



Next generation data storage system to support big data, IoT and machine learning at the Norwegian Meteorological Institute

Isak Buhl-Mortensen, Thomas N. Nipen, Ivar A. Seierstad, and Cristian Lussana
Norwegian Meteorological Institute, Oslo, Norway

The Norwegian meteorological institute (MET Norway) routinely collects and archives in-situ observations measured by conventional weather stations following the WMO standard. However, it is apparent that non-conventional observations, those shared by private companies and citizens, cannot be ignored. Over the last couple of years, the number of such observations has constantly risen. From the point of view of a national meteorological service, this data comes with a number of issues, such as insufficient metadata and a total lack of control on both the measurement practices applied and the instrumentation used. On the other hand, this large volume of non-conventional data (up to hundreds of observations per km square per minute) allows for the near-surface atmospheric state to be observed at an unprecedented level of detail thus opening new possibilities for disaster risk reduction and research in atmospheric sciences.

Redundancy is the key factor that helps transform otherwise unreliable data into usable data for national meteorological services. MET Norway has recently improved the temperature forecasts on Yr.no by introducing amateur station data into the processing chain. Yr.no has millions of users per week so this improvement is beneficial for a large community. This has required a tailored system based on two aspects: (1) distributed storage and (2) data quality control.

We present our plans for a distributed database for mass storage and analysis of in-situ data. This storage backend will lay the foundation for products based on Big data, IoT and machine learning. To match a constant increase in data load, it becomes necessary to scale out and embrace the nature of distributed systems - a constant compromise between performance (availability) and information consistency. We favor availability at the expense of eventual consistency (milliseconds). For transactions that require higher consistency, distributed database management systems like Cassandra (C*) allow clients to specify the level of consistency. C* also supports ordered columns and a time-window compaction strategy, making C* performant for time series data. In terms of redundancy, a C* cluster employs leaderless replication and therefore has no single point of failure. This makes C* popular - it is the main technology behind Netflix's time series storage solution of customer viewing history.

C* may sound perfect for time series data, however there is a need for other access patterns, and unlike a relational database this comes at a cost, denormalization - SQL-like operations such as joins are not supported. To overcome these constraints, we denormalize the data into gridded, time series, and point cloud representations. For the point cloud representation we are currently testing a model that distributes data across the cluster via geohash, and allows for data selection within a geofence.