Time Series DB for IoT

Choosing the Right IoT Data Platform
On-Premises and Azure

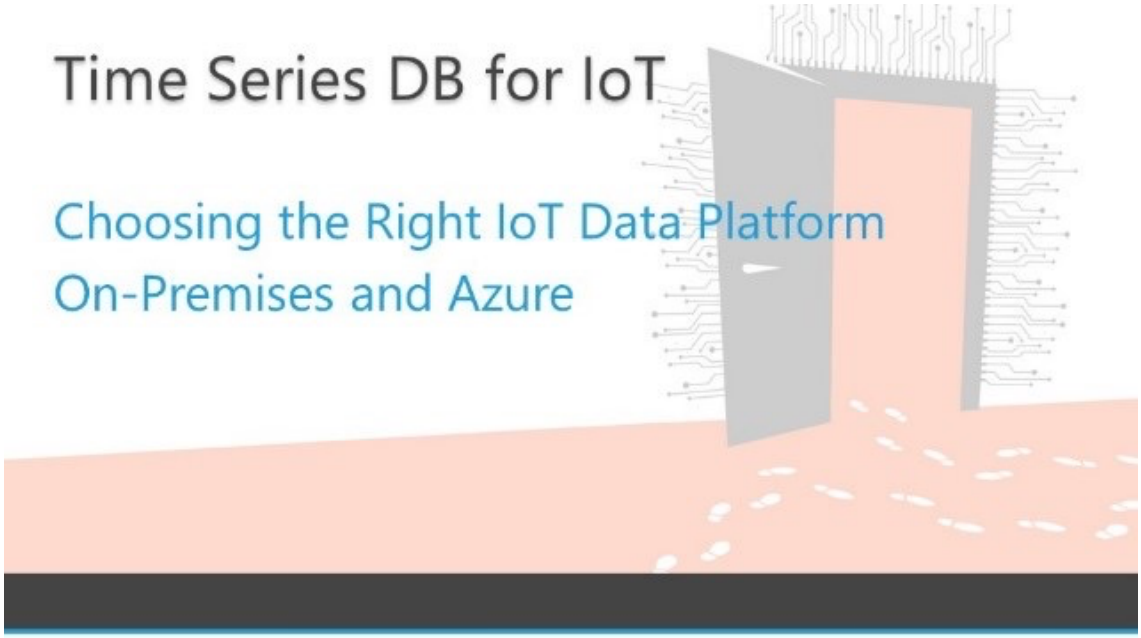

September 12 2017 IoT Summit Sofia, Bulgaria

About me

- **Project Manager @** ICB SOFTWARE & CONSULTING
  - 15 years professional experience
  - .NET Web Development MCPD
- **External Expert Horizon 2020**
- **External Expert Eurostars & IFD**
- **Business Interests**
  - Web Development, SOA, Integration
  - Security & Performance Optimization
  - IoT, Computer Intelligence
- **Contact**

ivelin.andreev@icb.bg
www.linkedin.com/in/ivelin
www.slideshare.net/ivoandreev

Microsoft CERTIFIED Professional Developer

European Commission

eurostars™

Innovationsfonden

2017

Agenda

Time Series

- Why Time Series
- Choosing the Right IoT Data Platform
- InfluxDB vs Competitors
- Key Concepts
- Demo

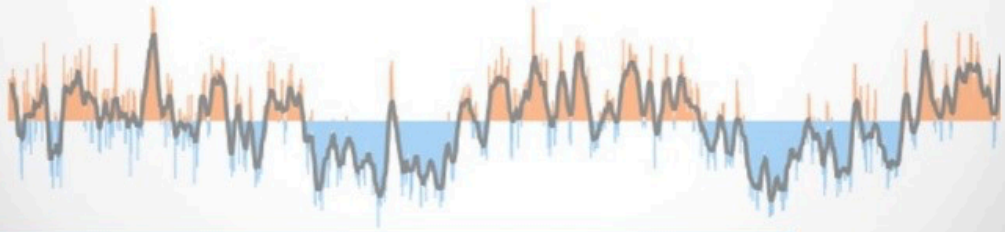

IoT Data

- **IoT** (buzzword of the year)
  - 30 billion "things" by 2020 (Forbes)
- **Top IoT Industries**
  - Manufacturing
  - Healthcare
- **What are the benefits from IoT?**
- **The discussion has shifted**
  - How to make IoT work?
  - How to gain insight on hidden relations?
  - How to get actionable results?

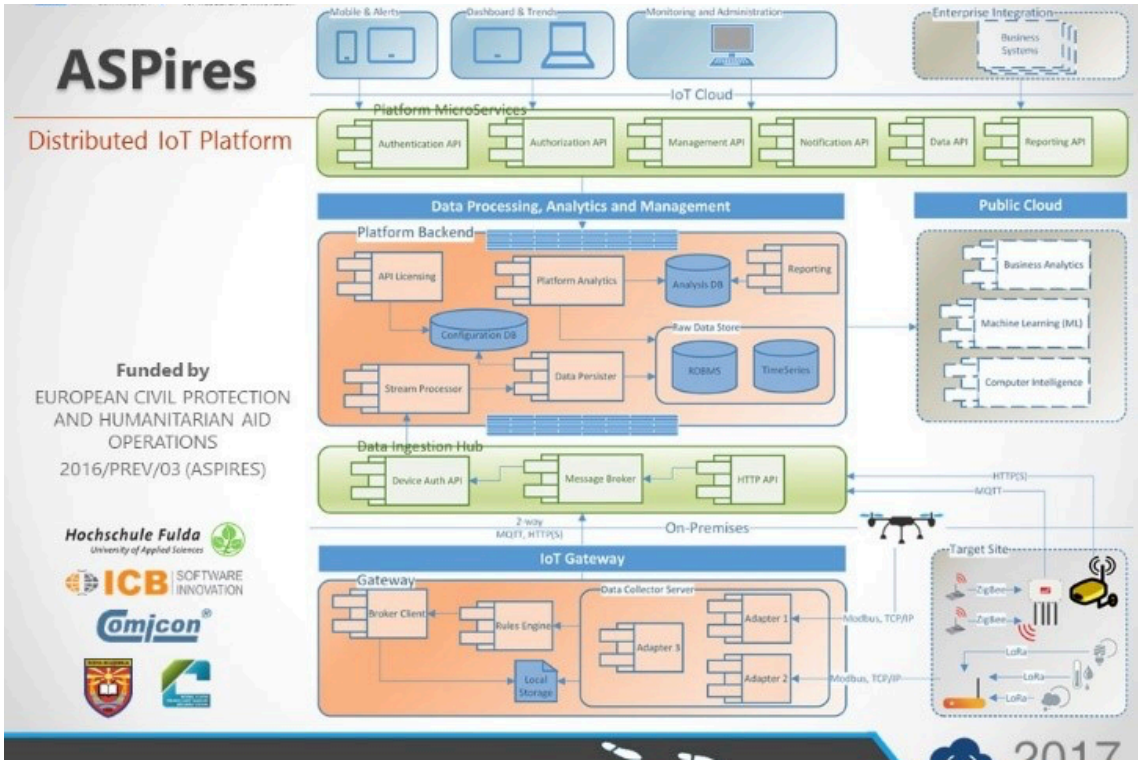
ASPires
Distributed IoT Platform
Funded by
EUROPEAN CIVIL PROTECTION AND HUMANITARIAN AID OPERATIONS
2016/PREV/03 (ASPIRES)
Hochschule Fulda
ICB SOFTWARE INNOVATION
Comicon
Mobile & Alerts
Dashboard & Trends
Monitoring and Administration
Enterprise Integration
Business Systems
IoT Cloud
Platform MicroServices
Authentication API
Authorization API
Management API
Notification API
Data API
Reporting API
Data Processing, Analytics and Management
Public Cloud
Platform Backend
API Licensing
Platform Analytics
Analysis DB
Reporting
Configuration DB
Raw Data Store
Stream Processor
Data Persister
RDBMS
TimeSeries
Business Analytics
Machine Learning (ML)
Computer Intelligence
Data Ingestion Hub
Device Auth API
Message Broker
HTTP API
HTTP(S)
MQTT
2-way MQTT, HTTP(S)
On-Premises
IoT Gateway
Gateway
Broker Client
Rules Engine
Local Storage
Data Collector Server
Adapter 1
Adapter 3
Adapter 2
Modbus, TCP/IP
Target Site
ZigBee
LoRa
2017

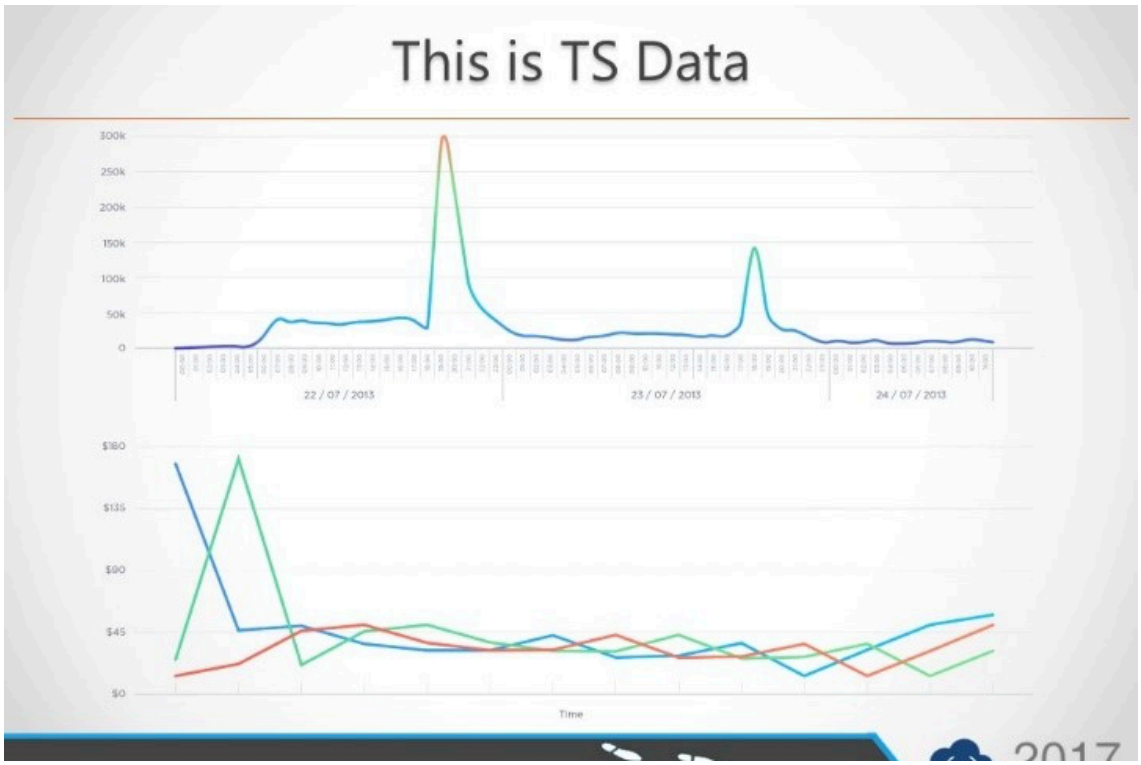
This is TS Data
300k
250k
200k
150k
100k
50k
0
22 / 07 / 2013
23 / 07 / 2013
24 / 07 / 2013
$180
$135
$90
$45
$0
Time
2017

IoT Data is Time Stamp Data

- **Series**
  - Identified by source name (i.e. sensor ID) and metric name (i.e. temp)
  - Consists of ordered {time, value} measurements
  - Regular and Irregular

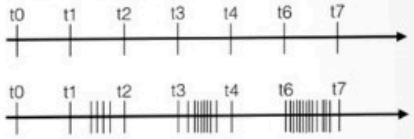

- **Time Series DB**
  - Optimized for TS Data
- **Process Historian – more than TS DB**
  - Interfaces to read data from multiple data sources
  - Render graphics for meaningful points
  - Statistical process control
  - Redundancy and high availability

And this is NOT

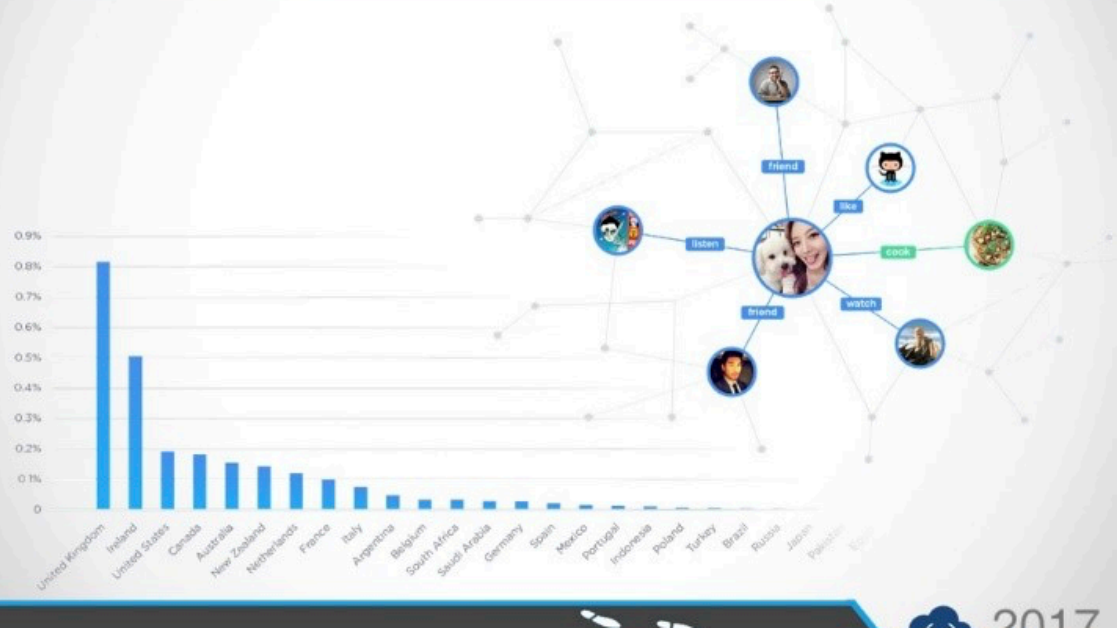

TS Data Characteristics

- **Writes**
  - 95%-99% of all operations
  - Streaming live data from multiple devices
  - Typically sequential appends
- **Updates** to modify values are rare
- **Deletes** are bulk on large ranges *(days, months, years)*
- **Queries**
  - Typically sequential
  - Concurrent reads are common
- **Performance issues** are typically I/O bound
  - Caching does not work well for BigData
  - Systems are typically distributed by design

Credits: Baron Schwartz

https://www.xaprb.com/blog/2014/06/08/time-series-database-requirements/

Data Platform Services for IoT Data

- Data Aggregation Services
  - Collect, normalize, aggregate metrics and events data
- Data Storage Services
  - Distributed by nature
  - High write load, fast retrieval, efficient compression
- Analytics and Visualization Services
- **Technologies not designed for the use case**
  - Relational DB Engines
  - Columnar or key-value databases (HBase, Cassandra)
  - Search engine (Elasticsearch)
  - Document-oriented (MongoDB, DocumentDB)
  - First gen. TS DB (Graphite, OpenTSDB)

Validation Setup

InfluxDB	v1.3.0 single node	**SQL Server**	2016 Standard
CPU: Intel Core i7 3.40GHz (quad, 8MB cache)			
Memory: 16GB DDR3 1333MHz			
OS: Windows 10 Enterprise, 64 bit			
SSD: Samsung 850 EVO 250GB			
HDD: Seagate 7200RPM, 16MB cache, 500GB			

InfluxDB schema:

Name	Type	
tagID	Int	Tag Key
value	Int	Field Key
time	DateTime	TimeStamp

SQL Server schema:

Name	Type	Indexed
ID (PK)	int	Yes
TagID	int	Yes
Value	int	No
Timestamp	DateTime	Yes

When TS overperform RDBMS

- **Target scenarios**
  - High I/O rate
  - Number of tags
  - Volume of data
  - Aggregation of irregular data
  - Compression & De-duplication
- **Requires a learning**
  - Do you expect that many data?
  - Do you need to plot?
  - Do you need to aggregate?

Write Performance (writes/sec)

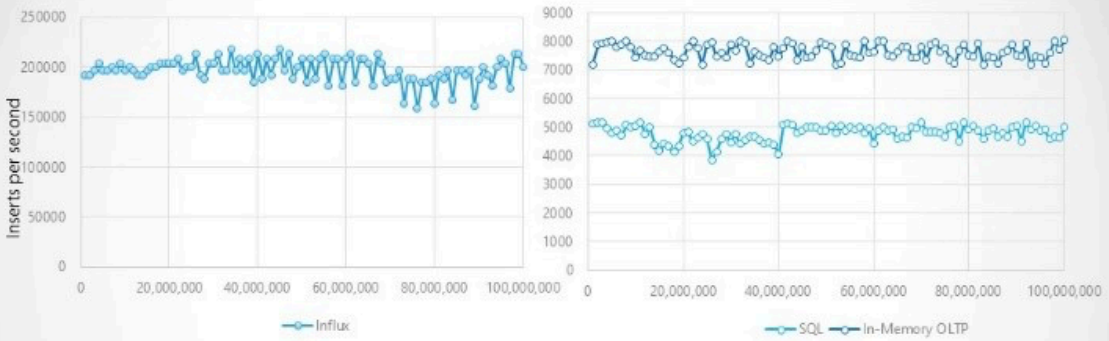

- Influx average write speed of 200,000 writes/second
- SQL Server average write speed of 5,000-8,000 writes/second
- Memory optimized tables increases write performance by 1.4x
- Write trend unaffected by the amount of data (in this volumes)
- Influx outperforms SQL Server 40x

InfluxDB vs SQL Server 2016

Credits: Angelin Nedelchev (ICB)

2017

Disk Storage

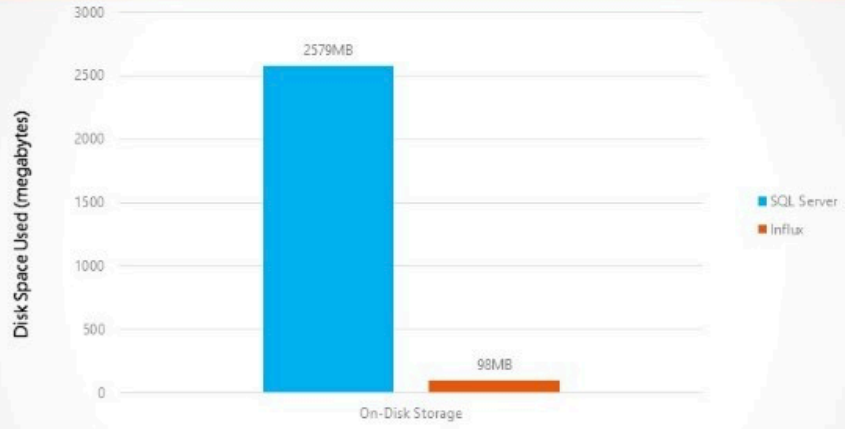

- 27% improvement is SQL Server page compression is enabled (1883MB)
- Influx has 19x-26x less disk requirements for the same functionality

2017

Read Performance (queries/sec)

Benchmark query:

Aggregate samples of 10,000 points, average over 1,000 runs.

- SQL Server mean query response time – 78ms (13 queries/sec)
- Influx mean query response time – 1,3ms (769 queries/sec)
- Influx has 59x better query throughput than SQL Server

Database as a Service in Azure

- **Typical Azure IoT Services** (Before 2017.04.20)

- **With Time Series Insigths** (beta)

- **Other**
  - CosmosDB with MongoDB (SLA 10ms reads, 15ms writes)

Read Performance (rows at a time)

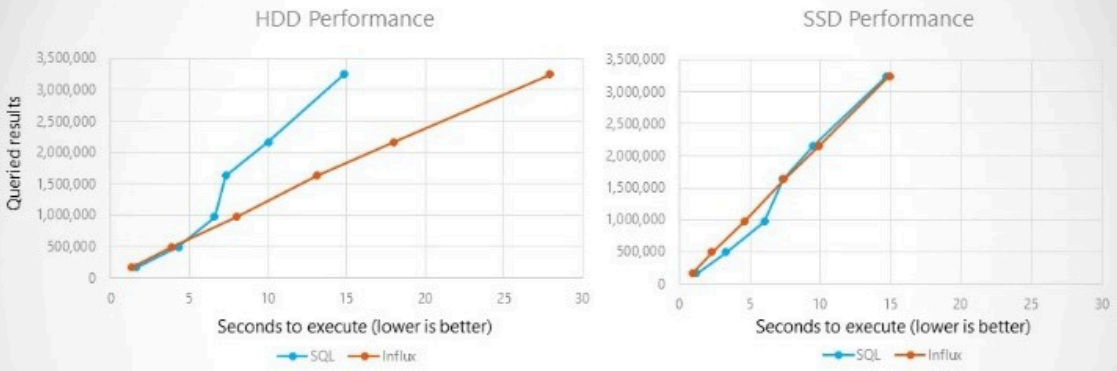

- SQL Server relatively unaffected by SSD (due to caching)
- Influx performance improves up to 100% on SSD, chunk size 10'000
- SQL Server is up to 2x better on HDD (Influx is better up to 600K)
- Relatively equal on SSD (Influx is better up to 1.6M)

Azure Stream Insights

Fully managed analytics, storage and visualization service

Benefits

- Reduced number of services used
- Monitor IoT solutions
- Event source from IoT Hub and Event Hub (only)
- Visualize and analyze data at large scale
- Root cause analysis and anomaly detection
- APIs available for management of raw data

Cons

- Pricey (€126.53, 1M events - €1,138.50 month/unit, 10M events)
- Specific use cases
- Limited capabilities

InfluxDB vs NonSQL **for TS Data**

- InfluxDB vs MongoDB
  - WRITE: 27x greater
  - STORAGE: 84x less
  - QUERY: Equal performance
- InfluxDB vs Elasticsearch
  - WRITE: 8x greater
  - STORAGE: 4x less
  - QUERY: 3.5x – 7.5x faster
- InfluxDB vs OpenTSDB
  - WRITE: 5x greater
  - STORAGE: 16.5x less
  - QUERY: 4x faster
- InfluxDB vs Cassandra
  - WRITE: 5.3x greater
  - STORAGE: 9.3x less
  - QUERY: up to 168x faster
- InfluxDB vs DocumentDB
  - More popular (Google Trends)
  - Cloud and on-premises
  - No external dependencies
  - Aggregations

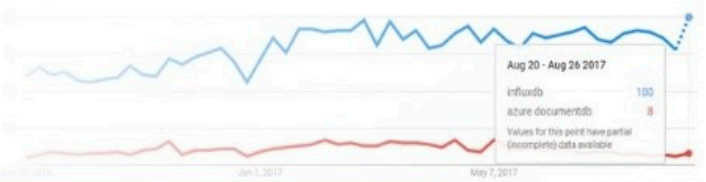

2017

InfluxDB

- **Open-source distributed TS database**
- **Key Features**
  - Easy setup, no external dependencies, implemented in Go
  - Runs on Linux, Windows, OS X
  - Supports .NET, Java, JS, R, PHP, Python, Ruby, Go, Node.js
  - Comprehensive documentation
  - Scalable, highly efficient
  - REST API (JSON)
  - SQL-like syntax
  - On-premise and cloud
- **Top ranked TS DB**

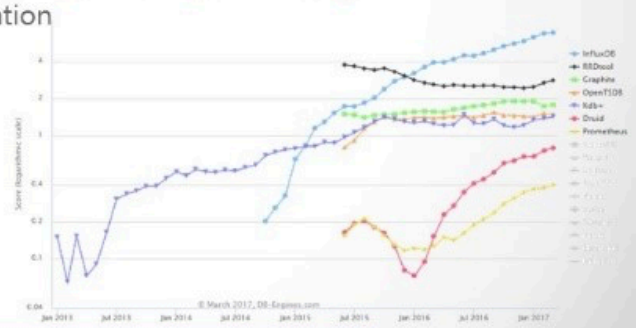

2017

Key Concepts

Term	Description
Measurement	Container (Table)
Point	Single record for timestamp
Field Set	Required; Not-indexed
Field key	Define what is measured
Field value	Actual measured value (string, bool, int64, float64)
Tag Set	Metadata about the point Optional; Indexed; Key-value;
Tag key	Unique per measurement
Tag value	Unique per tag key
Series	Points with common tag set

- Aggregation functions
- Retention policies
- Downsampling
- Continuous queries

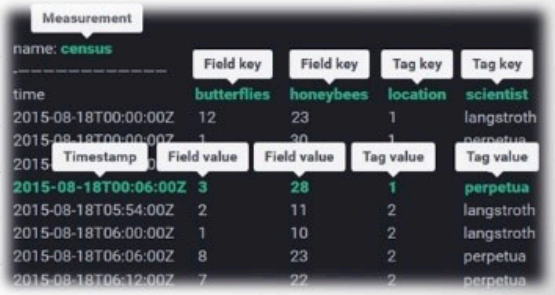

InfluxData Advantages

- **End-to-end solution**
- **Flexible tag support**
  - Graphite, RRD (no tags at all)
  - OpenTSDB, Kairos (tag number)
- **Regular & Irregular TS**
  - Graphite, OpenTSDB
- **Multiple Data Types**
- **High Performance**
  - Hbase, Elastic, Cassandra
- **Functionality**
  - Hbase, Cassandra

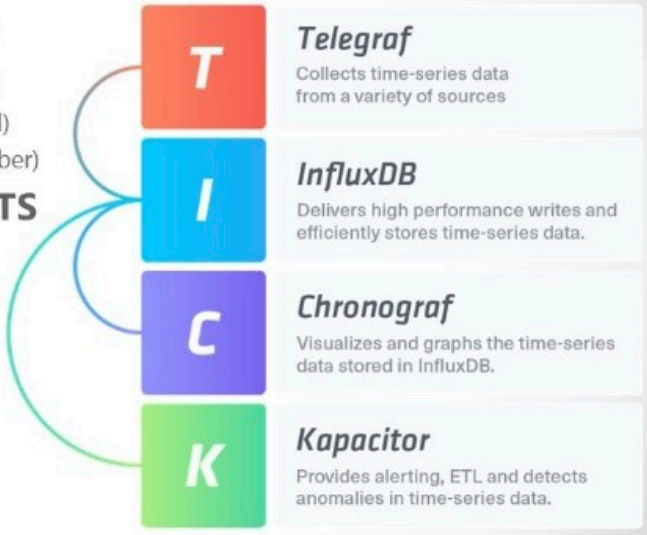

Scalability & Sizing Guidelines

- **Single node or cluster**
  - Single node is open source and free
- **Recommendations**

Load	Resources	Writes/Sec	Moderate Queries/Sec.	Unique Series
Low	Cores: 2-4; RAM: 2-4 GB	0 - 5K	0 - 5	0 – 100K
Moderate	Cores: 4-6; RAM: 8-32 GB	5K - 250K	5 - 25	100K - 1M
High	Cores: 8+; RAM: 32+	250K – 750K	25 - 100	1M - 10M

- **Query complexity (Moderate)**
  - Multiple functions, few regular expressions
  - Complex GROUP BY clause or sampling over weeks
  - Runtime 500ms – 5sec

InfluxDB & Grafana

Functions

Aggregations	Selectors	Transformations	Predictors
COUNT()	BOTTOM()	CEILING()	HOLT_WINTERS()
DISTINCT()	FIRST()	CUMULATIVE_SUM()	
~~INTEGRAL()~~	LAST()	DERIVATIVE()	
MEAN()	MAX()	DIFFERENCE()	
MEDIAN()	MIN()	ELAPSED()	
MODE()	PERCENTILE()	FLOOR()	
SPREAD()	SAMPLE()	~~HISTOGRAM()~~	
STDDEV()	TOP()	MOVING_AVERAGE()	
SUM()		NON_NEGATIVE_DERIVATIVE()	

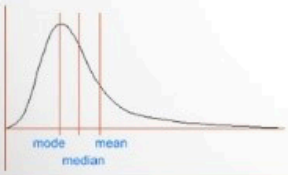

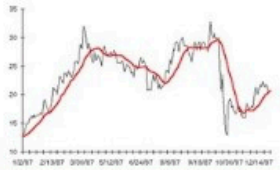

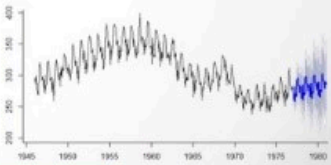

Takeaways

- **Why Time Series Matters for IoT**
  - https://www.influxdata.com/_resources/**webinar-time-series-monitoring-metrics-real-time-analytics-iotsensor-data/**
- **InfluxDB**
  - https://www.influxdata.com/**resources/**
  - https://docs.influxdata.com/influxdb/v1.3/**guides/hardware_sizing/**
  - https://docs.influxdata.com/influxdb/v1.3/**concepts/glossary/**
  - https://docs.influxdata.com/influxdb/v1.3/query_language/**schema_exploration/**
- **Grafana Plugins**
  - https://**grafana.com**/plugins
- **Time Series Insights**
  - https://docs.microsoft.com/en-us/azure/time-series-insights/**time-series-insights-overview**


Thanks to our Sponsors:
Global Sponsor:
Microsoft
With the support of:
epam
ICB
SOFTWARE
INNOVATION