Distributed Calculus with Hadoop MapReduce inside Orange Search Engine

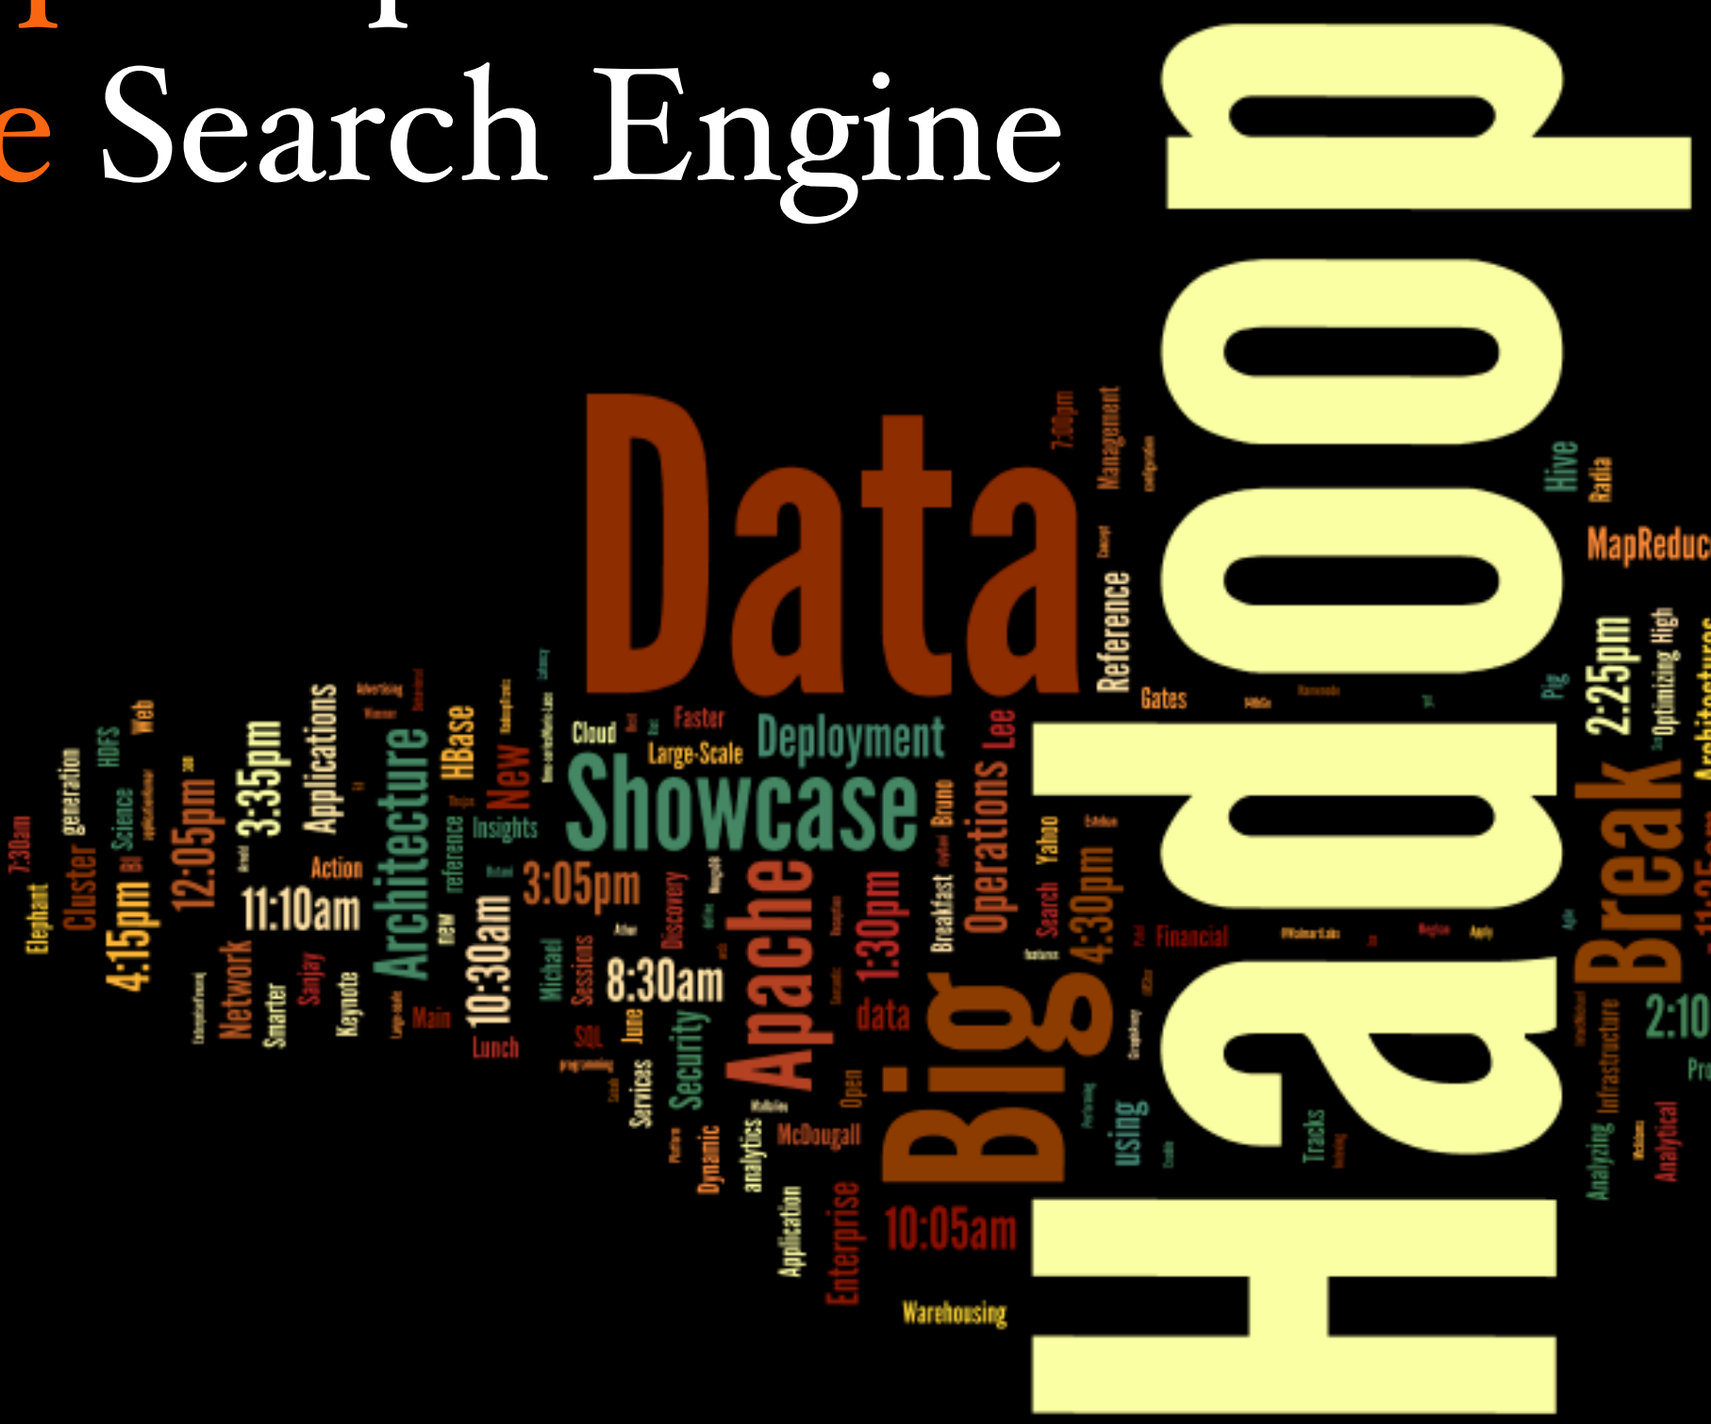

SophiaConf2012

olivier.varene@orange.com

What is Big Data ?

olivier.varene@orange.com

$ 5 billions (2012)

to

$ 50 billions

(by 2017)

Forbes

olivier.varene@orange.com

«Big Data is the new definitive source of competitive advantage across all industries»

Jeff Kelly

olivier.varene@orange.com

Product success

«The days are over when you build a product once and it just works.

You have to take ideas, test them, iterate them, use data and analytics to understand what works and what doesn't in order to be successful»

LivingSocial @PCWorld 2012

olivier.varene@orange.com

Big Data

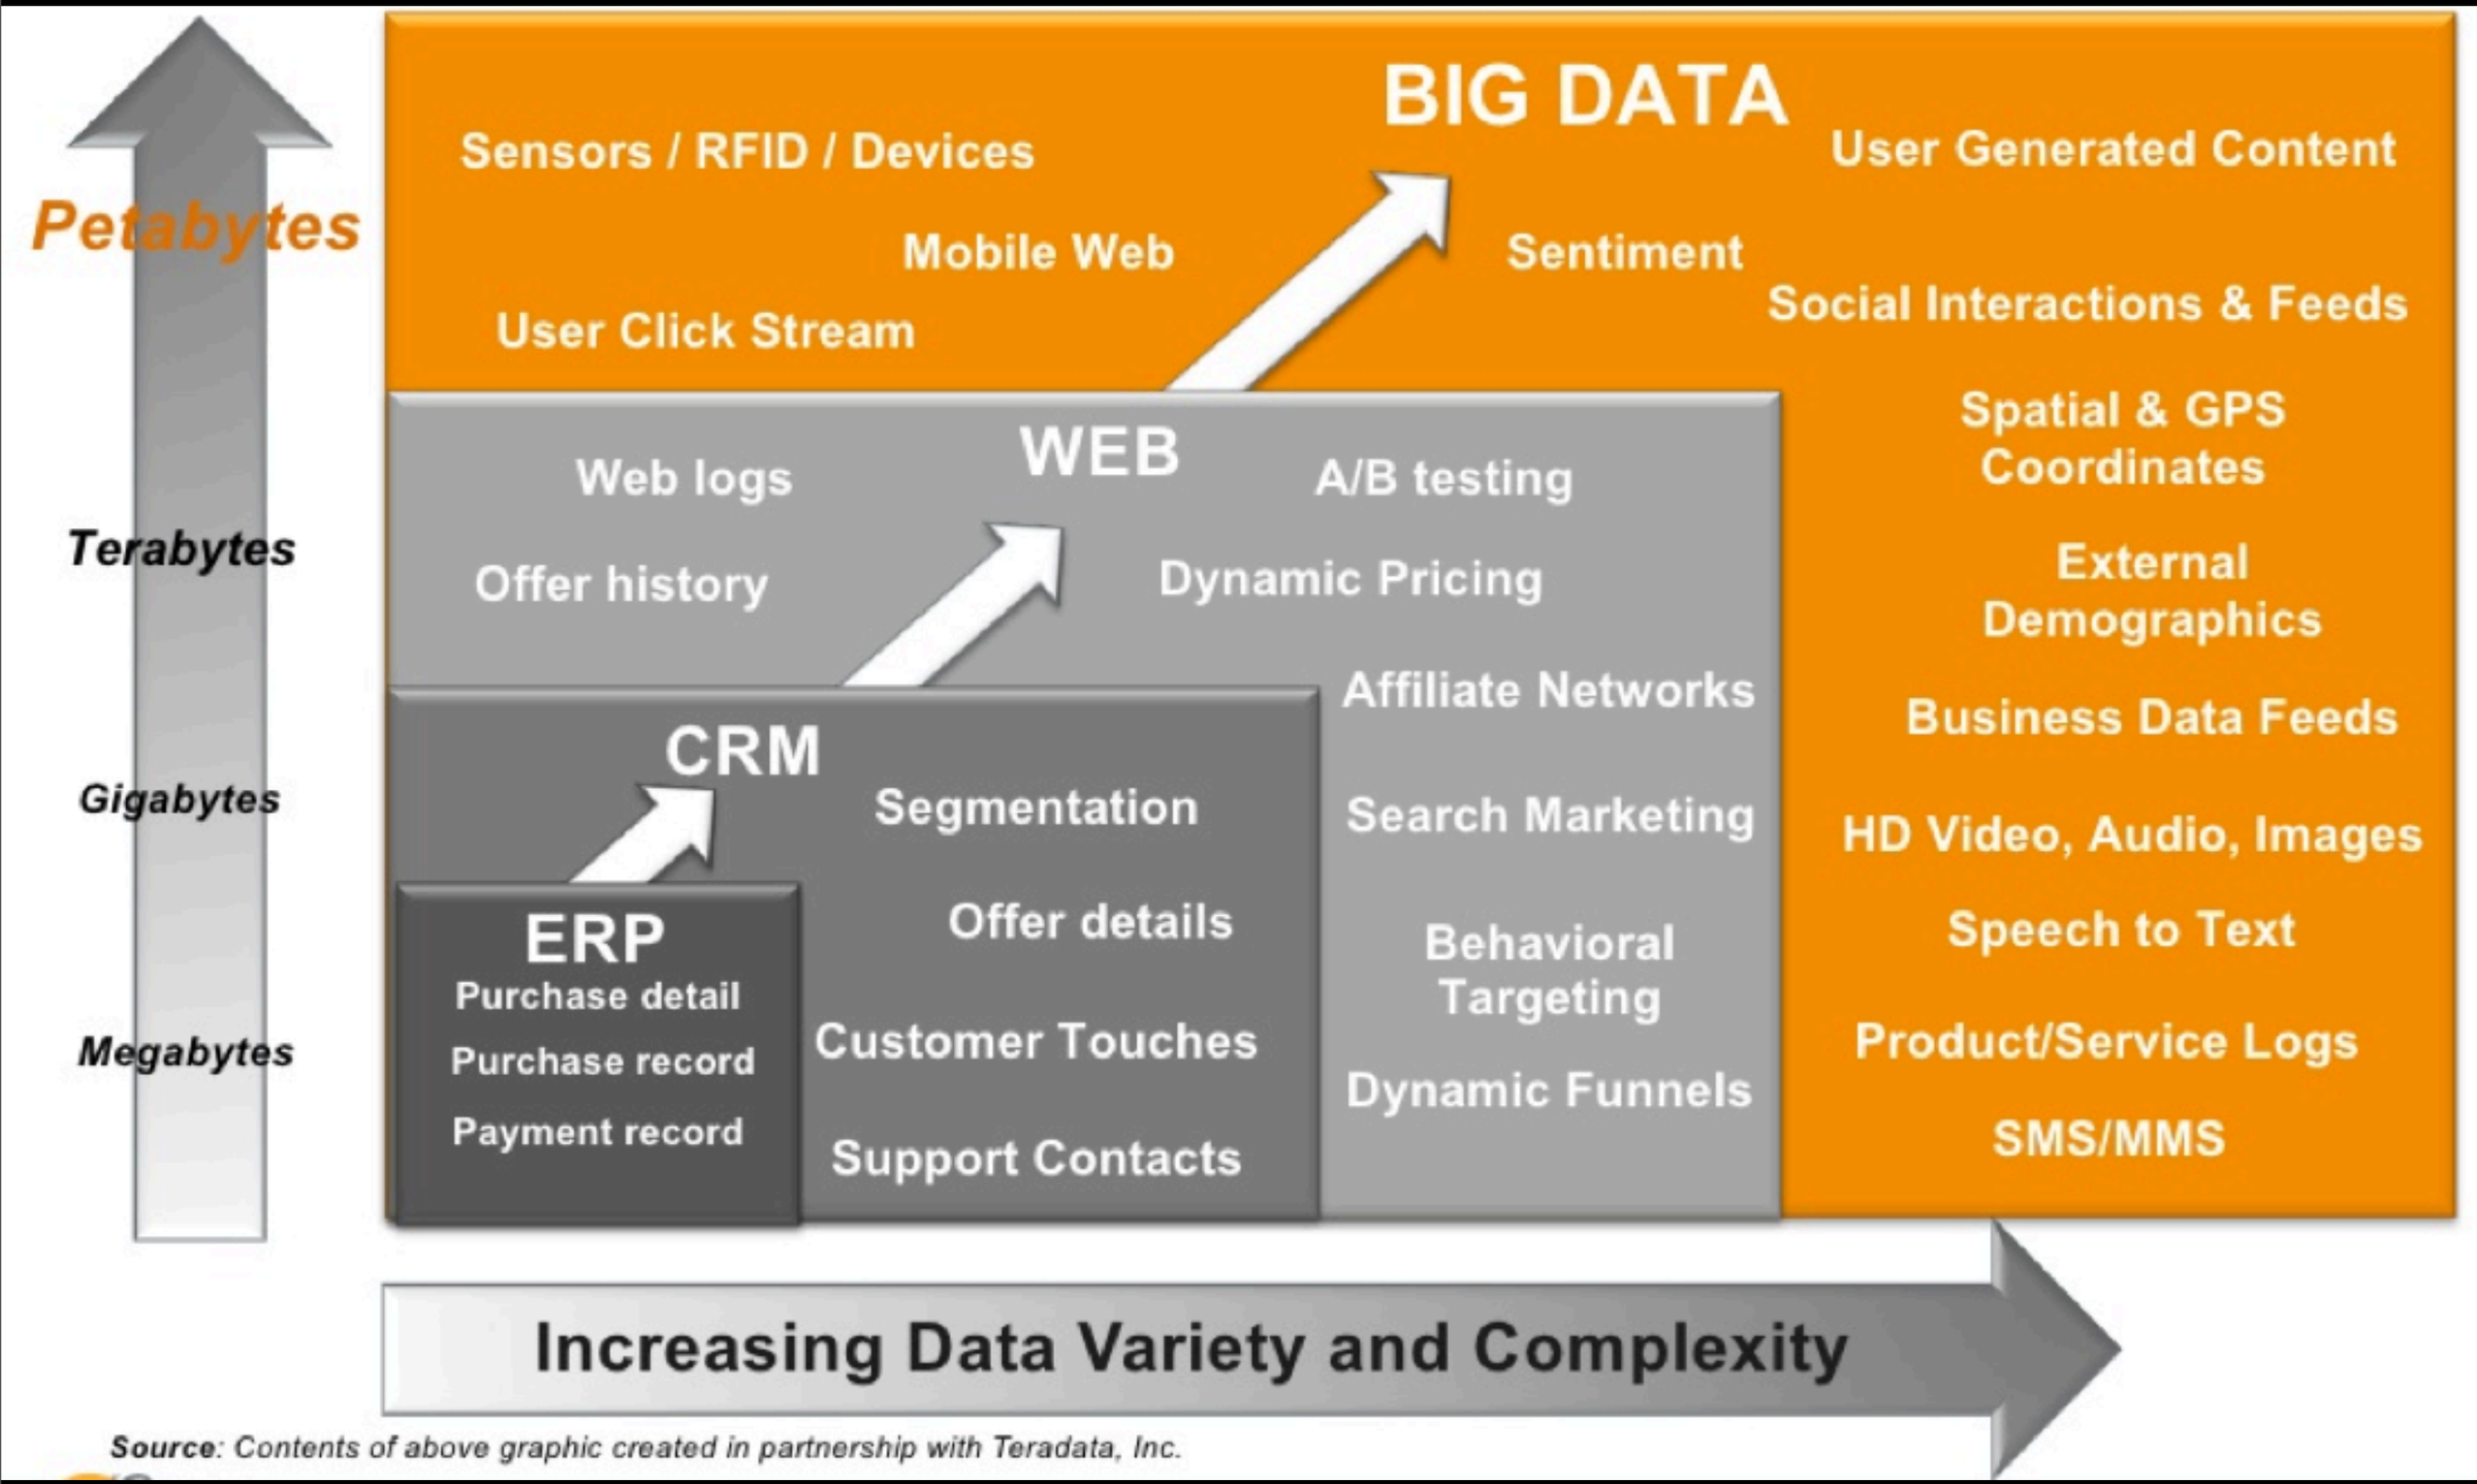

Source: *Contents of above graphic created in partnership with Teradata, Inc.*

What about Hadoop ?

olivier.varene@orange.com

Beliefs

« We believe that by 2015, more than half the world's data will be processed by Apache Hadoop »

HortonWorks @HadoopSummit 2012

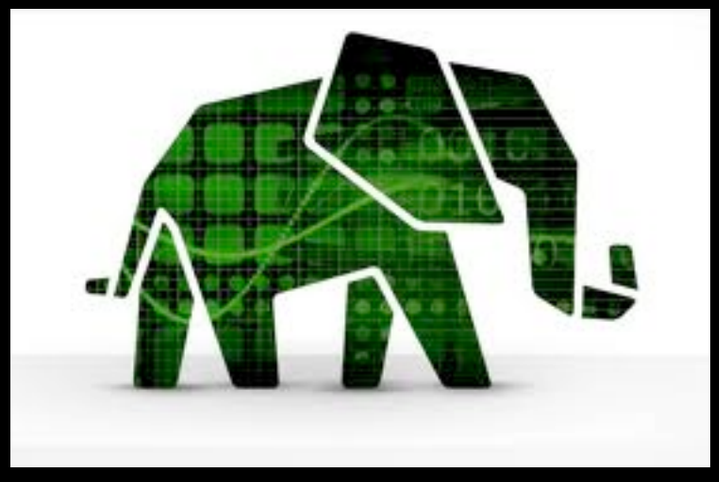

olivier.varene@orange.com

Actors

olivier.varene@orange.com

Hadoop eco-system

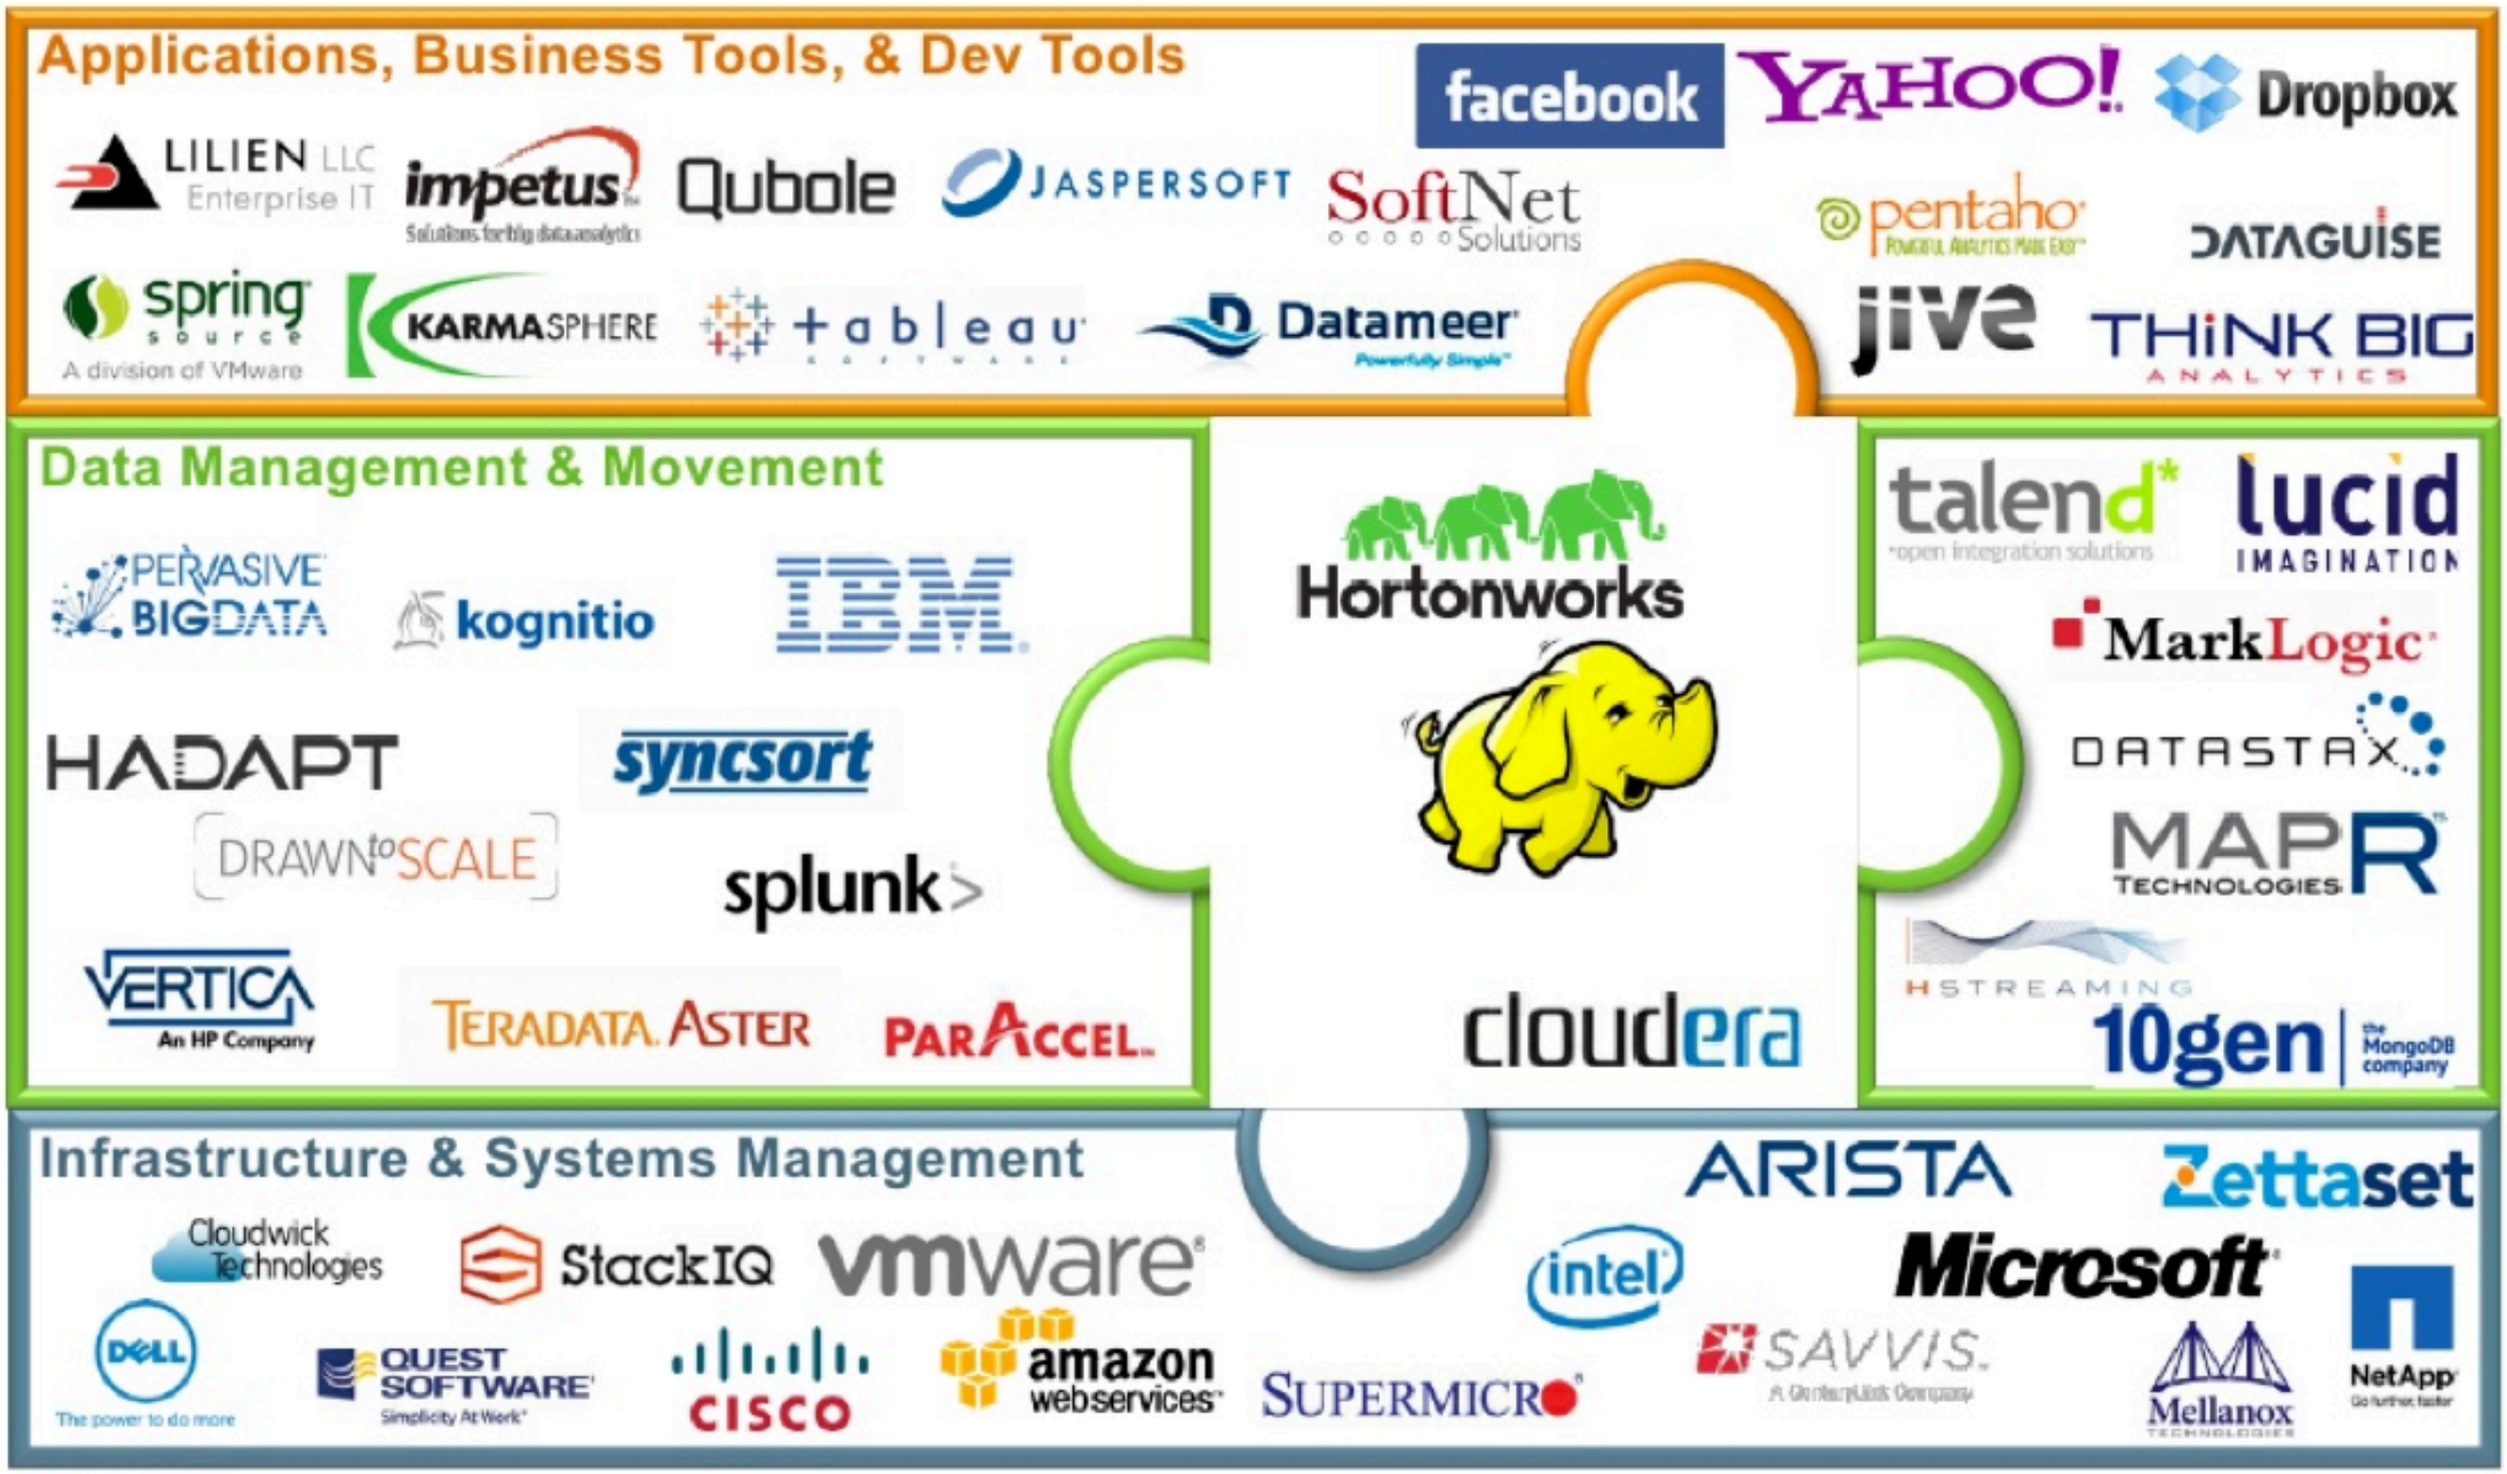

olivier.varene@orange.com

Context at Orange

(more than 2 years ago)

olivier.varene@orange.com

Orange Search Engine

http://www.orange.fr

http://www.lemoteur.fr

http://www.voila.fr

olivier.varene@orange.com

Search Engine Architecture

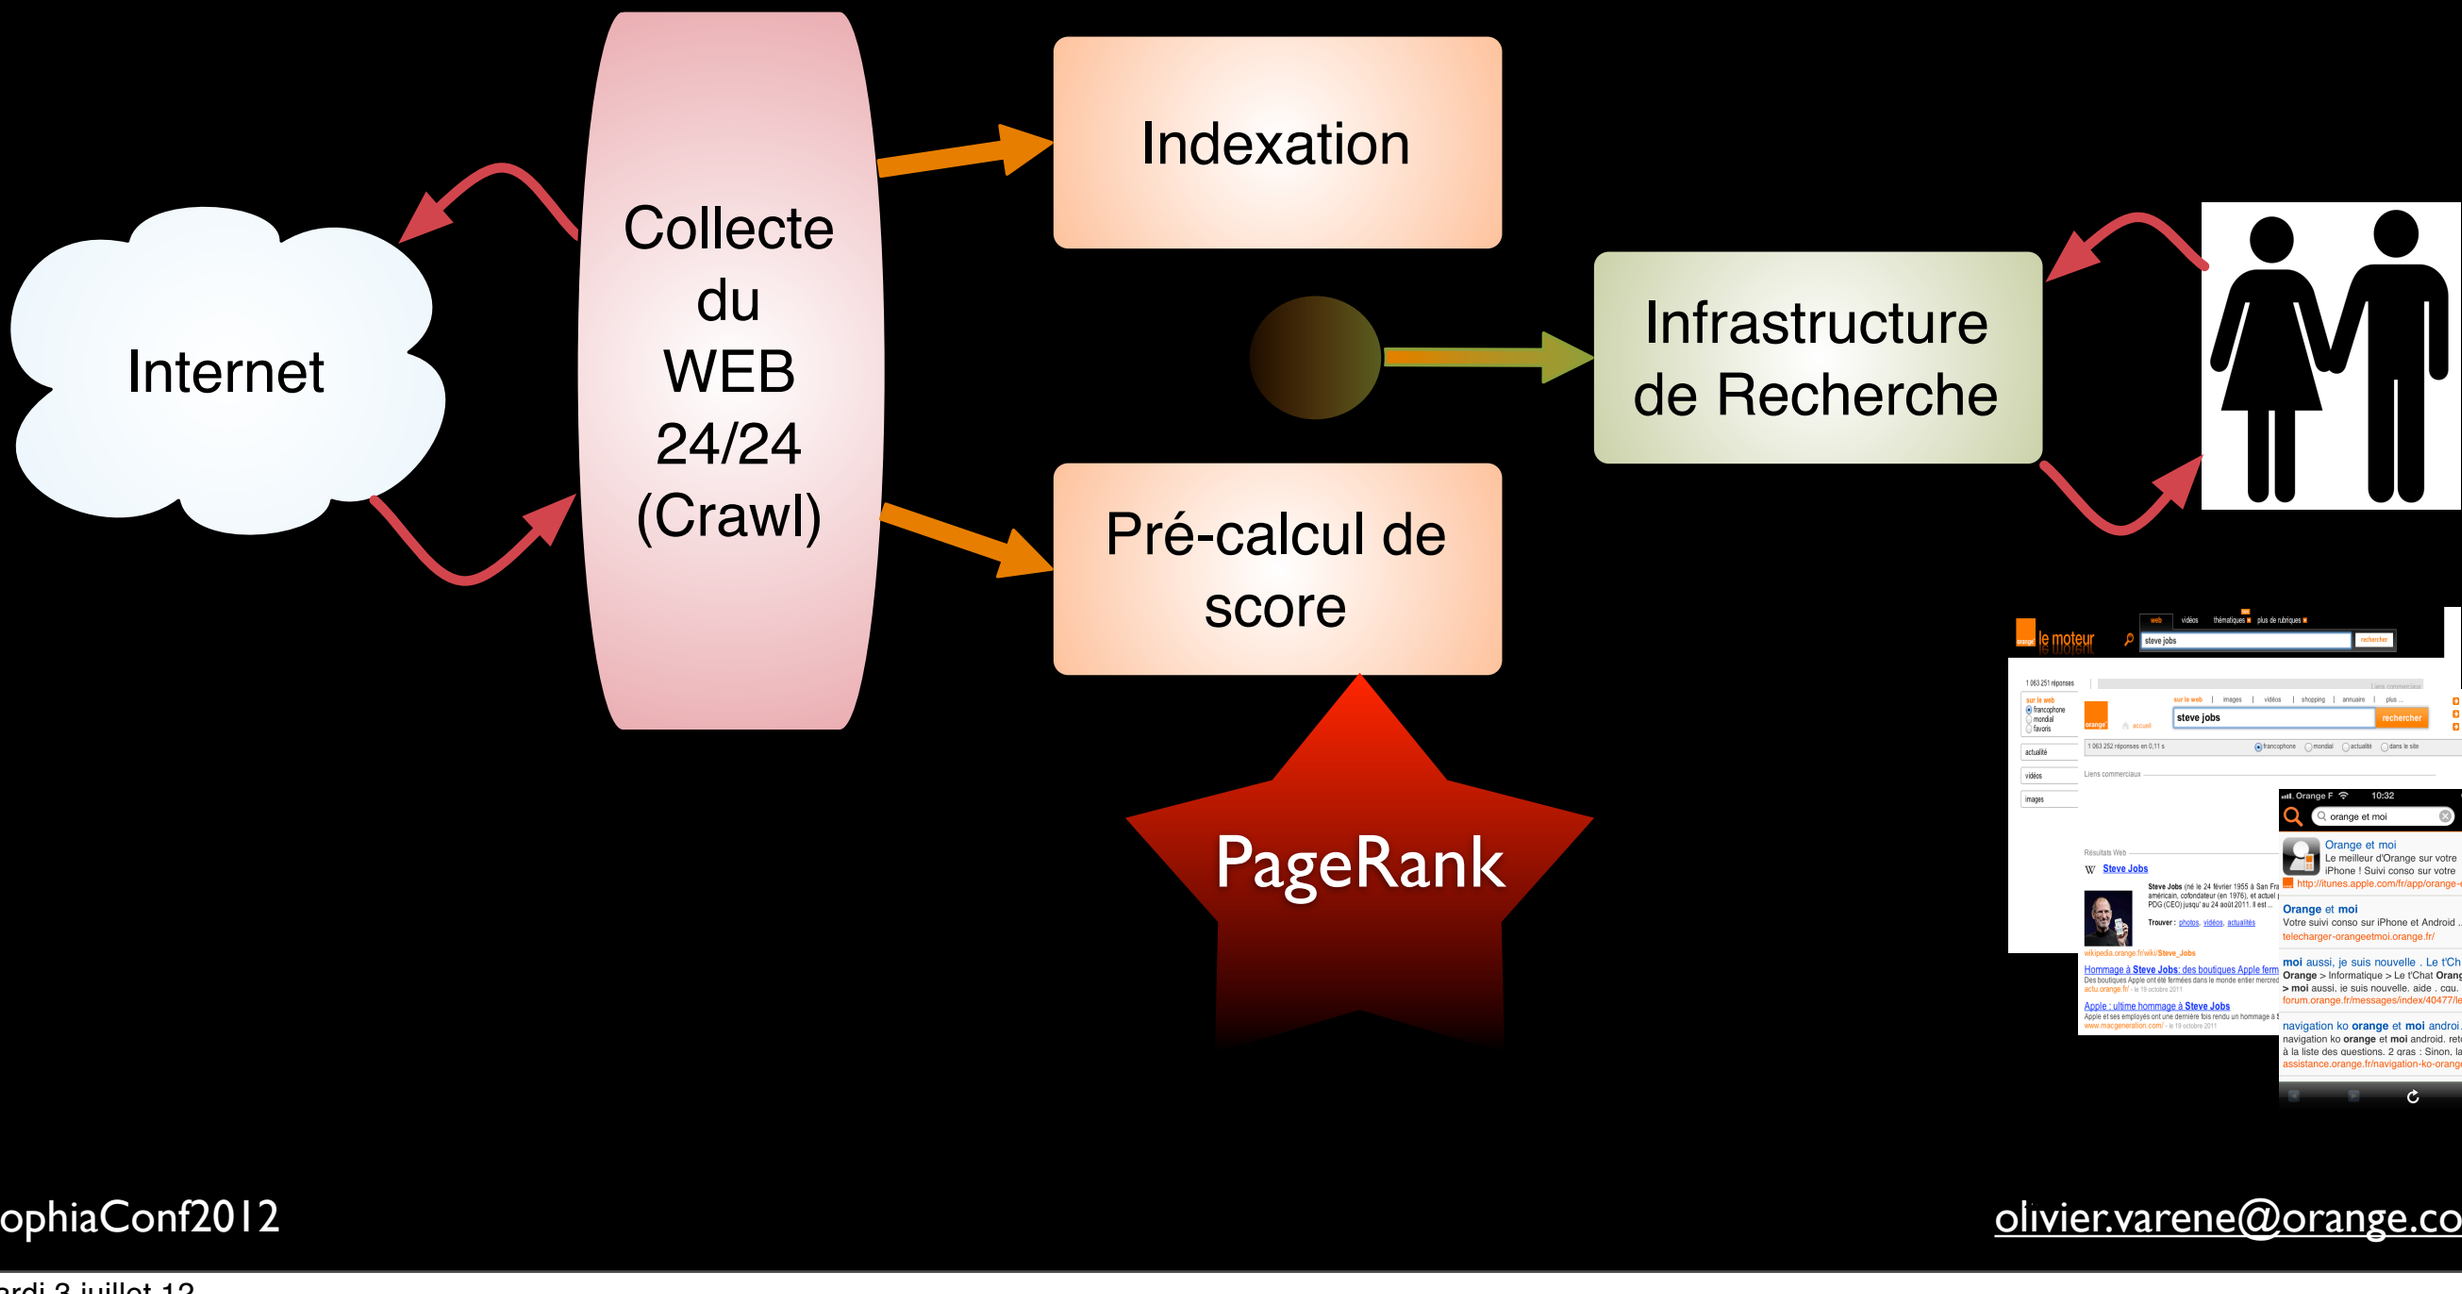

olivier.varene@orange.com

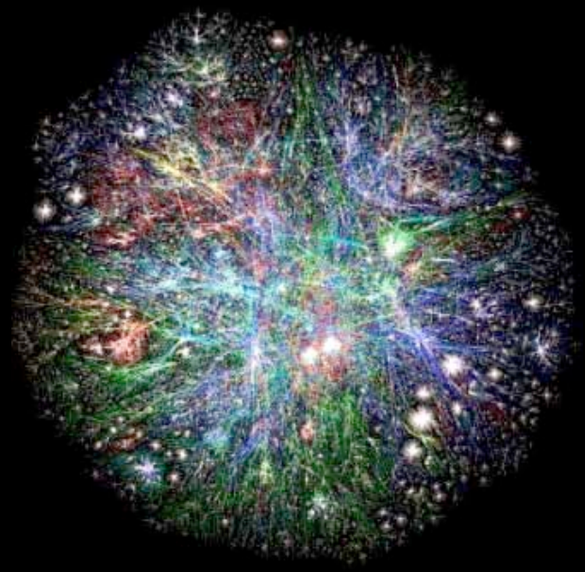

Main issue

- PageRank calculus on billions nodes and 10s billions edges
- regularly failed ! (hardware ...)
- 4 to 8 weeks calculus
- unscalable
- failure rate aroud 80%
- One person full time to supervise

olivier.varene@orange.com

Answer

olivier.varene@orange.com

PageRank portable to Hadoop / MapReduce ?

- Simple programing model:

  Map(in_k,in_v) => list(out_k,intermed_v)
  Reduce(out_k,intermed_v) => list(out_v)

- Scalable
- Batch Processing
- YES !

olivier.varene@orange.com

Hadoop Axioms

- System shall manage and heal himself
- Performance shall scale linearly
- Compute shall move to data
- Modular and extensible

olivier.varene@orange.com

Our install ?

olivier.varene@orange.com

Our install

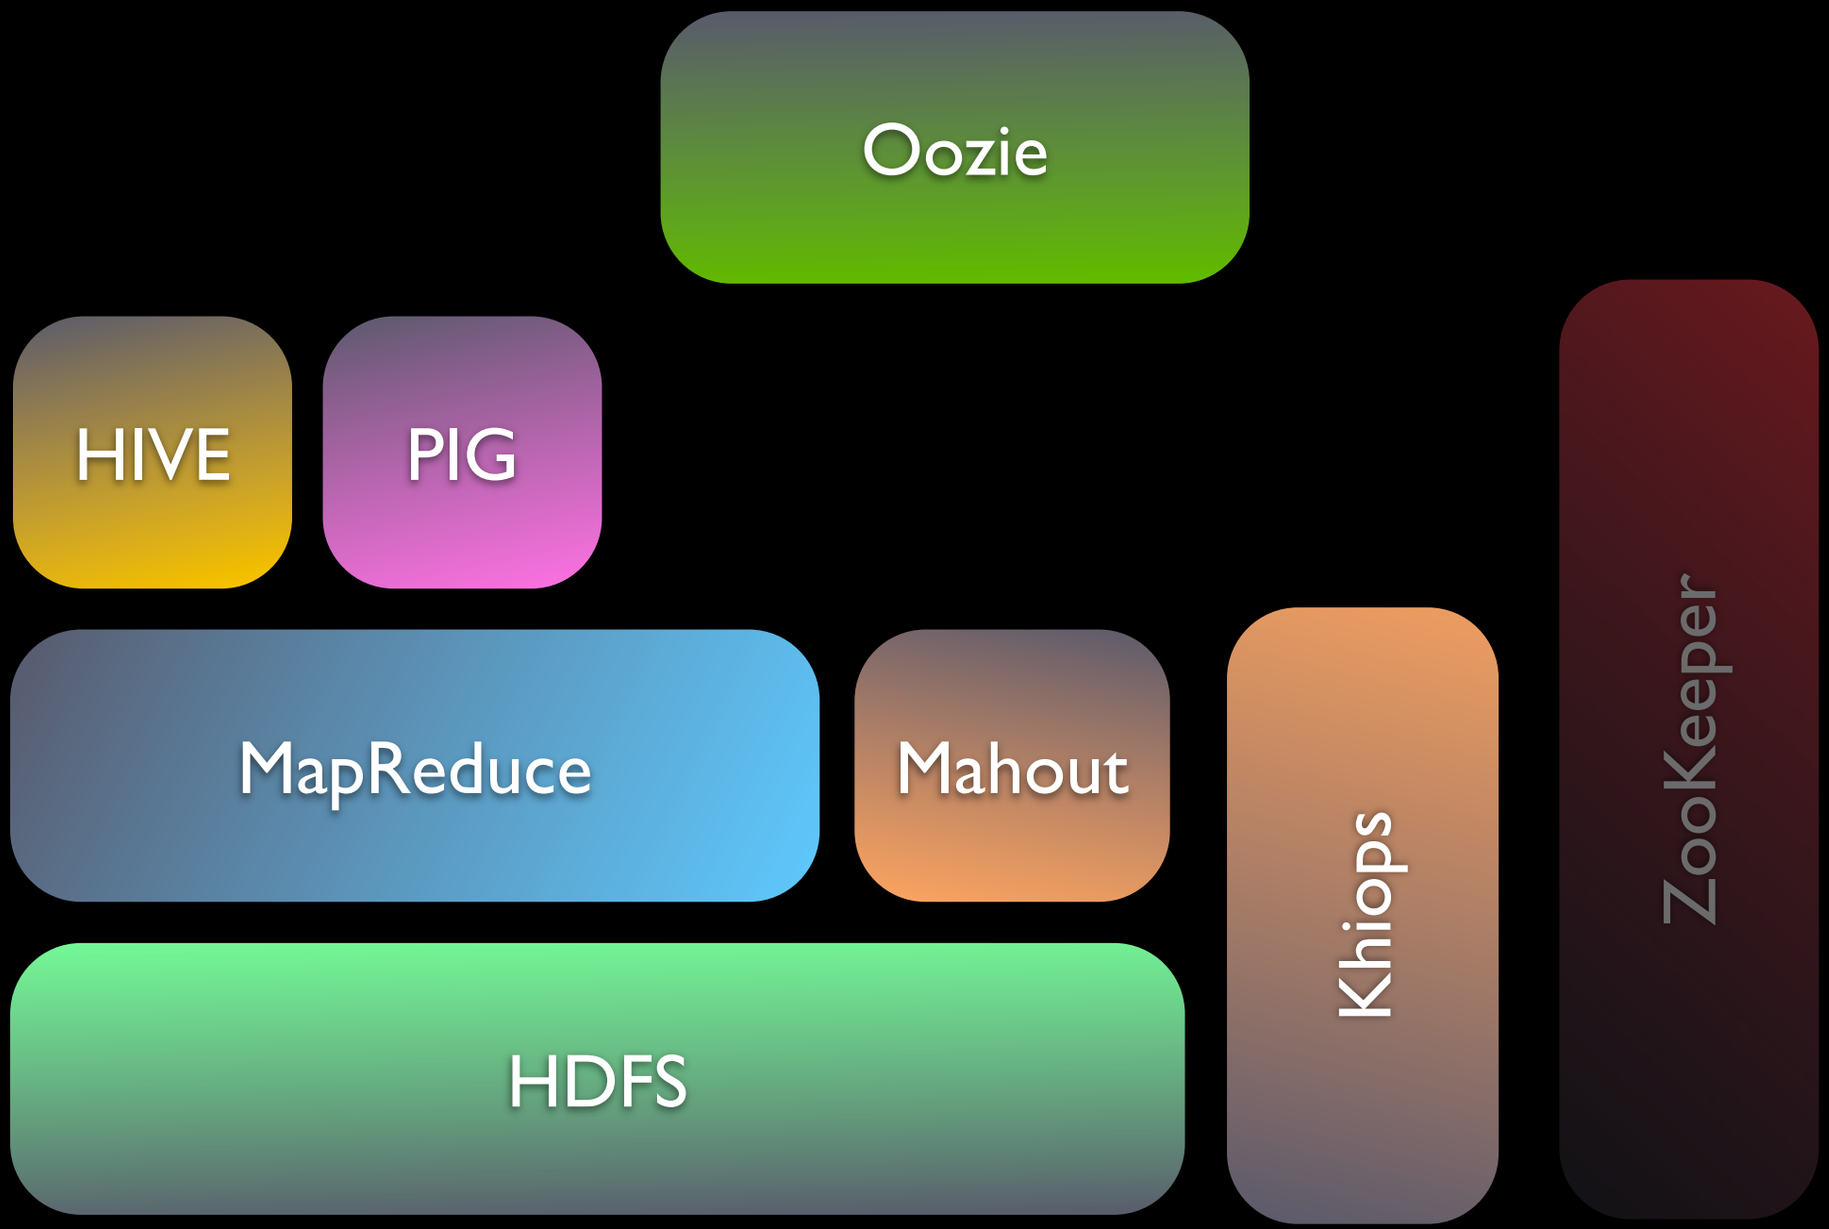

olivier.varene@orange.com

HDFS ?

olivier.varene@orange.com

Hadoop - HDFS

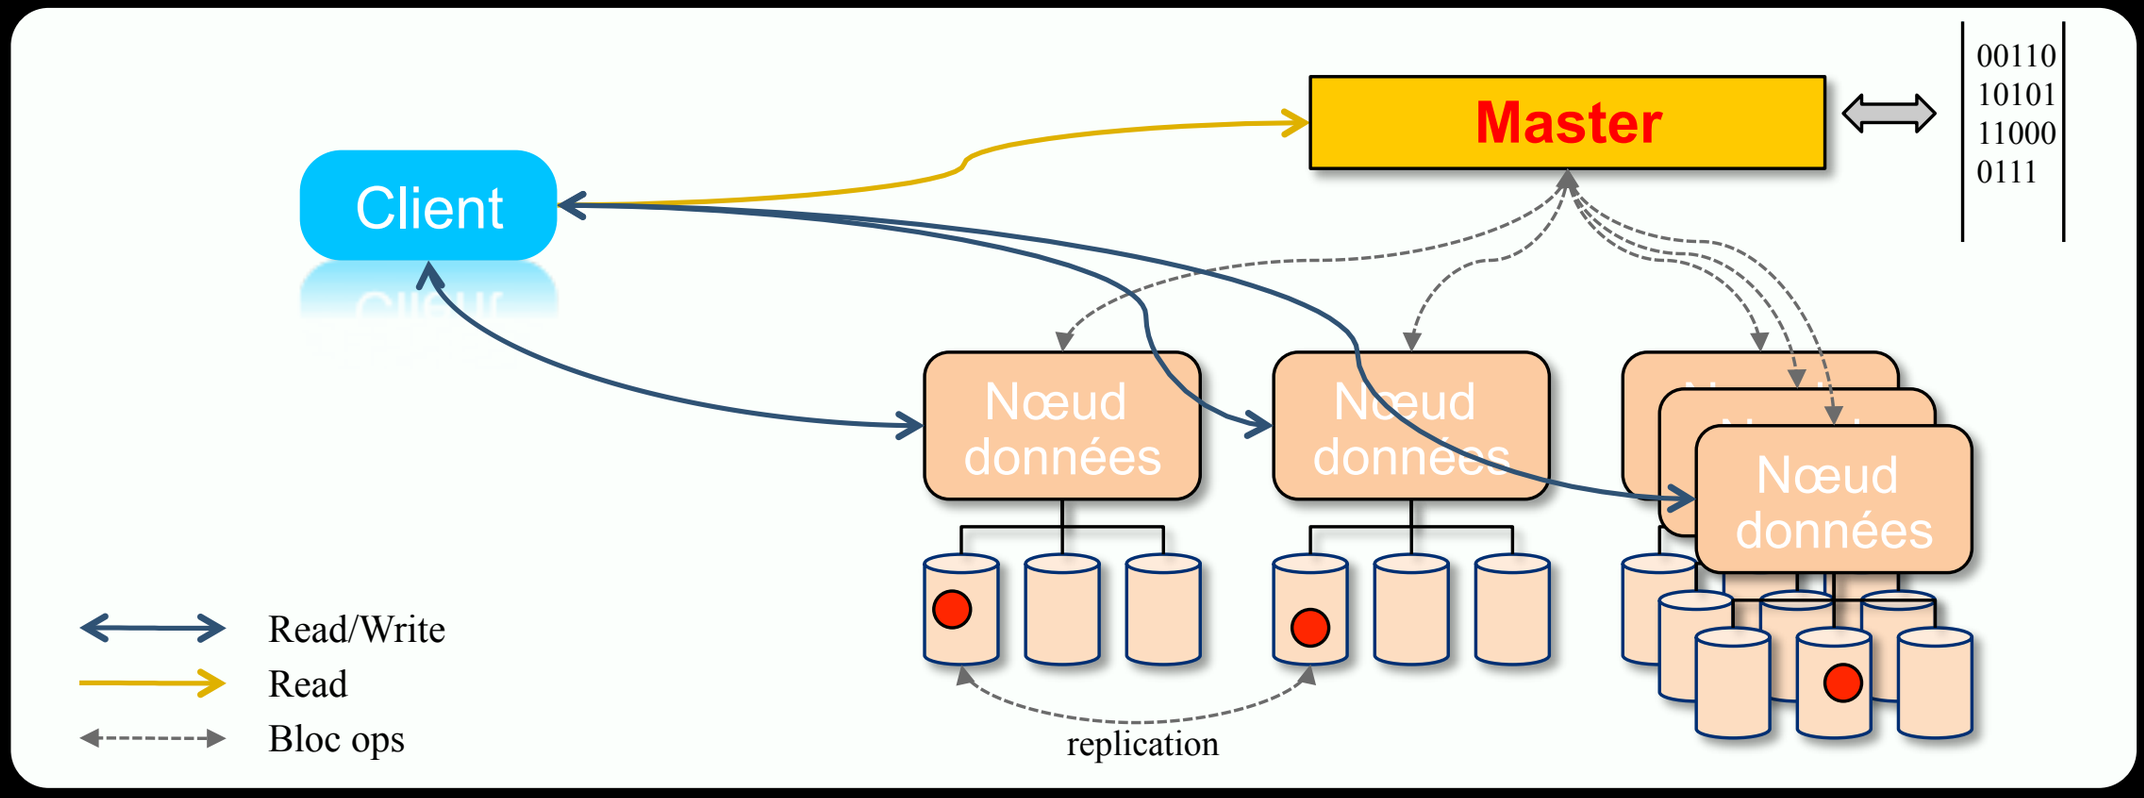

COTS - replication - big blocks
maximize throughput - Metadata in RAM

olivier.varene@orange.com

Map Reduce ?

olivier.varene@orange.com

MapReduce

cat | «your map» | sort -u | «your reduce»

Programming paradigm

olivier.varene@orange.com

MapReduce

cat | «your map» | sort -u | «your reduce»

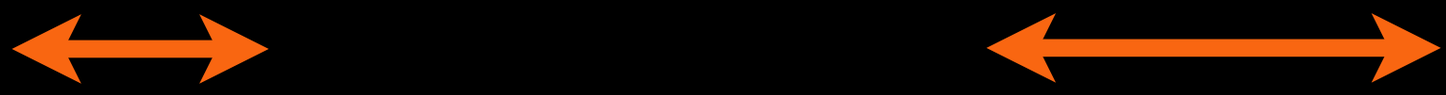

FrameWork

olivier.varene@orange.com

MapReduce

cat | «your map» | sort -u | «your reduce»

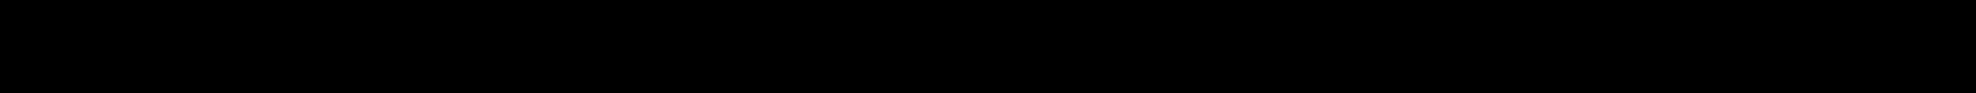

your Job

olivier.varene@orange.com

MapReduce Insides

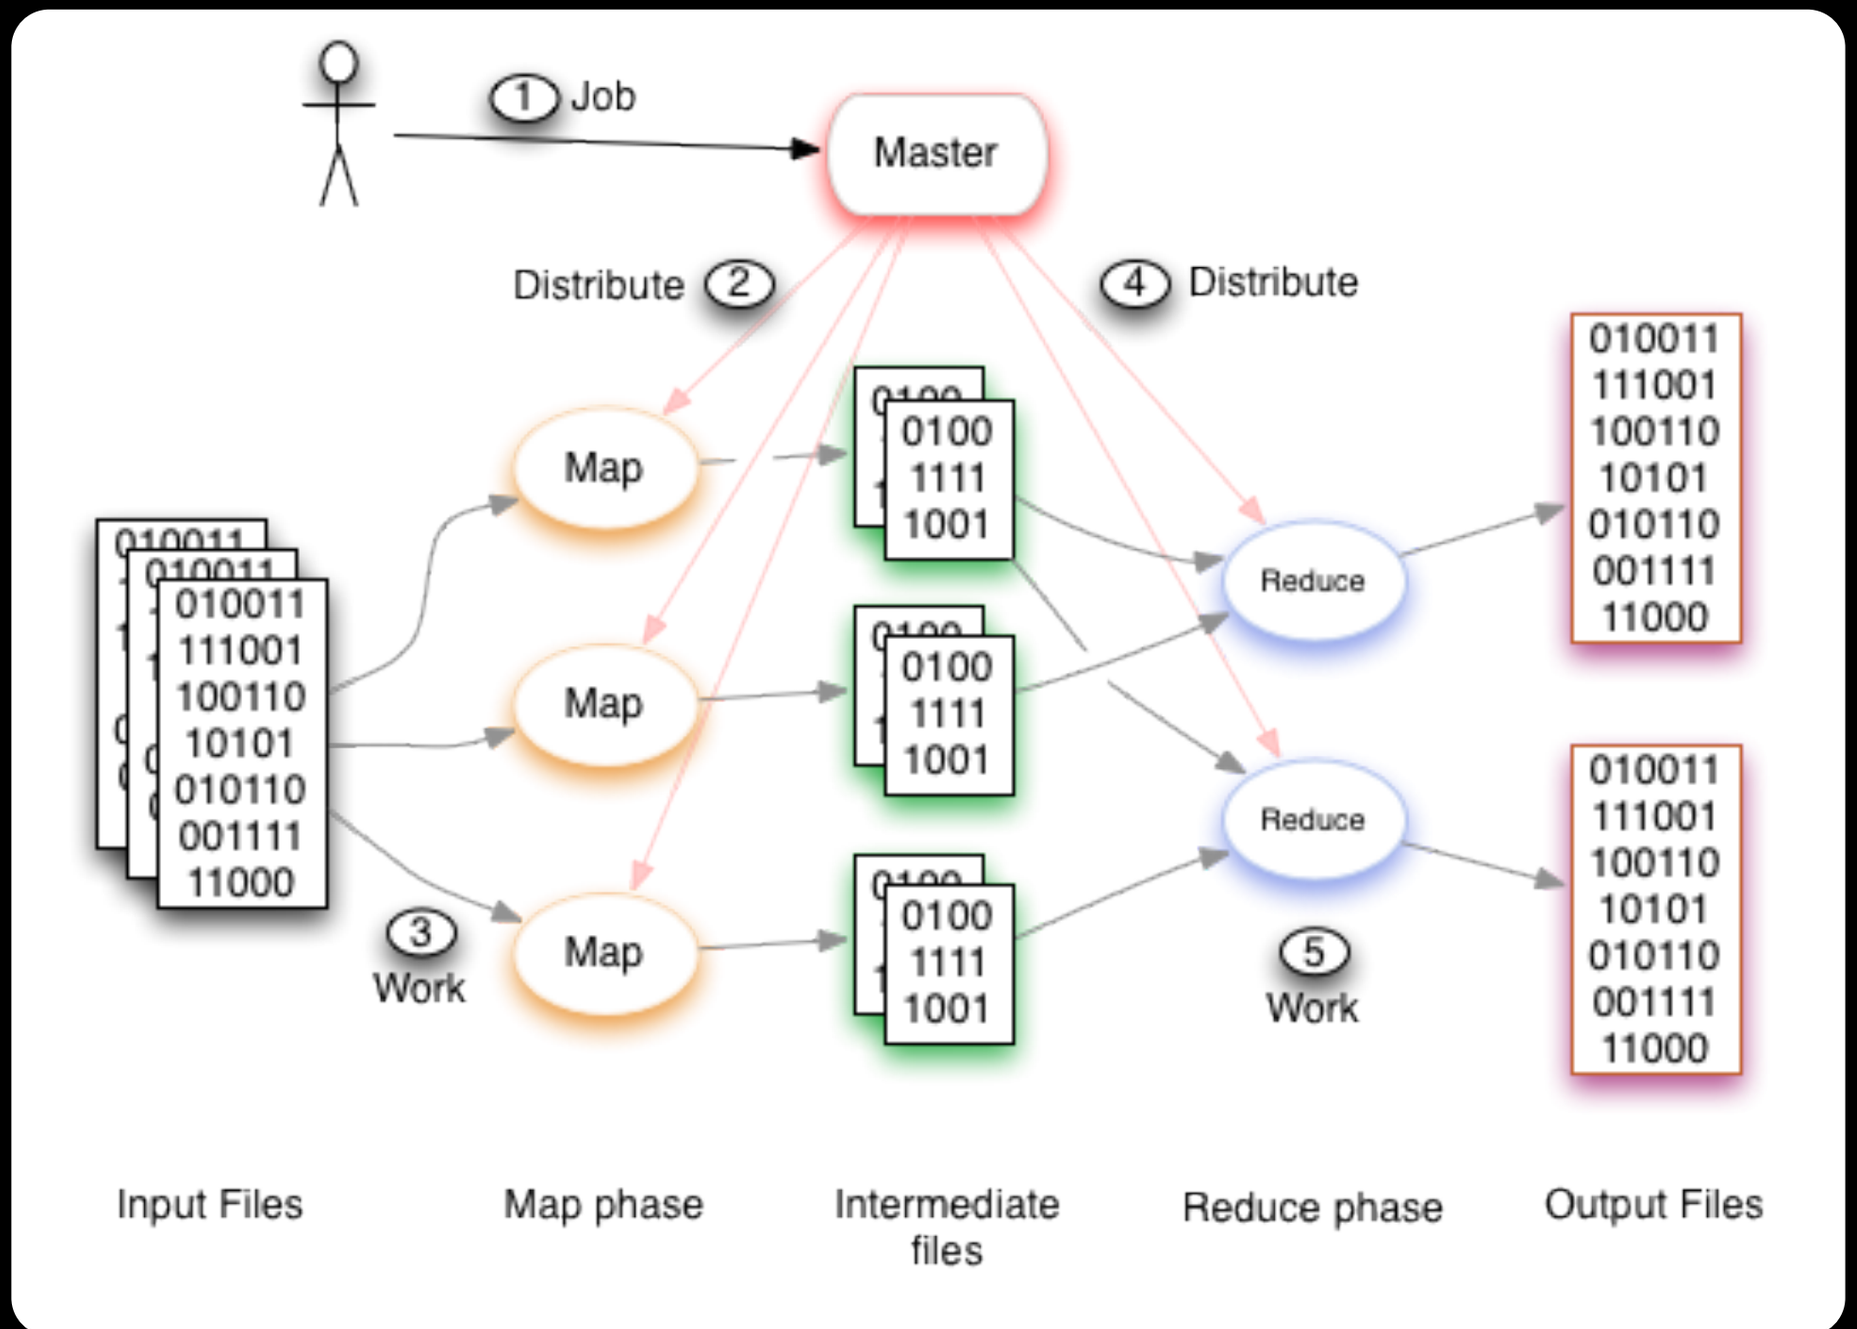

olivier.varene@orange.com

Interfaces

- Java API
- Pipes
- Streaming (python, perl, C/C++, ...)

olivier.varene@orange.com

PIG

- High level data analysis script language
- extensible via UDF
- Structure of a Pig script
  - load
  - filter
  - foreach | group by | join | your functions
  - order
  - store

olivier.varene@orange.com

HIVE

- High level SQL-like query and analysis language
- extensible via UDF
- Structure of a Hive script
  - create table
  - load data
  - select ... from ...
  - insert | group by | join

olivier.varene@orange.com

Application domains ?

olivier.varene@orange.com

Projects

- Scoring
- User profiling
- Log analysis and statistics
- ... and many others to come

olivier.varene@orange.com

ROI ?

olivier.varene@orange.com

ROI

- Lines Of Code **10X gain**
- Development Time **2X gain**
- IT cost **4X gain**

less bug, automatic, scalable ...

olivier.varene@orange.com

Perfect World ?

★ YES

- Run cost
- Development cost
- Scalable
- Stable
- Heterogenity

★ NO

- SPOF (almost solved)
- Fastidious debugging
- Localy non optimum
- mono-site

olivier.varene@orange.com

Thank you

Olivier Varene - Orange

olivier.varene@orange.com

@VareneO

thanks to

- Hadoop - Apache (http://hadoop.apache.org/)
- Khiops - Orange (http://www.khiops.com)
- Shauwn Connoly - HortonWorks (http://youtu.be/yPfysFAGv8s)
- Forbes - article (http://www.forbes.com/sites/siliconangle/2012/02/17/big-data-is-big-market-big-business/)
- Living Social (sentence)
- Terradata (Volumetry Graph)
- http://www.wordle.net/ (Words Cloud)

olivier.varene@orange.com