

The Internet-of-Things (IoT) is exploding: It is estimated that by 2020 there will be six devices online for every person on earth (source: Cisco IBSG). IoT devices generate data in a unique way. They post readings frequently—sometimes even multiple times in a second—and each of these readings has an associated timestamp. This never-ending stream of timestamped data is commonly referred to as a time-series. With more and more IoT devices being deployed, efficiently storing and querying time-series data has become a pressing problem.

Our project is based on two simple observations:

- 1.) While there are several databases optimized for storing relational data (Oracle), array data (TileDB), and key-value data (MongoDB), there are few databases optimized for storing time-series data.
- 2.) The storage of time-series data can be organized around a simple but useful heuristic: “New data is king.” In other words, most queries on time-series data (such as analytics and problem-diagnosis queries) rely more heavily on new data than old data. Therefore, a time-series storage manager should ensure fast access to the newest data.

Metronome is a database storage manager optimized for time series, with support for windowing¹, and providing fast access to the new data.

¹ Queries on time-series take a common form: You specify the predicate conditions and the range of times to consider. In other words, you specify the start timestamp and the end timestamp, and this range is referred to as a window. You always need to specify the window in a time-series query.