


# Time Series Databases

@flashpys




Map  
Intelligence  
Agency



# What are Time Series DBs?

- Great in storing time-value pairs
- Usually used for server system metrics
- Compression
- Time Bucketing
- Calculation



# What TSDBs offer

- Efficient storage
- "Data-gathering Tool"
- Alert management
- Query language
- Cloud Hosting
- Clustering\*



# Overview

Most popular TSDB  
Designed as TSDB  
Great ecosystem



Java World



Well-known  
Great ecosystem



Prometheus

Built on top of PG  
Great documentation  
Works with PostGIS



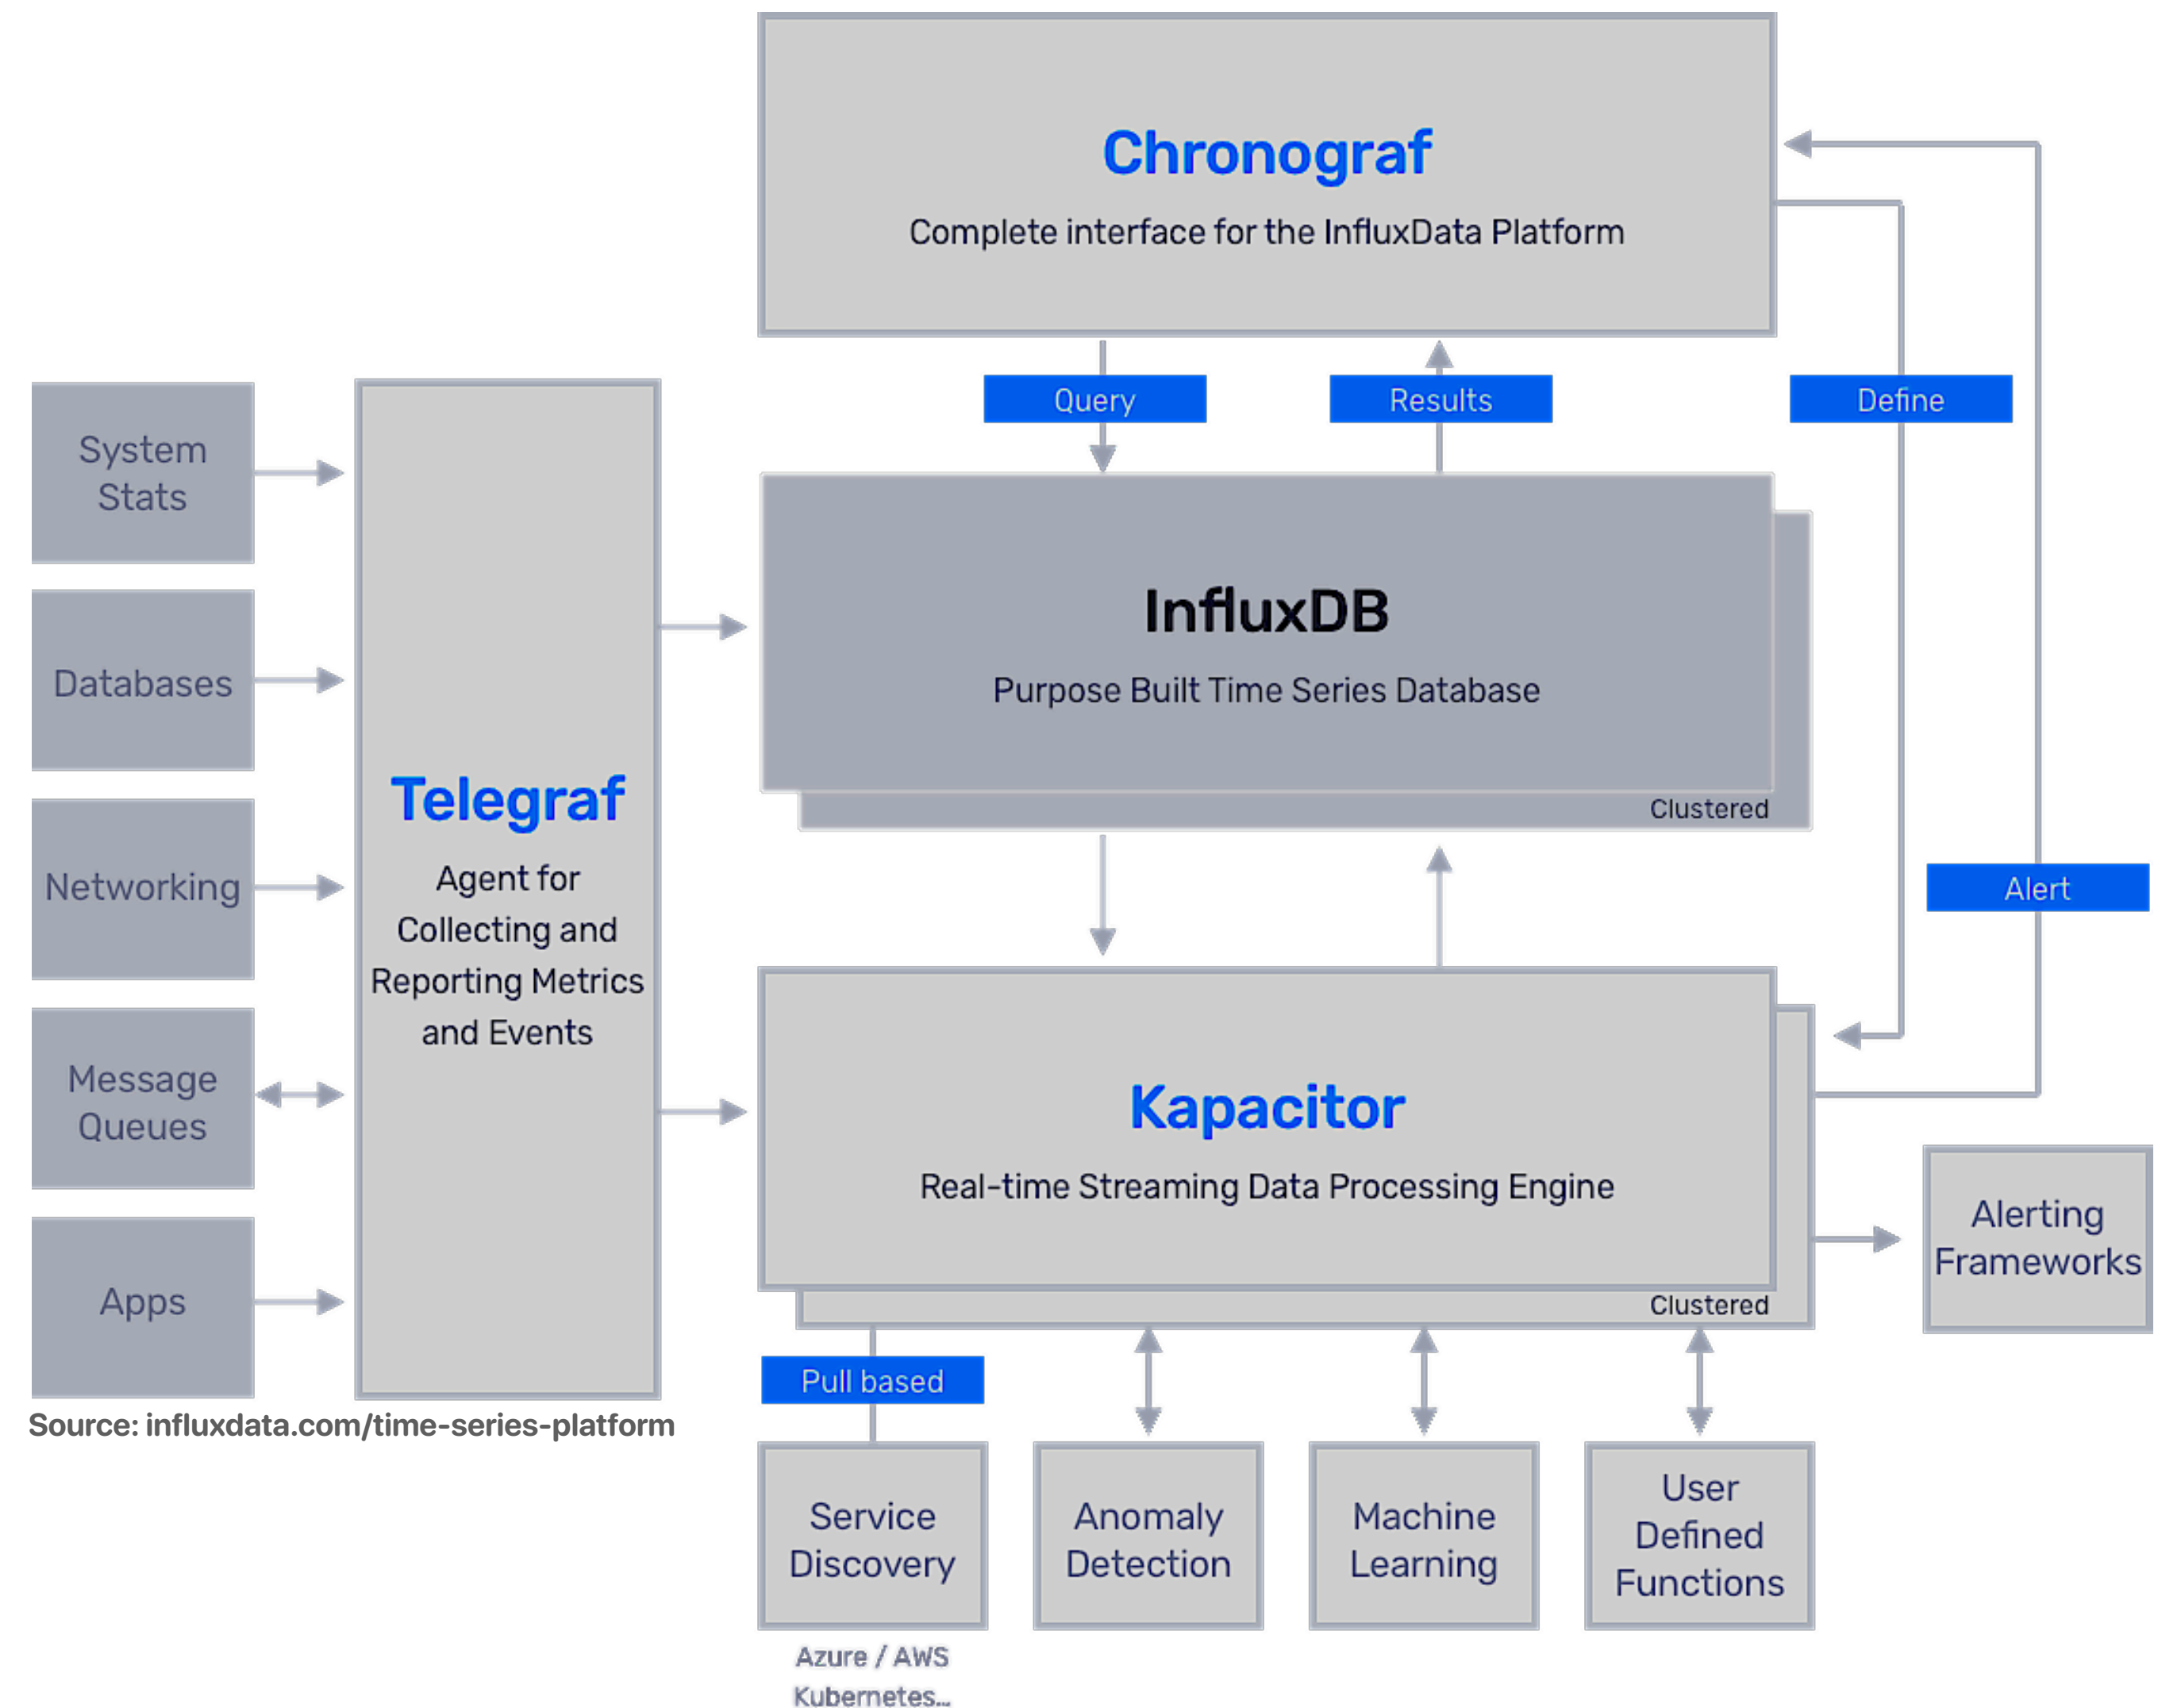
TimescaleDB

Serverless TSDB  
GraphQL included



# InfluxDB

- TICK-Stack
- Complete Time Series Ecosystem
- Quite easy to set up



# Flux

- "Data scripting language"
- Good standard library
- First of its kind

```
1 dataStream1 = from(bucket: "bucket1")
2   |> range(start: -1h)
3   |> filter(fn: (r) =>
4     r._measurement == "network" and
5     r._field == "bytes-transferred"
6   )
7
8 dataStream2 = from(bucket: "bucket1")
9   |> range(start: -1h)
10  |> filter(fn: (r) =>
11    r._measurement == "httpd" and
12    r._field == "requests-per-sec"
13  )
14
15 join(
16   tables: {d1:dataStream1, d2:dataStream2},
17   on: ["_time", "_stop", "_start", "host"]
18 )
```




# TimescaleDB

- "Only" a Postgres extension
- Works flawlessly, almost invisibly
- Keeps all the good Postgres stuff
  - All of SQL
  - Tooling
  - Clustering



Hypertable



Chunk

# Visualization

-  Grafana
- Data source plugins for nearly every database
- Great visualization

