

# Software Engineering Intern, 2021

## Timescale

As you might have read, TimescaleDB is a time-series database, which is often used in application domains like IoT (the Internet of Things), where thousands or hundreds of thousands of sensors may be sending a deluge of data.

Your task is to make TimescaleDB gobble up data fast!

It's a fairly straight-forward task:

1. Download and install TimescaleDB from <http://docs.timescale.com/latest/getting-started/installation>
2. Create a table with as simple schema: A time value (stored a timestamp without timezone information), and an integer value (which can always be sent as "0").
3. Insert data into TimescaleDB.

For this puzzle, we'd like you to try to try to maximize the ingest rate of TimescaleDB, i.e., the number of writes per second that you can achieve, when ingesting **10 million rows**. For your answer, please submit the following:

1. A write-up of your approach, i.e., what you did, and the reasoning behind this approach and why you think it maximized ingest performance.

2. Any code you had to write code to achieve this (either github repo, single file, or a zip/tar ball)
3. The performance you achieved. Please submit this as a graph showing the ingest rate per second over the lifetime of your 10M-row experiment.

For a bonus, ingest 250 million rows of data into Timescale. Add a brief discussion of any additional steps you took during this experiment, as well as the resulting graph.

As background reading, some of the following resources may help:

- <https://blog.timescale.com/time-series-data-why-and-how-to-use-a-relational-database-instead-of-nosql-d0cd6975e87c>
- <http://docs.timescale.com/latest/getting-started/setup>
- [http://docs.timescale.com/latest/api/api-timescaledb#create\\_hypertable](http://docs.timescale.com/latest/api/api-timescaledb#create_hypertable)