influxdata®

Leveraging the Apache Flight Python Client and InfluxDB

Anais Dotis Georgiou

Anais Dotis-Georgiou
Developer Advocate

LinkedIn

influxdata®

Agenda

- Introduction to InfluxDB and Time Series Data
- Commitment to Open Data Architecture with the FDAP Stack
- Leveraging the Arrow Flight Client InfluxDB v3 Python Client Library
- Projects that Leverage th

influxdata®

Introduction to InfluxDB and Time Series Data

Time Series Data

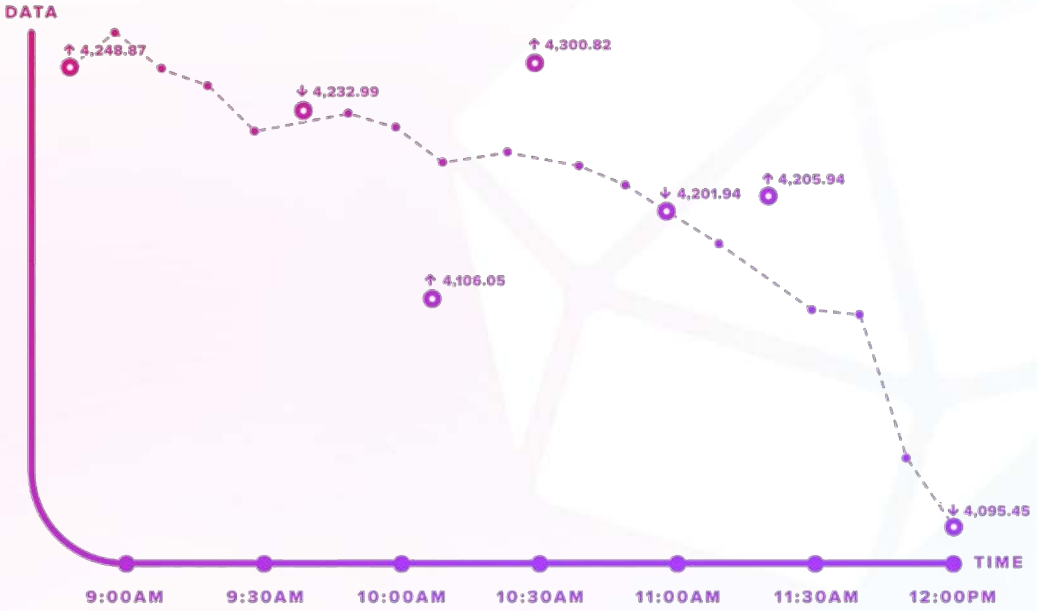

Time Series Data

influxdata®

Time Series Data Types

Metrics

Measurements at **regular** time intervals

Events

Measurements at **irregular** time intervals

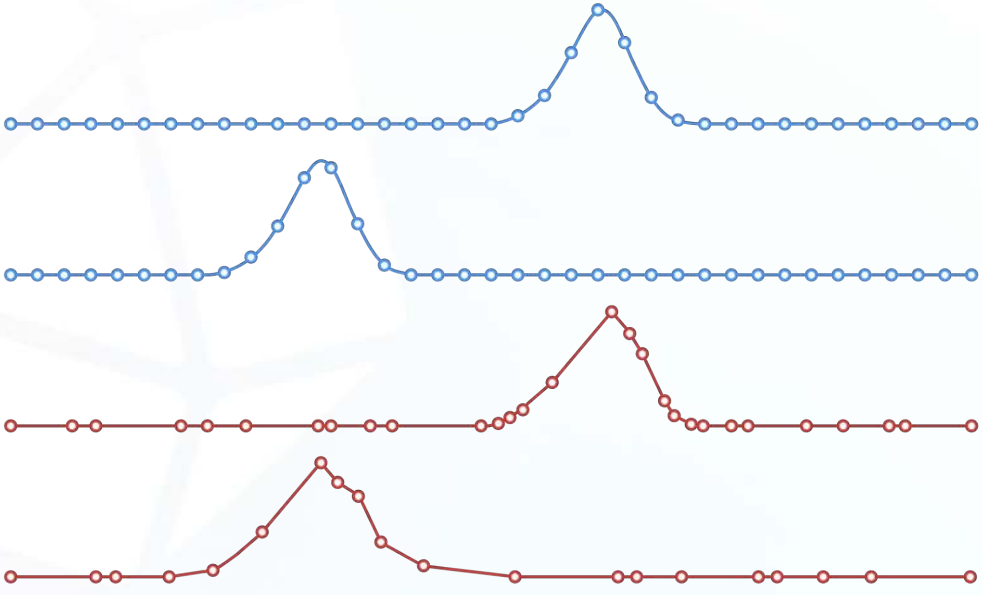

influxdata®

Time Series Databases

Time Series Data

High write throughput

Efficient Queries Over Time Ranges

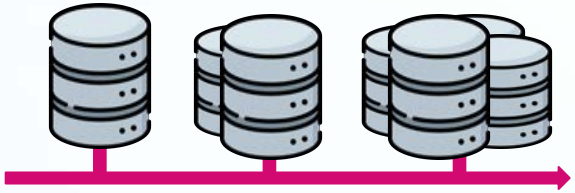

Scalability and Performance

InfluxDB 3.0

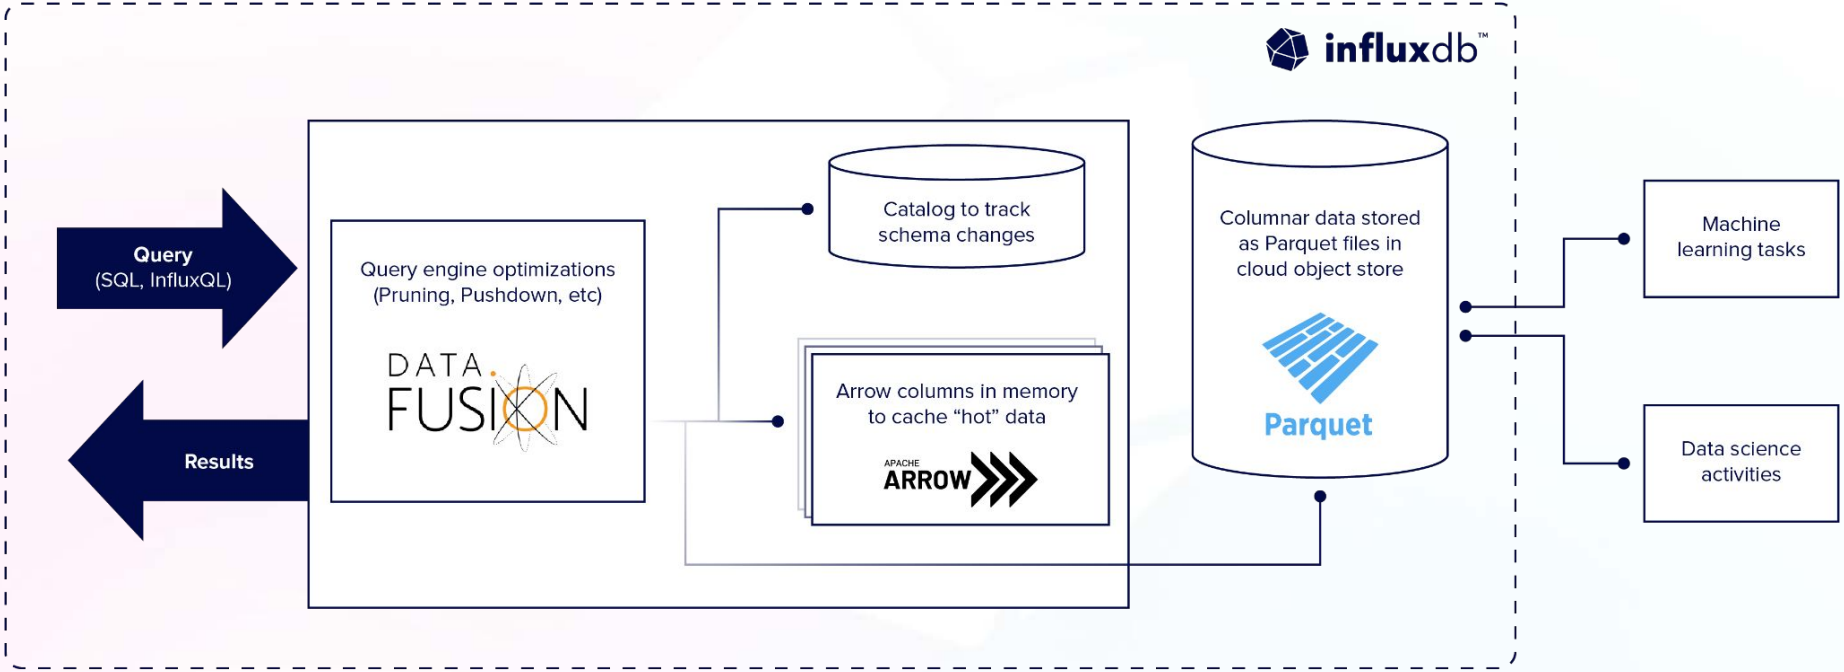

influxdata®

Some of our customers

IoT monitoring	DELAIR, INTESENS The connected maintenance, ARABLE, Greensmith, TESLA, ASML, ecoult energy storage solutions, NEAS ENERGY, aquicore, BBOXX, APMT, WayKonect
Software	COMCAST, coupa, robinhood, cruise, SAP, BBC NEWS, BARCLAYS, TD Bank, MuleSoft, ORACLE, cisco, wayfair, PayPal, VERISIGN, Capital One, SIEMENS

 |

influxdata

Ingest Benchmark

Time Range: 21.5 hours

InfluxDB OSS 1.8 | InfluxDB 3.0

Data Ingest Performance: Points Per Second

Higher = Better

Ingest Vector	InfluxDB OSS 1.8	InfluxDB 3.0
Total Telegraf Clients	100	4,500 x45
Rows / Hour	8,859,600	329,206,806
Rows / Second	2,461	91,446
Values / Second	162,264	4,310,764

influxdata®

Commitment to Open Data Architecture with the FDAP Stack

Data Storage

InfluxDB is a database purpose-built for handling time series data at massive scale for real-time analytics.

Developers can ingest, store, and analyze all types of time series data; metrics, events, traces in a single platform. Designed to handle high-speed, high-volume, and high-cardinality data.

influxdata®

InfluxDB 3.0

Schema on write

Write and query millions of rows per second

Single datastore for all time series data (metrics, logs, and traces)

SQL, InfluxQL Support

 |

influxdata®

Advantages of Columnar Data Storage (sidebar)

Sidebar–Advantages of Columnar Data Storage

```
measurement1,tag1=tagvalue1 field1=1i timestamp1
measurement1,tag1=tagvalue2 field1=2i timestamp2
measurement1,tag2=tagvalue3 field1=3i timestamp3
measurement1,tag1=tagvalue1,tag2=tagvalue3 field1=4i,field2=true timestamp4
measurement1, field1=1i timestamp5
```


influxdata®

Sidebar–Advantages of Columnar Data Storage

Name: measurement1					
field1	**field2**	**tag1**	**tag2**	**tag3**	**time**
1i	null	tagvalue1	null	null	timestamp1
2i	null	tagvalue2	null	null	timestamp2
3i	null	null	tagvalue3	null	timestamp3
4i	true	tagvalue1	tagvalue3	tagvalue4	timestamp4
1i	null	null	null	null	timestamp5

influxdata®

Sidebar–Advantages of Columnar Data Storage

1i	2i	3i	4i	1i
null	null	null	true	null
tagvalue1	tagvalue2	null	tagvalue1	null
null	null	tagvalue3	tagvalue3	null
null	null	null	tagvalue4	null
timestamp1	timestamp2	timestamp3	timestamp4	timestamp5

```
1i, 2i, 3i, 4i, 1i ;
null, null, null, true, null ;
tagvalue1, tagvalue2, null, tagvalue1, null ;
null, null, null, tagvalue3, tagvalue3, null ;
null, null, null, tagvalue4, null ;
timestamp1, timestamp2, timestamp3, timestamp4, timestamp5 .
```


influxdata®

Leveraging the Arrow Flight Client InfluxDB v3 Python Client Library

Arrow Flight Client

```
pip install pyarrow
```

Library Import

```
from pyarrow.flight import FlightClient, Ticket
import json
```

Initialization

```
host = "us-east-1-2.aws.cloud2.influxdata.com"
database = 'my_db'
sql = "SELECT * from my_table"

client = FlightClient(f"grpc+tls://{host}:443")
ticket_data = {
    "namespace_name": "my_db",
    "sql_query": sql,
    "query_type": "sql",
}
```

Query

```
ticket_bytes = json.dumps(ticket_data)
ticket = Ticket(ticket_bytes)
flight_reader = client.do_get(ticket)
arrow_table = flight_reader.read_all()
data_frame = arrow_table.to_pandas()
print(data_frame.to_markdown())
```

influxdata®

Arrow Flight SQL Client

```
pip install flightsql-dbapi
```

Library Import

```
from flightsql import FlightSQLClient
```

Initialization

```
client =
FlightSQLClient(host='cluster-id.influxdb.io',
    token='DATABASE_TOKEN',
    metadata={'database': 'DATABASE_NAME'},
    features={'metadata-reflection': 'true'})
```

Query

```
info = client.execute("SELECT * FROM home")
ticket = info.endpoints[0].ticket
reader = client.do_get(ticket)
table = reader.read_all()

print(table)
```


influxdata®

InfluxDB v3 Python Client Library

```
from influxdb_client_3 import InfluxDBClient3
```

Library Import

```
host = “eu-central-1-1.aws.cloud2.influxdata.com”
org="6a841c0c08328fb1"
token = “”
database = “database”

client = InfluxDBClient3(
    token=token,
    host=host,
    org=org)
```

Initialization

```
sql = '''SELECT * FROM table'''
table = client.query(query=sql, language='sql',
mode='all')
print(table)
```

Query

Projects that Leverage the Arrow Flight

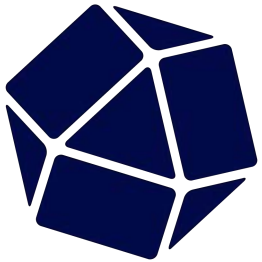

+


Grafana

Grafana Flow

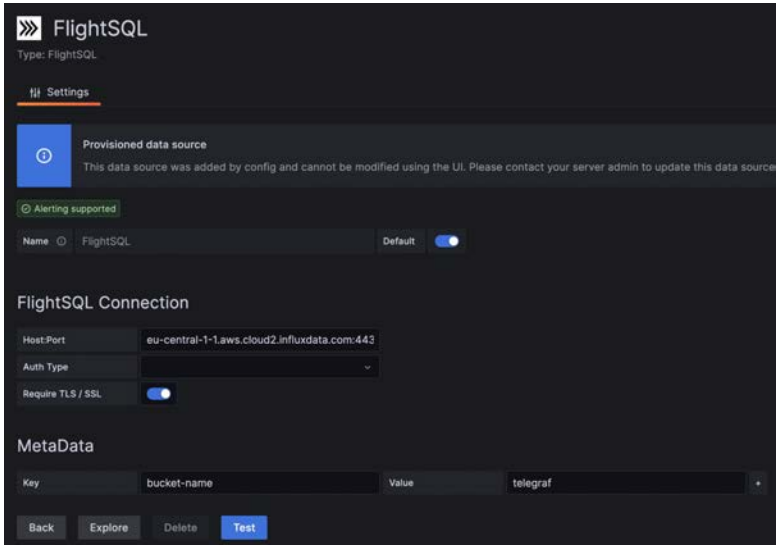

Datasource

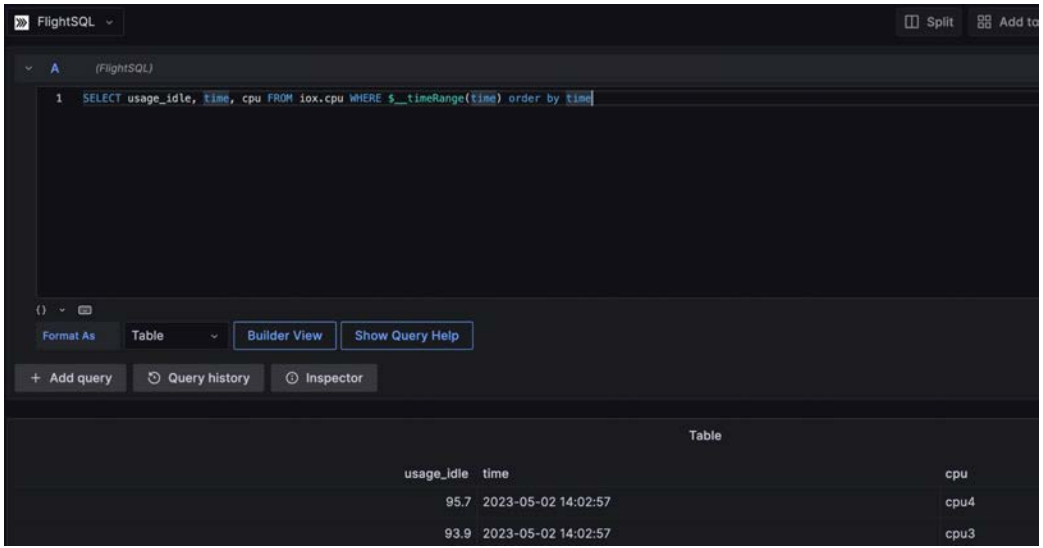

Explore

 |

influxdata®

Try it yourself - Grafana Quick Starts

main | 1 branch | 0 tags | Go to file | Add file | Code

Jayclifford345 Update README.md — 18d0b6c now — 10 commits

emergency-generator	added new example	last month
plugins	updated readme	3 months ago
system-monitoring	updated docker-compose	3 months ago
vision-ai-pipeline	added vision ai	2 months ago
.gitignore	initial commit	3 months ago
README.md	Update README.md	now

README.md

InfluxDB 3.0 Quick Starts

Welcome to the "InfluxDB IOx Quick Starts" repository! This repository is dedicated to providing easy to follow tutorials on how to integrate InfluxDB IOx with Grafana and Superset. InfluxDB IOx is a powerful time-series database solution and when paired with Grafana or Superset, both popular open-source visualization platform, you can create beautiful and informative dashboards to better visualize and understand your data.

This repository will cover the basics of setting up InfluxDB IOx and Grafana or Superset, how to connect them and how to query data from InfluxDB IOx. Whether you are new to InfluxDB IOx or an experienced user, this repository is designed to help you get up and running with InfluxDB IOx, Grafana and Superset quickly and easily.

https://github.com/InfluxCommunity/InfluxDB-3-Quick-Starts

 |

influxdata®

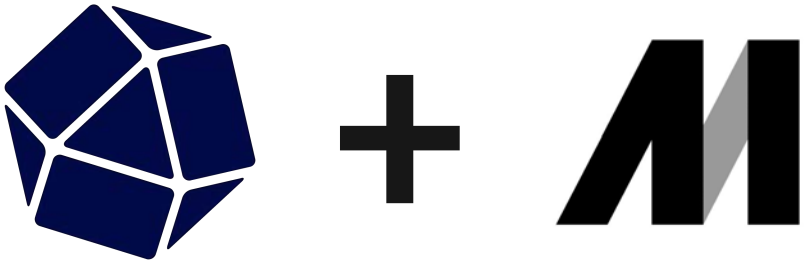

🚀 Mage & InfluxDB - Anomaly Detection

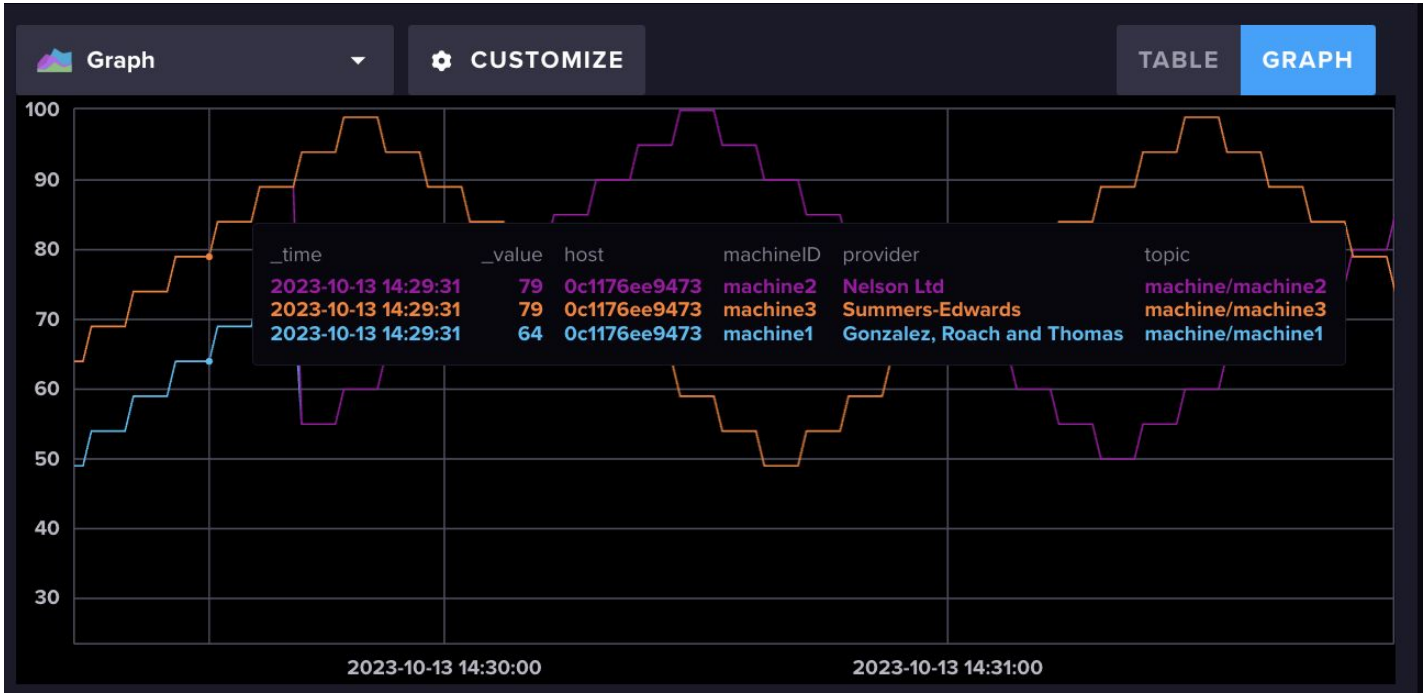

 |

influxdata®

Try it yourself

https://github.com/InfluxCommunity/Mage_Demo

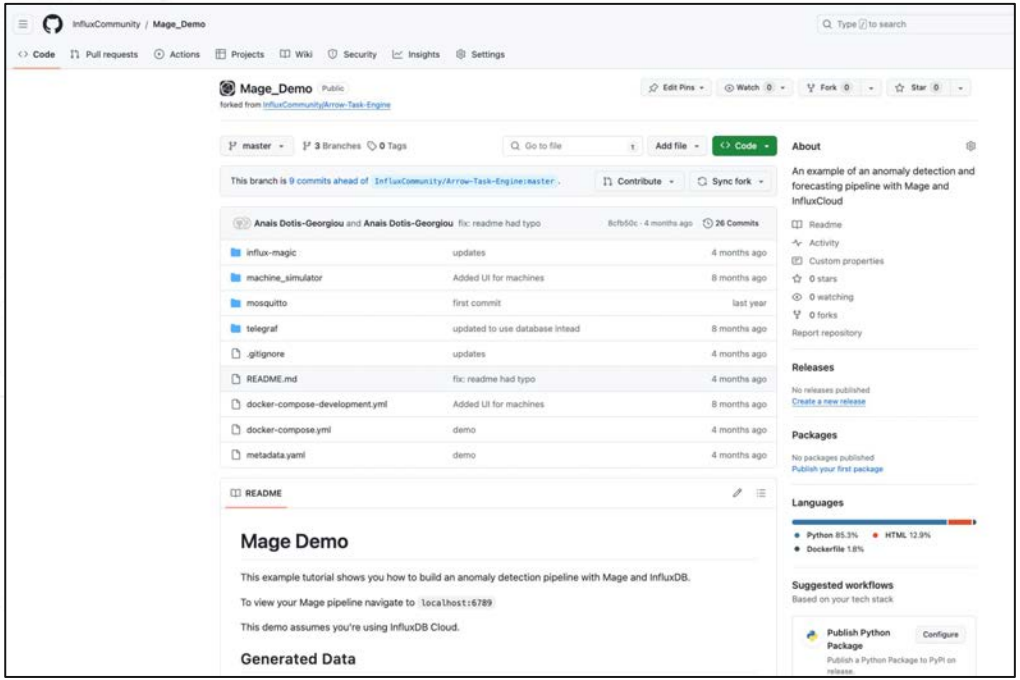

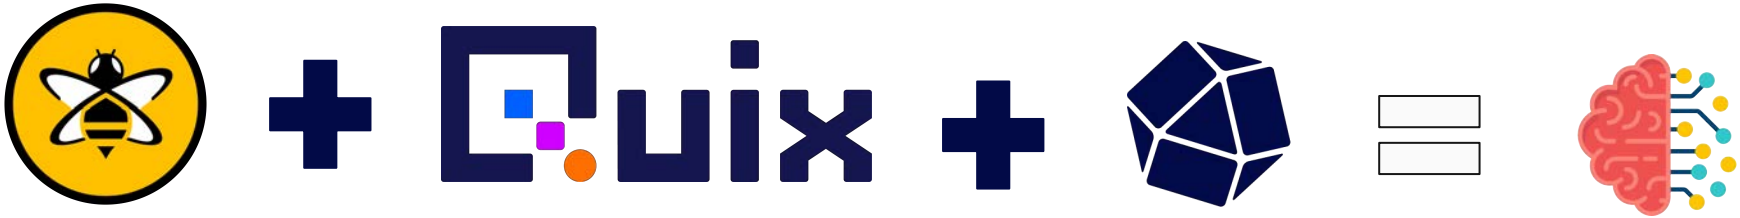
uix

Solution Architecture

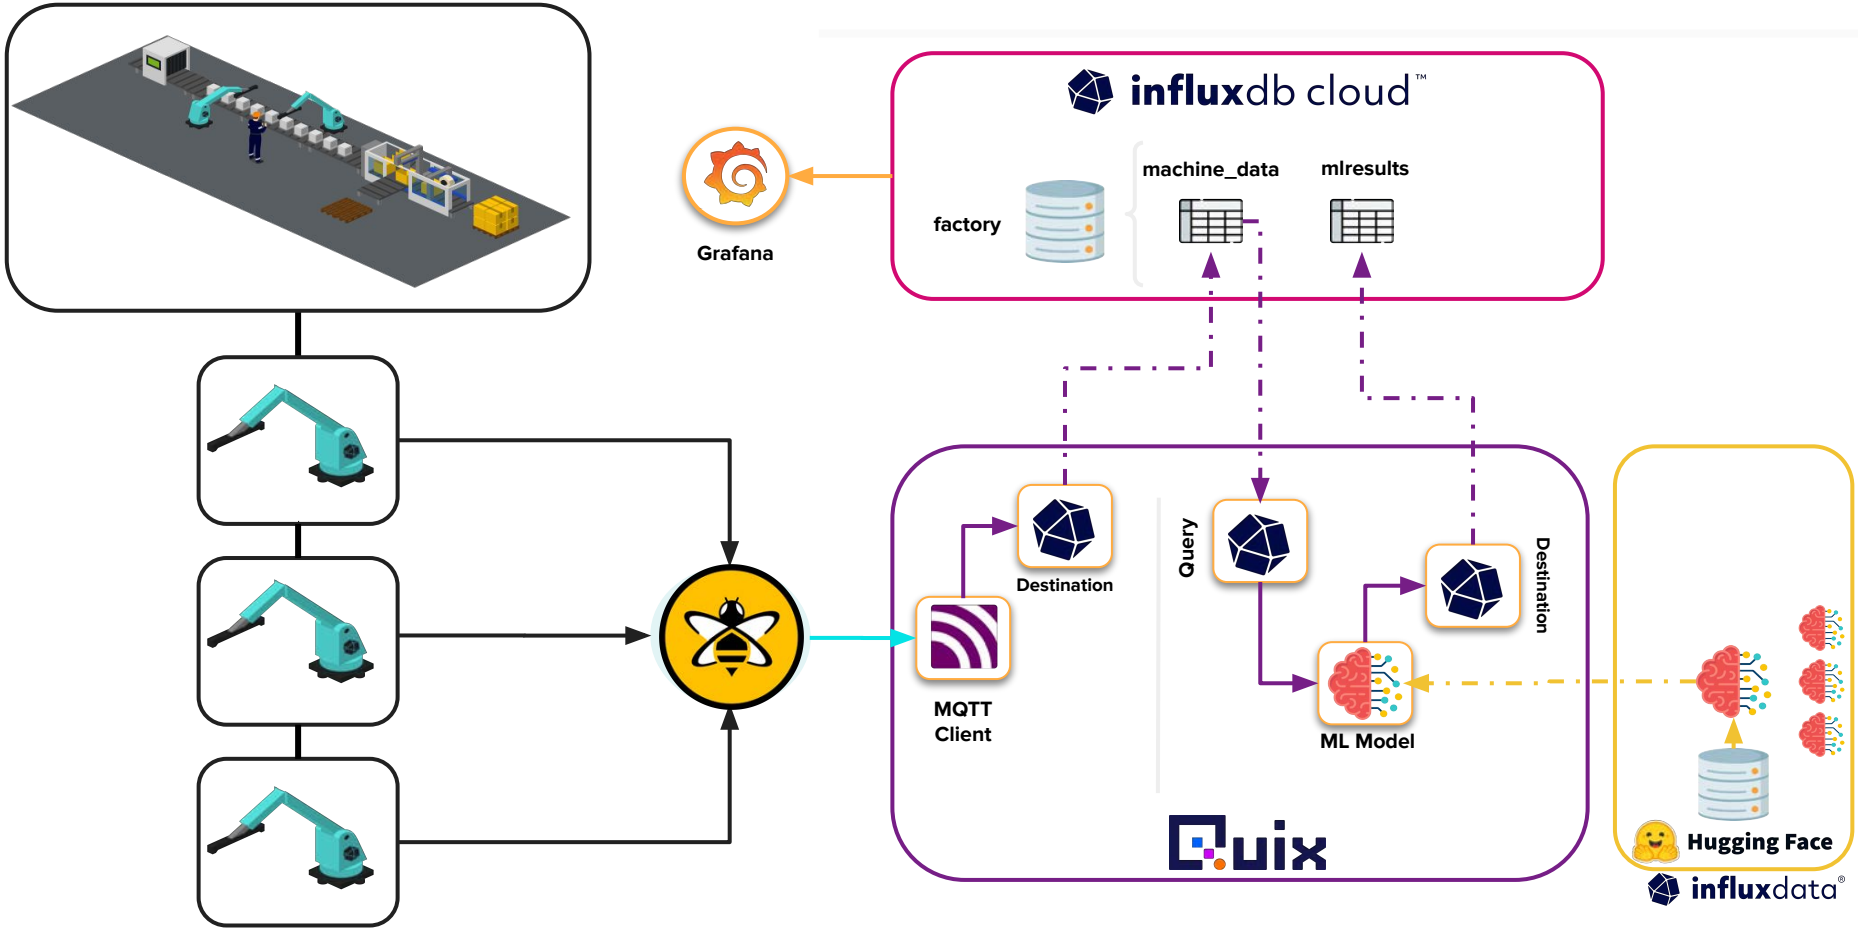

Try it yourself

https://github.com/InfluxCommunity/quix-anomaly-detection-example

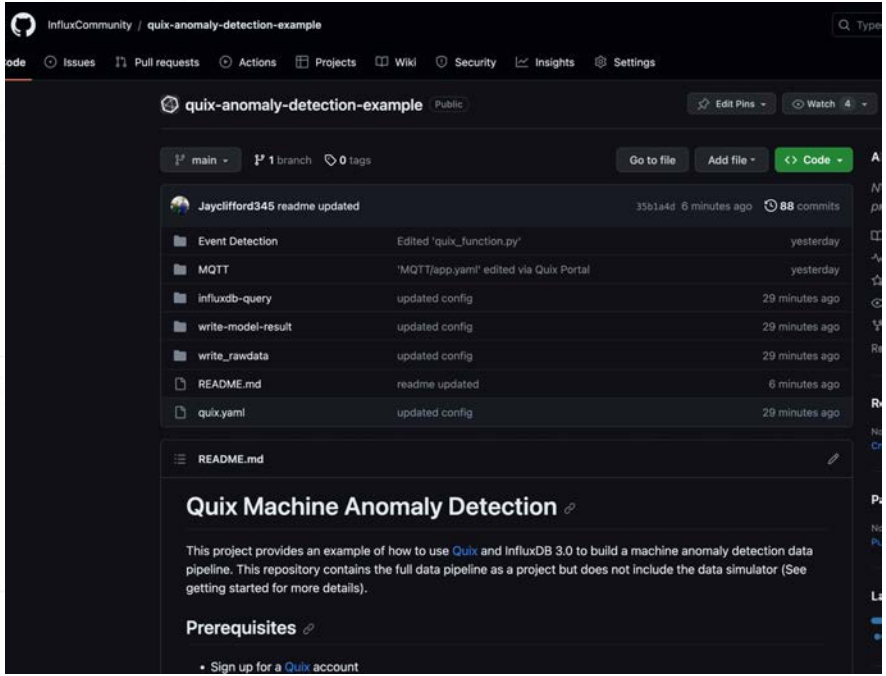

Join the InfluxDB Community

Sign up

Influxdata.com

Get InfluxDB

Via cloud marketplace

aws

Google Cloud

Microsoft Azure

Learn

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

influxdata®

Thank you

influxdata®

THANK YOU