

# IoT

# Cassandra Data Modeling

November 2016

# IoT Energy





# Context

- Thousand industrial sites in US, Canada, Europe
- Thousand buildings and super markets
- Hundred thousand users
- Million captors
  - Data semi-structured or Un-structured
  - At least one data event each 2 seconds each captor
  - 50K records a day each captor
  - Billion records a day in total
- Real time monitoring
- Multi dimensions reporting
  - Daily
  - Monthly
  - Yearly
  - ...
- Data analytics



# Cassandra

- Lambda-architecture

- US Metro - SMACK Stack (Spark, Mesos, Akka, Cassandra and Kafka)
- Lambda-architecture for Twitter Sentiment Analysis
- Amazon Web Services Lambda

- Cassandra

- Best Practice of One Table Per Query
- De-normalized data models.
- Primary key with a hashing algorithm
  - Partitioning
  - Clustering
- Scalability
- Performance

Any questions?

