



[linkedin.com/in/simonaubury](https://linkedin.com/in/simonaubury)



[github.com/saubury/race-mapper](https://github.com/saubury/race-mapper)

# Did I beat Ben?

## Race tracking with Kafka KSQL, MQTT & Kibana





# Hello!

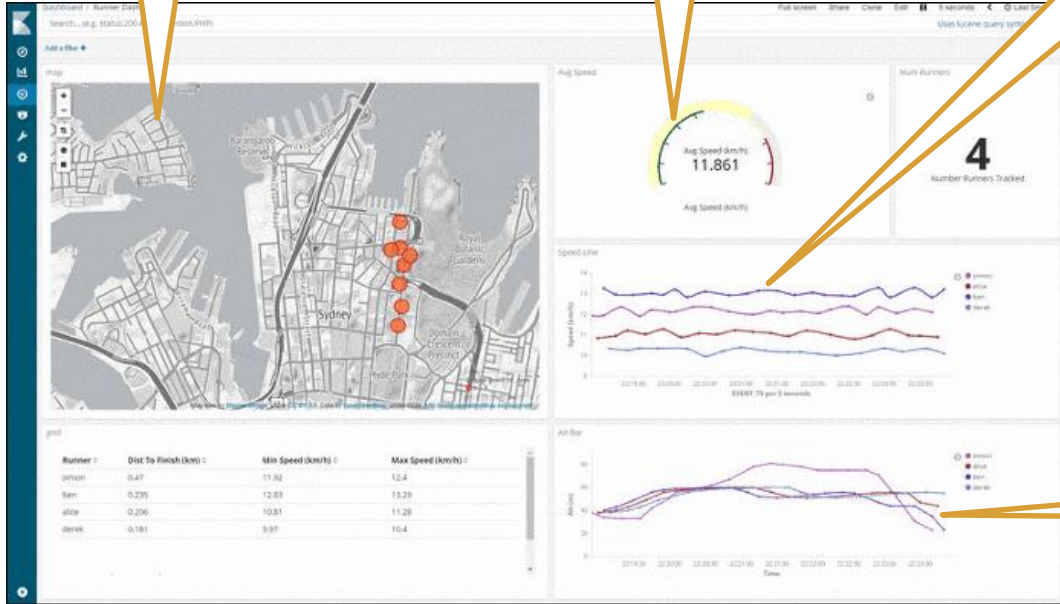
*I am* **Simon Aubury**

I am here because I love streaming & running

Map

Avg Speed

Runner pace




**Finish line ... end result**  
This is what we're building

Elevation




## The parts





## **OwnTracks** open-source phone app




- OwnTracks is an open-source app
  - iOS and Android - <https://owntracks.org>
- Sends latitude, longitude, altitude to an MQTT server.
- Phone location is a great proxy for runner location!





## What is MQTT




- MQTT is lightweight TCP/IP protocol
  - Small footprint
  - Low power
- MQTT acts more like a key/value store
  - whereas Kafka is a complete streaming platform.





## Hosted MQTT - CloudMQTT

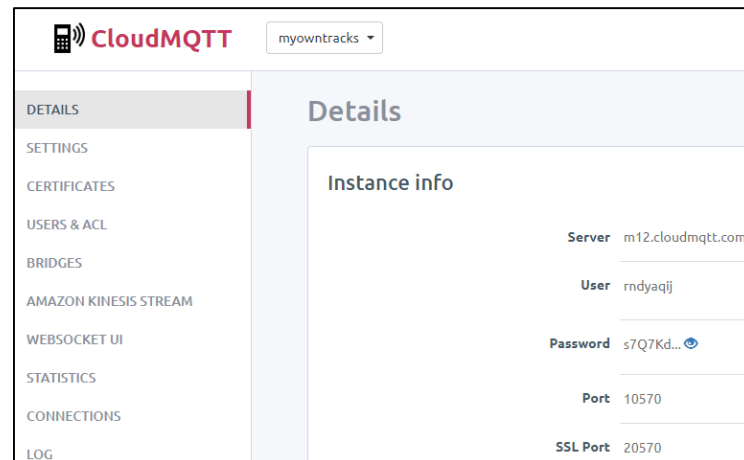


### CloudMQTT

- [www.cloudmqtt.com](http://www.cloudmqtt.com)
- Free tier just fine


### Hosted MQTT server

- Always up
- Exposed end-points

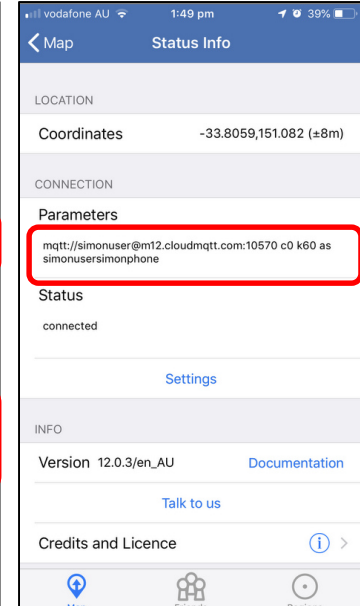
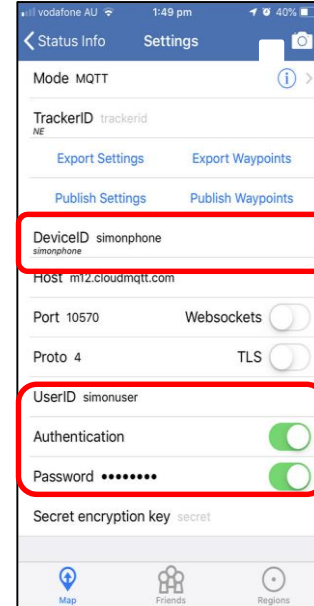




## Setup OpenTracks




- Config Opentracks with CloudMQTT endpoint
- One MQTT Server
- Each runner has own logon
- Each device has own ID








## Kafka connect




- Kafka Connect is a framework for streaming data between Apache Kafka and other data systems
  - Kafka connect **sources** – from something into Kafka
  - Kafka connect **sinks** – from Kafka into something






## MQTT to Kafka with **Kafka Connect**



```
{  
  "connector.class" : "io.confluent.connect.mqtt.MqttSourceConnector",  
  "mqtt.server.uri" : "tcp://m1234.cloudmqtt.com:10570",  
  "mqtt.password" : "SomePassword",  
  "mqtt.username" : "kafkauser",  
  "mqtt.topics" : "owntracks/#",  
  "kafka.topic" : "data_mqtt",  
  "key.converter": "org.apache.kafka.connect.storage.StringConverter",  
  "value.converter":  
    "org.apache.kafka.connect.converters.ByteArrayConverter",  
  "tasks.max" : "1",  
  "confluent.topic.bootstrap.servers" : "kafka:29092",  
  "confluent.topic.replication.factor" : "1"  
}
```



## Do we have data in topic?




```
kafka-console-consumer \  
--bootstrap-server localhost:9092 \  
--property print.key=true --property key.separator=" : " \  
--topic data_mqtt
```


```
owntracks/simonuser/simonphone : {"lon":151.061, "lat":-33.812,  
"batt":35,"acc":8,"p":100.546,"vel":53,"vac":3,"conn":"m","tst":1559019380,"  
tid":"NE","_type":"location","alt":45}
```



## KSQL **create stream**



```
CREATE STREAM realtime_runner_location_stream  
  (lon DOUBLE, lat DOUBLE, alt INTEGER)  
WITH (kafka_topic = 'data_mqtt', value_format='JSON');
```




## KSQL parse location

```
CREATE STREAM runner_location with (value_format='JSON')
AS
select split(rowkey, '/') [2] as who
, tst as event time
, cast(lat as varchar) || ',' || cast(lon as varchar) as LOCATION
, alt
from realtime_runner_location_stream;
```




## KSQL league table



```
CREATE TABLE runner_status with (value_format='JSON')
AS
select split(rowkey, '/') [2] as who
, min(vel) as min_speed
, max(vel) as max_speed
, min(GEO_DISTANCE(lat, lon, -33.87014, 151.211945, 'km')) as dist_to_finish
, count(*) as num_events
from realtime_runner_location_stream WINDOW TUMBLING (size 5 minute)
group by split(rowkey, '/') [2]
```




## Kafka To Elastic & Kibana



- Now I have 2 derived topics
  - runner-location — a stream of runner's and their location
  - runner-status — a table of runner's speed and remaining distance
- I want to send these topics into Kibana



## Kibana using the Elastic Kafka Connect Sink Connector

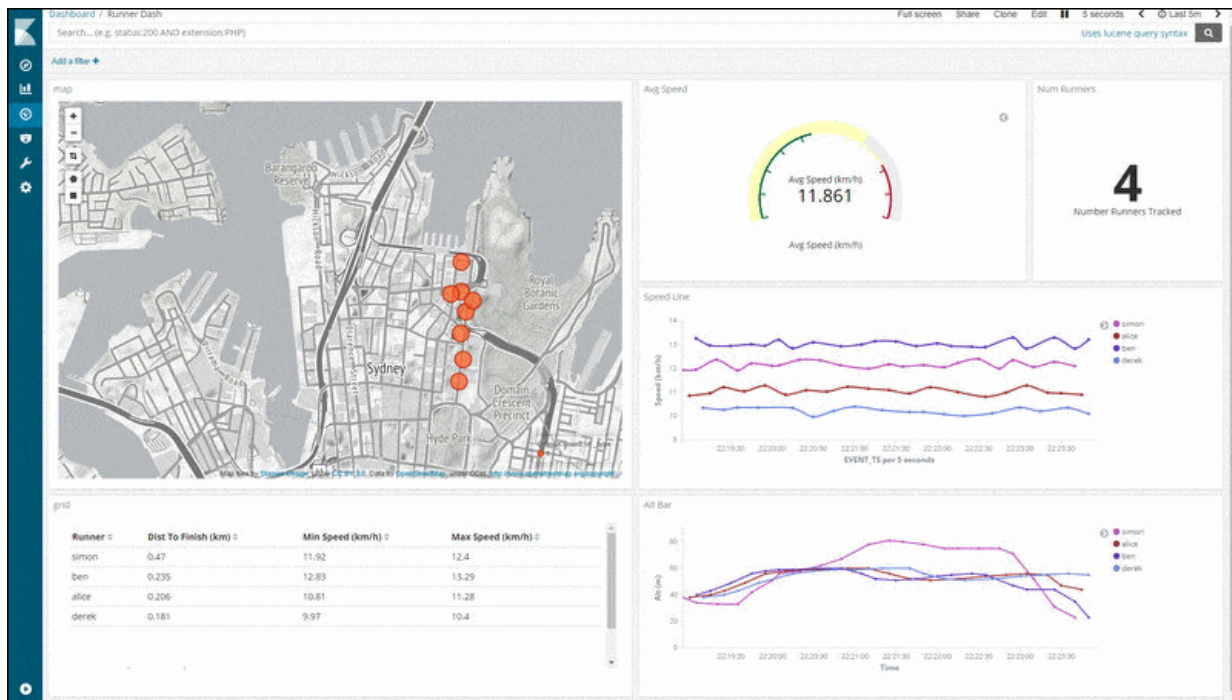


```
{
  "connector.class": "io.confluent.connect.elasticsearch.ElasticsearchSinkConnector",
  "type.name": "type.name=kafkaconnect",
  "topics": "RUNNER STATUS",
  "transforms": "ExtractTimestamp",
  "key.ignore": "true",
  "schema.ignore": "true",
  "transforms.ExtractTimestamp.type": "org.apache.kafka.connect.transforms.InsertField$Value",
  "topic.index.map": "RUNNER_STATUS:runner_status",
  "value.converter.schemas.enable": "false",
  "value.converter": "org.apache.kafka.connect.json.JsonConverter",
  "connection.url": "http://elasticsearch:9200",
  "key.converter": "org.apache.kafka.connect.storage.StringConverter",
  "transforms.ExtractTimestamp.timestamp.field": "EVENT_TS"
}
```





# Final Kibana dashboard





**Did I Beat Ben?**

*no ...*





But I did beat my sister!



# Thanks!

Any **questions** ?



[linkedin.com/in/simonaubury](https://www.linkedin.com/in/simonaubury)



[github.com/saubury/race-mapper](https://github.com/saubury/race-mapper)

*Presentation template by SlidesCarnival*