DigBigData

Apache Cassandra & Friends

About Me

- DigBigData Founder & CTO
- Working with Cassandra since 2010
- Apache Cassandra MVP for 2014
- Systems Engineer at heart

About this Presentation

- We began to see patterns in the problems we were asked to solve
- We wanted to work with a predictable technology stack
- Choice of technologies is based on our real world experience
- We are Cassandra evangelists but for good reason

DigBigData Ethos

- Choose the right tool for the right job
- Know as best you can what you are getting yourself into
- Research!
- Model!
- Simulate!
- Evaluate!

Apache Cassandra

- Our NoSQL database of choice
- Because :
  - Time series data!
  - High availability
  - Almost Linear Scalability
  - Reliable multi-DC data distribution without the headaches
  - No single point of failure
  - Very active development by a talented team

Apache Cassandra

- What you don't get:
  - Free text searching, faceting and indexing
  - Sort by value (for Top N)
  - Aggregate queries such as GROUP BY, AVG, MIN, MAX
  - Partial Row Caching
  - Analytics toolkit out of the box
- Not included for very good reasons!

Lambda Architecture

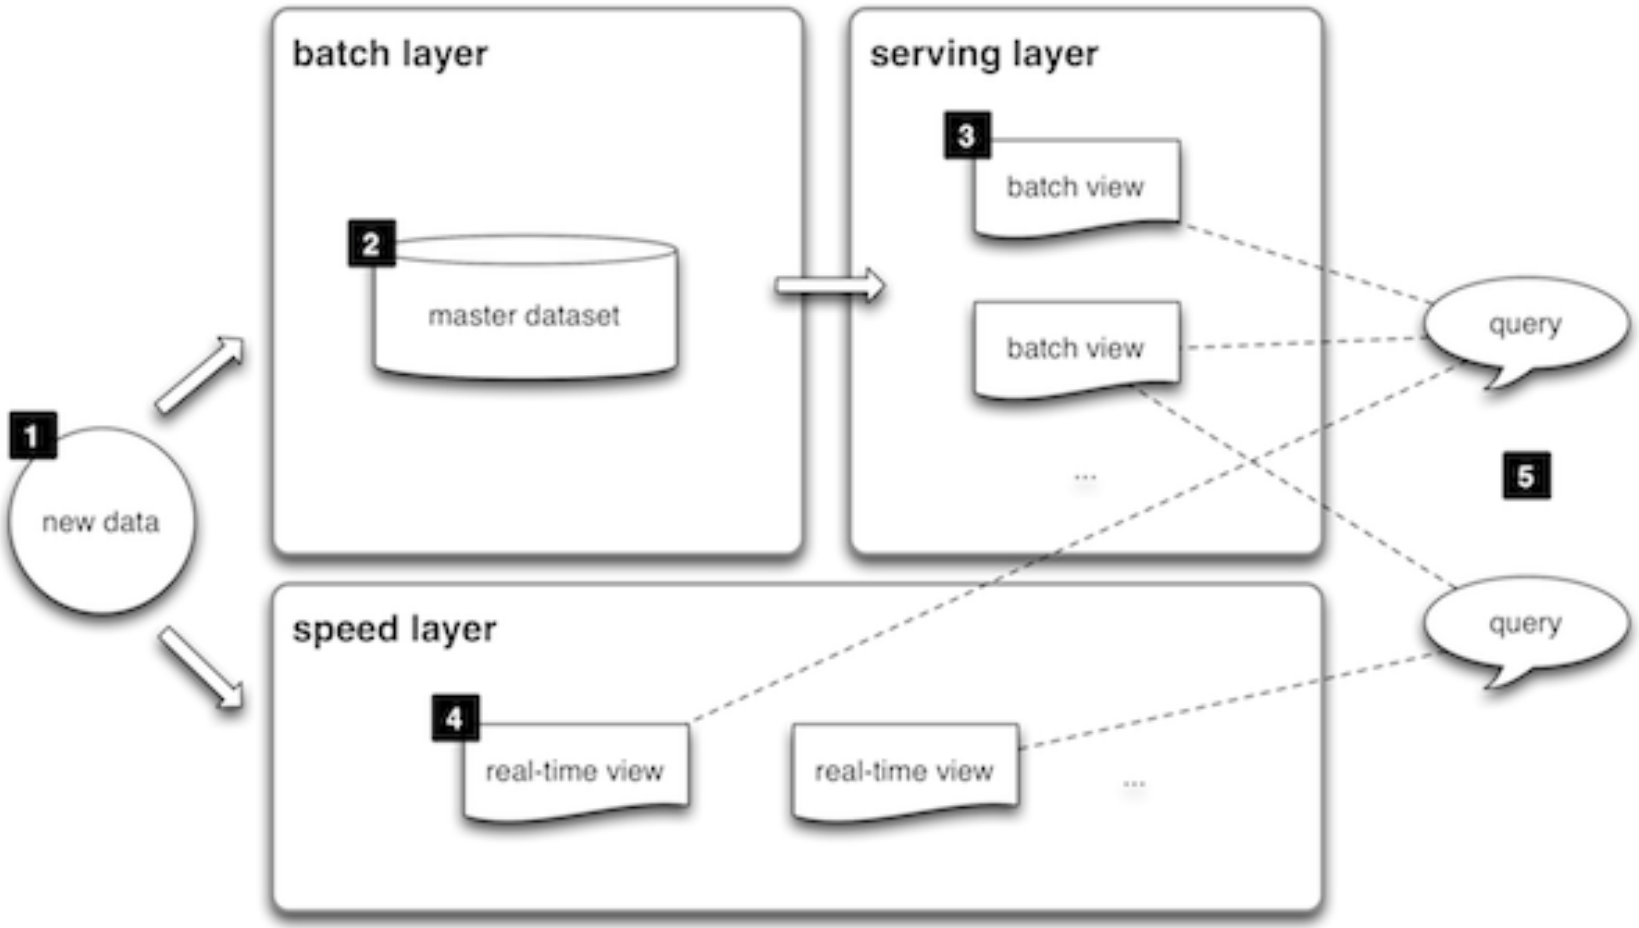

Responsiveness

- Today we want to make decisions in real-time or near real-time.
- Tomorrow isn't good enough.
- Smarter ways of batching and analyzing data sets are emerging to facilitate this (e.g. Spark, Impala, Drill etc.)

Lambda Revisited

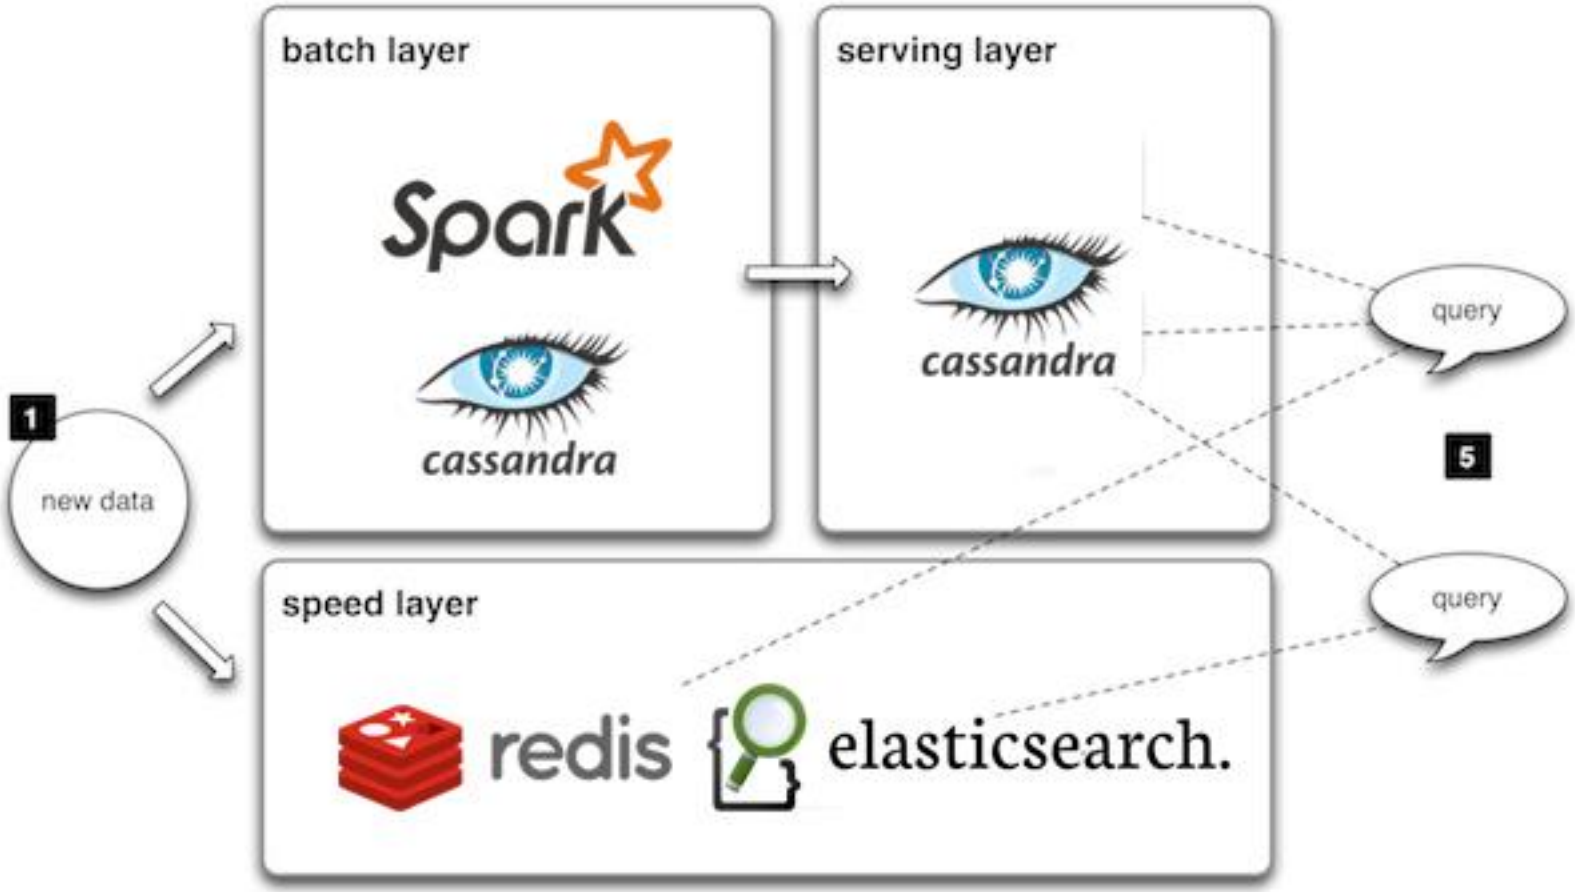

Every Layer is a Speed Layer!

- Cassandra for consolidated storage & source of truth
- Spark Streaming for real-time streaming calculations
- Spark / Shark for Analytics
- Redis for key / value multi-purpose cache and sort by value
- Elasticsearch for real-time search providing free text, multi-lingual search, faceting etc.

More on Spark

- Top Level Apache Project (v1.0.2 current version)
- 100x faster than Hadoop
- MapReduce in memory or 10x faster if on disk
- Possible to easily create RDDs from CQL queries
- Libraries available for easy interaction with Cassandra (DataStax Spark Driver / Calliope)

More on Spark

- Via MLlib we can execute complex machine learning algorithms such as:
  - **Basic Statistics**
  - Support Vector Machines
  - Logistic Regression
  - **Classification / Naïve Bayes**
  - **k-means clustering**
  - Recommendation via alternating least squares

More on Spark Streaming

- Spark Streaming support also included in the Cassandra driver
- Allows for windowed analytics based on intervals
- Can mix batch and streamed analytics, good for historical data + latest data operations

More on Shark

- With Shark we can achieve super fast in-memory queries on subsets of data in Cassandra
- Effectively all the features of Hive running on RDD not HDFS
- Uses HiveQL queries
- Mllib algorithms available in HiveQL also
- CqlStorageHandler provided to read RDD from Cassandra or read SSTables directly
- https://github.com/richardalow/cassowary

Other Possibilities

- Table joins over Cassandra tables
- Average, min, max, sorting, grouping, percentiles, map/reduce in near realtime
- Aggregations with write-back to Cassandra for time series
- Cassandra as “fact” store

More on REDIS

“Redis is an open source, BSD licensed, advanced key-value store. It is often referred to as a data structure server since keys can contain strings, hashes, lists, sets and sorted sets.”

More on REDIS

- Good for…
  - Sorting sets & lists (Can solve retaining Top N at a given time)
  - Pubsub messaging
  - (more) Accurate counters
  - Merging sets
  - Transactions!
- Works in memory, can serve data fast based on key
- Could use shared resources on Cassandra nodes (could populate most recent data via triggers API

More on ElasticSearch

- Distributed real-time search engine
- Built from the ground up for reliability and scalability
- Supports lots of other features as well free text search
  - Spatial
  - Query by arbitrary fields
  - Facets
- Multi-lingual query support
- Great operational tools included for index management

More on Elasticsearch

- Although external to Cassandra it can provide rich query capabilities over the same data
- Simplify Data Models in Cassandra to maximise storage
- Separate read and write workloads (read from ES, write to Cassandra)
- Some integration for **Storm** for writing records to elastic search and Cassandra as data enters the system

Conclusions

- Cassandra excellent as source of truth (redundant, performant & highly available)
- Can logically split analytics and operational clusters without adding operational complexity (using virtual data centers)
- Spark / Shark can operate on data loaded from Cassandra via CQL and persist data back
- Integration in early phase but will be comprehensive v. soon

Conclusions

- Redis great for sorted sets and retaining real-time Top N
- Redis also good for partial row caching
- Elasticsearch for free text search, faceting and indexing

Conclusions

- Lastly…

- It is all open source software!

Questions

- ??