



Time Series DB

Should I use one in my application architecture?



Anais Dotis-Georgiou

Developer Advocate



LinkedIn



Contents

Who is InfluxDB?

What is Time Series?

Why do we need Time Series Databases?

How do I know if a Time Series database is right for my application?

Influx Data - At a glance

FOUNDED	2013 San Francisco HQ, 180 FTE's, 67+ in technical functions
FOCUS	Where developers build real-time applications for IoT, Analytics and Cloud native services
DIFFERENCE	One platform; one API across Multiple Clouds and On-Prem Ingest, query, story using common tools regardless of architecture
OSS FOUNDATION	1300+ Customers and 736,000 daily active OSS deployments; Google , Cisco, SAP, Comcast, Tesla, Siemens, PTC, Honeywell, JP Morgan Chase
BUSINESS MODEL	PLG Driven Usage and Subscription Model Pay for what you use; Pay how you want. Credit card, cloud provider, annual contract

You have probably
used InfluxDB



TESLA



nest™





HOME

SEARCH

WATCHLIST

ORIGINALS


MOVIES

SERIES

DTVE



STAR WARS THE MANDALORIAN



Disney


PIXAR

MARVEL

STAR
WARS

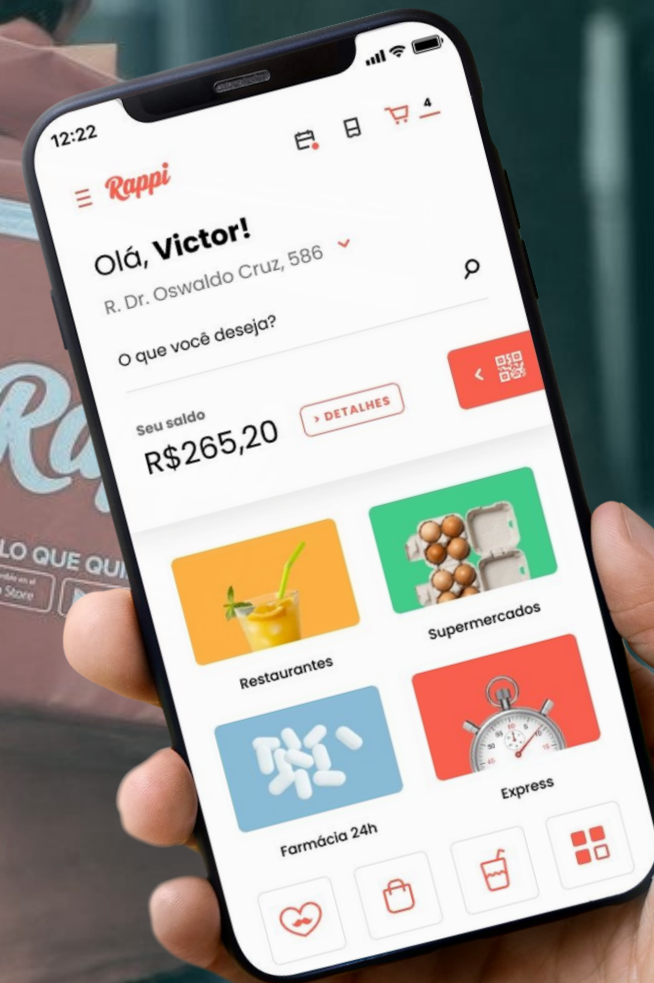
NATIONAL
GEOGRAPHIC

Originals





Rappi



What is Time Series Data?

Time series data is a collection of observations obtained through repeated measurements over time

Time Series Example - Weather Station

December 27, 2022


Time	Temperature	Dew Point	Humidity	Wind	Speed	Gust	Pressure	Precip. Rate.	Precip. Accum.	UV	Solar
12:04 AM	69.7 °F	66.0 °F	88 %	ENE	1.0 mph	1.7 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:09 AM	69.6 °F	66.1 °F	89 %	ENE	1.0 mph	1.4 mph	29.97 in	0.00 in	0.00 in	0	0 w/m ²
12:14 AM	69.6 °F	66.4 °F	89 %	ENE	1.0 mph	1.4 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:19 AM	69.7 °F	66.6 °F	90 %	ENE	0.9 mph	1.2 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:24 AM	69.7 °F	66.6 °F	90 %	ENE	0.5 mph	1.1 mph	29.97 in	0.00 in	0.00 in	0	0 w/m ²
12:29 AM	69.7 °F	66.6 °F	90 %	East	0.7 mph	0.9 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:34 AM	69.7 °F	66.6 °F	90 %	East	0.9 mph	1.4 mph	29.97 in	0.00 in	0.00 in	0	0 w/m ²
12:39 AM	69.9 °F	66.8 °F	90 %	East	0.8 mph	1.4 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:44 AM	70.0 °F	66.9 °F	90 %	East	0.4 mph	0.9 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:49 AM	70.1 °F	67.0 °F	90 %	East	0.3 mph	0.4 mph	29.97 in	0.00 in	0.00 in	0	0 w/m ²
12:54 AM	70.1 °F	67.0 °F	90 %	ENE	0.4 mph	0.7 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²
12:59 AM	70.2 °F	67.1 °F	90 %	East	0.3 mph	0.6 mph	29.96 in	0.00 in	0.00 in	0	0 w/m ²

Time Series Example - Weather Station

Summary December 27, 2022


	High	Low	Average
Temperature	70.2 °F	67.1 °F	68.6 °F
Dew Point	67.1 °F	62.6 °F	64.9 °F
Humidity	92 %	83 %	88 %
Precipitation	0.00 in	--	--

	High	Low	Average
Wind Speed	3.1 mph	0.0 mph	1.4 mph
Wind Gust	4.5 mph	--	1.9 mph
Wind Direction	--	--	ENE
Pressure	29.97 in	29.91 in	--



Time Series Example - Predictive Maintenance

X-Axis




Time Series Example - Health Charting


Time	Pulse	Blood Pressure	Temperature
8:00am	65	120/80	98.6
8:30am	68	110/70	98.5
9:00am	70	112/72	98.5
9:30am	110	98/50	101.5
10:00am	82	120/75	99.2



Metrics and Events



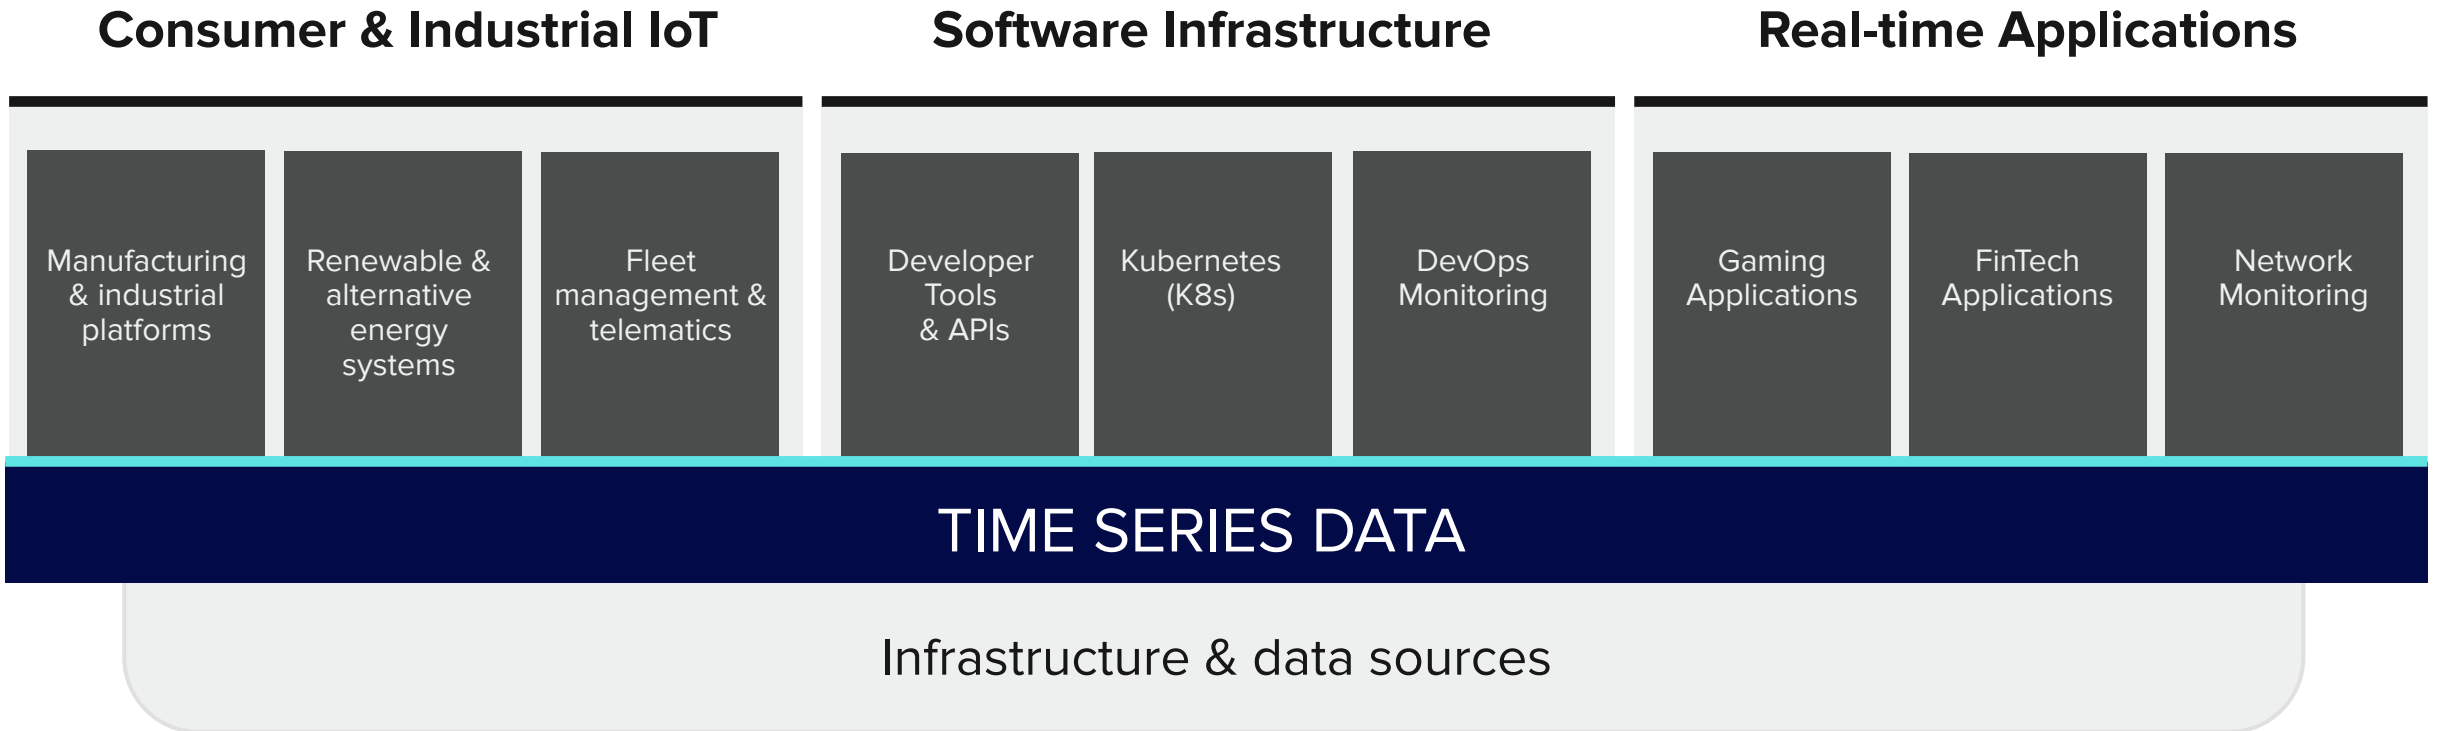
Metrics occur on regular time intervals



Events occur at specific points in time

Why do we need time series databases?

Time series data is advancing in many areas



Additional Data Sources: Computer/Device

Human Generated

- Account Updates
- Billing Updates
- Order Records



Transaction Oriented
Scales with business

Computer/Device Generated

- Log Events
- Device Status
- Sensor Status



Time Oriented
Scales with many factors

Scalability

Interval	Number per second	Number of generated records per day
Minute	.0166666666	1,440
Second	1	86,400
Millisecond	1 thousand	86.4 Million
Microsecond	1 million	86.4 Billion
Nanosecond	1 billion	86.4 Trillion

1 Device - at various intervals

→ Time Series databases are designed to be able to scale to a large volume of ingestion data

Queryability & Query Performance

Time Series Databases are typically organized by time

Time Series Database are optimized for fast retrieval of blocks of time series data



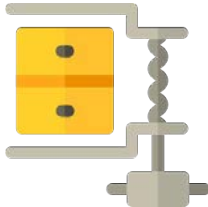

Summarization/Down Sampling

Weather Station Example

- Time Series Data can typically be summarized to reduce its size - often referred to as downsampling
- With Time Series, this can often happen at the database engine level without the need for a lot of additional coding

Today	Tomorrow
1 Sample Temperature Every 5 minutes	Yesterday's High & Low & Average Temperature
288 Records	1 Record

Data Storage


Time Series Database Feature	Benefit
DownSampling/Summarization	Reduces Number of Records retained, reducing storage size
<div data-bbox="236 558 524 605">Compression</div> <div data-bbox="1052 611 1258 815"></div>	The nature of time series data typically makes it very highly compressible ; this often gives time series databases a big advantage on needed storage size for the same amount of data
<div data-bbox="236 882 738 929">Data Retention Policies</div> <div data-bbox="1070 955 1240 1159"></div>	Time Series Databases typically let you set a retention period for data and data that falls outside that retention period is automatically removed; based on the organization of data by time, this mass deletion of data doesn't have negative impact on your database like it can with some databases



How do I know if Time Series is
right for my application?

Does my data have a time
element?


Do I care about changes in my
data over time?



How much data will I be working
with?

Will there be a need to do any
analytics with my data in the
future?

Am I concerned with Storage
Costs?



Am I concerned about application
performance with time series
data?

Summary: There are Great Options Out There

RELATIONAL

- Orders
- Customers
- Records



DOCUMENT

- High throughput
- Large document



SEARCH


- Distributed search
- Logs
- Geo



TIME SERIES

- Events, metrics, time stamped
- for IoT, analytics, cloud native

Time series is fastest growing data category by far



source: DB Engines

Getting started with InfluxDB

Sign up

Influxdata.com

Get InfluxDB

Via cloud marketplace



Google Cloud



Microsoft
Azure

Learn

- ✓ Self-service content
- ✓ Documentation
- ✓ InfluxDB University



THANK YOU

InfluxData



influxdata.com