



Benchmarking InfluxDB vs. MongoDB for Time Series Data, Metrics & Management

TECHNICAL PAPER

External Contributors

Vlasta Hajek Senior Software Engineer, Bonitoo

Ales Pour Engineer, Bonitoo

Ivan Kudibal Engineering Manager, Bonitoo

January 2019

An abstract graphic at the bottom of the page consisting of numerous thin, light blue lines that form a complex, wavy pattern. The lines are connected by small blue dots, creating a sense of data flow and connectivity. A prominent feature is a bright yellow circular highlight at the top right, from which several lines radiate downwards, suggesting a focal point or a specific data point of interest.

| | |
|---|-----------|
| Summary | 3 |
| Introduction | 3 |
| Why time series? | 4 |
| Test design | 4 |
| About InfluxDB | 4 |
| InfluxDB version tested: v1.7.2 | 4 |
| About MongoDB | 4 |
| MongoDB version tested: v4.0.4 | 4 |
| Comparison-at-a-glance | 5 |
| Overview | 6 |
| The dataset | 6 |
| Overview of the parameters for the sample dataset | 6 |
| Test methodology | 7 |
| Write performance | 7 |
| On-disk storage requirements | 8 |
| Query performance | 9 |
| Testing hardware | 10 |
| User experience comparison | 10 |
| Overview | 10 |
| Mental models | 10 |
| System configuration | 11 |
| Inapplicable advice from MongoDB documentation | 11 |
| Schema design | 11 |
| Compression and disk usage | 13 |
| Ad-hoc query composition | 13 |
| Go language support | 14 |
| User experience conclusion | 14 |
| About InfluxData | 15 |
| InfluxDB documentation, downloads & guides | 15 |
| What is time series data? | 15 |
| What is a time series database? | 15 |

Summary

In the course of this benchmarking paper, we looked at the performance of InfluxDB and MongoDB performance across three vectors:

- Data ingest performance - measured in values per second
- On-disk storage requirements - measured in bytes
- Mean query response time - measured in milliseconds

The benchmarking tests and resulting data demonstrated that InfluxDB outperformed MongoDB in data ingestion, on-disk storage, and query performance by a significant margin. Specifically:

- InfluxDB outperformed MongoDB by **3.9x** when it came to data ingestion.
- InfluxDB outperformed MongoDB by delivering **16.6x** better compression.
- InfluxDB outperformed MongoDB by up to **16%** when measuring query performance.

Introduction

In this technical paper, we'll compare the performance and features of InfluxDB and MongoDB for common [time series](#) workloads, specifically looking at the rates of data ingestion, on-disk data compression, and query performance. This data should prove valuable to developers and architects evaluating the suitability of these technologies for their use case. Specifically, the time series data management use cases involving building [DevOps Monitoring](#) (Infrastructure Monitoring, Application Monitoring, Cloud Monitoring), [IoT Monitoring](#), and [Real-Time Analytics](#) applications.

Our goal with this benchmarking test was to create a consistent, up-to-date comparison that reflects the latest developments in both InfluxDB and MongoDB. Periodically, we'll re-run these benchmarks and update this document with our findings. All of the code for these benchmarks is available on [GitHub](#). Feel free to open up issues or pull requests on that repository or if you have any questions, comments, or suggestions.

Why time series?

Time series data has historically been associated with applications in finance. However, as developers and businesses move to instrument more in their servers, applications, network and the physical world, time series is becoming the de facto standard for how to think about storing, retrieving, and mining this data for real-time and historical insight. To learn more about why you should insist on using a purpose-built, time series backend versus attempting to retrofit a document, full-text, or RDBMS to satisfy your use case, check out the [“Why Time-Series Matters for Metrics, Real-Time and IoT/Sensor Data”](#) technical paper.

Test design

About InfluxDB

InfluxDB version tested: v1.7.2

InfluxDB is an open-source time series database written in Go. At its core is a custom-built storage engine called the Time-Structured Merge (TSM) Tree, which is optimized for time series data. Controlled by a custom SQL-like query language named InfluxQL, InfluxDB provides out-of-the-box support for mathematical and statistical functions across time ranges and is perfect for custom monitoring and metrics collection, real-time analytics, plus IoT and sensor data workloads.

About MongoDB

MongoDB version tested: v4.0.4

MongoDB is an open-source, document-oriented database, colloquially known as a NoSQL database, written in C and C++. Though it's not generally considered a true Time Series Database per se, its creators often promote its use for [time series workloads](#). It offers modeling primitives in the form of timestamps and bucketing, which give users the ability to store and query time series data. Please note that this paper does not look at the suitability of InfluxDB for workloads other than those that are time series-based. InfluxDB is not designed to satisfy document storage use cases and therefore those will not be explored in this paper. For these use cases, we recommend sticking with MongoDB or similar NoSQL databases.

Comparison-at-a-glance

| | InfluxDB | MongoDB |
|-------------------|---|---|
| Description | Database designed for time series, events and metrics data management | Scalable, document-oriented NoSQL database |
| Website | https://influxdata.com/ | https://www.mongodb.com/ |
| GitHub | https://github.com/influxdata/influxdb | https://github.com/mongodb/mongo |
| Documentation | https://docs.influxdata.com/influxdb/latest/ | https://docs.mongodb.com/manual/ |
| Initial release | 2013 | 2009 |
| Latest release | v1.7.2, November 2018 | v4.0.4, November 2018 |
| License | Open Source, MIT | Open Source, AGPL |
| Language | Go | C/C++ |
| Operating Systems | Linux, OS X | Linux, OS X, Windows |
| Data Access APIs | HTTP Line Protocol, JSON, UDP | JSON, BSON |
| Schema | Schema-free | Schema-free |

Overview

In building a representative benchmark suite, we identified the most commonly evaluated characteristics for working with time series data. As we'll describe in additional detail below, we looked at performance across three vectors:

1. Data ingest performance - measured in values per second
2. On-disk storage requirements - measured in bytes
3. Mean query response time - measured in milliseconds

CONCLUSION:

InfluxDB outperformed MongoDB in write throughput, on-disk compression, and query performance.

The dataset

For this benchmark, we focused on a dataset that models a common DevOps monitoring and metrics use case, where a fleet of servers are periodically reporting system and application metrics at a regular time interval. We sampled 100 values across 9 subsystems (CPU, memory, disk, disk I/O, kernel, network, Redis, PostgreSQL, and Nginx) every 10 seconds. For the key comparisons, we looked at a dataset that represents 100 servers over a 24-hour period, which represents a relatively modest deployment.

Overview of the parameters for the sample dataset

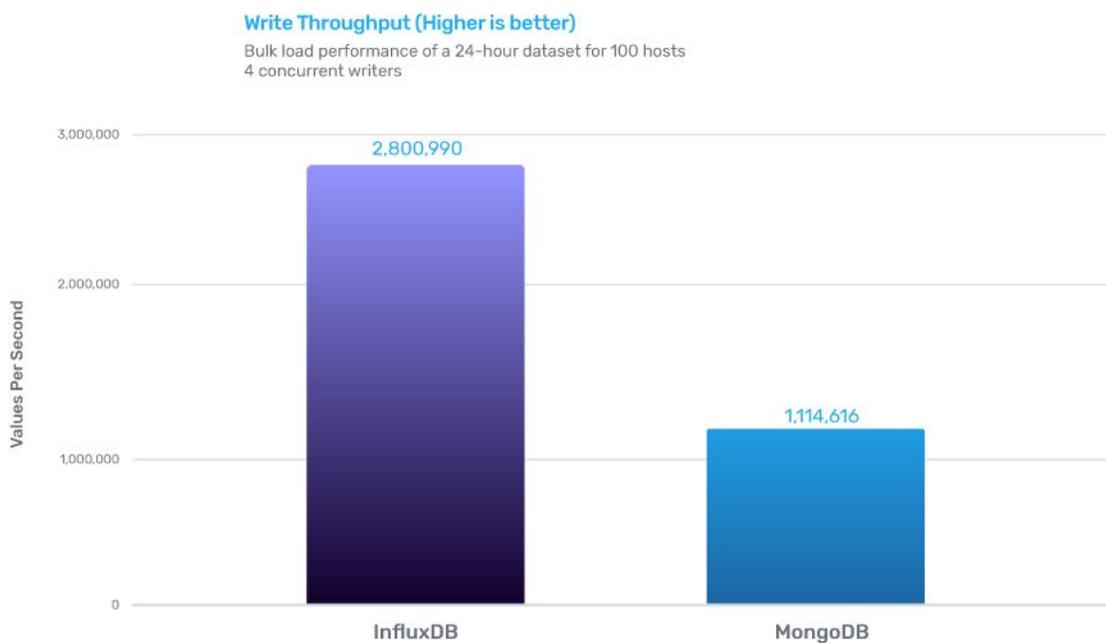
| | |
|----------------------------|------------|
| Number of servers | 100 |
| Values measured per server | 100 |
| Measurement interval | 10s |
| Dataset duration(s) | 24h |
| Total values in dataset | 87,264,000 |

This is only a subset of the entire benchmark suite, but it's a representative example. At the end of this paper we will discuss other variables and their impacts on performance. If you're interested in additional details, you can read more about the testing methodology on [GitHub](#).

Test methodology

Write performance

To test write performance, we concurrently batch loaded the 24-hour dataset with 16 worker threads (to be able to compare to the other databases tests). We found that the average throughput of MongoDB was **1,114,616 values per second**. The same dataset loaded into InfluxDB at a rate of **2,800,990 values per second**, which corresponds to approximately **2.4x faster ingestion** by InfluxDB. (Remember: the concurrency for this test was 16, with 100 hosts reporting.)



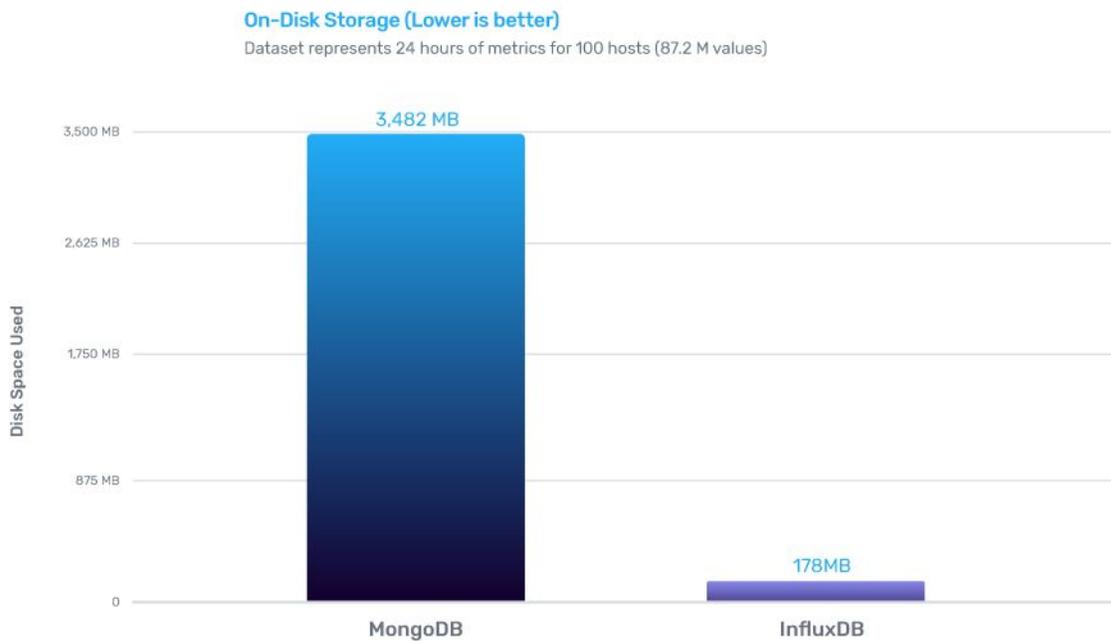
This write throughput stays relatively consistent across larger datasets (i.e. 48 hours, 72 hours, 96 hours).

CONCLUSION:

InfluxDB outperformed MongoDB by 2.4x when evaluating data ingestion performance.

On-disk storage requirements

For the same 24-hour dataset outlined above, we looked at the amount of disk space used after writing all values and allowing each database's native compaction process to finish. We found that the dataset required **3.4 GB** of storage for MongoDB. The same dataset required only **178 MB** for InfluxDB, corresponding to **20x better compression** by InfluxDB. This results in approximately **2.15 bytes per value** for InfluxDB and **43 bytes per value** for MongoDB.



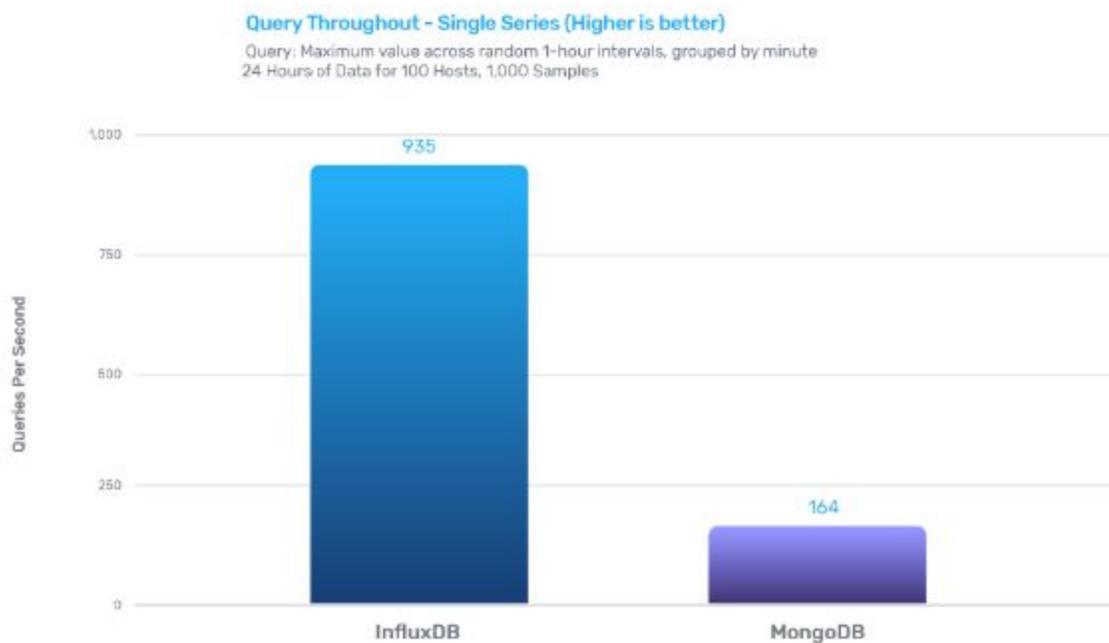
Largely, the additional storage requirement for MongoDB comes from its design as a document-oriented database. Without the time series-optimized compression for values and timestamps that InfluxDB has with its TSM storage engine, MongoDB ends up requiring considerably more space on disk to store the same dataset.

CONCLUSION:

InfluxDB outperformed MongoDB by delivering 20x better on-disk compression.

Query performance

To test query performance, we chose a query that aggregates data for a single server over a random 1-hour period of time, grouped into one-minute intervals, potentially representing a single line on a visualization, a common DevOps monitoring and metrics function. Querying an individual time series is common for many IoT use cases as well.



To reduce variability, the query times were averaged over 1,000 runs. With single worker thread, we found that the mean query response time for MongoDB was **6.09ms (164.2 queries/sec)**. The same query took an average of **1.07ms (934.5 queries/sec)** on InfluxDB, demonstrating approximately **5.7x higher performance**. As concurrency increased, the performance difference remained relatively consistent.

CONCLUSION:

InfluxDB outperformed MongoDB by delivering up to 5.7x faster query performance.

Testing hardware

All of the tests performed were conducted on two virtual machines in AWS infrastructure, running Ubuntu 16.04 LTS. We used instance type r4.4xlarge (Intel Xeon E5-2686 v4 2.3GHz, 16 vCPU, 122 GB RAM, 1x EBS Provisioned 6000 IOPS SSD 250GB) for a database server and c5.2xlarge instance type (Intel Xeon E5-2666v3 2.9GHz, 8 vCPU, 15 GB RAM) for a client host with the data load and query clients.

User experience comparison

Overview

MongoDB is a general-purpose document store. MongoDB is intended to store "schema-less" data, in which each object may have a different structure. In practice, MongoDB is typically used to store large, variable-sized payloads represented as JSON or BSON objects.

Both because of MongoDB's generality, and because of its design as a schema-less datastore, MongoDB does not take advantage of the highly-structured nature of time series data. In particular, time series data is composed of tags (key/value string pairs) and sequences of time-stamped numbers (which are the values being measured). As a result, MongoDB must be specifically configured to work with time series data.

In comparison, InfluxDB is a special-purpose time series database. Thus, it automatically takes advantage of the structure of time series data.

Mental models

MongoDB is not out-of-the-box ready for time series usage. It requires users to be conversant in its particular lexicon. The vocabulary of a MongoDB system is focused on "collections",

"documents", "arrays", and "indexes". None of these map directly to time series data. In all cases, engineering tradeoffs must be made to achieve that mapping.

In contrast, InfluxDB directly uses time series concepts like "measurements", "tags", and "values". This is because InfluxDB is specifically designed to make it easy to store and analyze time series data. As such, InfluxDB is much simpler to think about when working in the domain of time series data. The cognitive burden on an InfluxDB user is much lower than when using MongoDB.

System configuration

Upon startup, the MongoDB process advises the server administrator to disable transparent hugepages. We configured our machines to match this recommendation.

Inapplicable advice from MongoDB documentation

There are many official blog posts and documentation articles that attempt to recommend the best ways to use MongoDB for time series data. However, once the full time series storage problem is considered, those recommendations no longer apply.

For example, we commonly saw advice to use a multi-level data storage strategy, in which each MongoDB document stores an array of point data. The stated goal of this design is to choose a trade-off between the number of separate documents in a collection, the raw data size, and the index size. Specifically, each document would represent a predetermined span of time (typically 60 seconds). Within that document, the array would be pre-allocated to represent at most "one point per second".

There are numerous problems with this approach:

1. It does not permit variable-rate data, which is common in realistic workloads.
2. It does not permit storing different tags with each value, which is critical for performing ad-hoc querying.
3. It requires application-specific tuning, which is brittle and stressful to maintain.

Schema design

Designing a time series application with MongoDB requires significant up-front engineering investment. The decisions made in the planning stage of a MongoDB-based application will

have long-lasting impact on what can be done with the data. For example, some faster configurations store tag data in separate collections, which can make it awkward or impossible to perform ad-hoc querying.

The choices to be made include:

- How many MongoDB collections to use?
- How will numeric values be stored (all as float, or some as integers and some as floats)?
- How much to rely on MongoDB's generic compression, versus use an application-specific optimization?
- To save disk space, should tags be normalized, which introduces coordination problems and is not an idiomatic MongoDB design?
- Is the write-throughput benefit of using multiple collections worth the complexity?
- If using many collections, how to create them without race conditions? Another approach could be to store each measurement in separate collection. Such design would then face limitations such as lack of querying/aggregation across multiple collections.

The schema that we developed is derived from first principles to most correctly and performantly map MongoDB to the time series use case, with particular attention paid to supporting ad-hoc querying. We rely on the copy-on-write and Snappy compression features of MongoDB's new WiredTiger storage engine. In MongoDB, our time series data uses one document per multiple values for a measurement with the same set of tags and timestamp, ie. it corresponds to a point in InfluxDB protocol. Each document contains the name of the measurement (e.g. "cpu" or "redis"), the set of tags as an array of objects (e.g. { "key": "hostname", "val": "host_42" }), the set of fields for given measurement (e.g. { "usage_user": 65, "usage_system": 7, ... }) as a subdocument, and the timestamp in nanoseconds. All points are stored in one collection. This design is maximally flexible at query time, allowing ad-hoc analysis similar to that of InfluxDB, but has the tradeoff of causing a large amount of disk space to be used. Also, in absolute terms, write throughput is low with higher cardinality.

After experimentation, we chose to use a single compound index, which has the following form: `[measurement, tags, timestamp_ns]`.

Notably, although MongoDB supports so-called "covered queries" to reduce disk seeks, they do not apply in the time series case because tag data requires a multikey index.

We chose one configuration in the space of possible MongoDB configurations. Much thought was put into this design, and we strove to make MongoDB perform the best way we knew how.

In contrast, InfluxDB required zero schema design.

Compression and disk usage

MongoDB uses a configurable compressor for raw collection data. In particular, the WiredTiger engine supports Snappy, which we used. Indexes are compressed with prefix compression. We checked our index statistics to verify that compression was being used. Unfortunately, these generic compression methods were not enough to mitigate the disk space needed by MongoDB. In our benchmarks, MongoDB required over an order of magnitude more storage space than InfluxDB. In contrast, InfluxDB uses compression algorithms that are tailored for time series data. This results in less disk space usage.

Ad-hoc query composition

MongoDB requires the user to construct short "aggregation" programs to analyze data. These are JSON documents that specify hybrid imperative/declarative computation over the data in a collection. For example, here is an example aggregation query for MongoDB:

```
db.point_data.aggregate(
  [
    {
      $match: {
        measurement: "cpu",
        timestamp_ns: { $gte: 1451607326000000000, $lt:
1451610926000000000},
        tags: {$in: [{key: "hostname", val: "host_42"}]},
      },
    },
    {
      $project: {
        _id: 0,
        time_bucket: { $subtract: ["$timestamp_ns", { $mod:
["$timestamp_ns", 60e9] } ] },
        fields: { $filter: { cond: { $eq: [ "$$field.key",
"usage_user" ] }, input: "$fields", as: "field" } },
        measurement: 1
      }
    },
    {
```

```

    $group : {
      _id : {time_bucket: "$time_bucket", tags: "$tags"},
      max_value: { $max: "$fields.val" }
    }
  },
  {
    $sort : { '_id.time_bucket': 1 }
  }
]
)

```

In contrast, the equivalent InfluxQL query is:

```

SELECT max(usage_user) from cpu where hostname = 'host_42' and time >=
'2016-01-01T00:15:26Z' and time < '2016-01-01T01:15:26Z' group by
time(1m)

```

Clearly, InfluxQL is more succinct. The reason for this is that InfluxQL is specifically designed for ad-hoc analysis of time series data.

Go language support

As of this writing, the canonical MongoDB driver for the Go language is `mgo`. Although easy to use, it imposes a mandatory serialization overhead on the user. In particular, it forces all queries to be serialized to BSON at query time, instead of ahead of time. This prevented us from squeezing out as much performance as we would have liked from the query benchmarking software. In contrast, InfluxDB uses a straightforward HTTP API that is amenable to ahead-of-time generation. Also, InfluxDB has a bulk protocol that allows clients to perform zero heap allocations on the write path when using a standard HTTP client library.

User experience conclusion

In summary, MongoDB is a general purpose document-oriented database, and InfluxDB is a special-purpose time series database. As a result, InfluxDB is orders of magnitude more convenient to use when storing and analyzing time series data.

About InfluxData

InfluxData is the creator of InfluxDB, the open source time series database. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by IoT devices, applications, networks, containers and computers. We are on a mission to help developers and organizations, such as Cisco, IBM, PayPal, and Tesla, store and analyze real-time data, empowering them to build transformative monitoring, analytics, and IoT applications quicker and to scale. InfluxData is headquartered in San Francisco with a workforce distributed throughout the U.S. and across Europe.

[Learn more.](#)

InfluxDB documentation, downloads & guides

[Download InfluxDB](#)

[Get documentation](#)

[Tutorials](#)

[Join the InfluxDB community](#)

What is time series data?

Time series data is nothing more than a sequence of values, typically consisting of successive measurements made from the same source over a time interval. Put another way, if you were to plot your values on a graph, one of your axes would always be time.

Time Series Databases are optimized for the collection, storage, retrieval and processing of time series data; nothing more, nothing less. Compare this to document databases optimized for storing JSON documents, search databases optimized for full-text searches or traditional relational databases optimized for the tabular storage of related data in rows and columns.

What is a time series database?

[Baron Schwartz](#) has outlined some of the typical characteristics of a purpose-built Time Series Database. These include:

- 90+% of the database's workload is a high volume of high-frequency writes

- Writes are typically appends to existing measurements over time
- These writes are typically done in a sequential order, for example: every second or every minute
- If a time series database gets constrained for resources, it is typically because it is I/O bound
- Updates to correct or modify individual values already written are rare
- Deleting data is almost always done across large time ranges (days, months or years), rarely if ever to a specific point
- Queries issued to the database are typically sequential per-series, in some form of sort order with perhaps a time-based operator or function applied
- Issuing queries that perform concurrent reads or reads of multiple series are common



799 Market Street
San Francisco, CA 94103
(415) 295-1901
www.InfluxData.com
Twitter: [@InfluxDB](https://twitter.com/InfluxDB)
Facebook: [@InfluxDB](https://facebook.com/InfluxDB)