Big Data

Dr. Wenceslao Palma
wenceslao.palma@pucv.cl

Big Data


www.theclinic.cl/2017/01/19/ma
MENÚ
THE CLINIC
Martin Hilbert, experto en redes digitales: "Obama y Trump usaron el Big Data para lavar cerebros"
77
COMENTARIOS
Daniel Hopenhayn | 19 Enero, 2017
Tags: big data, Estados Unidos, Martin Hilbert, obama, redes digitales, Trump


Página 22
"Hay gente que cree que no se ha encontrado el remedio para el cáncer porque perjudicaría a las farmacéuticas"


ciperchile.cl/2017/02/13/estudio
CIPER Chile
13 Febrero, 2017 13s comentarios
Estudio prevé que el 50% de los trabajadores chilenos será reemplazado por máquinas
Compartir
Publicar en Twitter
+1
Correo electrónico

Big Data: the four V's

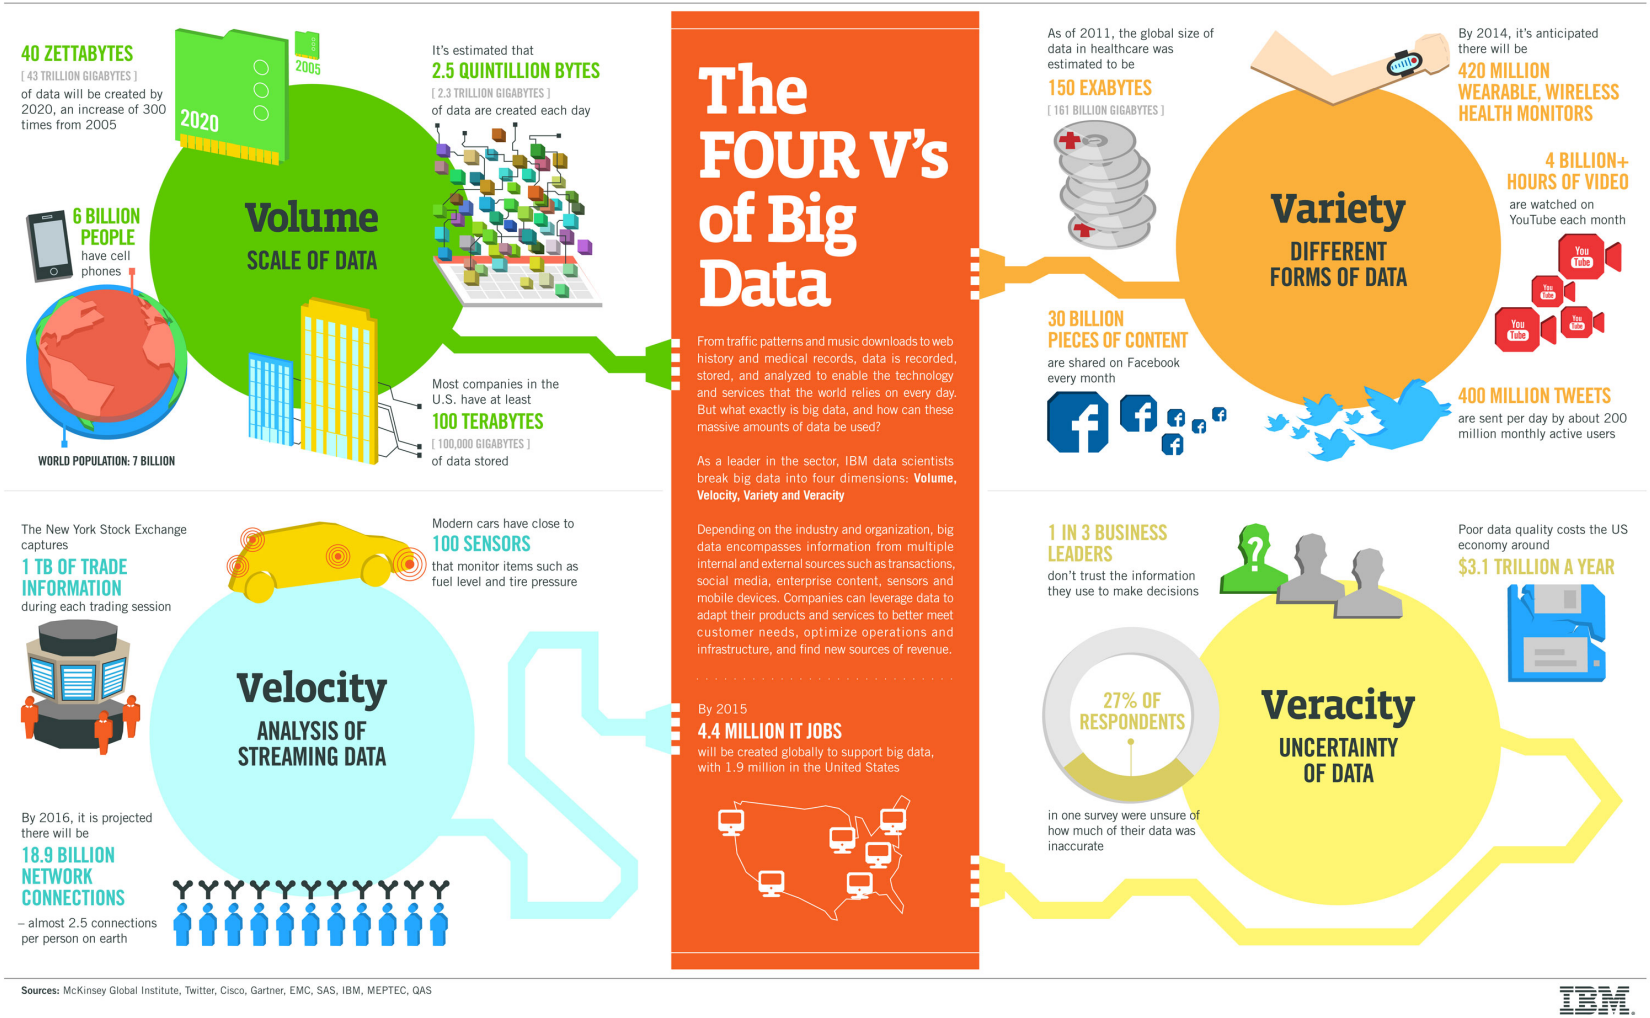

Big Data: the five V's

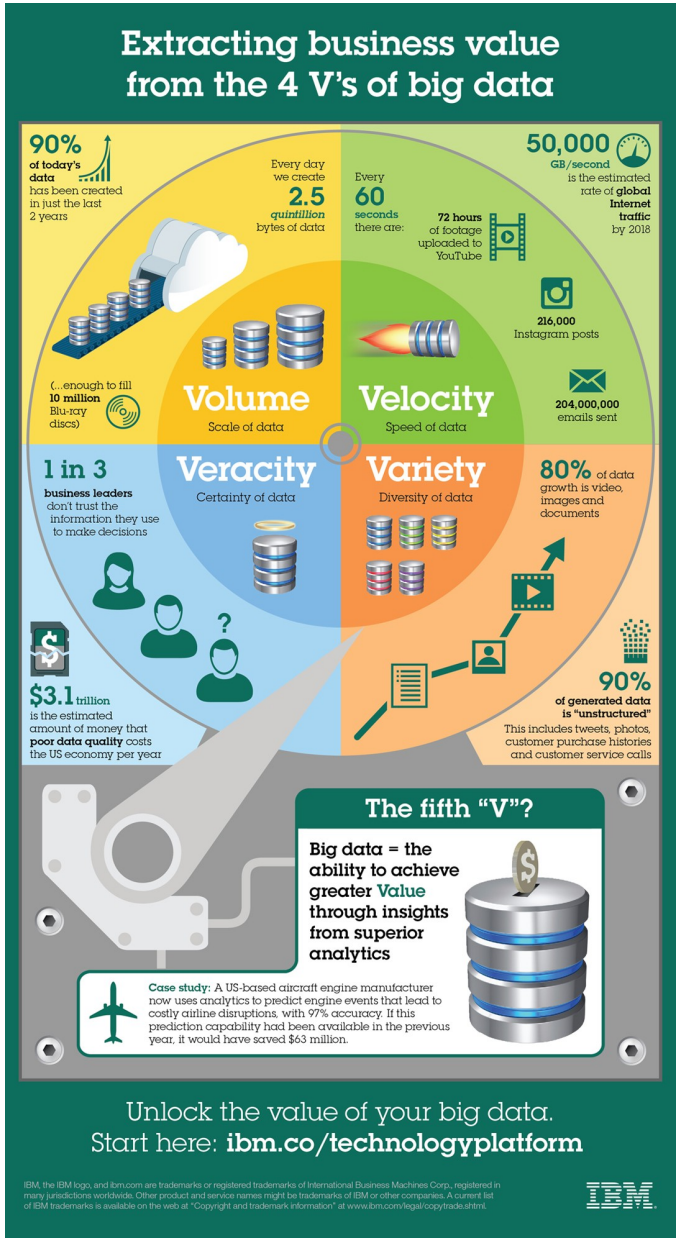

Big Data: How does it mean?

Big Data as stated by Cesar Hidalgo

- A lot of people actually are confused between Big Data and a lots of data.
- Big Data has to be big in three different ways: size, resolution and scope.

Big Data Mining: data science - business intelligence

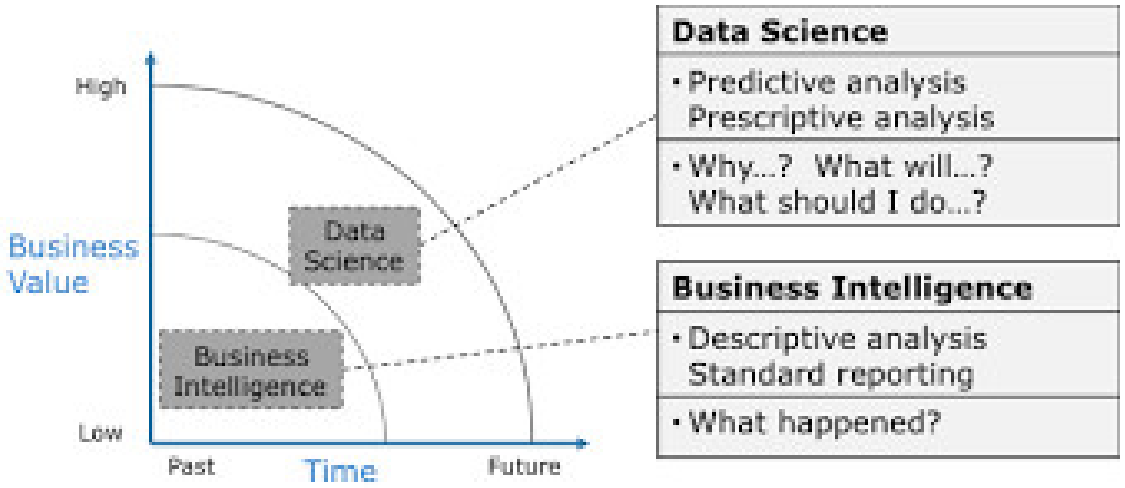

Big Data Mining: Google Flu Trends

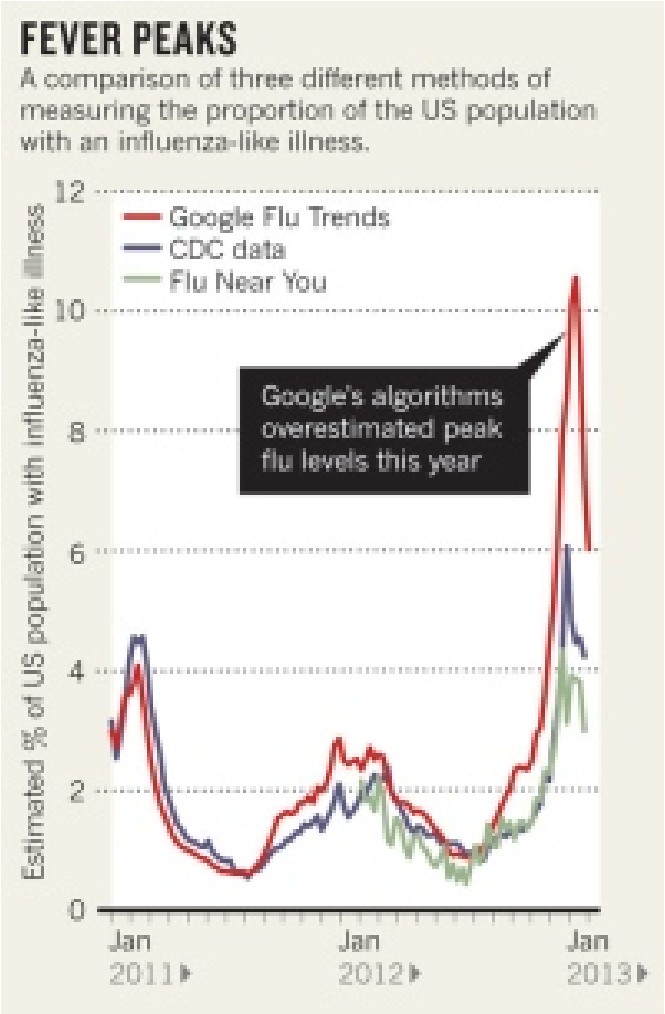

Big Data: how to tackle it?

- The only feasible approach to tackling large-data problems today is to divide and conquer.
- The general principles behind divide-and-conquer algorithms are broadly applicable to a wide range of problems in many different application domains.
- There are many issues that need to be addressed:
  - How to organize the data store?
  - How to break up a large problem into smaller tasks?
  - How to assign tasks across a potentially large number of computer nodes.?
  - How to coordinate synchronization among the different nodes?
  - How to share partial results from one node that is needed by another?
  - How do we accomplish all of the above in the face of software errors and hardware faults?

Big Data Mining: the hadoop ecosystem

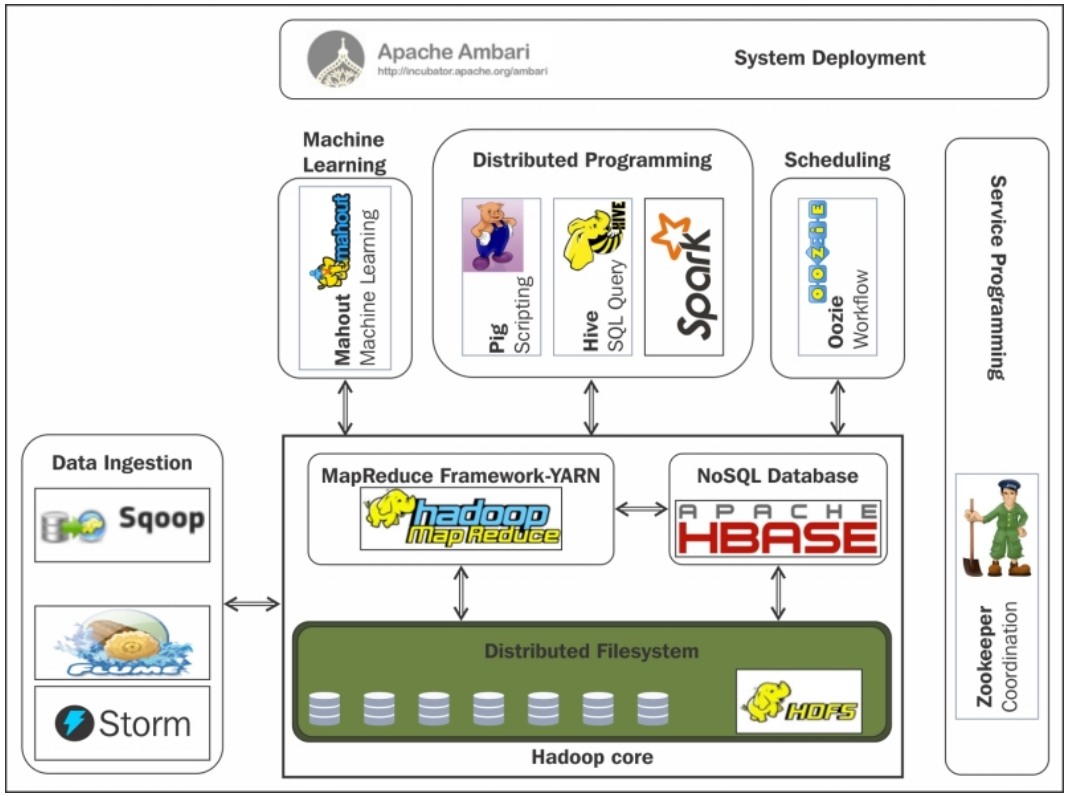

Big Data: Why worry about foundations?

- Chemistry before the tubes.

Big Data: Why worry about foundations?

- Chemistry before the tubes.
- Computer science before big data mining tools.

HDFS

Hadoop Distributed File System (HDFS) is the primary storage system used by Hadoop applications. HDFS creates multiple replicas of data blocks and distributes them on compute nodes throughout a cluster to enable reliable, extremely rapid computations.

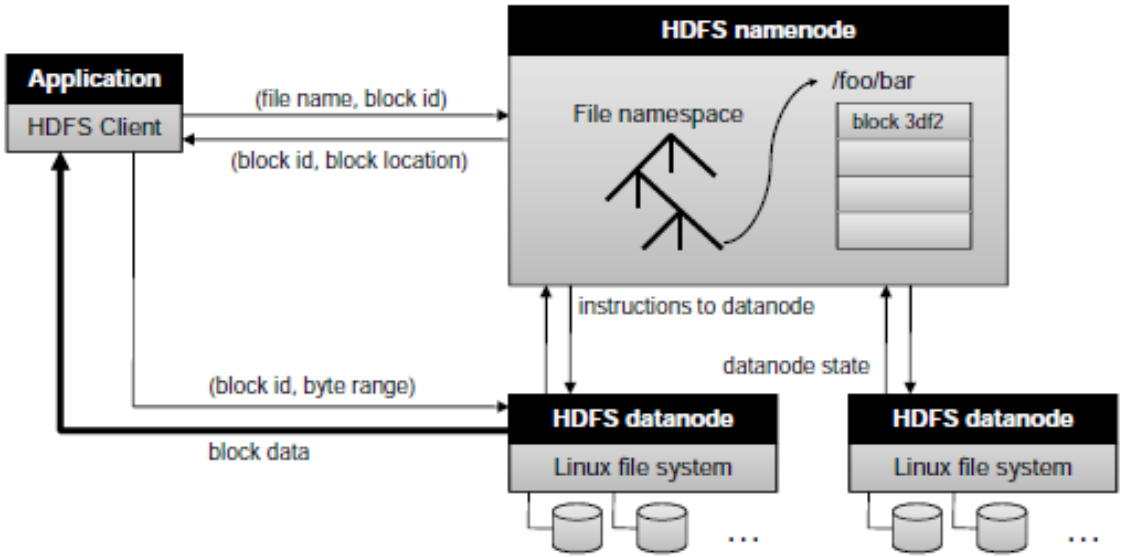

- Data replication makes the system more fault tolerant.
- Data replication provides scalability (w.r.t. data access).
- Data replication and partitioning provide high concurrency.

- Data replication makes the system more fault tolerant.
- Data replication provides scalability (w.r.t. data access).
- Data replication and partitioning provide high concurrency.

The problem with replication is that is hard to maintain data consistency over the time but in big data systems, the data is written once and the updates are stored as additional data sets over the time.

Big Data: MapReduce

MapReduce is a programming model for data processing introduced by Google (2004) to support parallel and fault-tolerant computations over large data sets on clusters of computers. It provides an abstraction that hides many system-level details from the programmer.

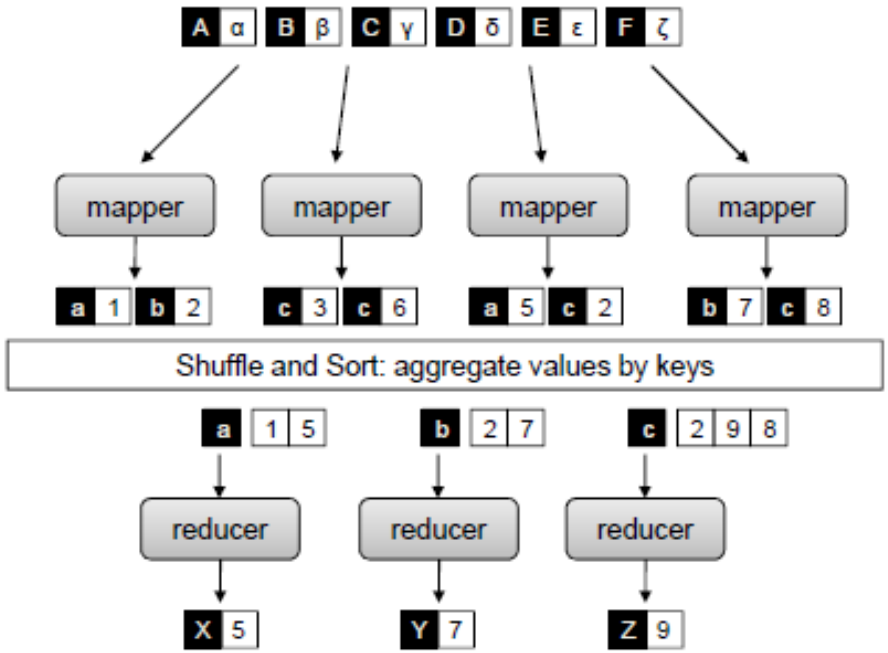

Big Data: Cloud Services

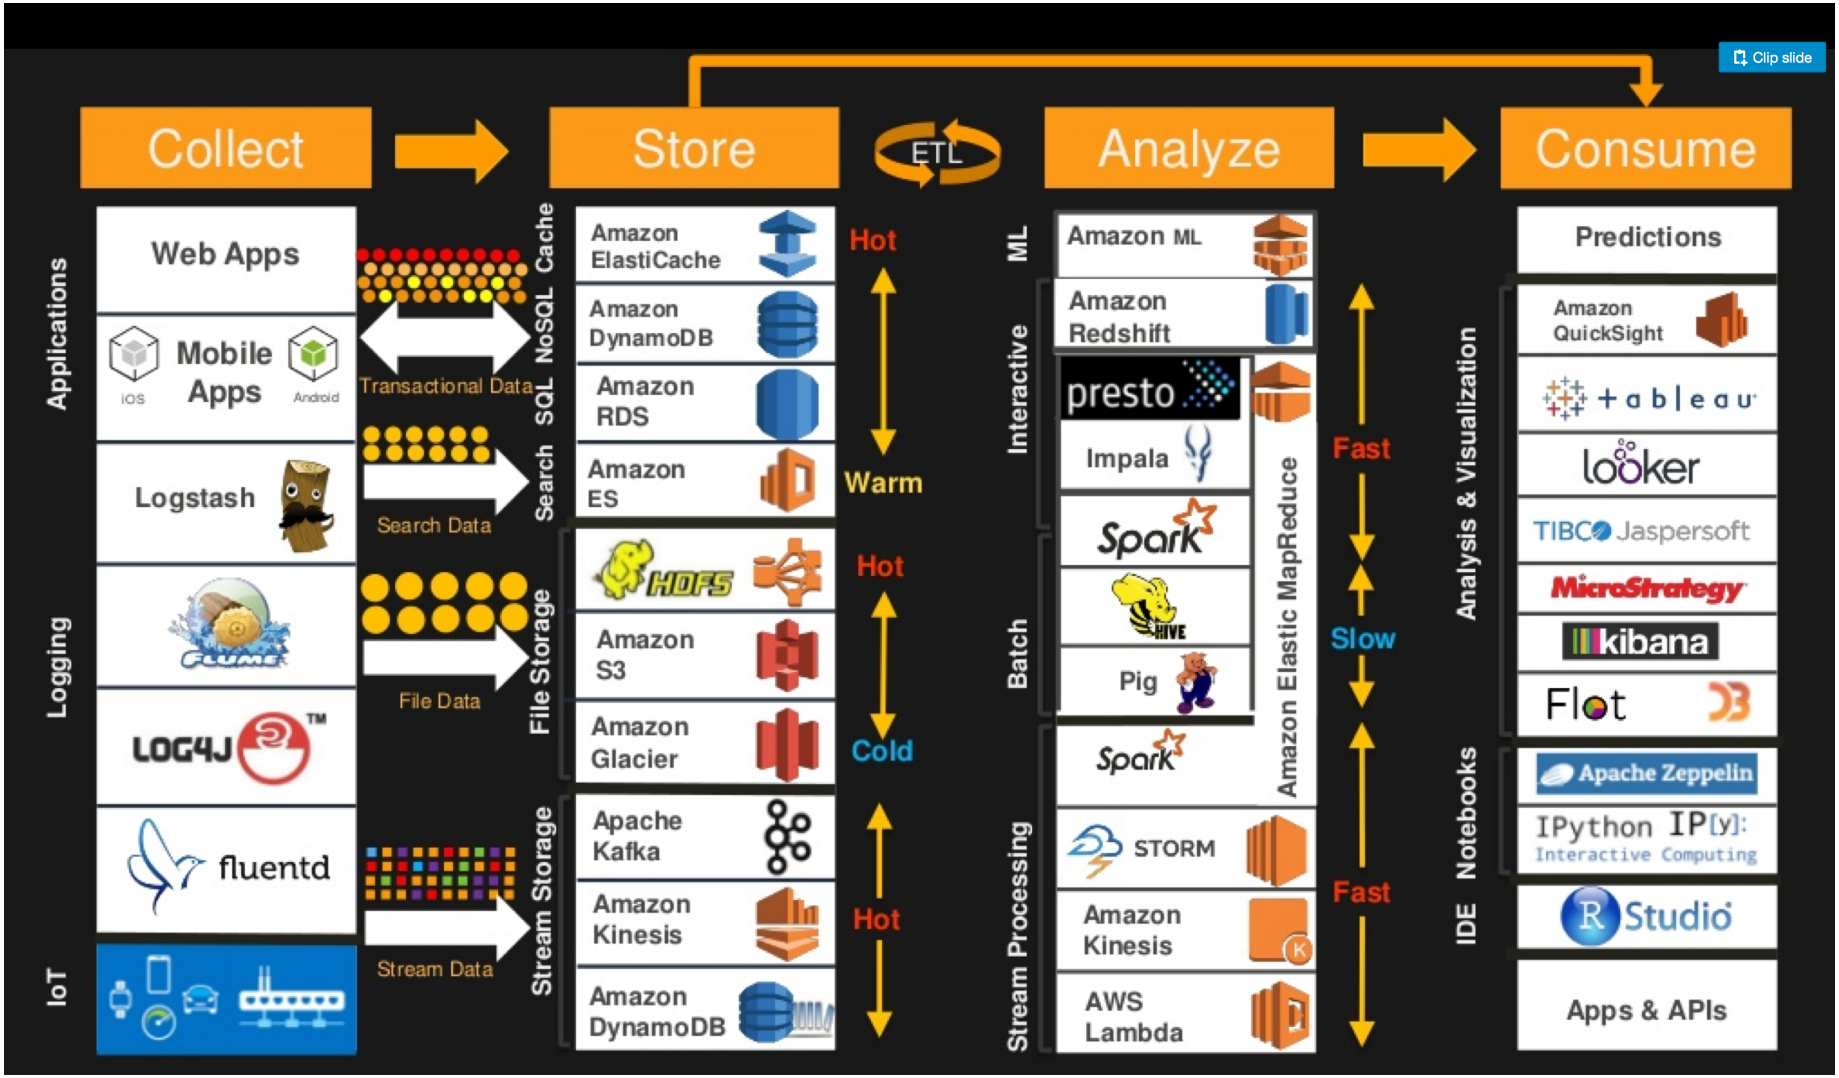

Big Data: Cloud Hosting

Big Data: Hadoop ecosystem providers

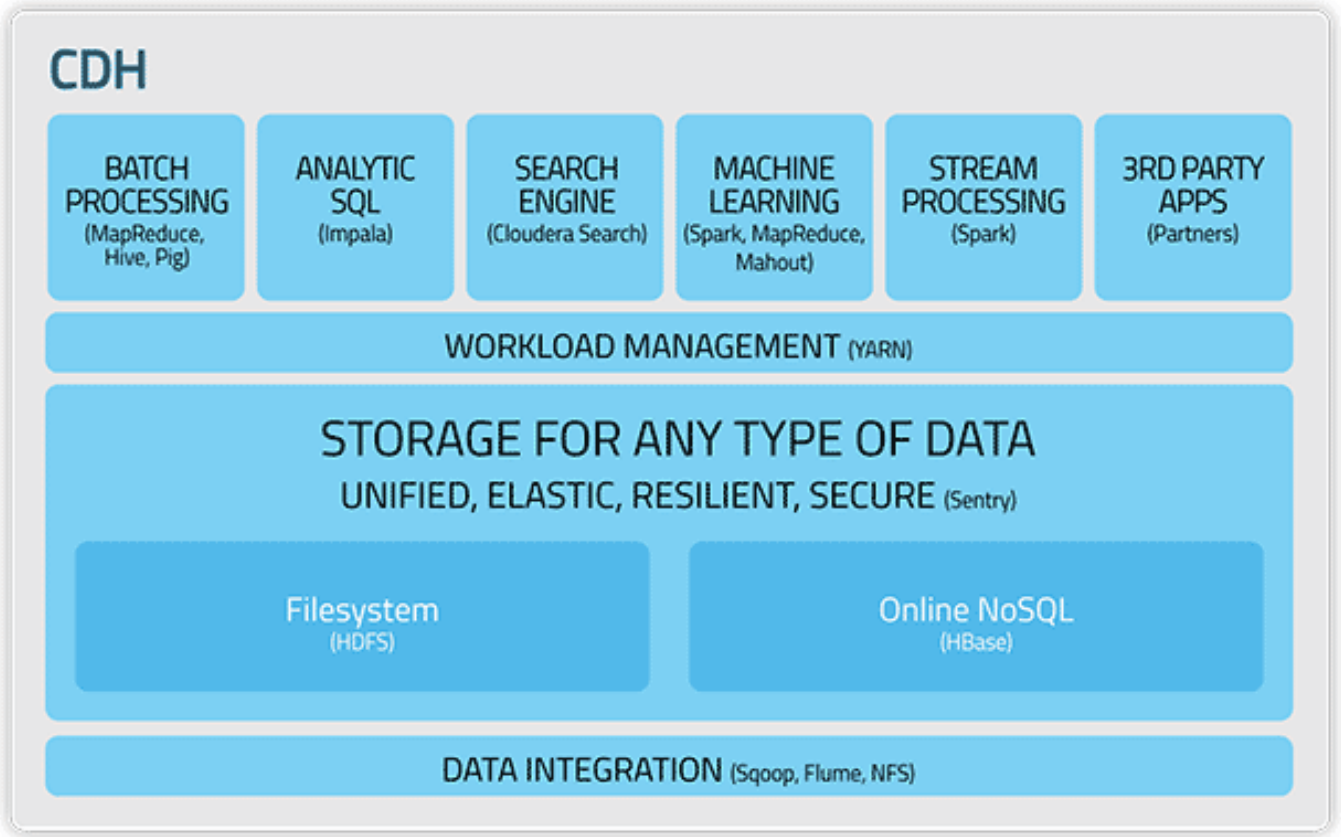

Big Data: Hadoop ecosystem providers

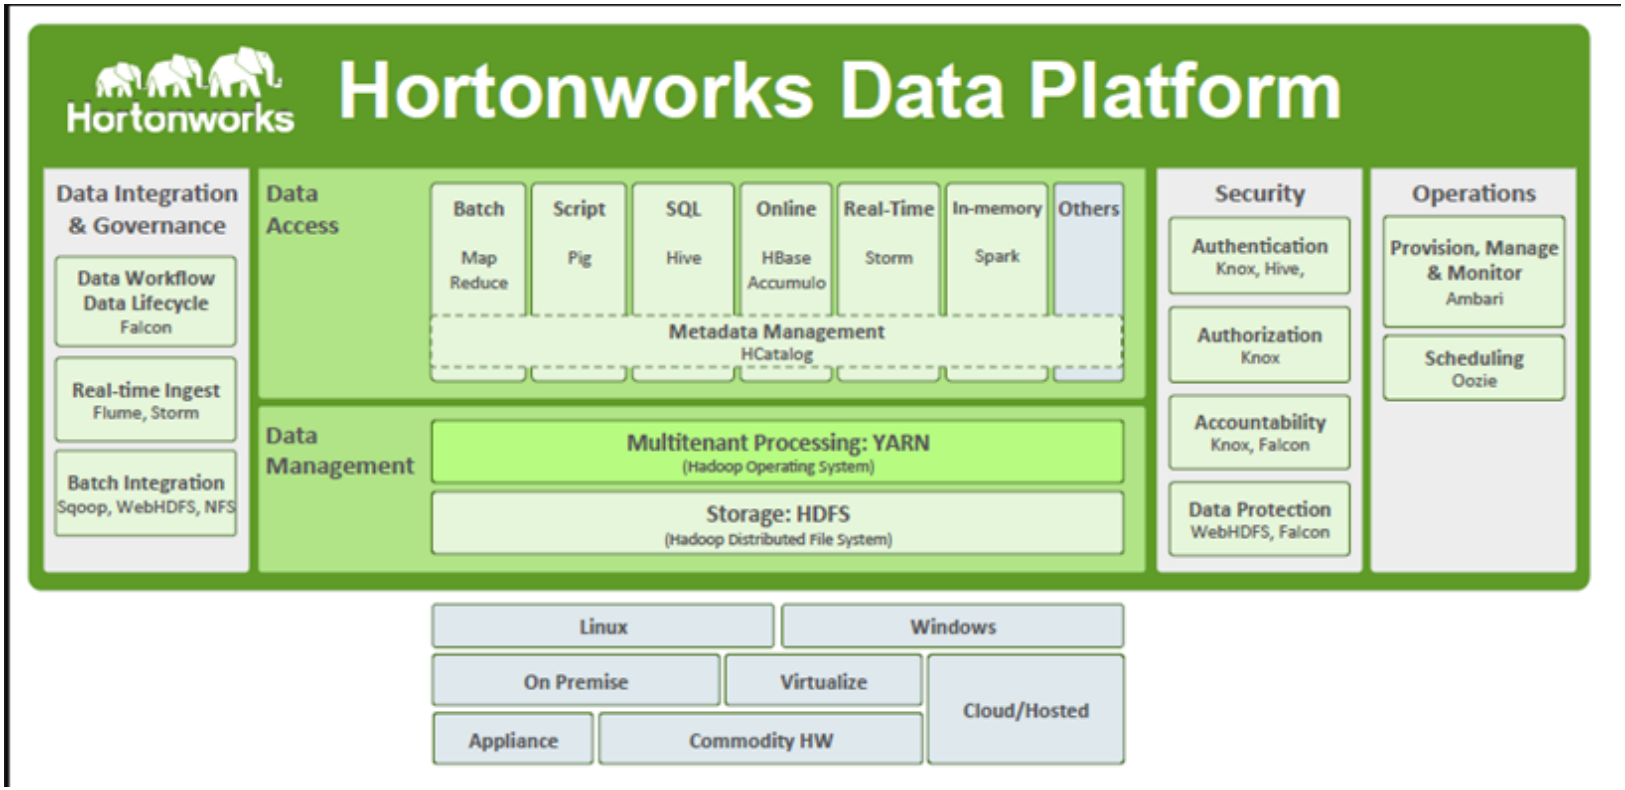

Big Data Mining: the unicorn