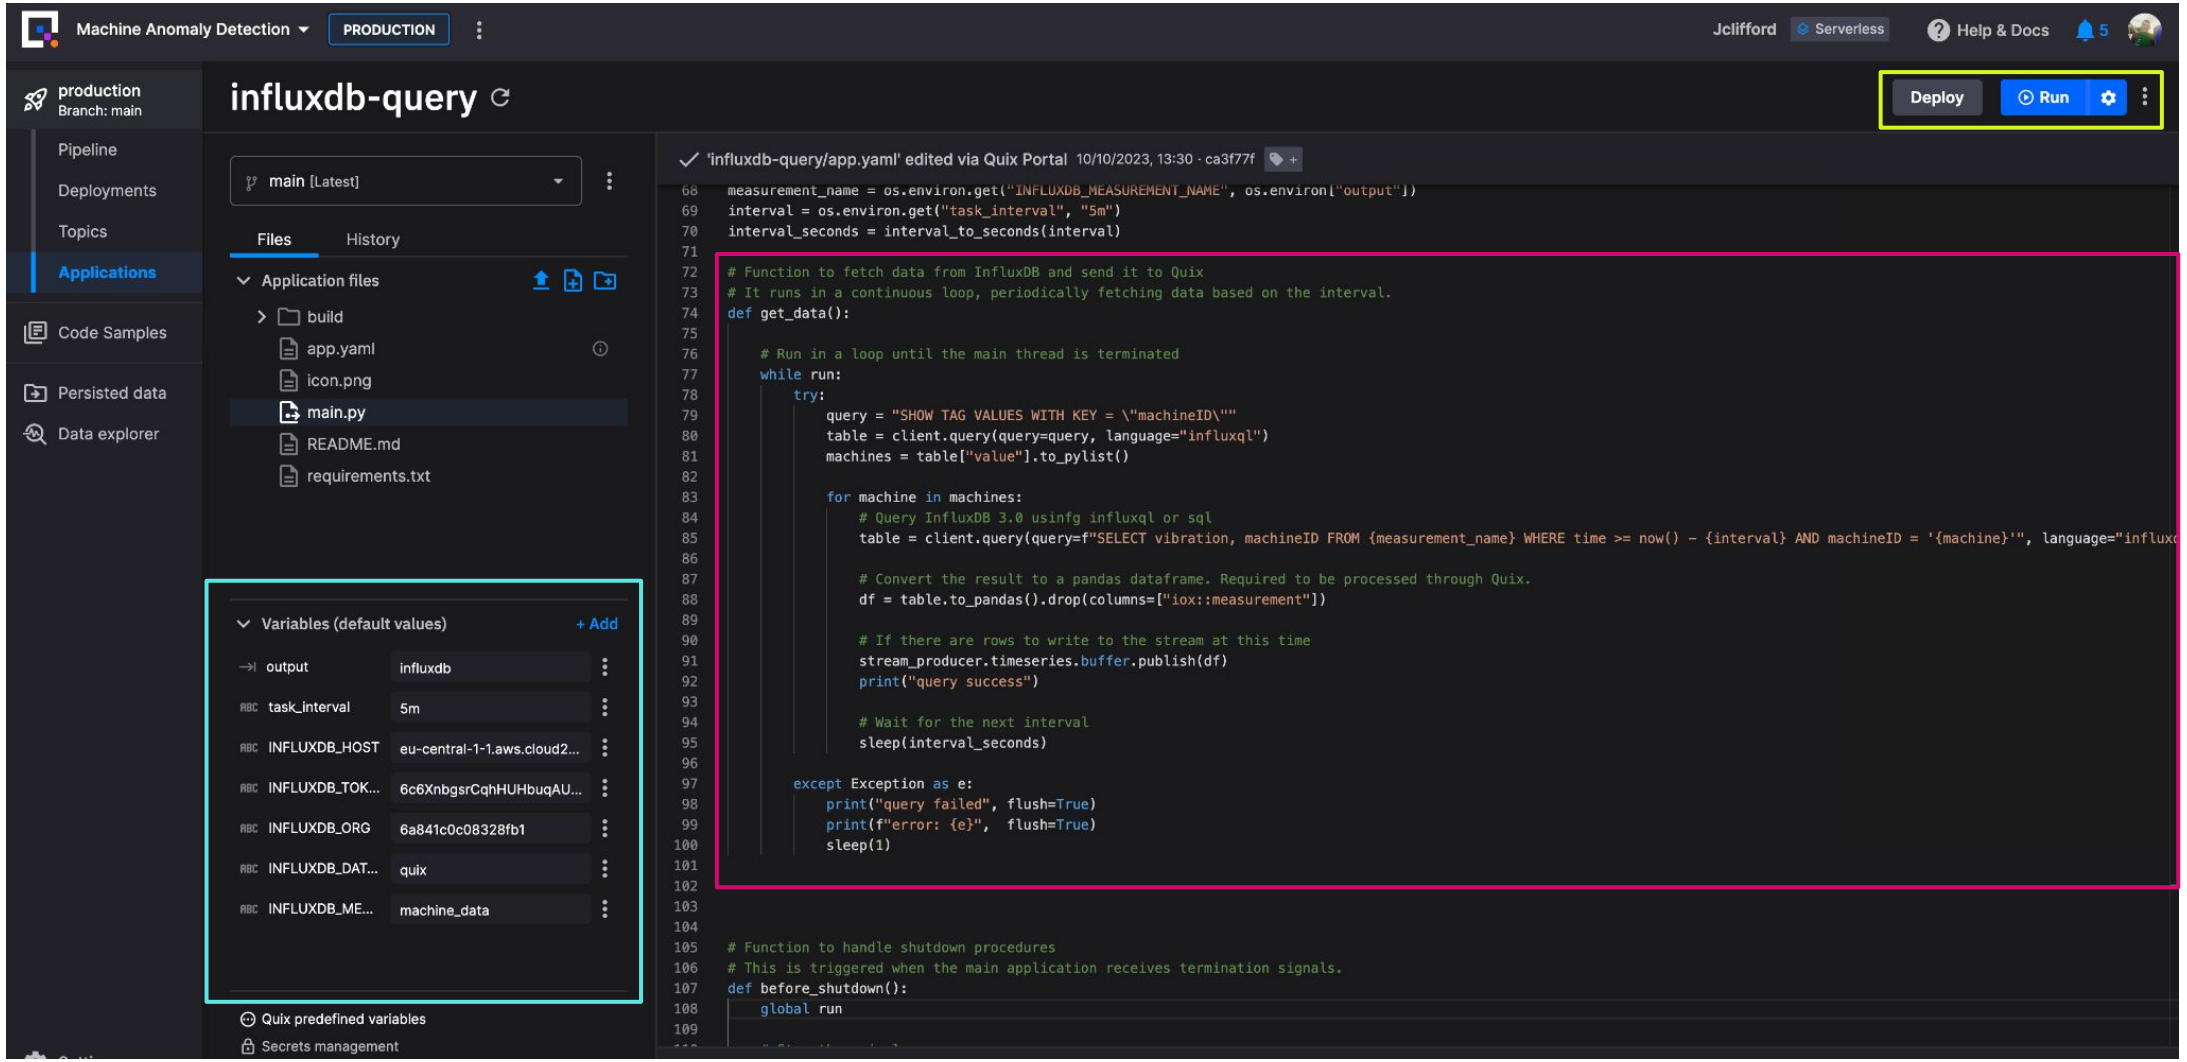

Quix | influxdata®

Quix → Model → Quix

'Event Detection/app.yaml' edited via Quix Portal 10/10/2023, 14:35 · 157f086

```python
import quixstreams as qx
from quix_function import QuixFunction
import os
from huggingface_hub import from_pretrained_keras

# Quix injects credentials automatically to the client.
# Alternatively, you can always pass an SDK token manually as an argument.
client = qx.QuixStreamingClient()

print("Opening input and output topics")
consumer_topic = client.get_topic_consumer(os.environ["input"], "default-consumer-group")
producer_topic = client.get_topic_producer(os.environ["output"])
model = from_pretrained_keras(os.environ["model"])


# Callback called for each incoming stream
def read_stream(consumer_stream: qx.StreamConsumer):

    # Create a new stream to output data
    producer_stream = producer_topic.get_or_create_stream(consumer_stream.stream_id + "vibration")
    producer_stream.properties.parents.append(consumer_stream.stream_id)

    # handle the data in a function to simplify the example
    quix_function = QuixFunction(consumer_topic, producer_stream, model)

    # React to new data received from input topic.
    consumer_stream.events.on_data_received = quix_function.on_event_data_handler
    consumer_stream.timeseries.on_dataframe_received = quix_function.on_dataframe_handler

    # When input stream closes, we close output stream as well.
    def on_stream_close(stream_consumer: qx.StreamConsumer, end_type: qx.StreamEndType):
        producer_stream.close()
        print("Stream closed:" + producer_stream.stream_id)

    consumer_stream.on_stream_closed = on_stream_close


# Hook up events before initiating read to avoid losing out on any data
consumer_topic.on_stream_received = read_stream

print("Listening to streams. Press CTRL-C to exit.")
```

Event Detection

main [Latest]

Files History

Application files

- build
- app.yaml
- main.py
- quix_function.py •
- README.md
- requirements.txt

Variables (default values) + Add

input	influxdb
output	detection-results
model	jayclifford345/vibration-a...
threshold	60

Edited 'quix_function.py' Commit

```python
        self.producer_stream.events.publish(data)

# Callback triggered for each new parameter data.
def on_dataframe_handler(self, stream_consumer: qx.StreamConsumer, df: pd.DataFrame):
    # Normalize the anomalous data
    print(df)
    df = df.set_index('timestamp')
    anom_data = df.drop(columns=['iox::measurement', 'machineID'])


    # Use the Autoencoder to predict on the anomalous data
    predictions = self.model.predict(anom_data)

    # Calculate reconstruction error
    mse = np.mean(np.power(anom_data.values - predictions, 2), axis=1)

    # Scale the MSE to a percentage
    min_mse = np.min(mse)
    max_mse = np.max(mse)
    mse_percentage = ((mse - min_mse) / (max_mse - min_mse)) * 100

    # Detect anomalies by comparing the scaled reconstruction error to some threshold
    threshold = float(os.environ["threshold"])  # Define a threshold value (in percentage)

    # Add 'is_anomalous' column to the DataFrame
    df['is_anomalous'] = mse_percentage > threshold
    df['mse_percentage'] = mse_percentage

    df = df.reset_index().rename(columns={'timestamp': 'time'})
    print(df)


    self.producer_stream.timeseries.buffer.publish(df)  # Send filtered data to output topic›
```


Quix | influxdata®

Quix ➔ InfluxDB

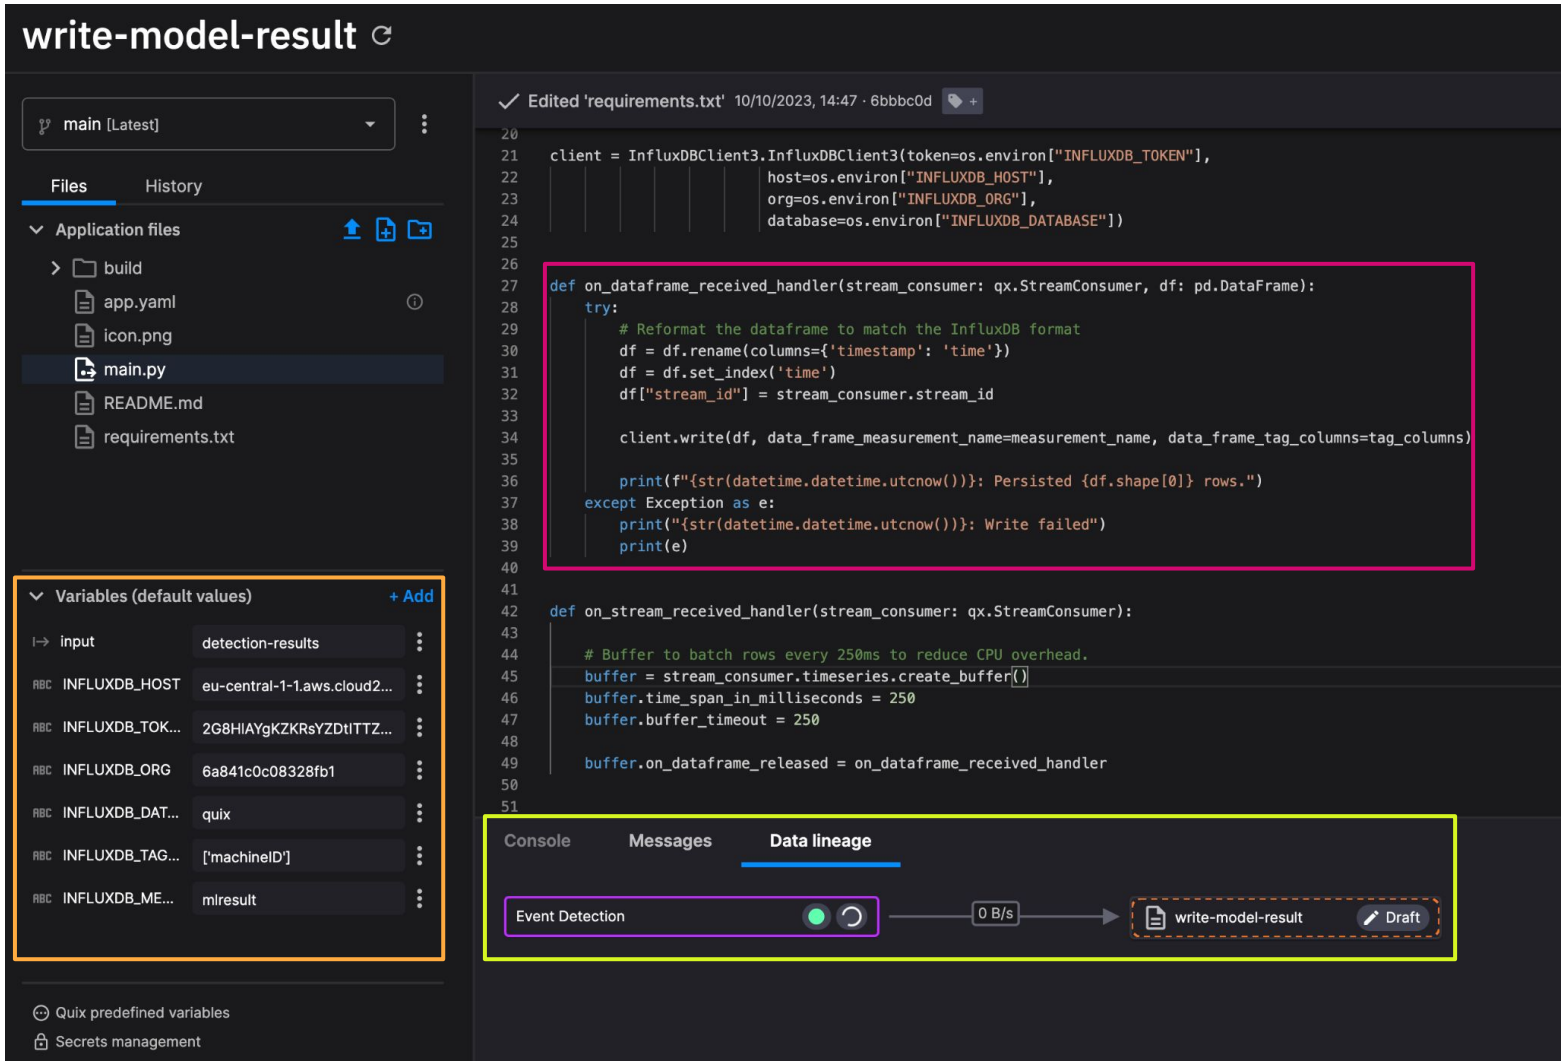

InfluxDB → Grafana

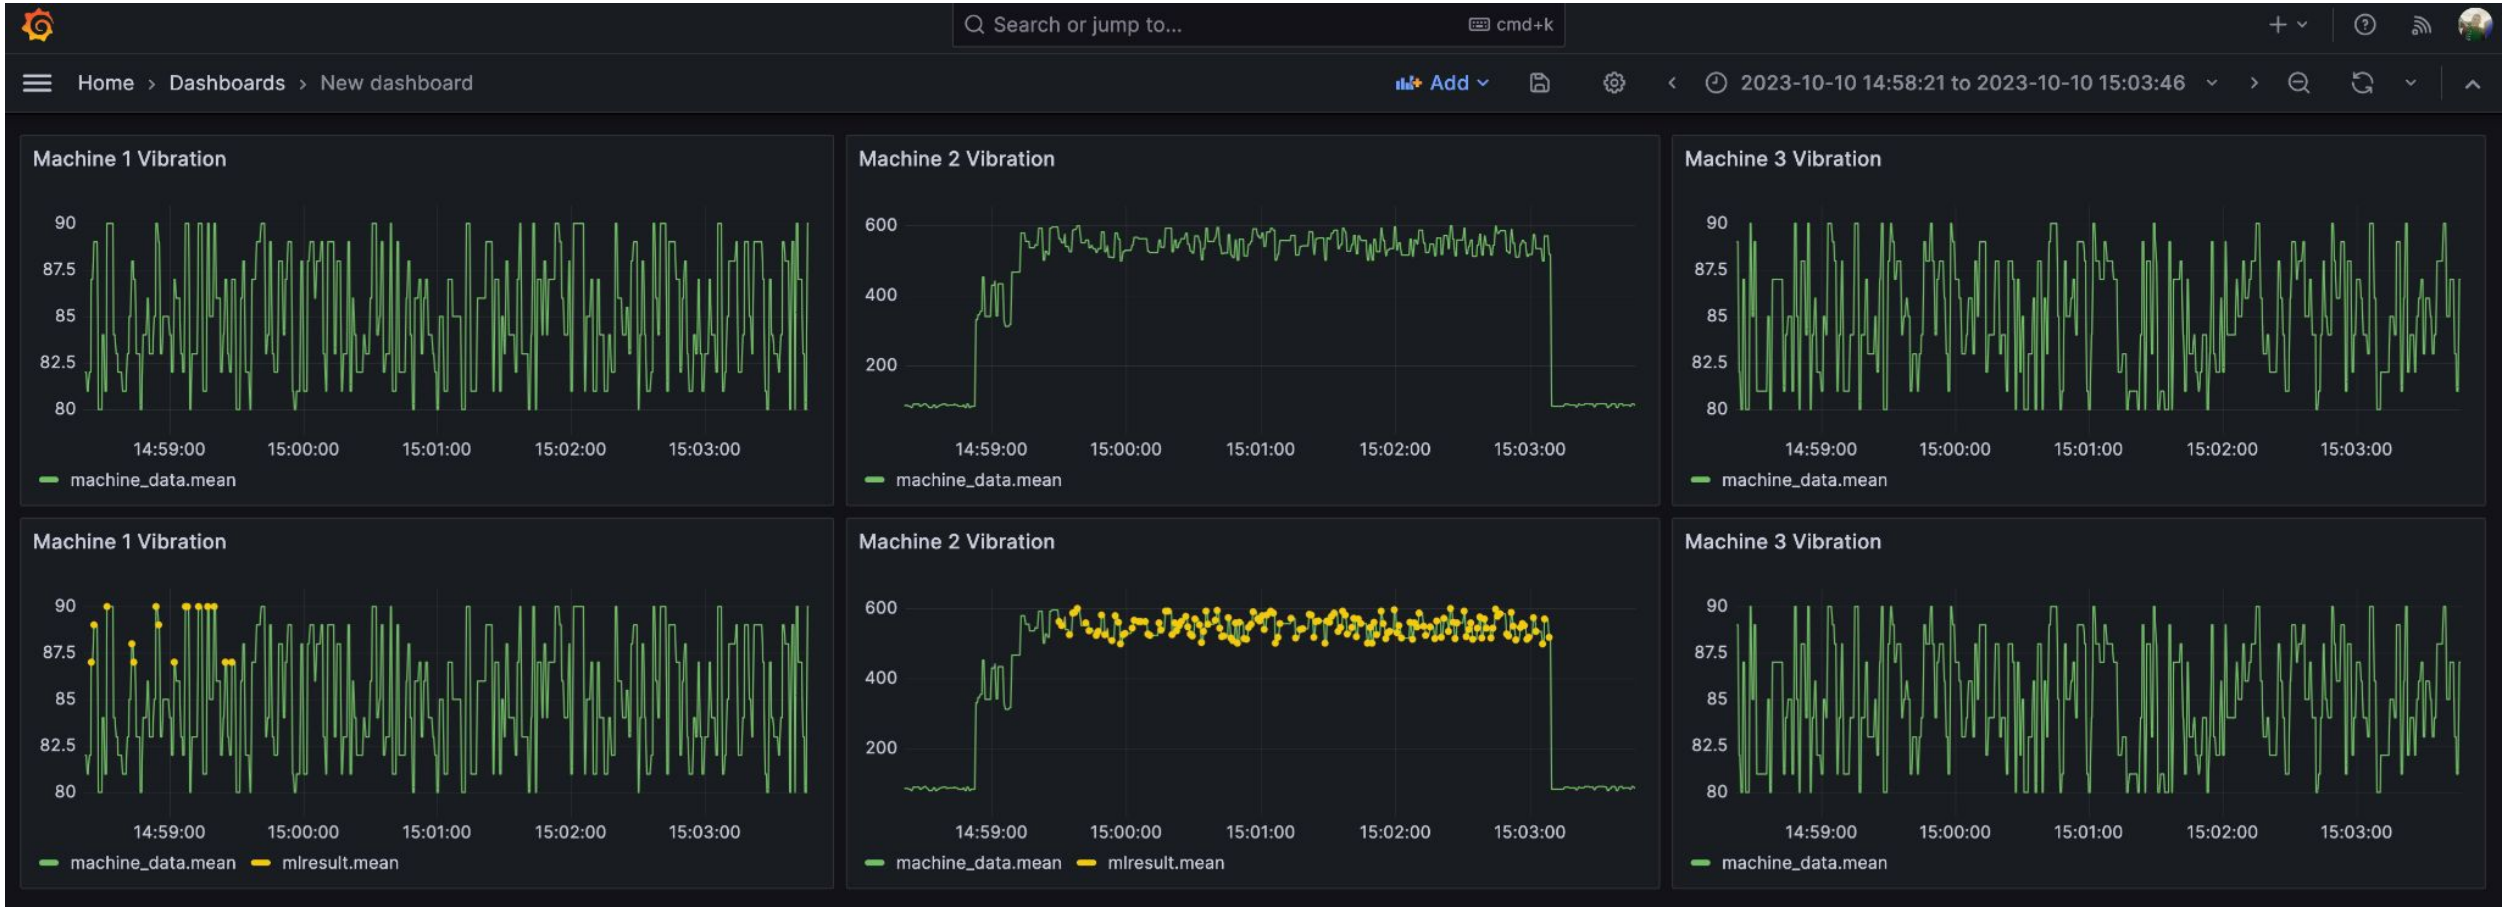

Quix | influxdata®