



NINTH ANNUAL LEADERSHIP EVENT

CYBER SECURITY

Security solutions through collaboration.[™] **SUMMIT**

October 28–30, 2019 | Minneapolis Convention Center

cybersecuritysummit.org | [#cybersummitmn](https://twitter.com/cybersummitmn)

Medical Device Security Logging: The Apache Software Foundation & Time Series Databases

Presented by:

Kevin A. McGrail

kmcgrail@InfraShield.com

About the Speaker



Kevin A. McGrail
Director, Business Growth InfraShield

Member of the Apache Software Foundation,
Release Manager for Apache SpamAssassin,
Director at the Dysautonomia Support Network,
Advisor to SecurityUniversity.edu and Virtru, and
Google G Suite TC, GDE & Ambassador.

<https://www.linkedin.com/in/kmcgrail>

The Apache Software Foundation



Who is the ASF?

The Apache Software Foundation is a 501(c)(3) Charity often referred to as just Apache or the ASF.

501(c)(3) Charity not a 501(c)(6) Trade Organization

We're known for the HTTP server and the Apache Software License.

What is the ASF's Mission?

To provide software for the public good.

***We do this by providing services and support
for many diverse software project
communities of individuals AT NO CHARGE.***

The Apache Software License

The ASLv2 is known for its permissive, business-friendly stance with patent grants and without copyleft provisions.

Powered by Apache

80% of the world's websites use our software

Every Smartphone in the world uses our software

Every plane in US airspace is tracked w/our software



Projects.Apache.org & The Incubator

There are currently 388 open source initiatives at the ASF:

201 committees managing 334 projects

5 special committees

49 incubating podlings



Quiz Break

Quiz - Part A

*\$s

Quiz - Part B

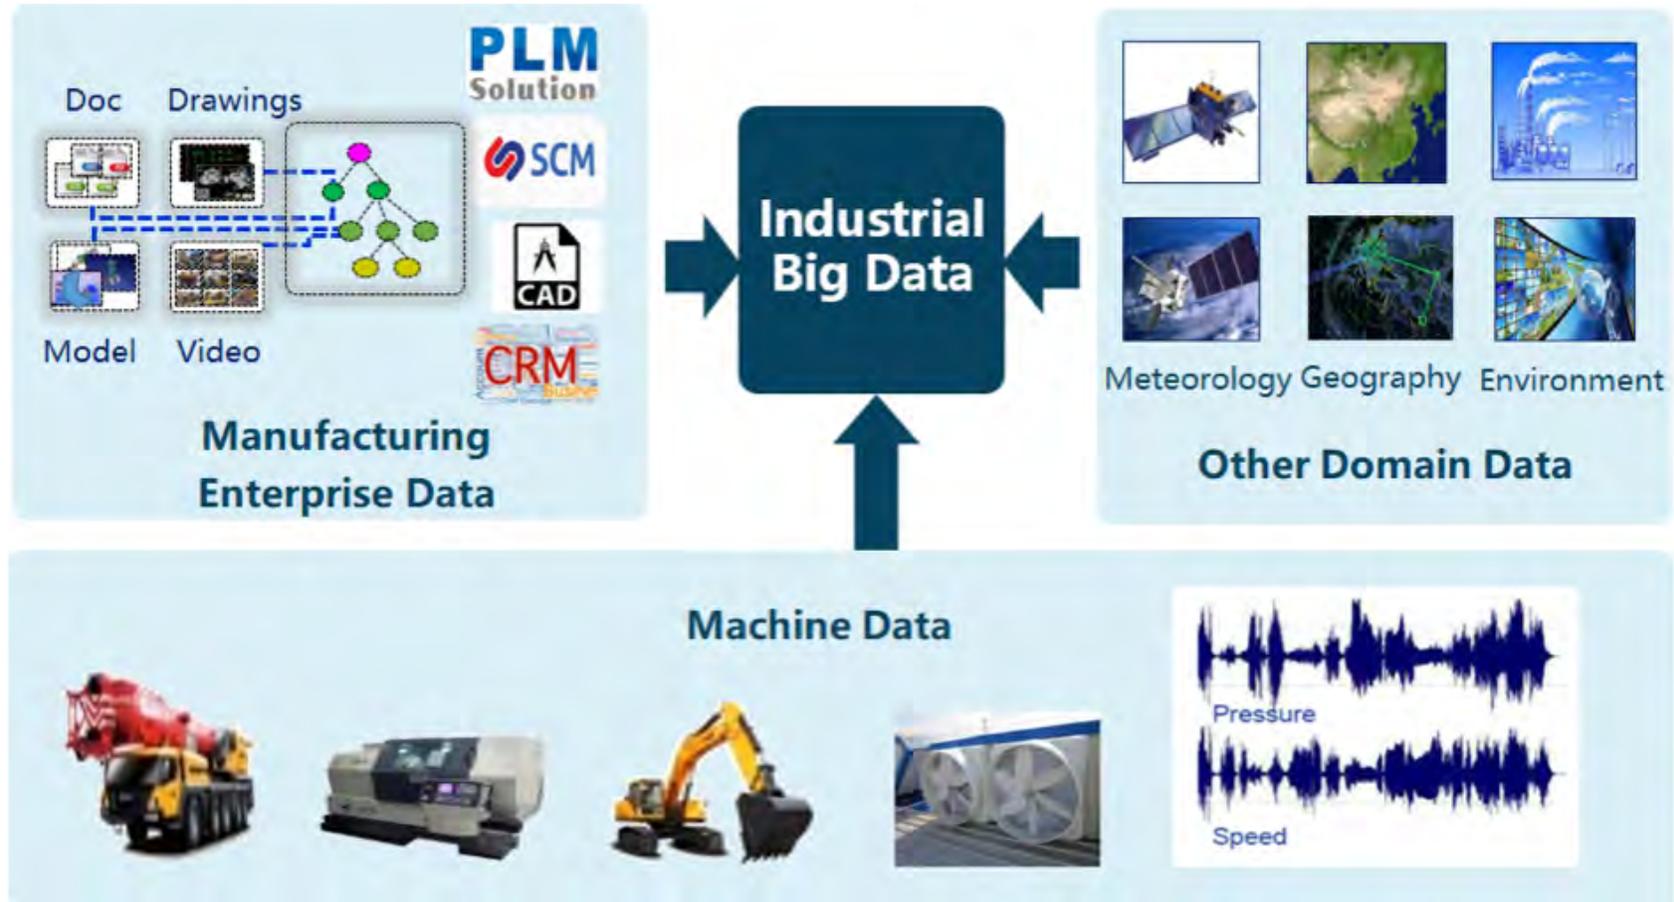
8-)}-----o

Quiz - Part C

8-) [> - = / / / >

What is a Time Series Database?

Machine Data



Big Data

“Big data is data sets that are so voluminous and complex that traditional data-processing application software are inadequate to deal with them.” Wikipedia

Apache has 48 projects under Big Data!

Apache Hadoop

Quiz: What's in a name?

Quiz: Why did they really call it Big Data?



Challenges with Machine Data

Produce Data 7*24
with High Frequency
and Large Volume



Big Data

50Hz, 500points/machine,
20K wind-turbines machines,
totally up to 500 million points/sec

□ Feature of Data Ingestion

- Out-of-order sometimes
- Sparse Table
(different machine has
different sensors)

□ Feature of Data Query

- Time Dimension is always
accessed
- Aggregation is the first-class
citizen
- Time-series-specific query
and analysis



□ Challenges

- Large Volume
- High Throughput
- Low Cost
(historical data)
- Low Latency for
Query
- Fast Aggregation
- Query-Analysis
hybrid workloads



The Solution

Apache IoTDB (Incubating)



Apache IoTDB Features



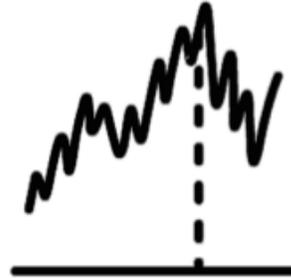
Persist data efficiently

- Millions points ingestion per sec per node
- Tens of millions of time series



Query data with low latency

- Efficiently filter data: millions of points per sec
- Aggregation: tens of ms latency on billions of points



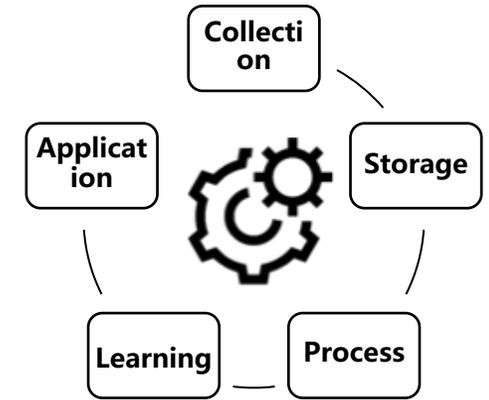
Exclusive operations of time series

- Segmentation
- Representation
- Subsequence matching
- Time-frequency transform
- Visualization



Integration with existing ecosystem

- Kafka
- MatLab
- Spark
- MapReduce
- Grafana



Cover the life cycle of data

- Connecting Edge to the Cloud
- Powerful query engine
- User Friendly analytics

How Does it Do This?

TsFile

Time series data files: high-tech write, high compression ratio, support for simple query

Simply put, **TsFile is a zip file for time series data.**

Suitable for embedded devices!

IoTDB

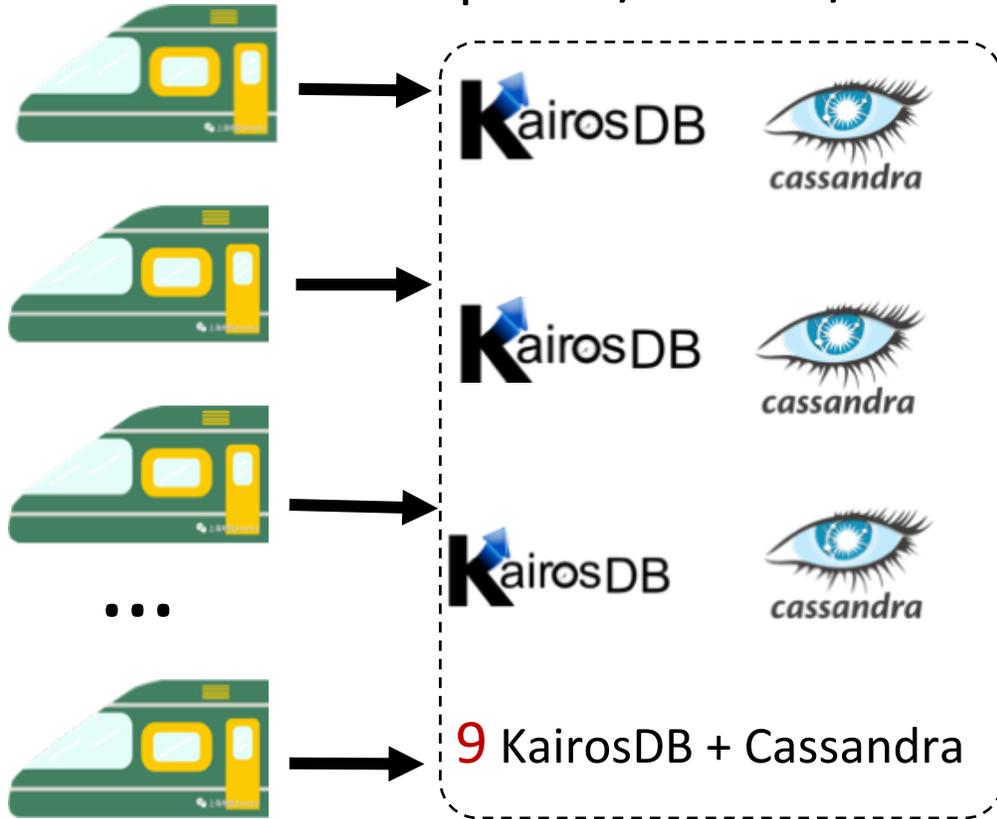
Efficiently operate on time series data from multiple TsFiles, including : CRUD and advanced queries like : max, min, avg and temporal alignment



Example: Shanghai Metro Monitoring

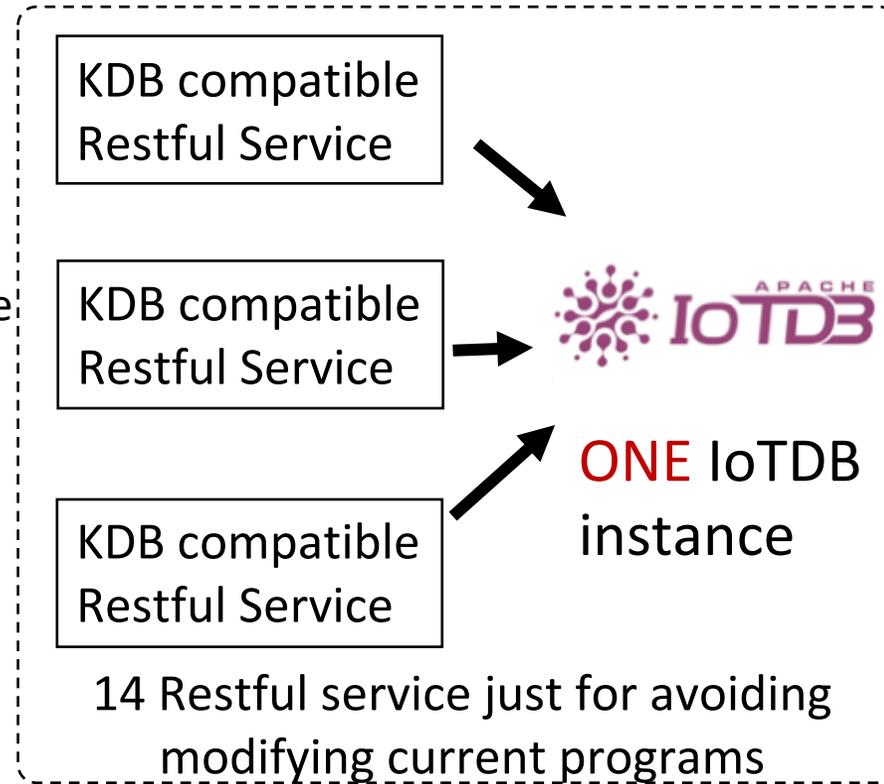
3200 points/500 ms/train

3200 points/200 ms/train



144 trains

upgrade
➔



300 trains

414 Billion
data points
per day
just using
ONE IoTDB
instance

$$300 \times 3200 \times 5 \times 86400 = 414,000,000,000$$

Join Apache IoTDB

- Mail list:
 - subscribe: dev-subscribe@iotdb.incubator.apache.org
 - discussion: dev@iotdb.apache.org
 - bug report: <https://issues.apache.org/jira/projects/IOTDB/issues/IOTDB>
- Website: <https://iotdb.apache.org>
- Ecosystem target:



IoTDB v0.8.0 is released! (the first Apache release version)

Thanks!

Image Credits:

KAM photo taken by Ted King, used with permission.

Company Logos are Brand Resources of their Respective Companies

Thanks to:

Jianmin Wang for lending some of his slides, Xiangdong Huang for his help, the School of Software at Tsinghua University, and everyone working on the Apache IoTDB project!



Kevin A. McGrail

www.linkedin.com/in/kmcgrail