cloudera

The Rise of Open Source Data Platforms: An Insider's View

Sept 30th 2017
DataCon.TW 2017
Jonathan Hsieh, Cloudera

Who Am I?

- **Tech Lead, Infrastructure @ Cloudera**
- **With Cloudera since 2009**
  - **Engineering Manager/Tech Lead** of Cloudera's HBase team
  - **Apache HBase** committer / PMC
  - **Apache Flume** founder / PMC
- **U of Washington:**
  - Research in Distributed Systems

Evolution

- Big Data
- Data Science
- Machine Learning
- What's Next

cloudera

cloudera
Ask Bigger Questions

The Rise of Open Source Big Data

Let's go back to 2006…

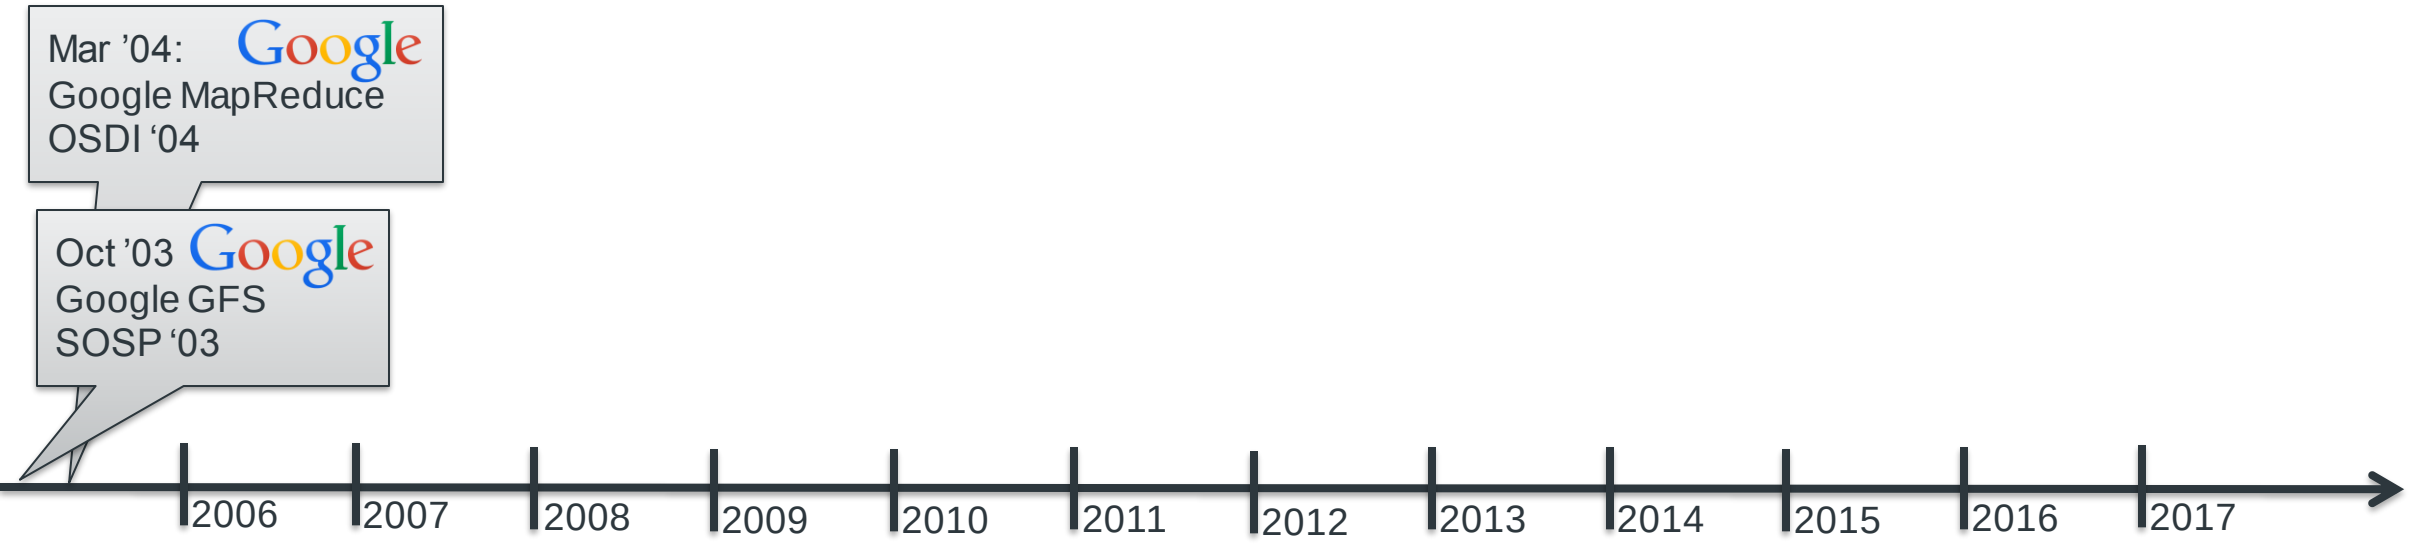

cloudera

Data Volumes were growing faster and faster

"Every two days we create as much information as we did from the dawn of civilization up until 2003."

-- Eric Schmidt
Google CEO, 2010

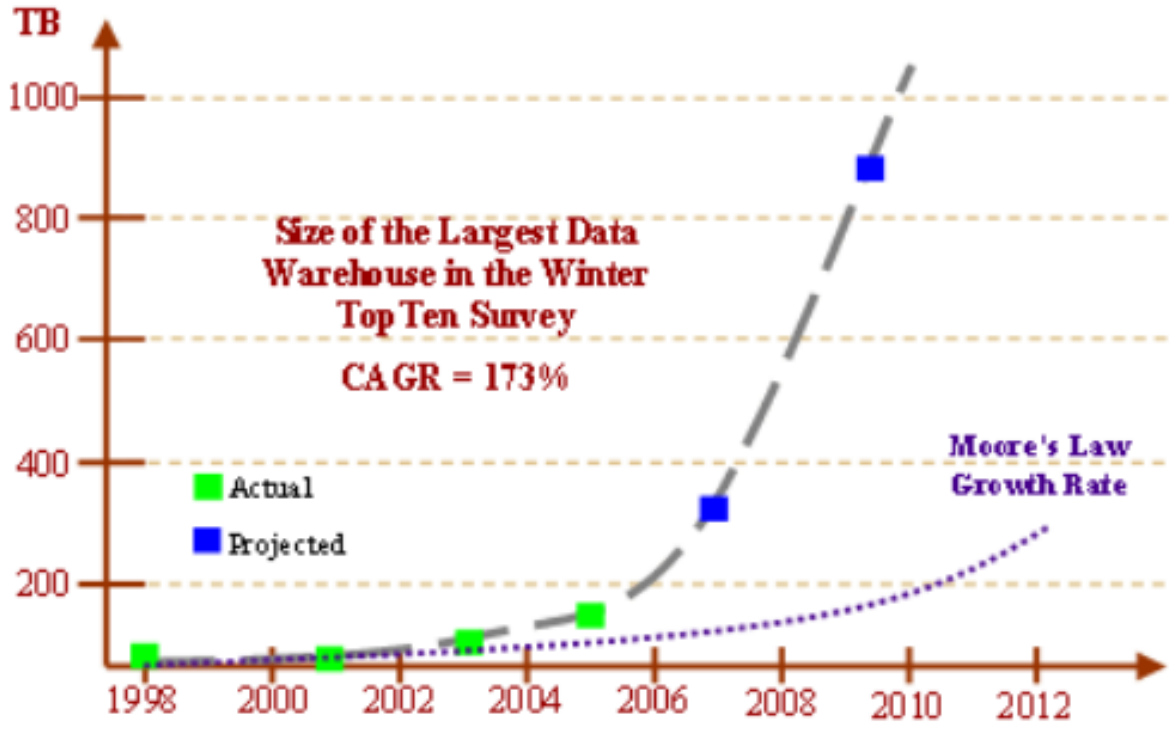

Source: Richard Winter, "Why are data warehouses growing so fast?" April 2008

cloudera

CPUs weren't getting any faster

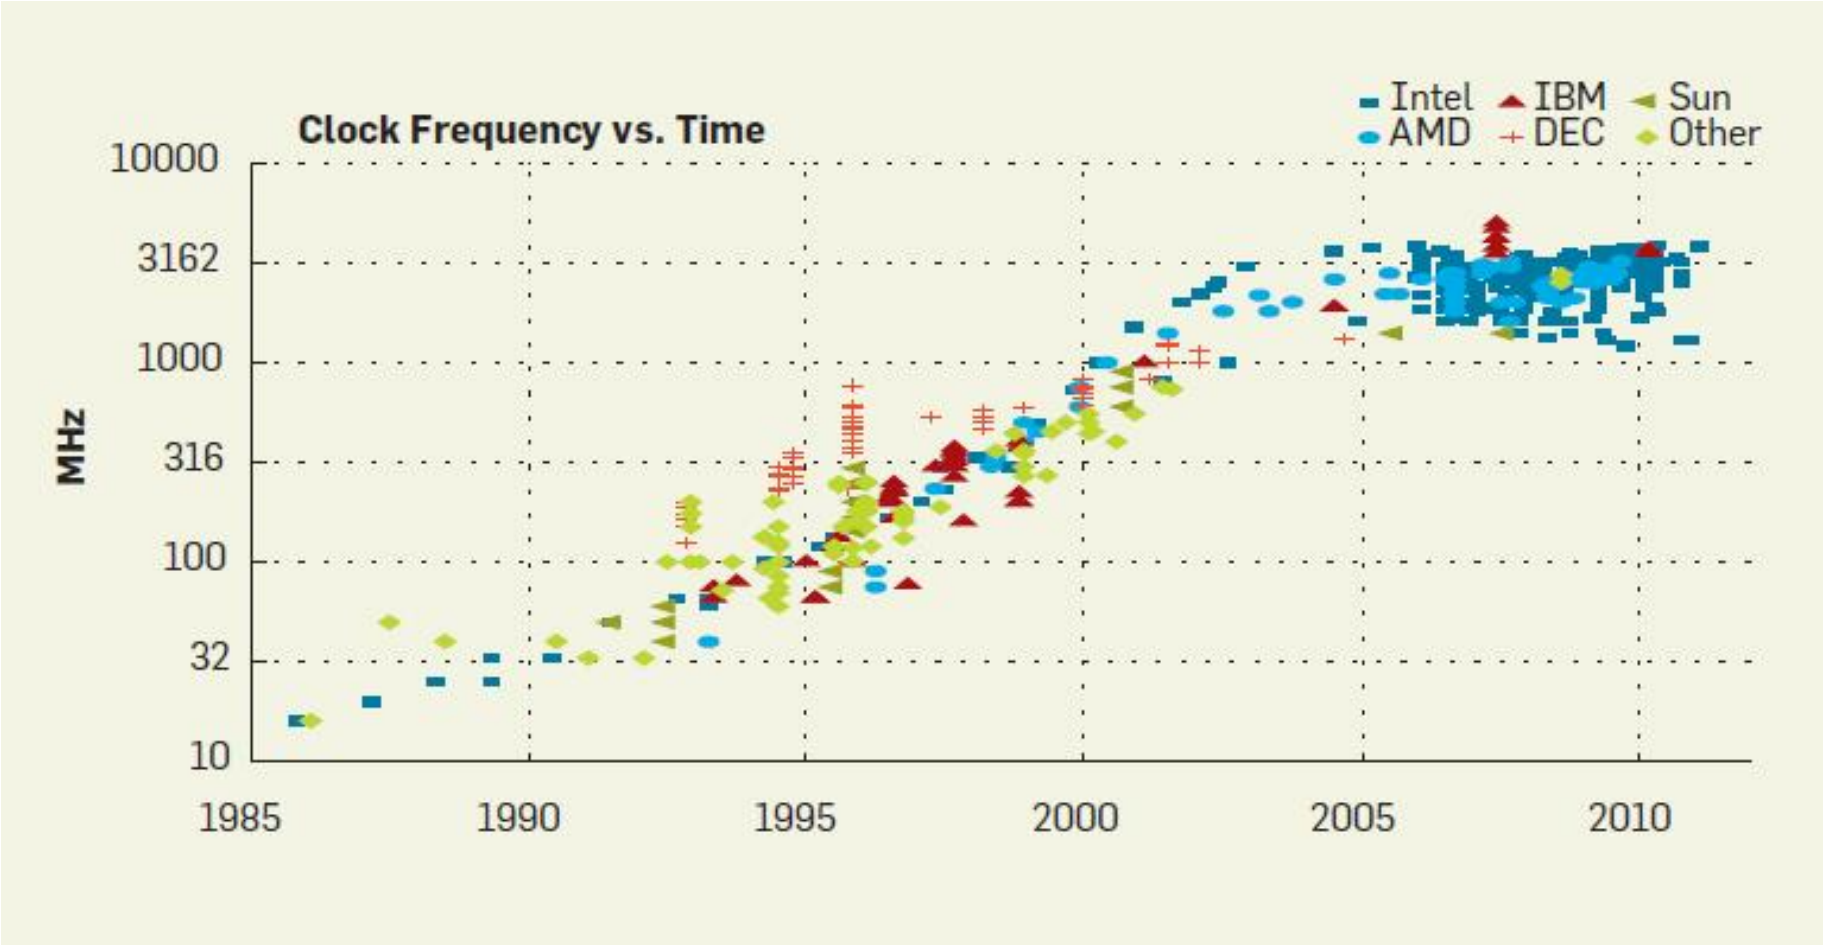

http://cacm.acm.org/magazines/2012/4/147359-cpu-db-recording-microprocessor-history/abstract

cloudera

Google's "Big Data" problem

- How to build an index of all the web's data?

- An algorithm:
  - Download the web (all of it)
  - Analyze it to build the index
    - Word count, page rank
    - Keep track of the links
  - Serve the index

cloudera

What is Big Data?

1. **Collect** the data
2. **Count** the data
3. **Report** the results

Approach: Use Distributed Systems

Instead of one expensive machine, we spread the work across many commodity machines

- **Horizontally Scalable**
  - More hardware means more capacity without diminishing returns
  - Deal with the **volume** of data

- **New Challenge: Reliability**
  - **Fault tolerant:** Tolerate failures, fail gracefully, have recovery mechanism
  - **Highly-Available:** Tolerate failures, recover immediately

cloudera

Distributed systems are hard!

Unicorns

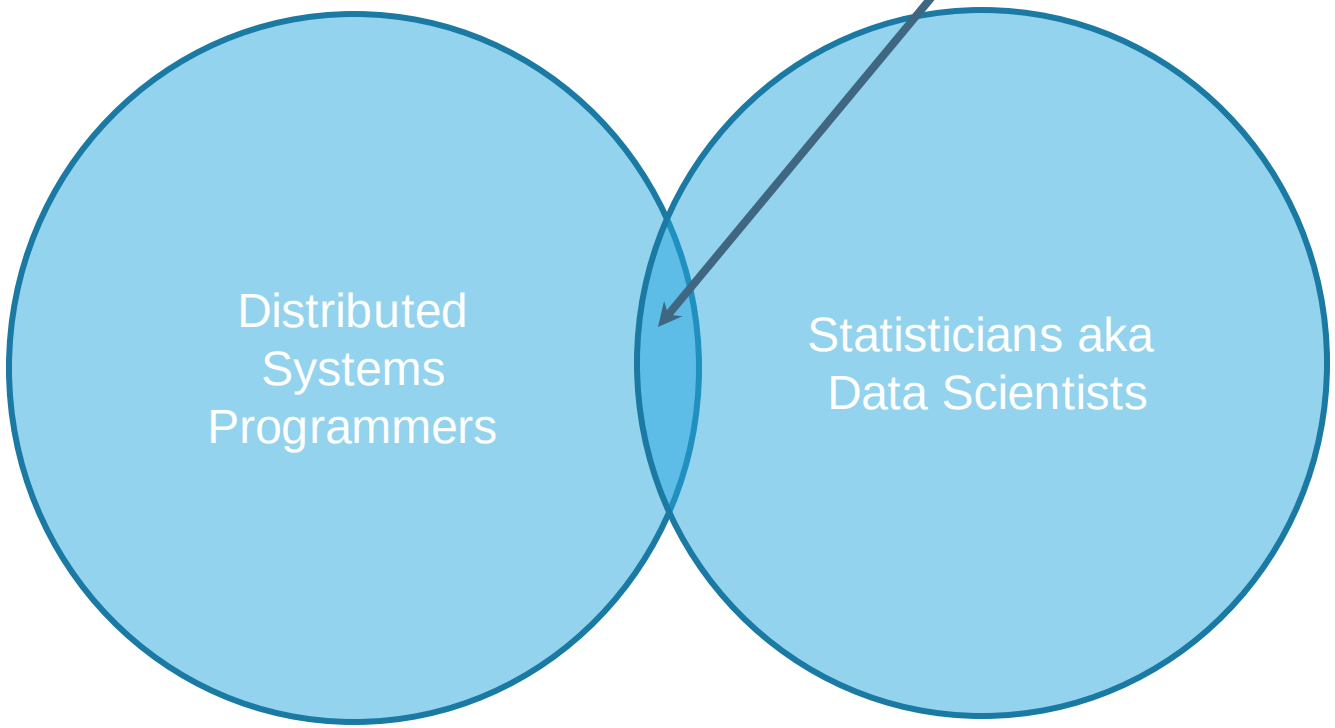

cloudera

Google built distributed storage and processing systems

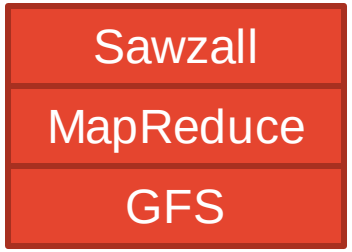

Storage:

- Google File System (GFS)
- Store PB's of data

Processing:

- MapReduce: Users specify a *map()* and a *reduce()* function
  - *map*: $K_1, V_1 \rightarrow list\ K_2, V_2$
  - *reduce*: K_2, $iter(V_2) \rightarrow list(K_3, V_3)$
- Sawzall: A DSL on top of MR

cloudera

Google built a Big Data Platform

Google

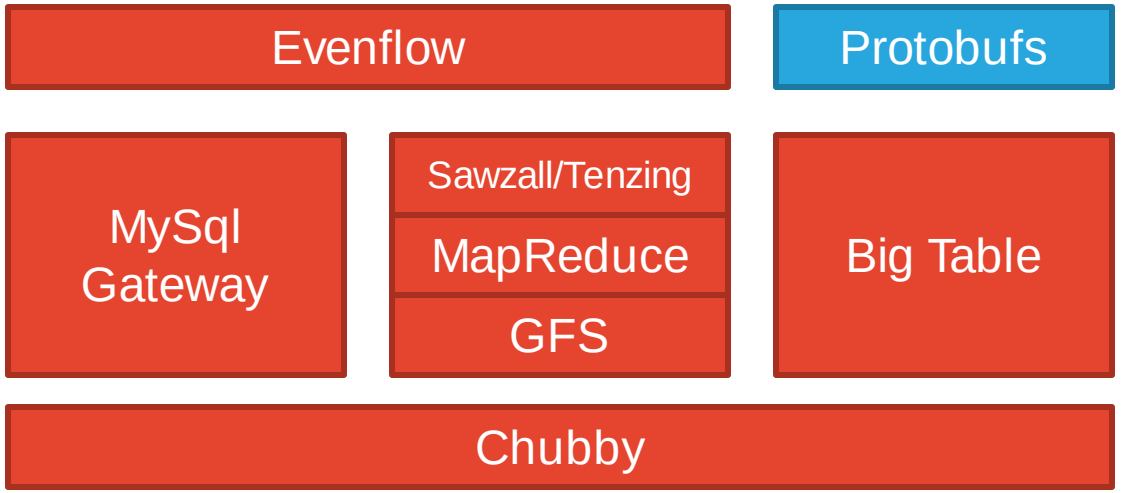

- Feeding data to GFS/MapReduce
- Serving data generated by it
- Coordinating the different systems
- Managing chains of jobs.
- Managing metadata

cloudera

Yahoo built a Big Data Platform

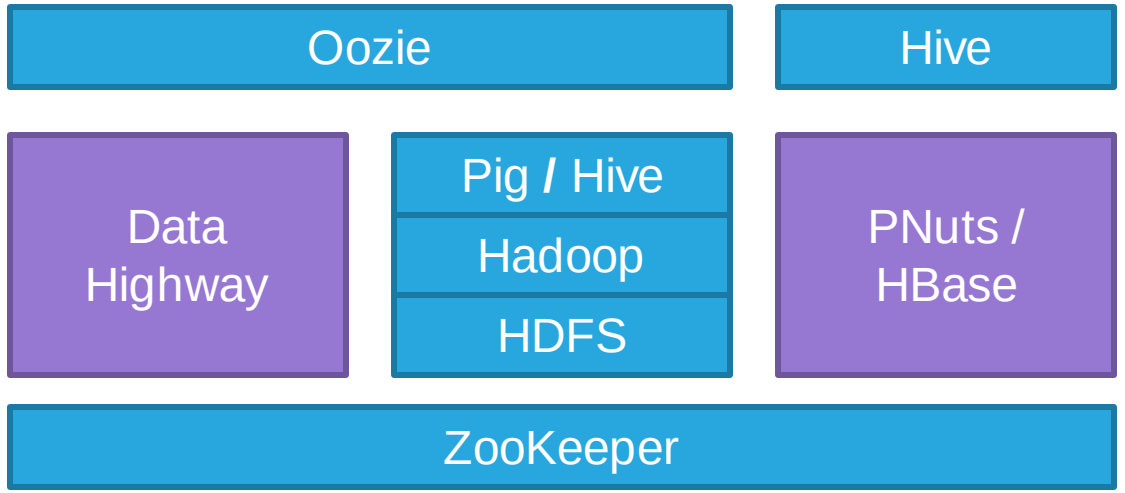

- Invested heavily in Hadoop (MapReduce clone) and HDFS (GFS clone)
- Created and open-sourced Oozie, Pig, and Zookeeper

Facebook built a Big Data Platform

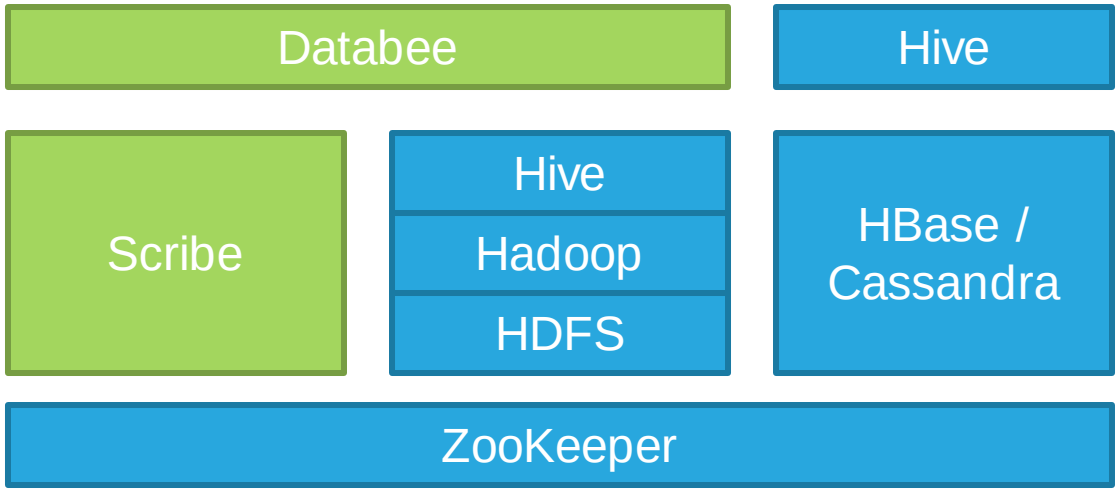

- Created and open-sourced Scribe, Hive, and Cassandra
- Invested heavily on HBase

Open Source Big Data Platform Timeline

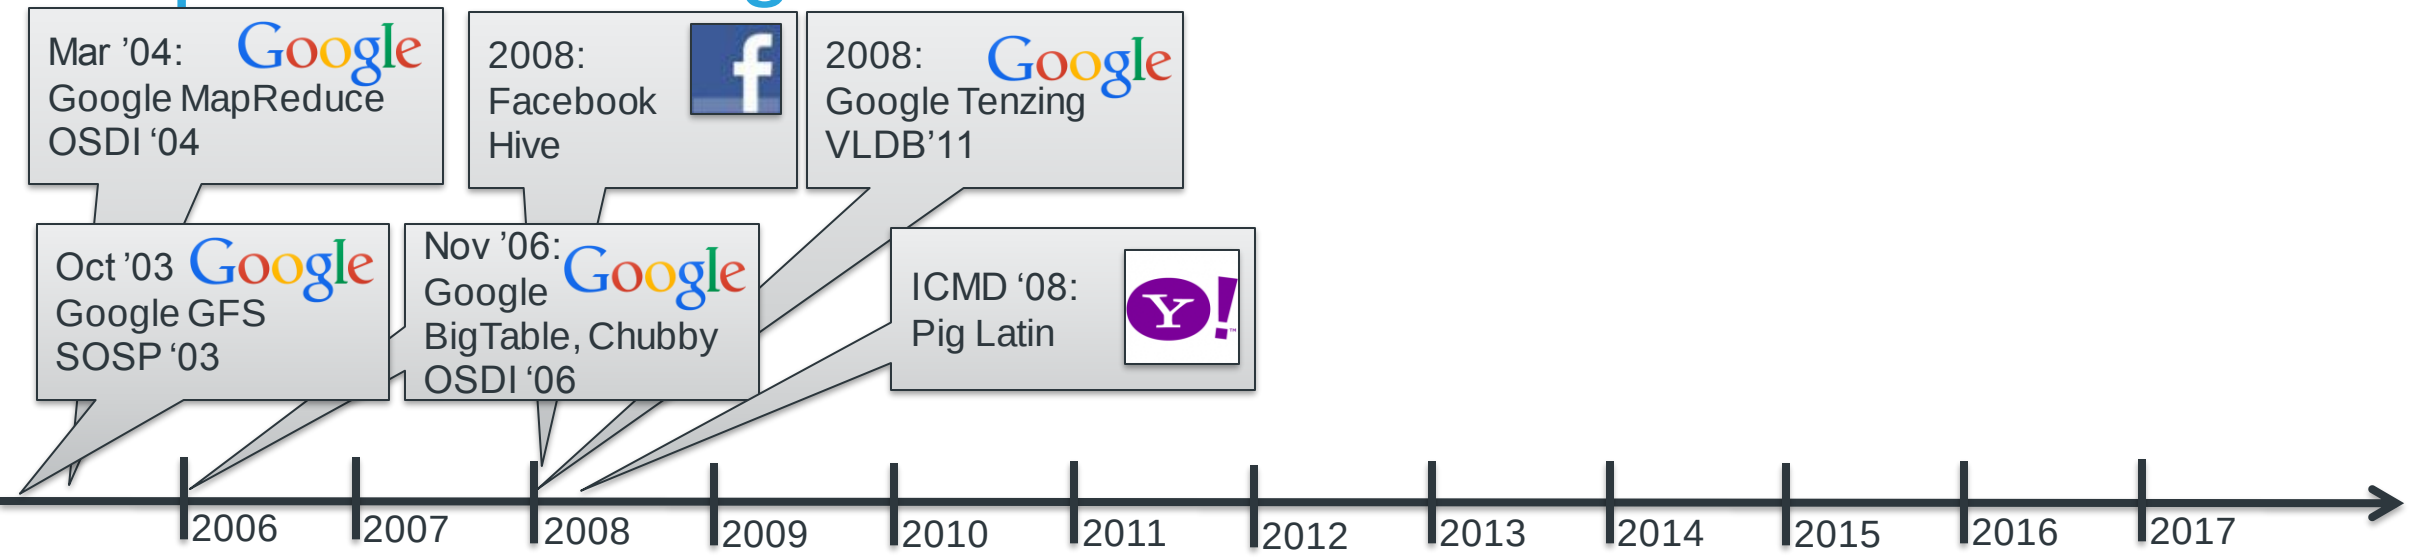

The future is already here — it's just not very evenly distributed.
-- William Gibson

cloudera

The Core Big Data Platform

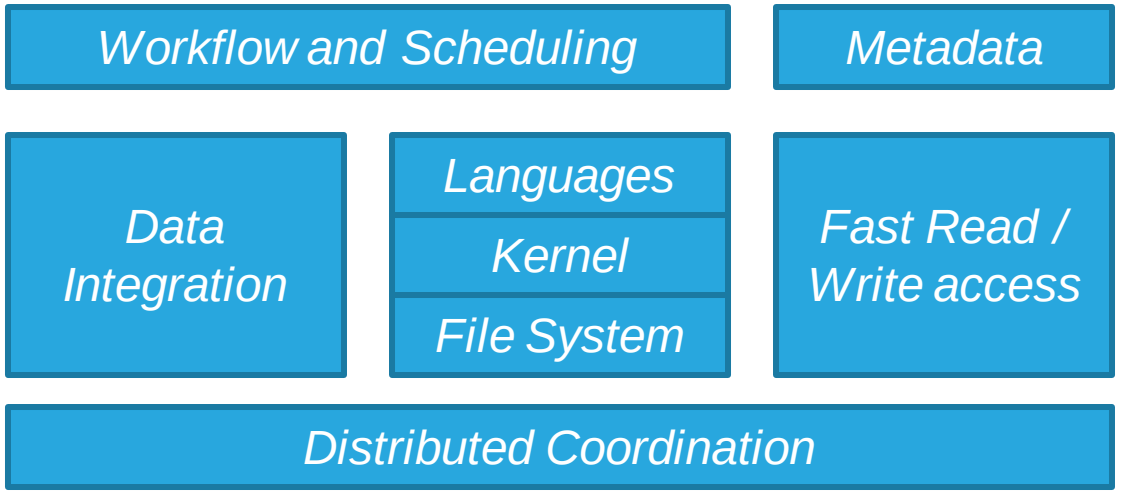

- Hadoop is an open MapReduce kernel
  - It is like the kernel of a Cluster Operating System.
- Each “friend” is a distributed service that has a similar workstation tool or abstraction.

cloudera

Cloudera built a Big Data Platform

cloudera

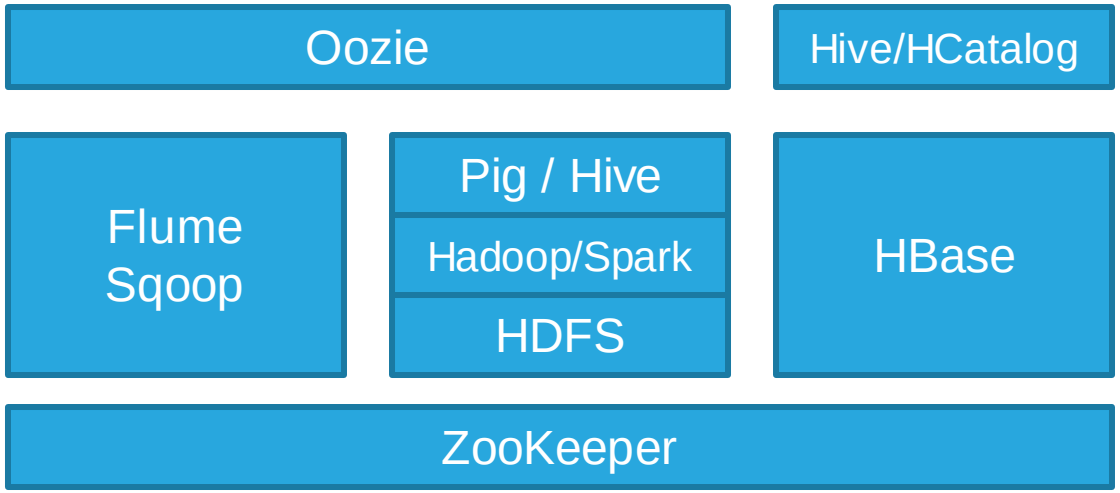

- Created and open-sourced Flume and Sqoop
- Core stack is free and available for all to use

Open Source Big Data Platform Timeline

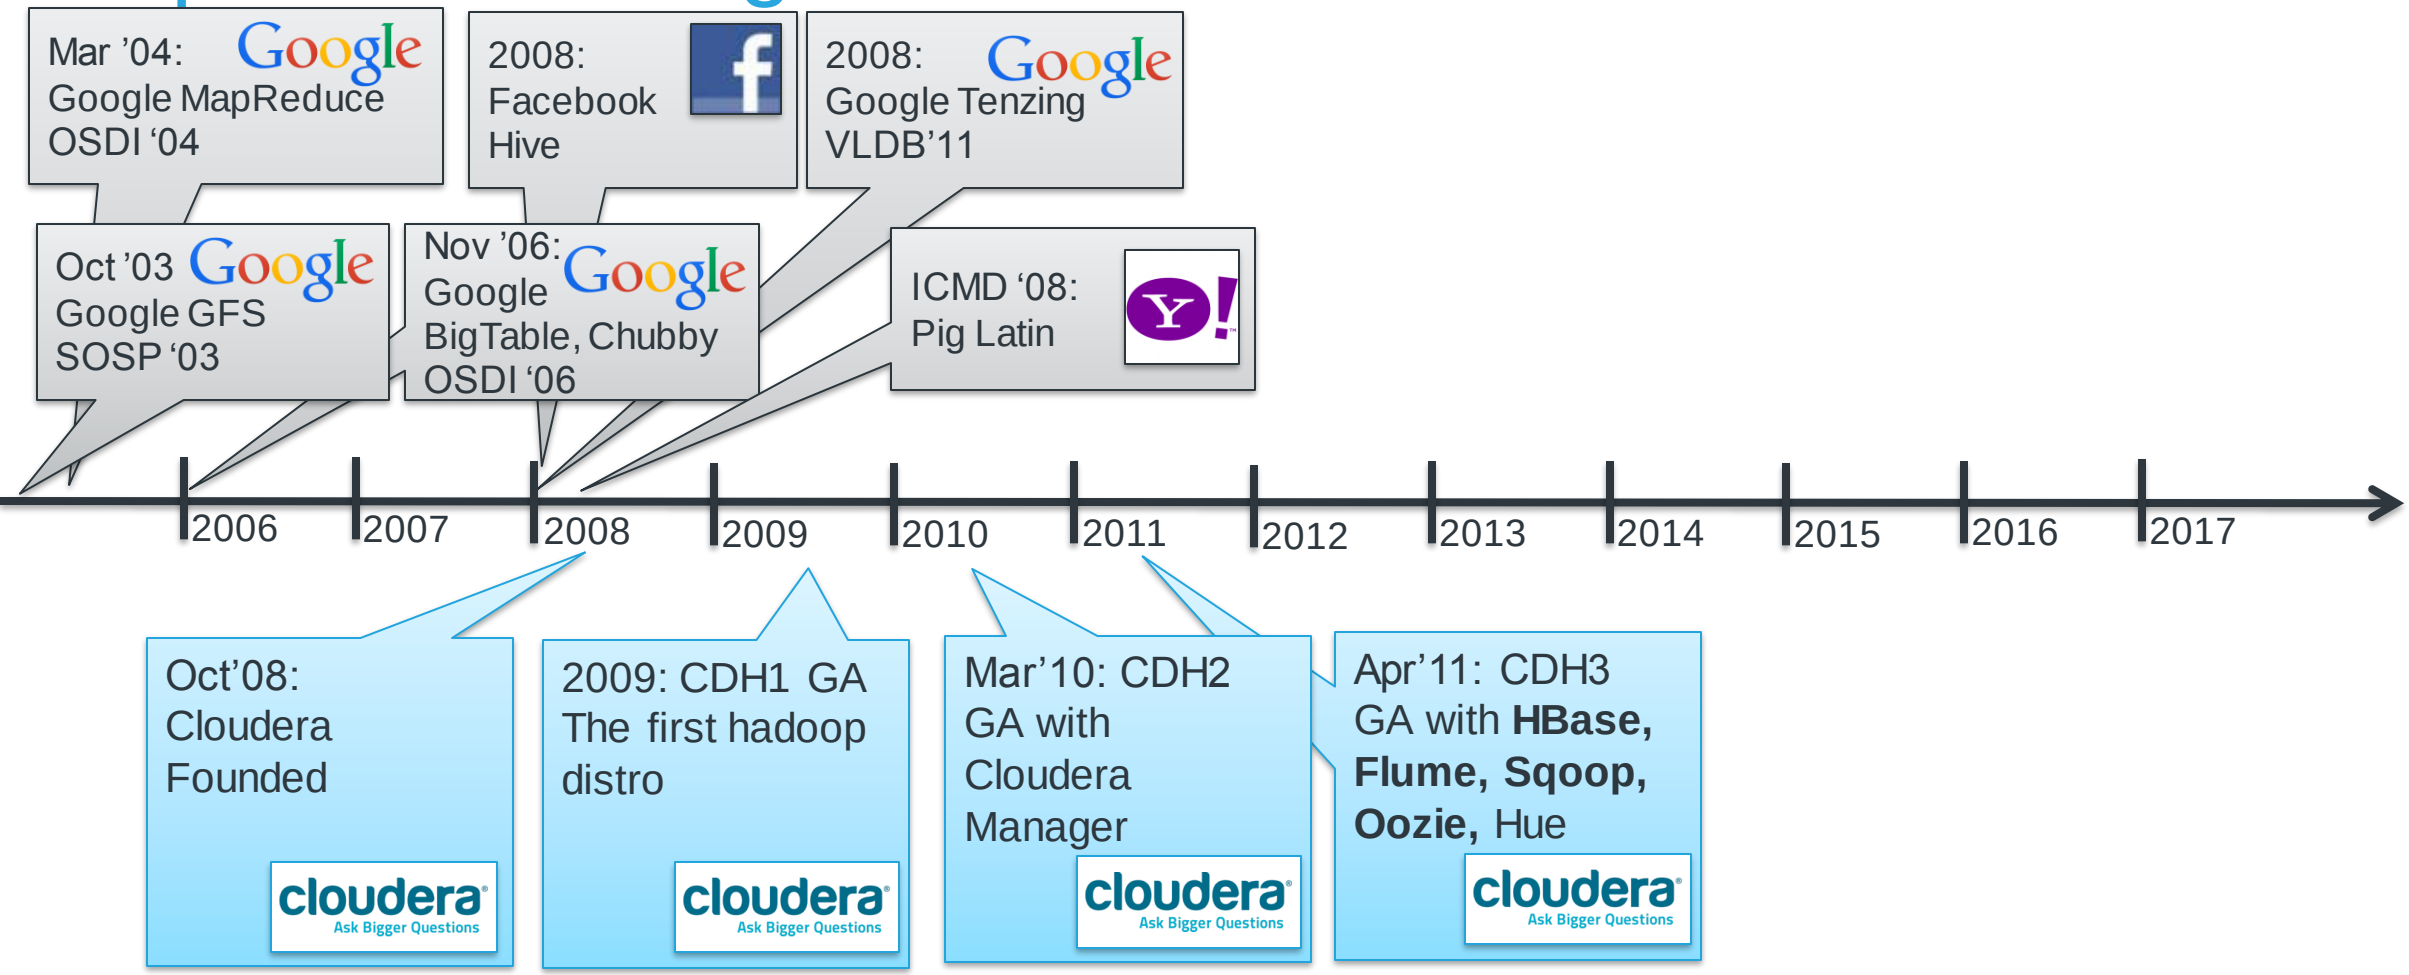

Functions move to Open Source

Function	Google	Yahoo!	Facebook	The Rest of Us
File Storage	GFS => Colossus	HDFS	HDFS	HDFS
Record storage (NoSQL)	BigTable => Megastore => Spanner	PNUTS => HBase	HBase	HBase Kudu
Batch processing	Google MapReduce	Hadoop MapReduce	Hadoop MapReduce	Hadoop MapReduce
Batch query	Sawzall, Tenzing, FlumeJava	Pig	Hive	Pig, Hive, Crunch
Resource Management	Borg => Omega	MR => Hadoop YARN	MR => Corona	Hadoop YARN
Ingest	EvenFlow Custom MySQL Proxy	Data Highway	Scribe / Calligraphus Custom MySQL proxy	Sqoop Flume
Coordination	Chubby	Zookeeper	Zookeeper	Zookeeper

cloudera

Limitations and next challenges

- High throughput, high latency batch processing with Map Reduce

- Deploy on prem
- Coupling storage and compute for **locality** is critical

cloudera

cloudera®
Ask Bigger Questions

The Rise of Data Science

Let's go back to 2012…

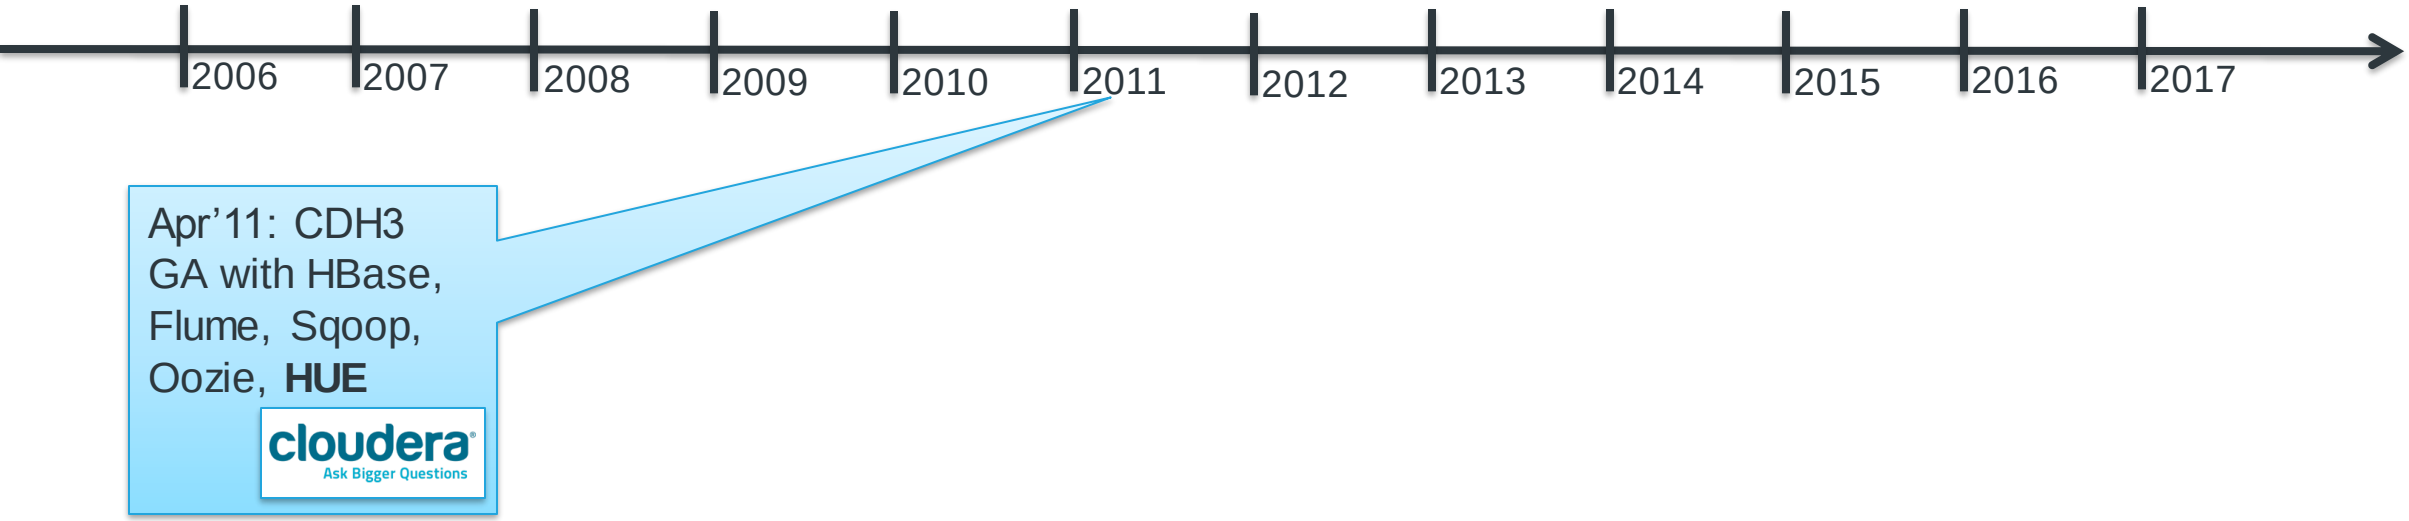

Rise of Mobile Data

- Apple iPhone released in 2007
- Google Android released in 2008
- By 2012, smart phone sales volume outgrew PCs

- The Internet and Social Networking moved to mobile

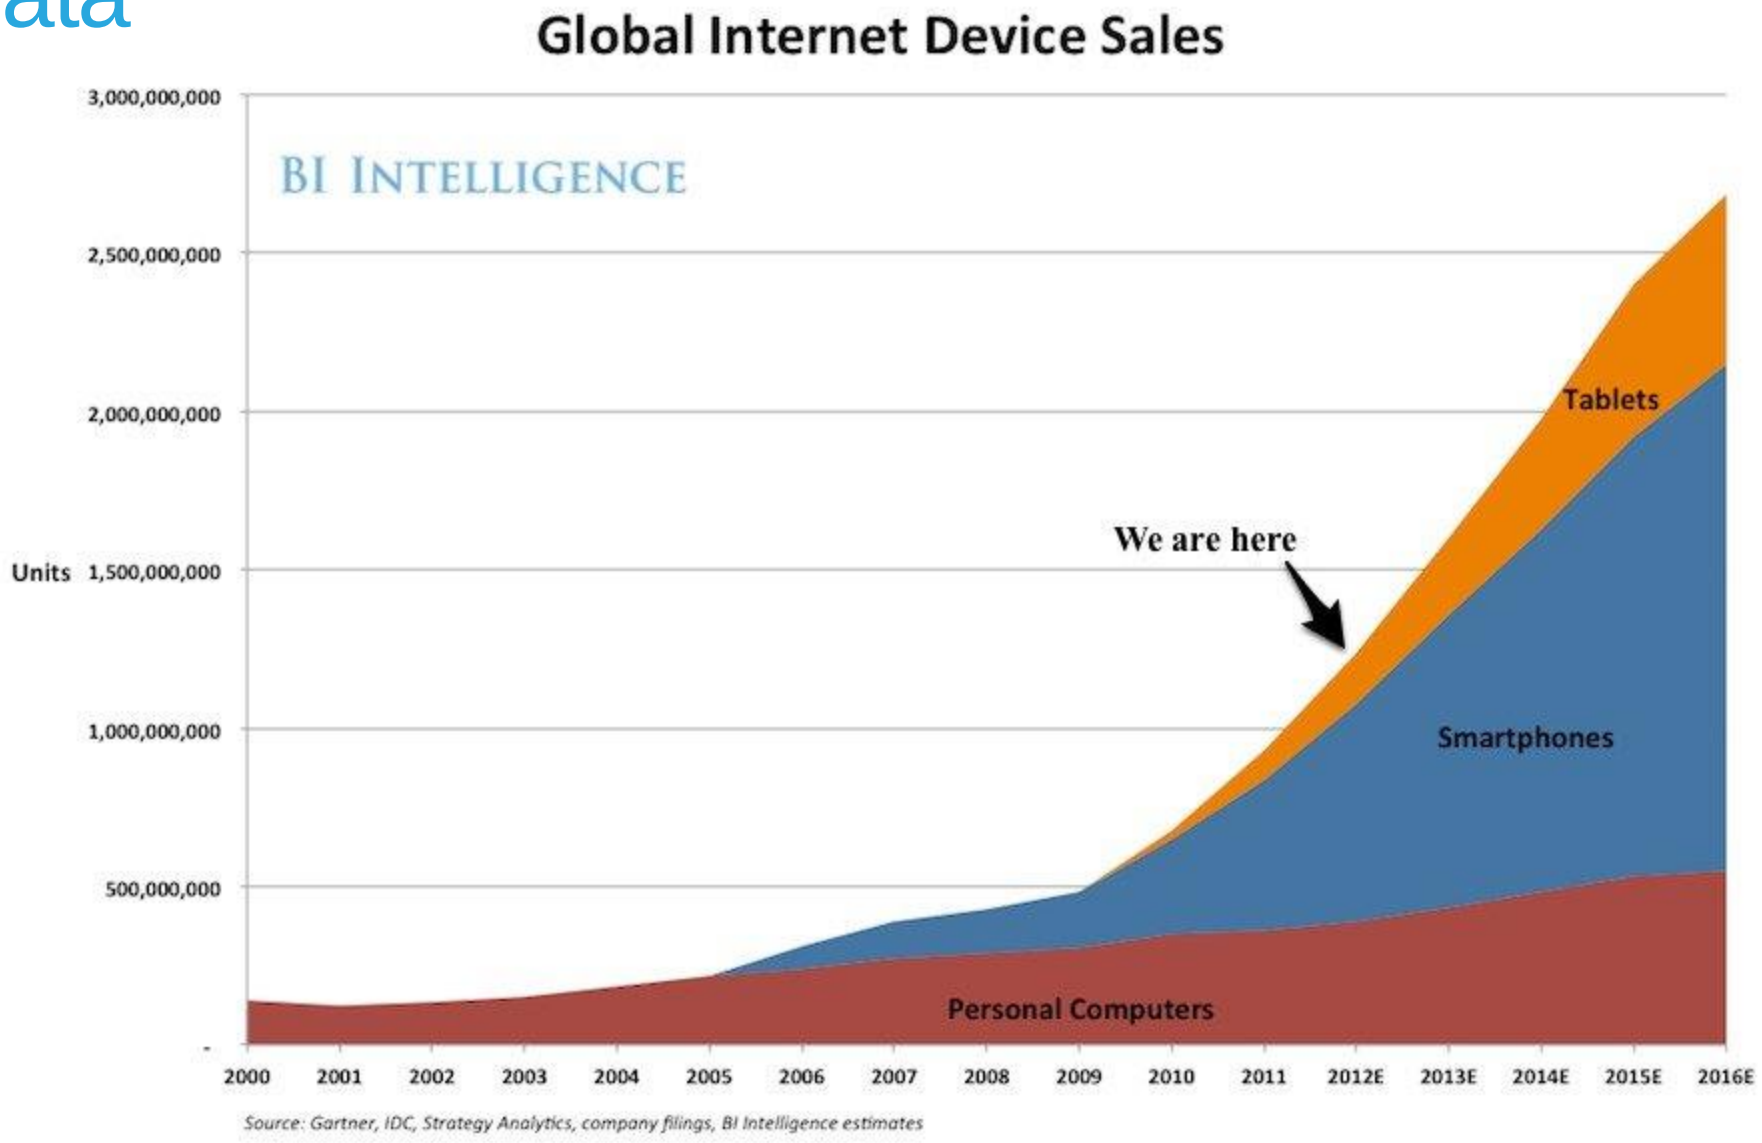

http://www.businessinsider.com/the-future-of-mobile-deck-2012-3

What is Big Data?

1. **Collect** the data
2. **Count** the data
3. **Report** the results

Big Data: Simplify and remove features to enable scaling

- Scalable and Fault tolerant first
- Exclude features to simplify
  - No transactions
  - No Schema
  - No Joins

This is still hard to do!

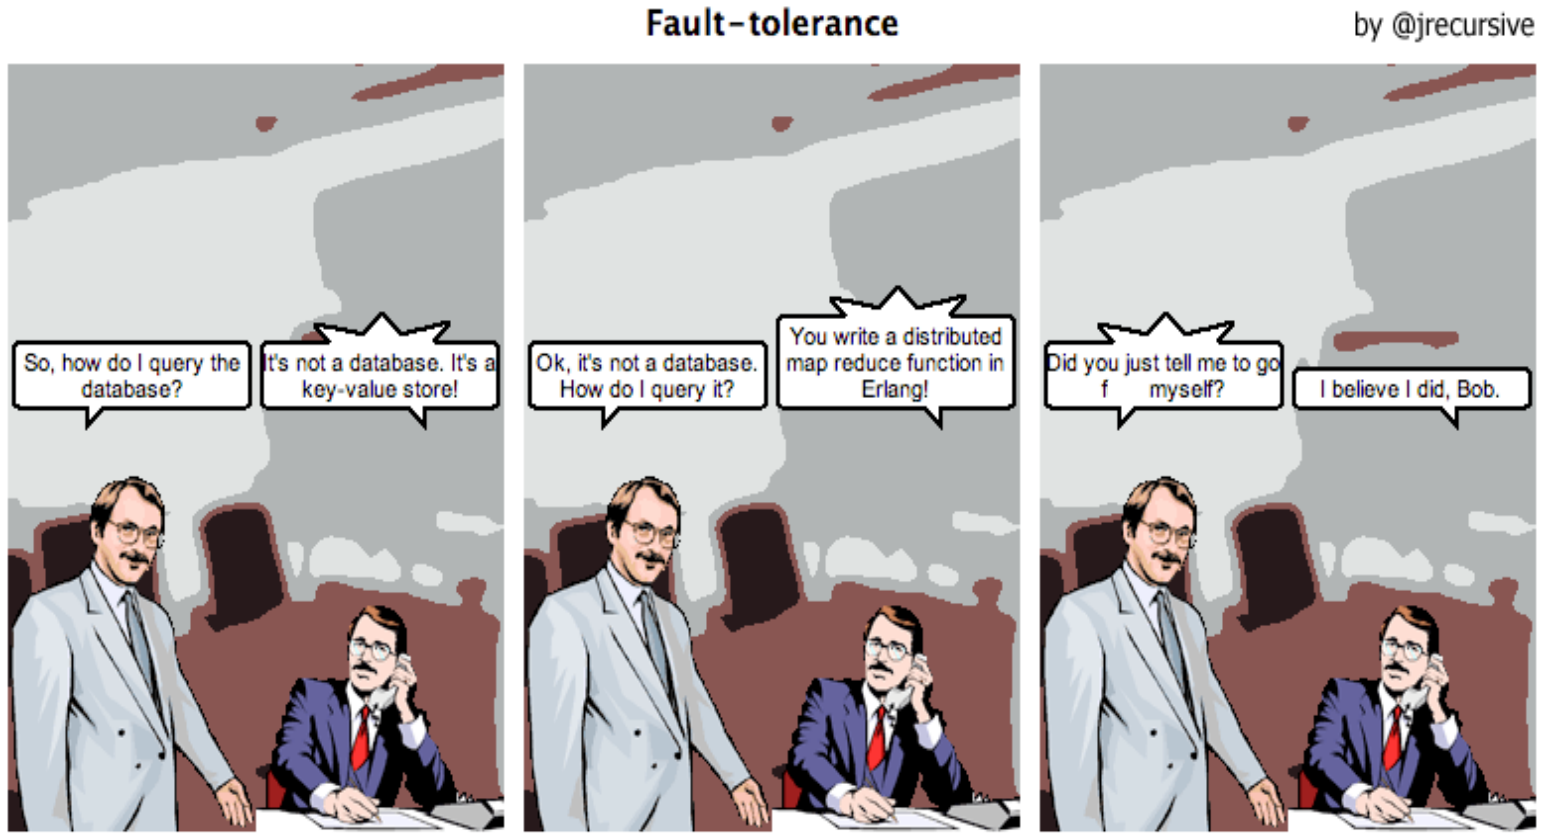

cloudera

Limitations and next challenges

- **High throughput, high latency batch processing** with MapReduce
- **Low-latency SQL** access with Impala
- **Full text search** with Solr
- **Incremental processing** with Spark
- **Near-real time stream processing** with Spark streaming
- Cluster is **secured** and **governed** for multitenancy
- Deploy on prem. Coupling storage and compute for **locality** is critical.
- Cloud is available for good for experimentation

cloudera

What is Big Data?

1. **Collect** the data
2. **Count** the data
3. **Report** the results

What is Data Science?

1. **Wrangle** the data
2. **Explore** the data
3. **Interact** with the results

Data Science and Analytics tools for your data

Interactive distributed SQL query engine

In-memory processing and stream processing

User friendly UX for queries

Fast storage for analytics

Incremental distributed full text search

R/Scala/Python Data Science Notebook for big data platforms (now CDSW)

cloudera

Sharing data requires security and governance

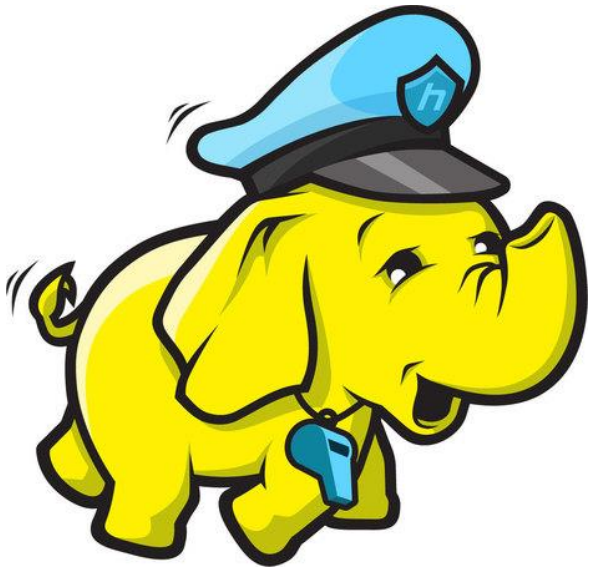

Kerberos Strong authentication
Throughout the data platform

Sentry™

Fine grained data
access control

xplain

Optimization based on
historical queries
(now Navigator Optimizer)

cloudera navigator

Lineage and auditing of data
operations.

Gazzang

Strong Encryption for data at rest
(now Navigator Encrypt + KeyTrustee)

cloudera

The Data Science and Analytics Stack

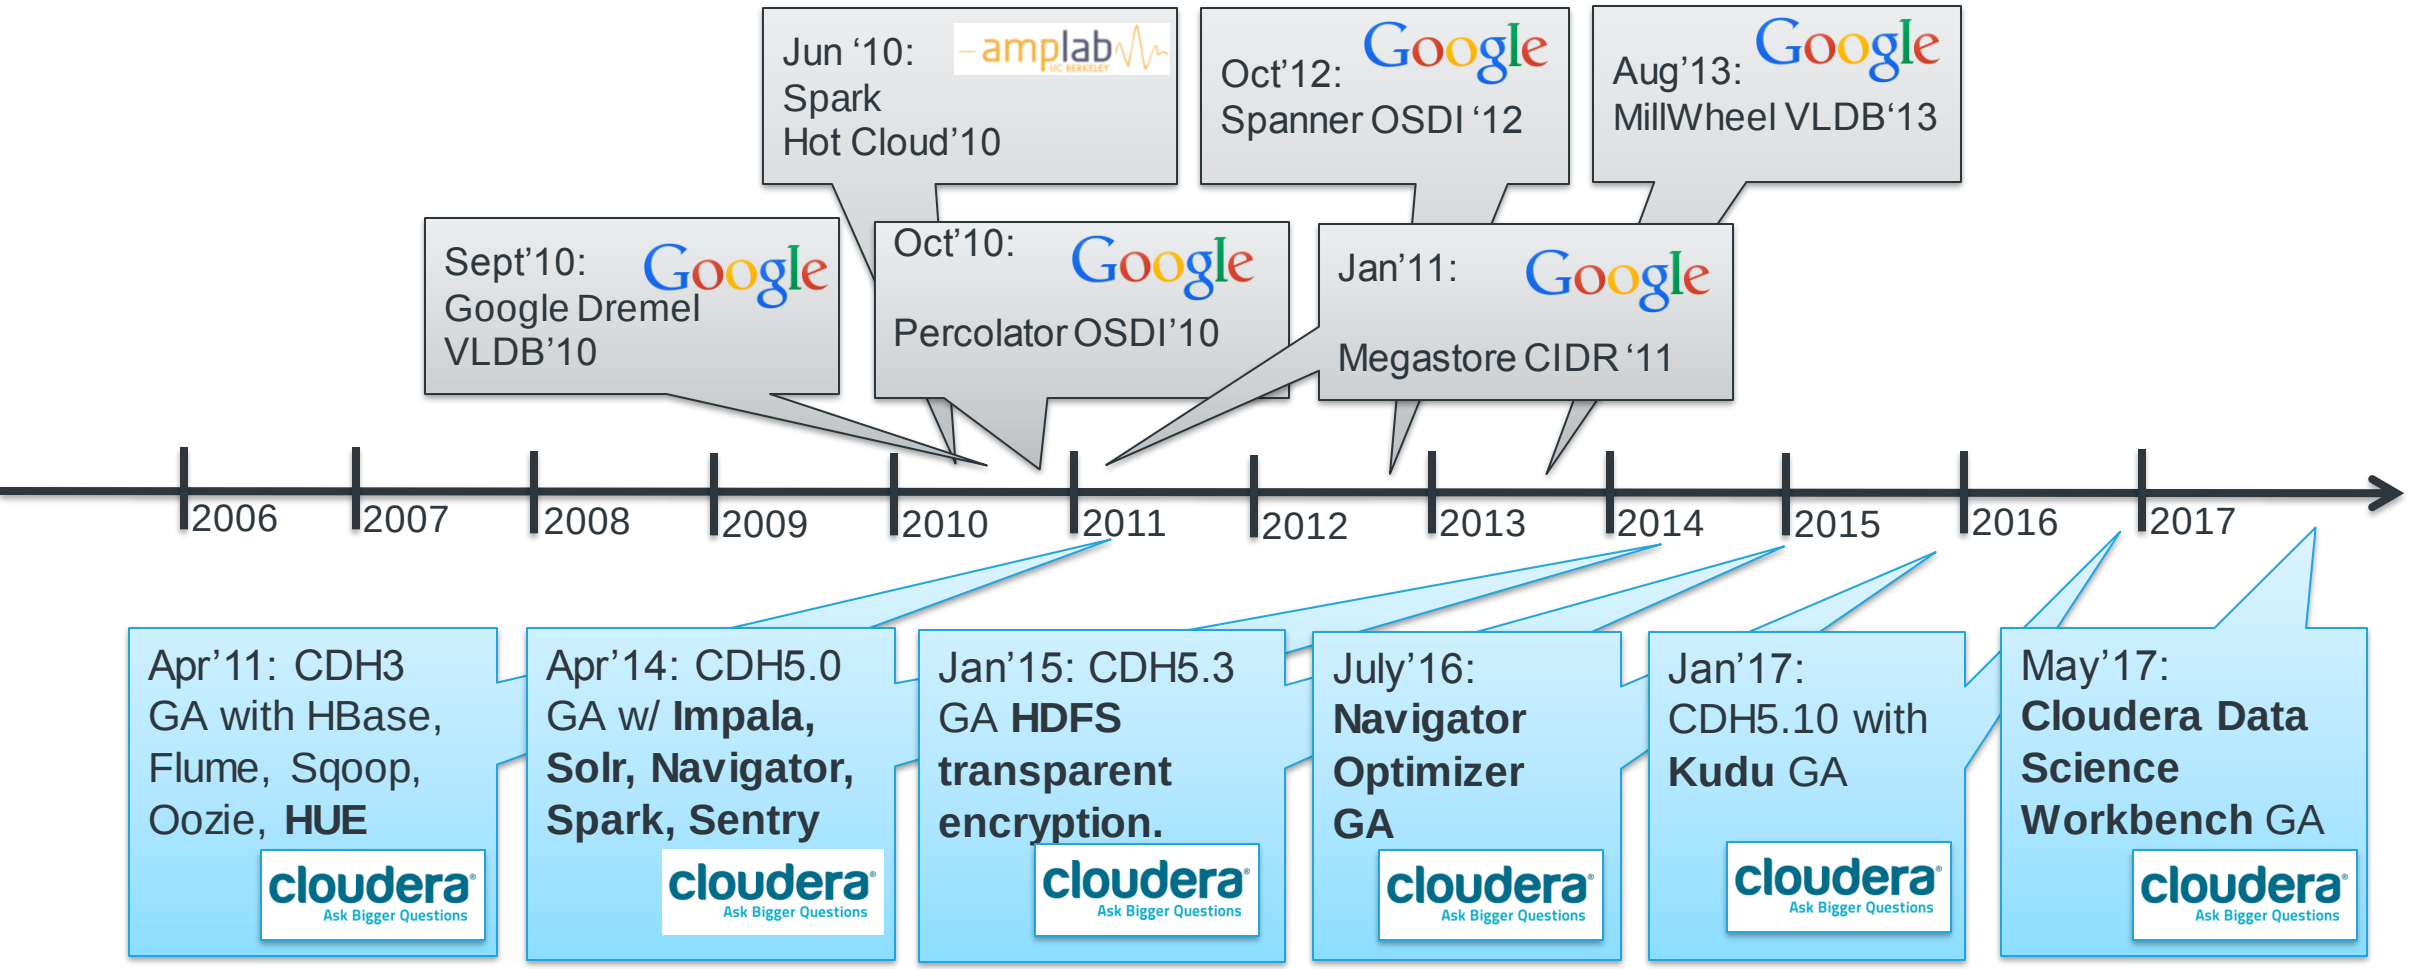

The Present: Enabling Machine Learning

Even More Data

- Sensors and IoT
  - Uber / Lyft / Didi
  - Fitbit / iWatch
  - Mobile payments: Paypal, Square, AliPay, WePay

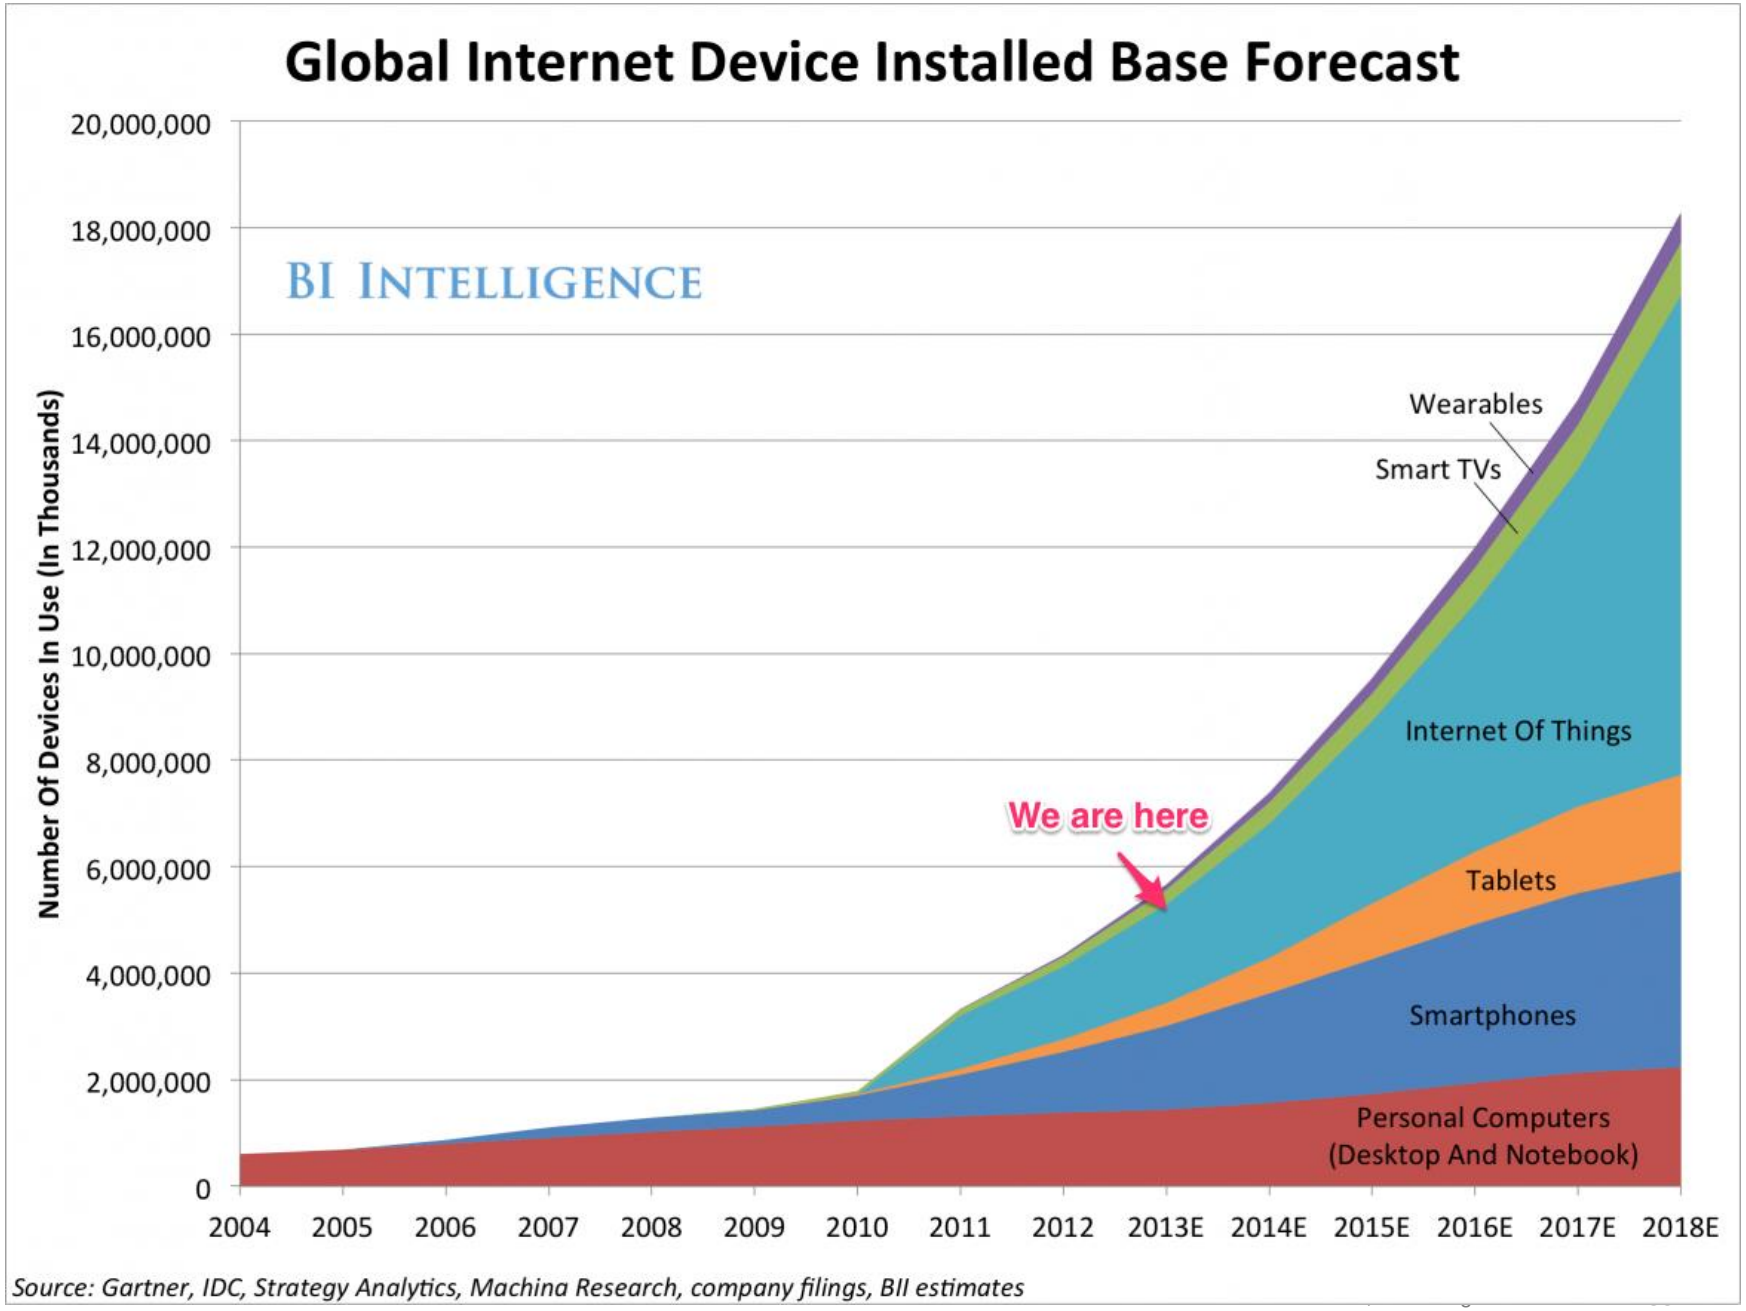

cloudera

What is Data Science?

1. **Wrangle** the data
2. **Explore** the data
3. **Interact** with the results

What is Machine Learning?

1. **Wrangle** the data
2. **Model** the data
3. **Predict** the results

More data usually beats better algorithms

- Use simple algorithms to learn about data from data
  - Supervised ML with training data
  - Unsupervised ML algorithms for clustering
  - Deep Neural nets that learn data using data.

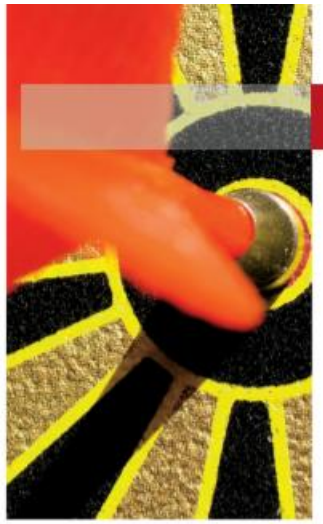

EXPERT OPINION

Contact Editor: **Brian Brannon,** bbrannon@computer.org

The Unreasonable Effectiveness of Data

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

Eugene Wigner's article "The Unreasonable Effectiveness of Mathematics in the Natural Sciences"[1] examines why so much of physics can be neatly explained with simple mathematical formulas behavior. So, this corpus could serve as the basis of a complete model for certain tasks—if only we knew how to extract the model from the data.

Learning from Text at Web Scale

The biggest successes in natural-language-related

cloudera

THE DATA SCIENCE HIERARCHY OF NEEDS

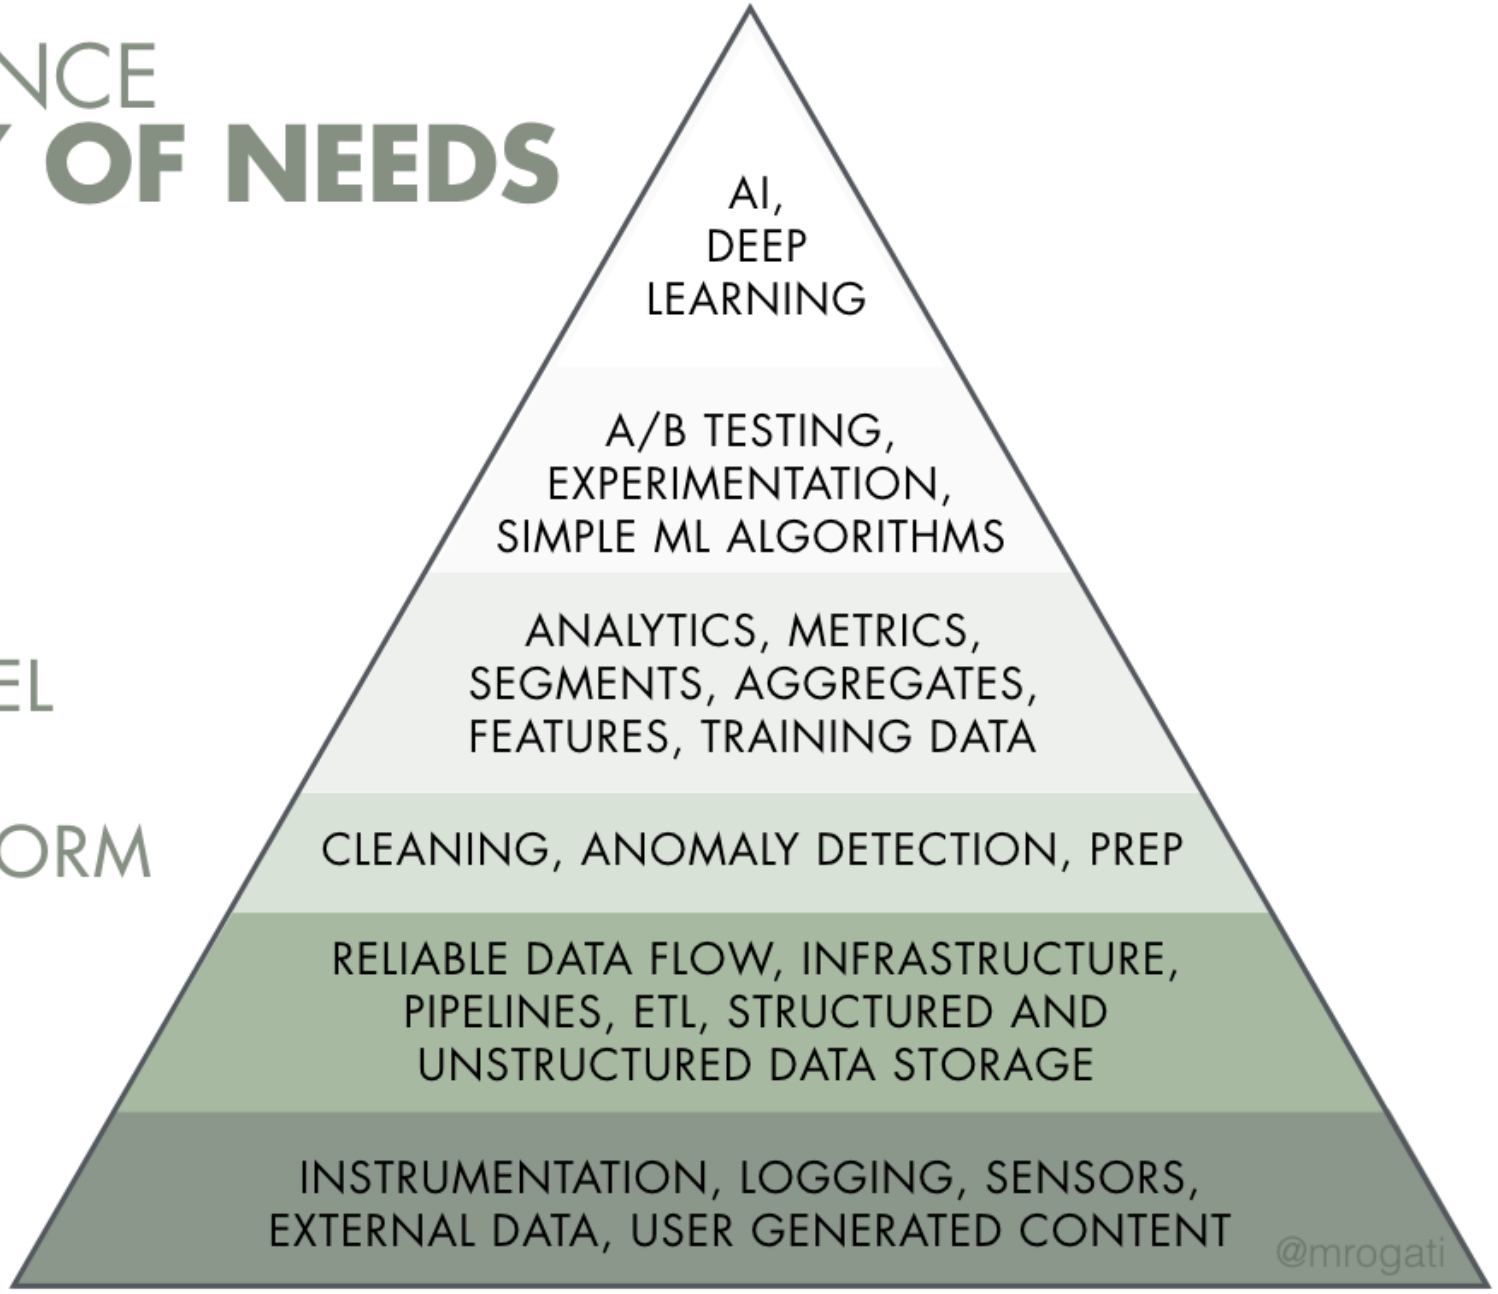

LEARN/OPTIMIZE

AGGREGATE/LABEL

EXPLORE/TRANSFORM

MOVE/STORE

COLLECT

https://hackernoon.com/the-ai-hierarchy-of-needs-18f111fcc007

What's Next

cloudera

Apache Hadoop

Interest by region Region

Big data

Interest by region Region

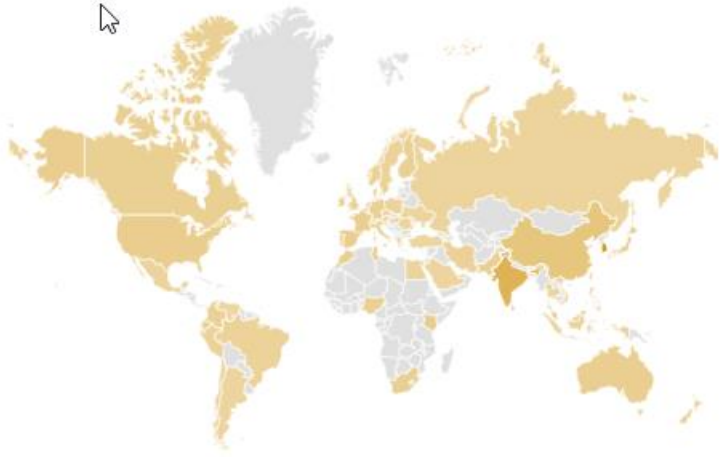

Machine learning

Interest by region Region

Data Science

Interest by region Region

Hardware comparison

2008

- CPU:
  - 8 Core 2GHz
- RAM:
  - 4GB
- NIC:
  - Dual 1GBe
- Disk:
  - 0.5 TB HDD

2017

- CPU: (4x)
  - 2x 14 Cores 2+Ghz
- RAM: (128x)
  - 512GB
- NIC: (5x)
  - 10GBe
- Disk: (250x)
  - 12-24x 6TB HDD (~144TB)
  - SSD System Disk

cloudera

From Cloudera's original pitch deck (Sept 2008)

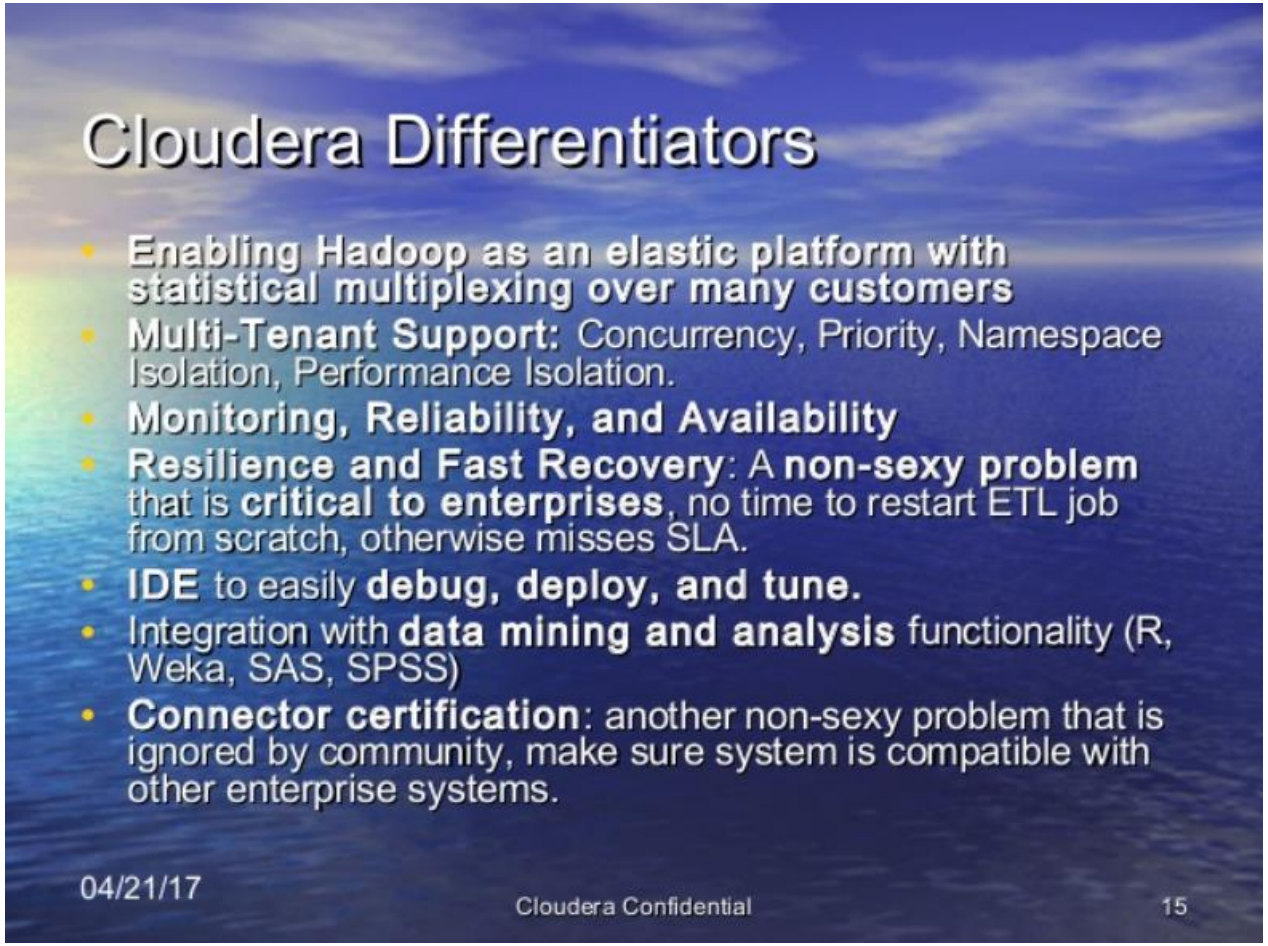

https://www.slideshare.net/AccelPartners/clouderas-original-pitch-deck-from-2008

cloudera

Data Platforms in the Cloud

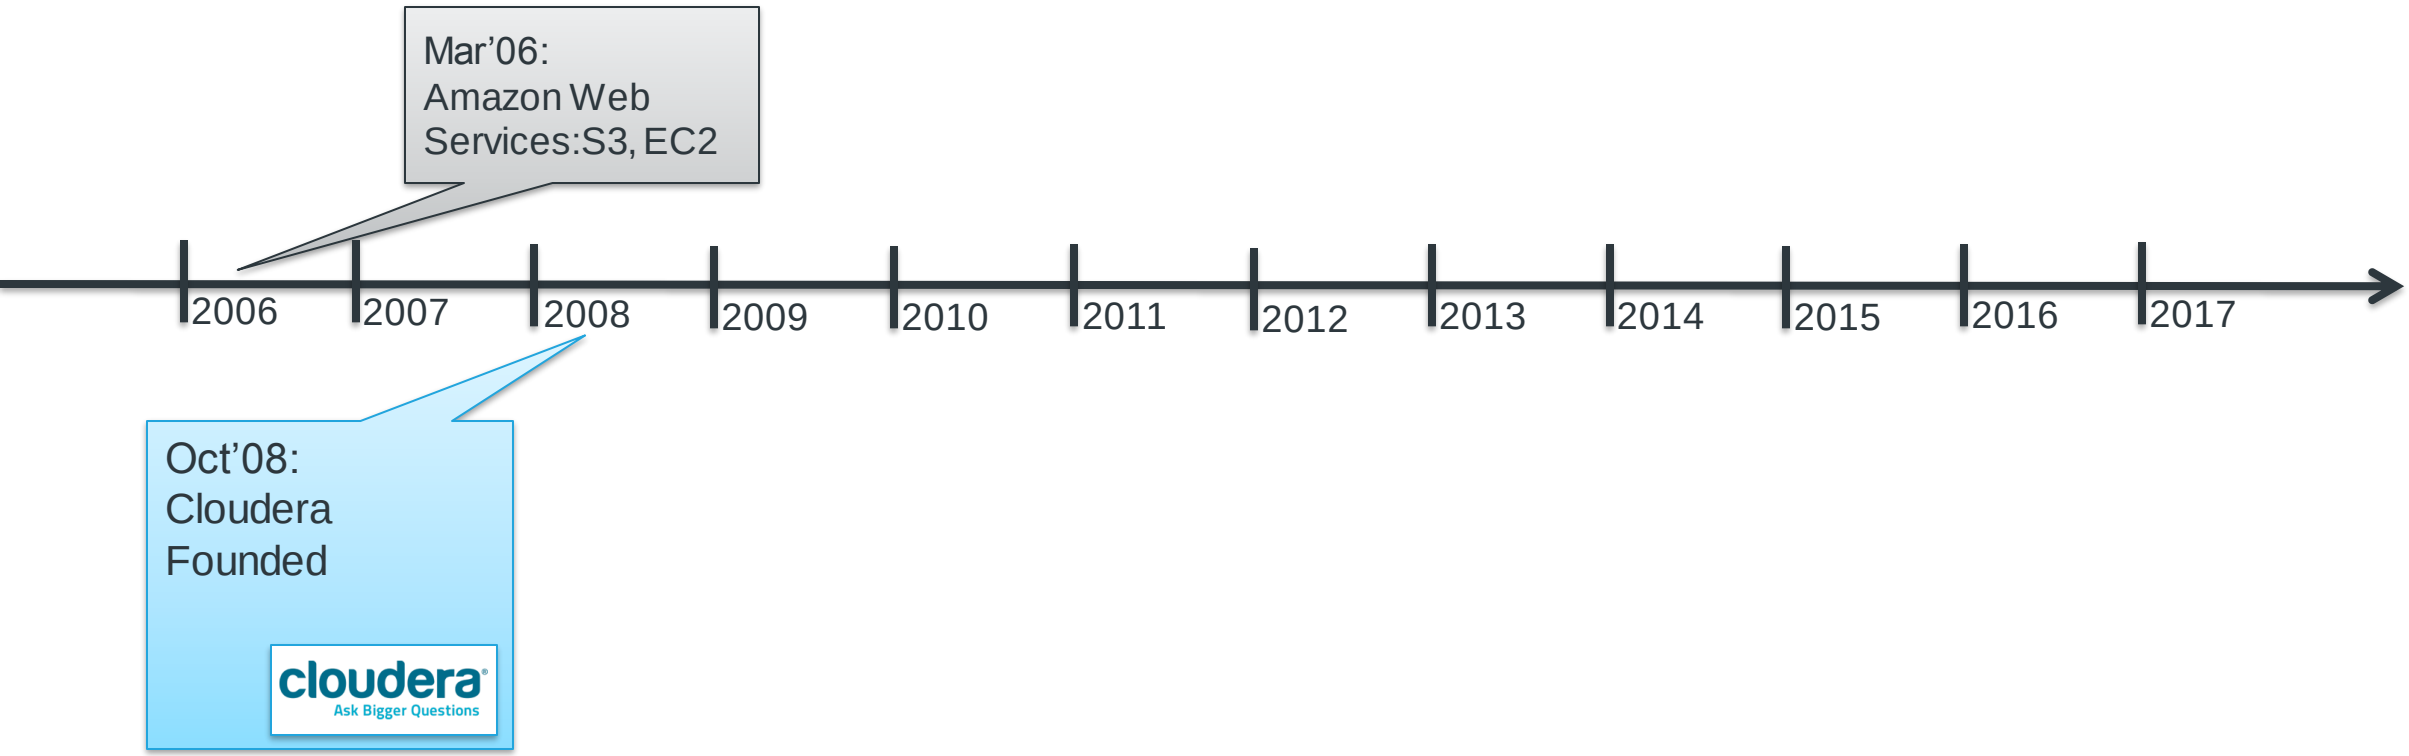

cloudera

The growth of cloud is now undeniable

- AWS revenues have a slow start
- But growth and competitor are here now.

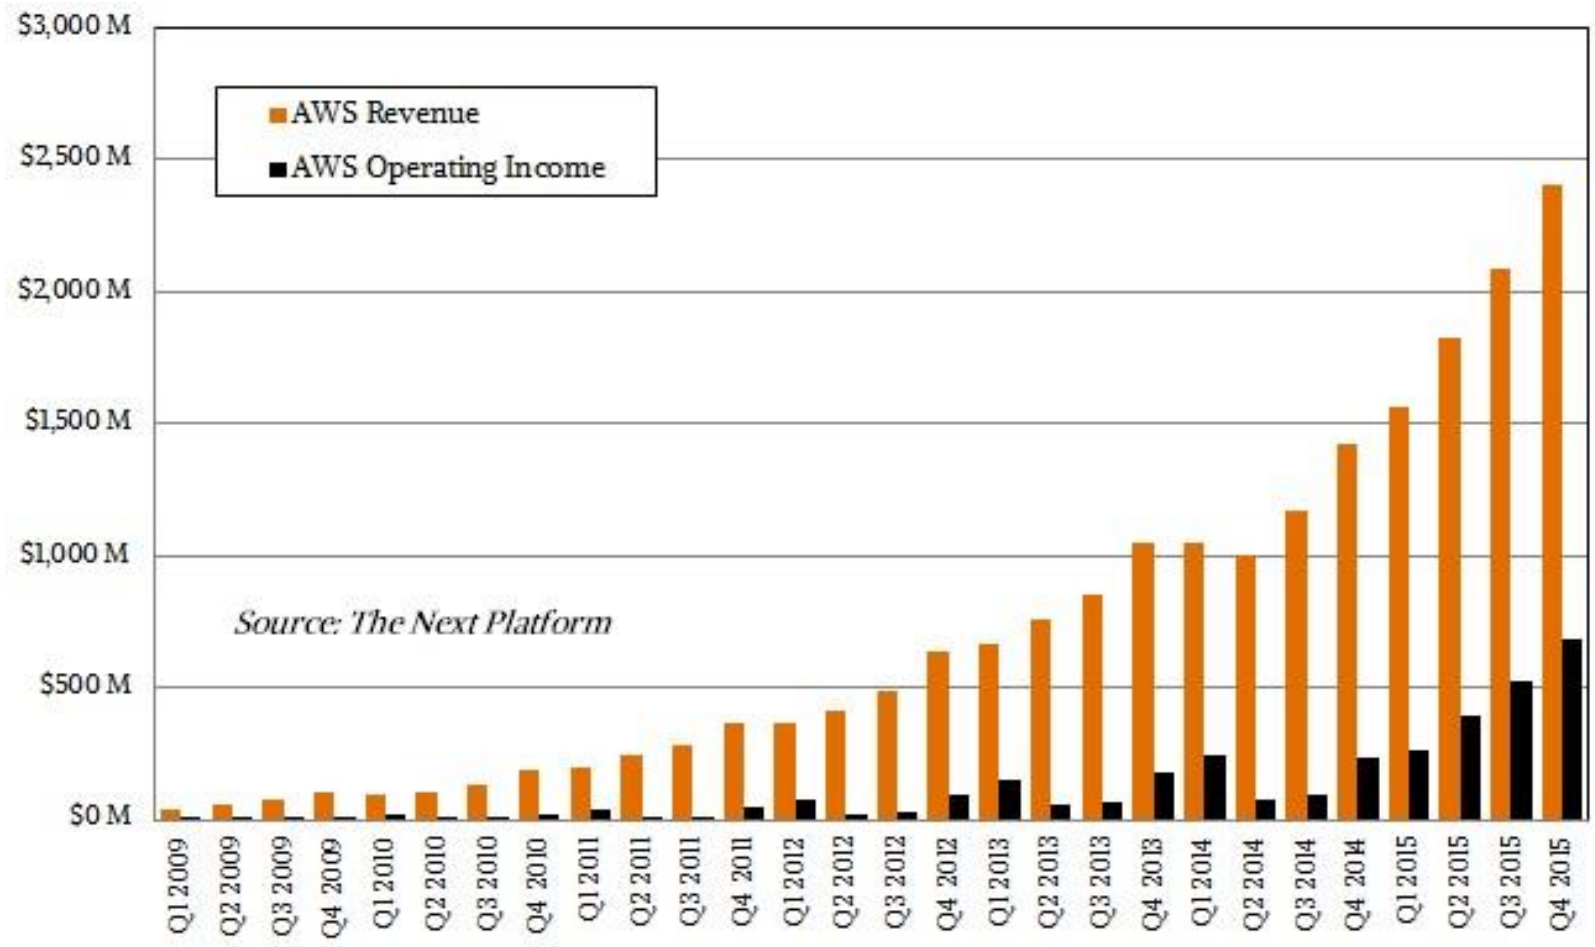

https://www.nextplatform.com/2016/02/01/how-long-can-aws-keep-climbing-its-steep-growth-curve/

cloudera

Data Platforms in the Cloud

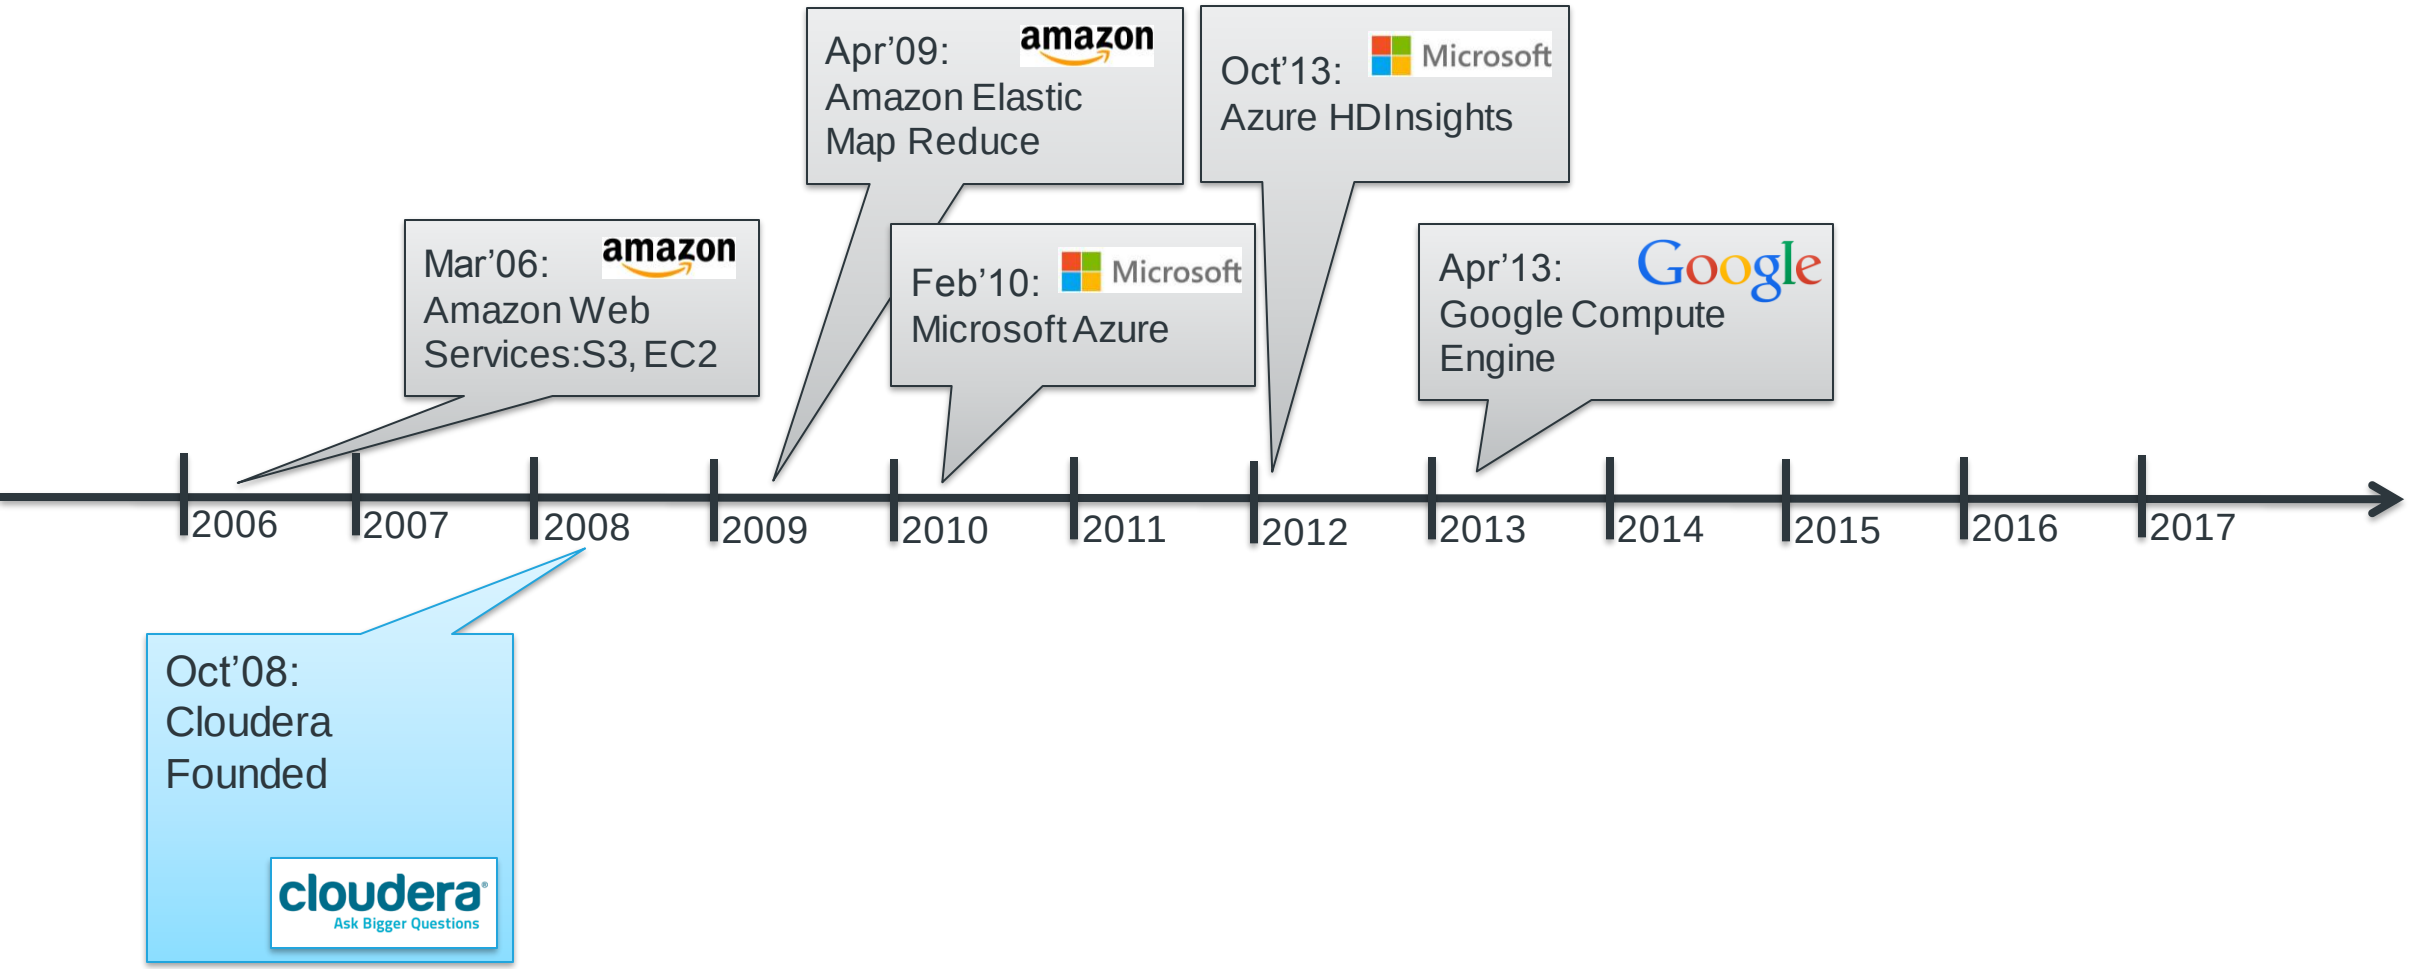

cloudera

New Assumptions with Data Platforms on Cloud

- On-prem Assumption
  - Machines are always on and added every few months. High Cap-Ex
  - Coupling Storage and compute for locality is critical for efficient performance
  - Deal with Services
  - Cost for maintaining services software and hardware health on Ops team
- Cloud Assumption
  - Machines are transient and elastic and added and removed every few hours. Low Cap-Ex
  - Decoupling storage from Compute is critical for efficiently lowering Op-Ex
  - Deal with Jobs or Requests
  - Cost for maintaining services software and hardware heal outsourced.

cloudera

Data Platforms in the Cloud

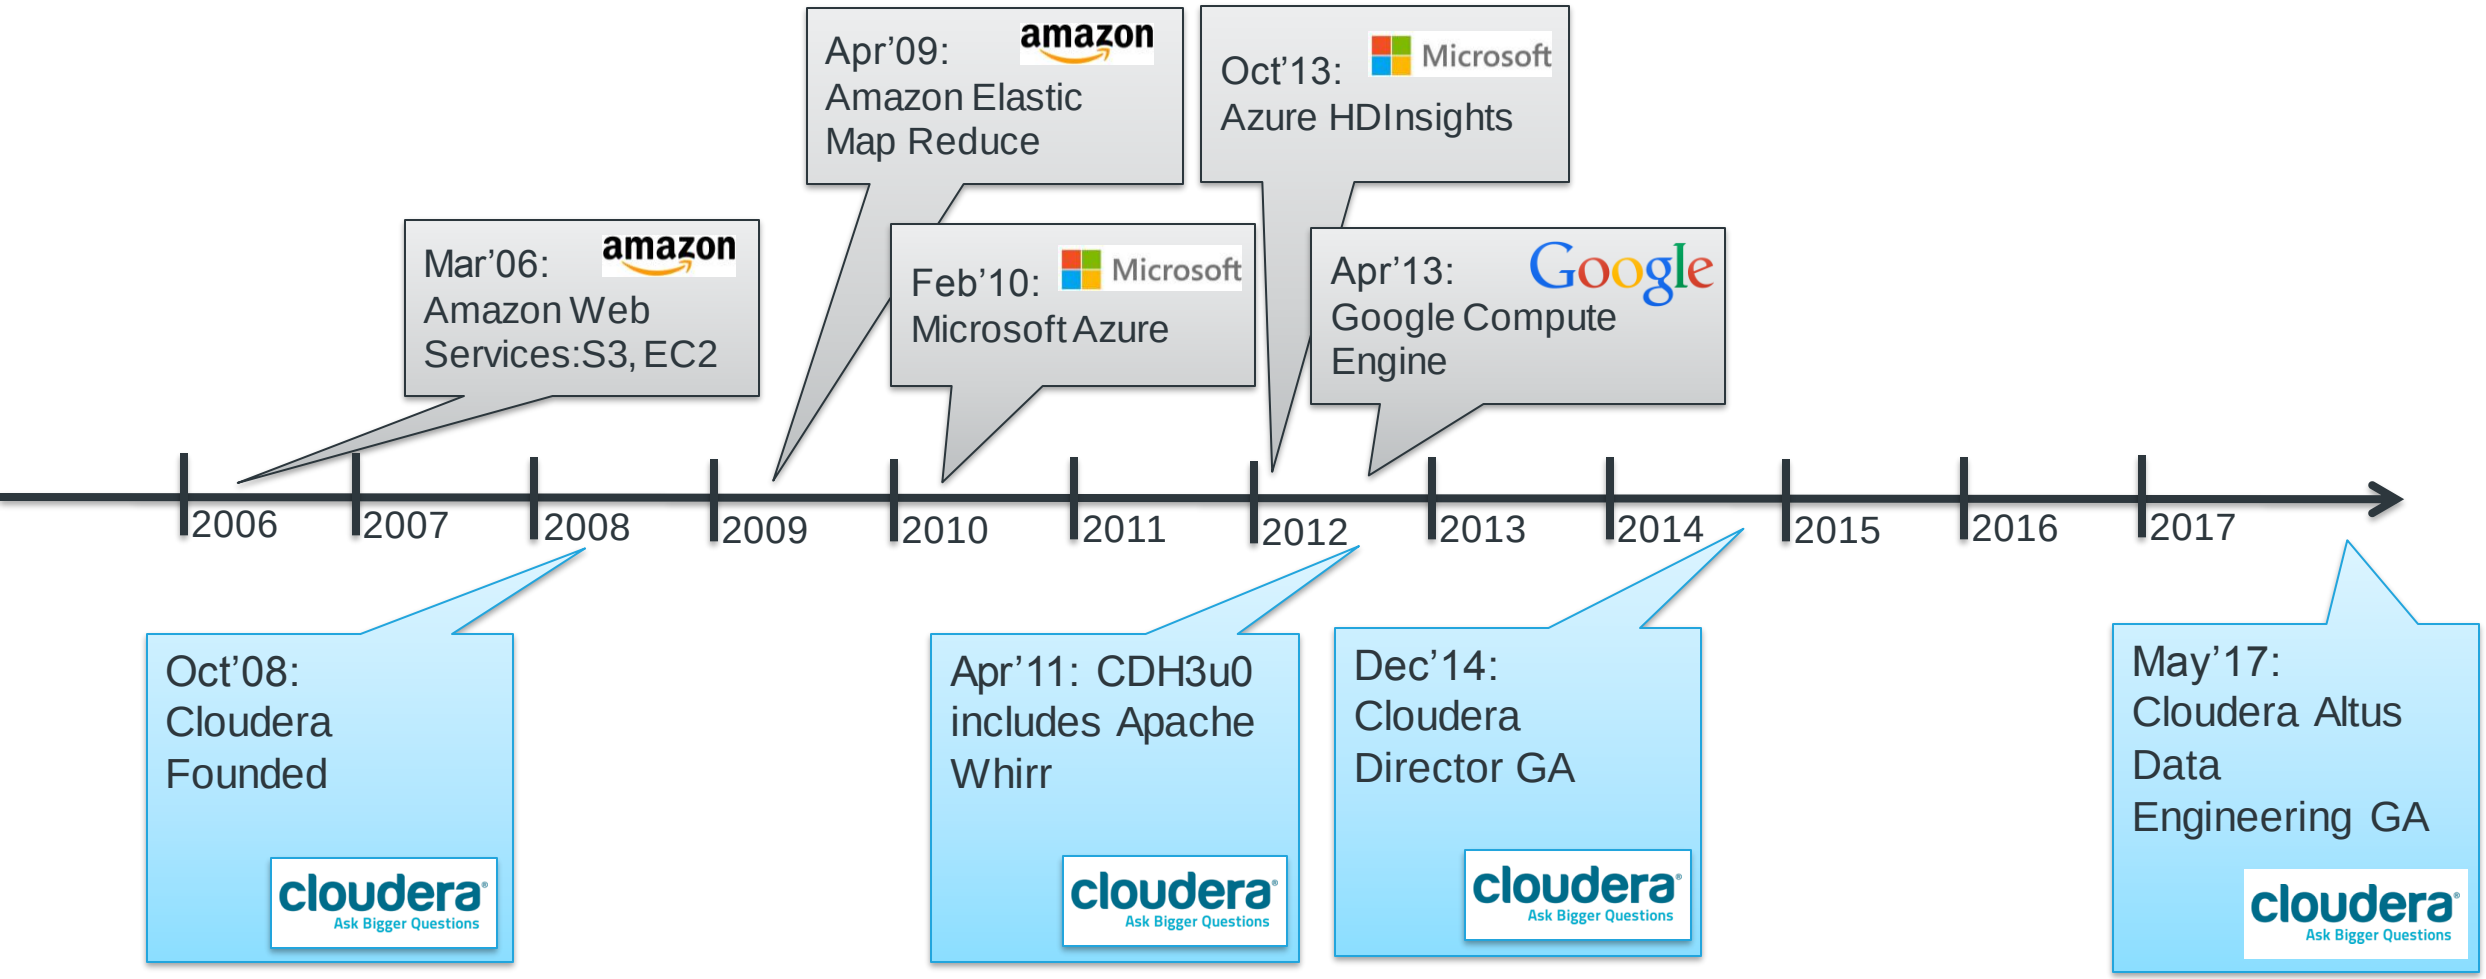

cloudera

The only way you can predict the future is to build it.
--Alan Kay

cloudera

Thank you!

cloudera