QCon

Please evaluate my talk via the mobile app!

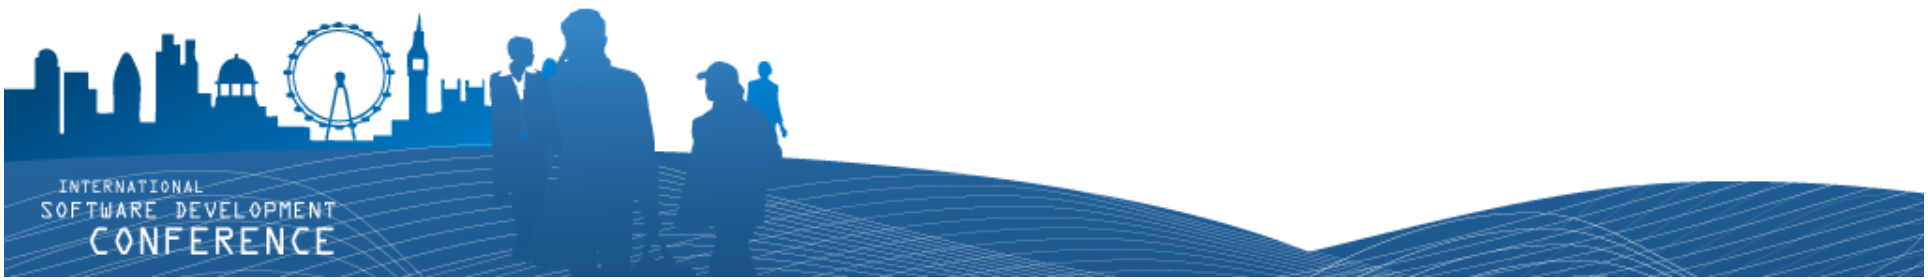

Hortonworks

Next Gen Hadoop

Gather around the campfire and I will tell you a good YARN

Akmal B. Chaudhri*
Hortonworks

*about.me/akmalchaudhri

My background

- **~25 years experience in IT**
  - Developer (Reuters)
  - Academic (City University)
  - Consultant (Logica)
  - Technical Architect (CA)
  - Senior Architect (Informix)
  - Senior IT Specialist (IBM)
  - TI (Hortonworks)
- **Worked with various technologies**
  - Programming languages
  - IDE
  - Database Systems
- **Client-facing roles**
  - Developers
  - Senior executives
  - Journalists
- **Broad industry experience**
- **Community outreach**
- **University relations**
- **10 books, many presentations**

Hortonworks

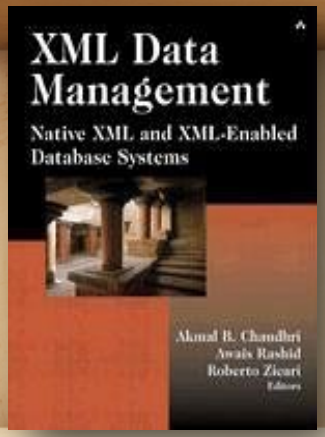
XML Data Management
Native XML and XML-Enabled Database Systems
Akmal B. Chaudhri
Awais Rashid
Roberto Zicari
Editors

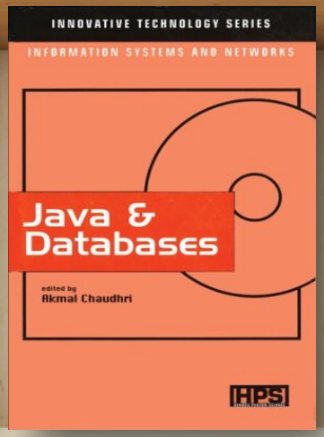
INNOVATIVE TECHNOLOGY SERIES
INFORMATION SYSTEMS AND NETWORKS
Java & Databases
edited by
Akmal Chaudhri
HPS

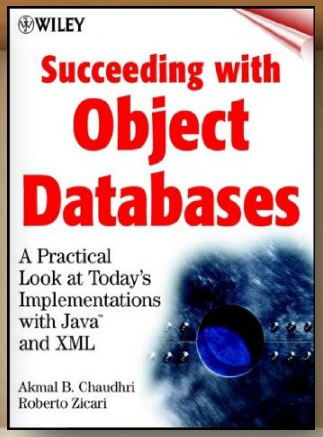
WILEY
Succeeding with Object Databases
A Practical Look at Today's Implementations with Java™ and XML
Akmal B. Chaudhri
Roberto Zicari

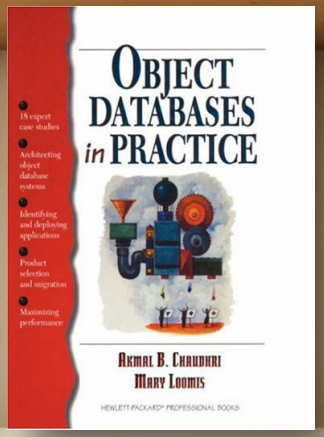
18 expert case studies
Architecting object database systems
Identifying and deploying applications
Product selection and migration
Maximizing performance
OBJECT DATABASES in PRACTICE
AKMAL B. CHAUDHRI
MARY LOOMIS
HEWLETT-PACKARD® PROFESSIONAL BOOKS

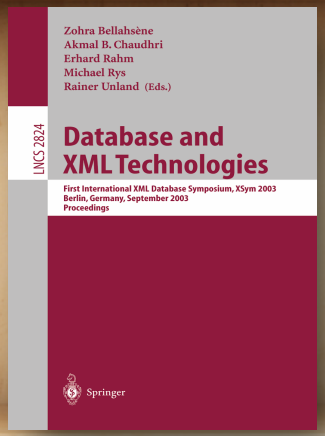
Zohra Bellahsène
Akmal B. Chaudhri
Erhard Rahm
Michael Rys
Rainer Unland (Eds.)
LNCS 2824
Database and XML Technologies
First International XML Database Symposium, XSym 2003
Berlin, Germany, September 2003
Proceedings
Springer

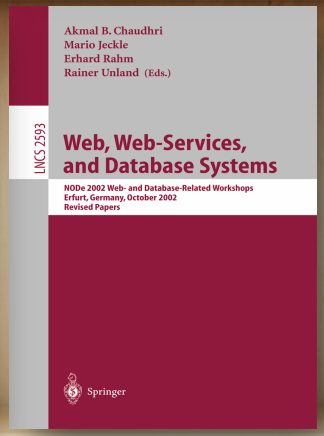
Akmal B. Chaudhri
Mario Jeckle
Erhard Rahm
Rainer Unland (Eds.)
LNCS 2593
Web, Web-Services, and Database Systems
NODe 2002 Web- and Database-Related Workshops
Erfurt, Germany, October 2002
Revised Papers
Springer

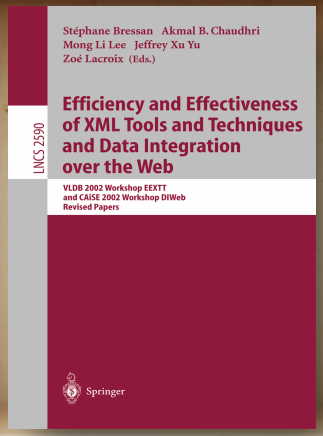
Stéphane Bressan Akmal B. Chaudhri
Mong Li Lee Jeffrey Xu Yu
Zoé Lacroix (Eds.)
LNCS 2590
Efficiency and Effectiveness of XML Tools and Techniques and Data Integration over the Web
VLDB 2002 Workshop EEXTT
and CAiSE 2002 Workshop DIWeb
Revised Papers
Springer

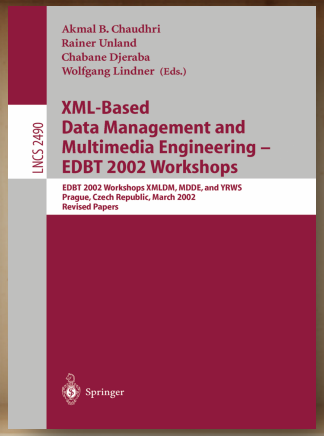
Akmal B. Chaudhri
Rainer Unland
Chabane Djeraba
Wolfgang Lindner (Eds.)
LNCS 2490
XML-Based Data Management and Multimedia Engineering – EDBT 2002 Workshops
EDBT 2002 Workshops XMLDM, MDDE, and YRWS
Prague, Czech Republic, March 2002
Revised Papers
Springer


XML Data Management
Native XML and XML-Enabled Database Systems
XML数据管理
纯XML和支持XML的数据库系统


事例で学ぶ
オブジェクトデータベース活用技法

Agenda

- Understanding Big Data
- Understanding Hadoop 2.0
- Hadoop Architecture Fundamentals
- The Future of Hadoop

Hortonworks

Understanding Big Data

What is Big Data?

1. In what timeframe do we now create the same amount of information that we created from the dawn of civilization until 2003?

   2 days

2. 90% of the world's data was created in the last (how many years)?

   2 years

3. What is 1024 petabytes also known as?

   An exabyte

4. What is the anticipated shortage in the U.S. of skilled workers with deep analytical skills by 2018?

   140,000 to 190,000

Source:
http://www.itbusinessedge.com/cm/blogs/lawson/just-the-stats-big-numbers-about-big-data/?cs=48051
http://techcrunch.com/2010/08/04/schmidt-data/

Hortonworks

The Three V's of Big Data

Variety	Unstructured and semi-structured data is becoming as strategic as the traditional structured data.
Volume	Data coming in from new sources as well as increased regulation in multiple areas means storing more data for longer periods of time.
Velocity	Machine data as well as data coming from new sources is being ingested at speeds not even imagined a few years ago.

Hortonworks

6 Key Hadoop DATA TYPES

1. **Sentiment**
   How your customers feel
2. **Clickstream**
   Website visitors' data
3. **Sensor/Machine**
   Data from remote sensors and machines
4. **Geographic**
   Location-based data
5. **Server Logs**
6. **Text**
   Millions of web pages, emails, and documents

Value

Hortonworks

Changes in Analyzing Data

Big data is fundamentally changing the way we analyze information.

- Ability to analyze vast amounts of data rather than evaluating sample sets.
- Historically we have had to look at causes. Now we can look at patterns and correlations in data that give us much better perspective.

Hortonworks

The Need for Hadoop

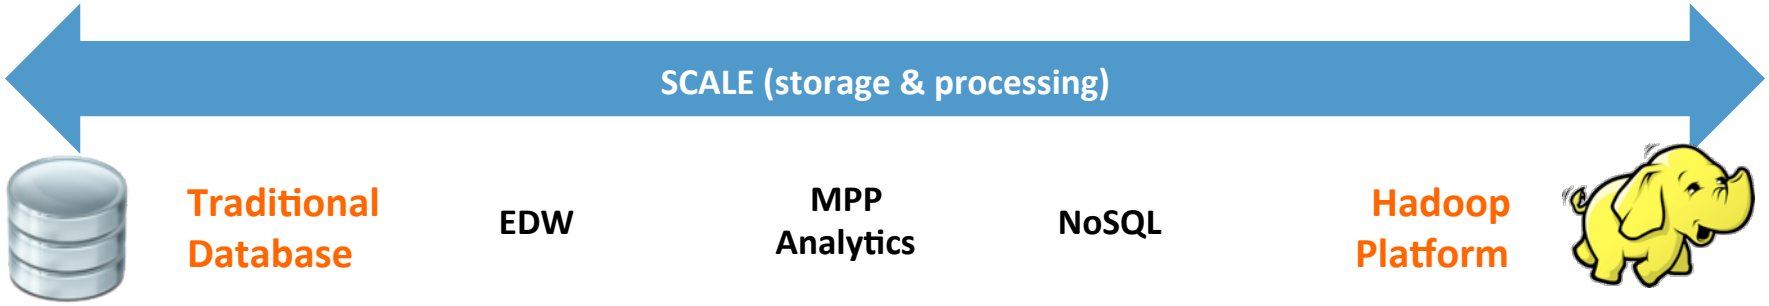

- Store and use all types of data
- Process all the data
- Scalability
- Commodity hardware

Hortonworks

Hadoop as a Data Factory

- A role Hadoop can play in an enterprise data platform is that of a data factory

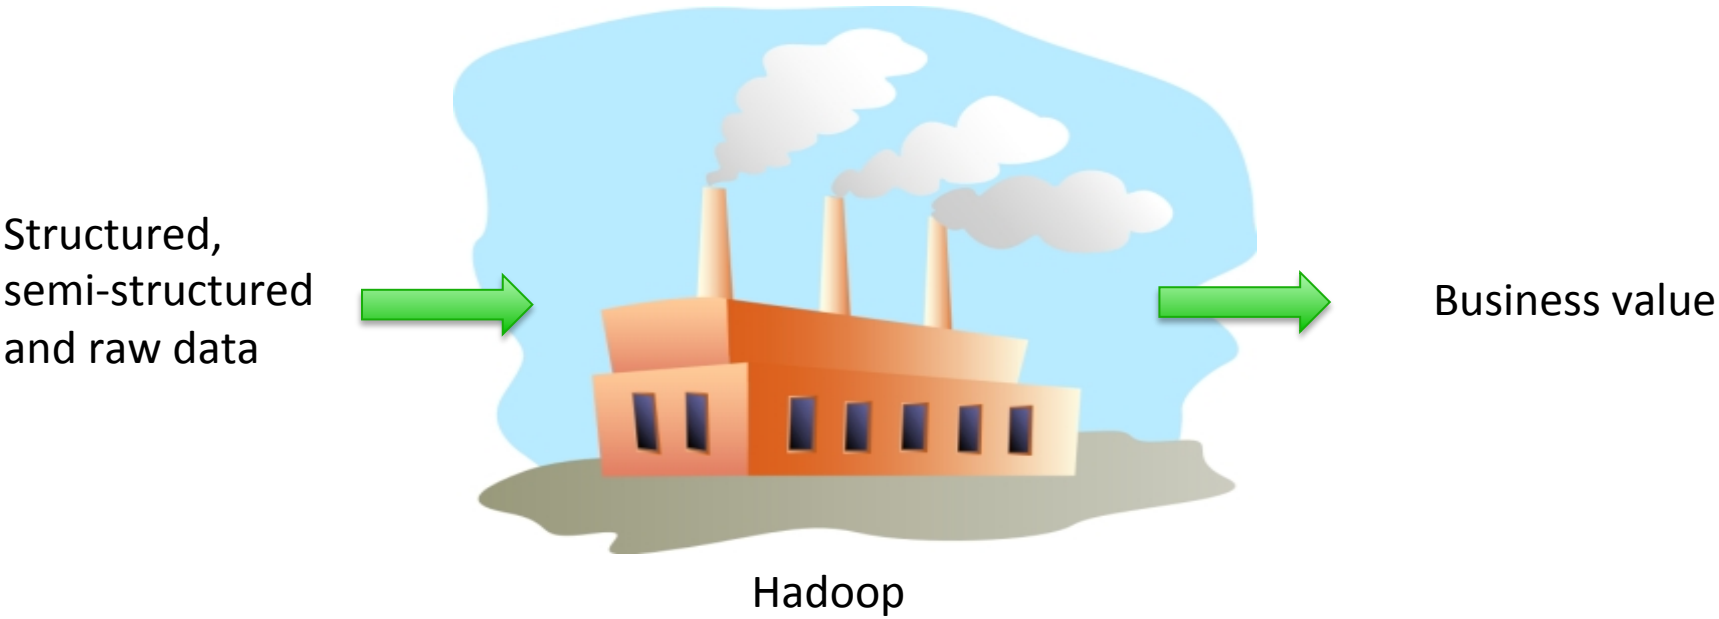

Hortonworks

Integrating Hadoop

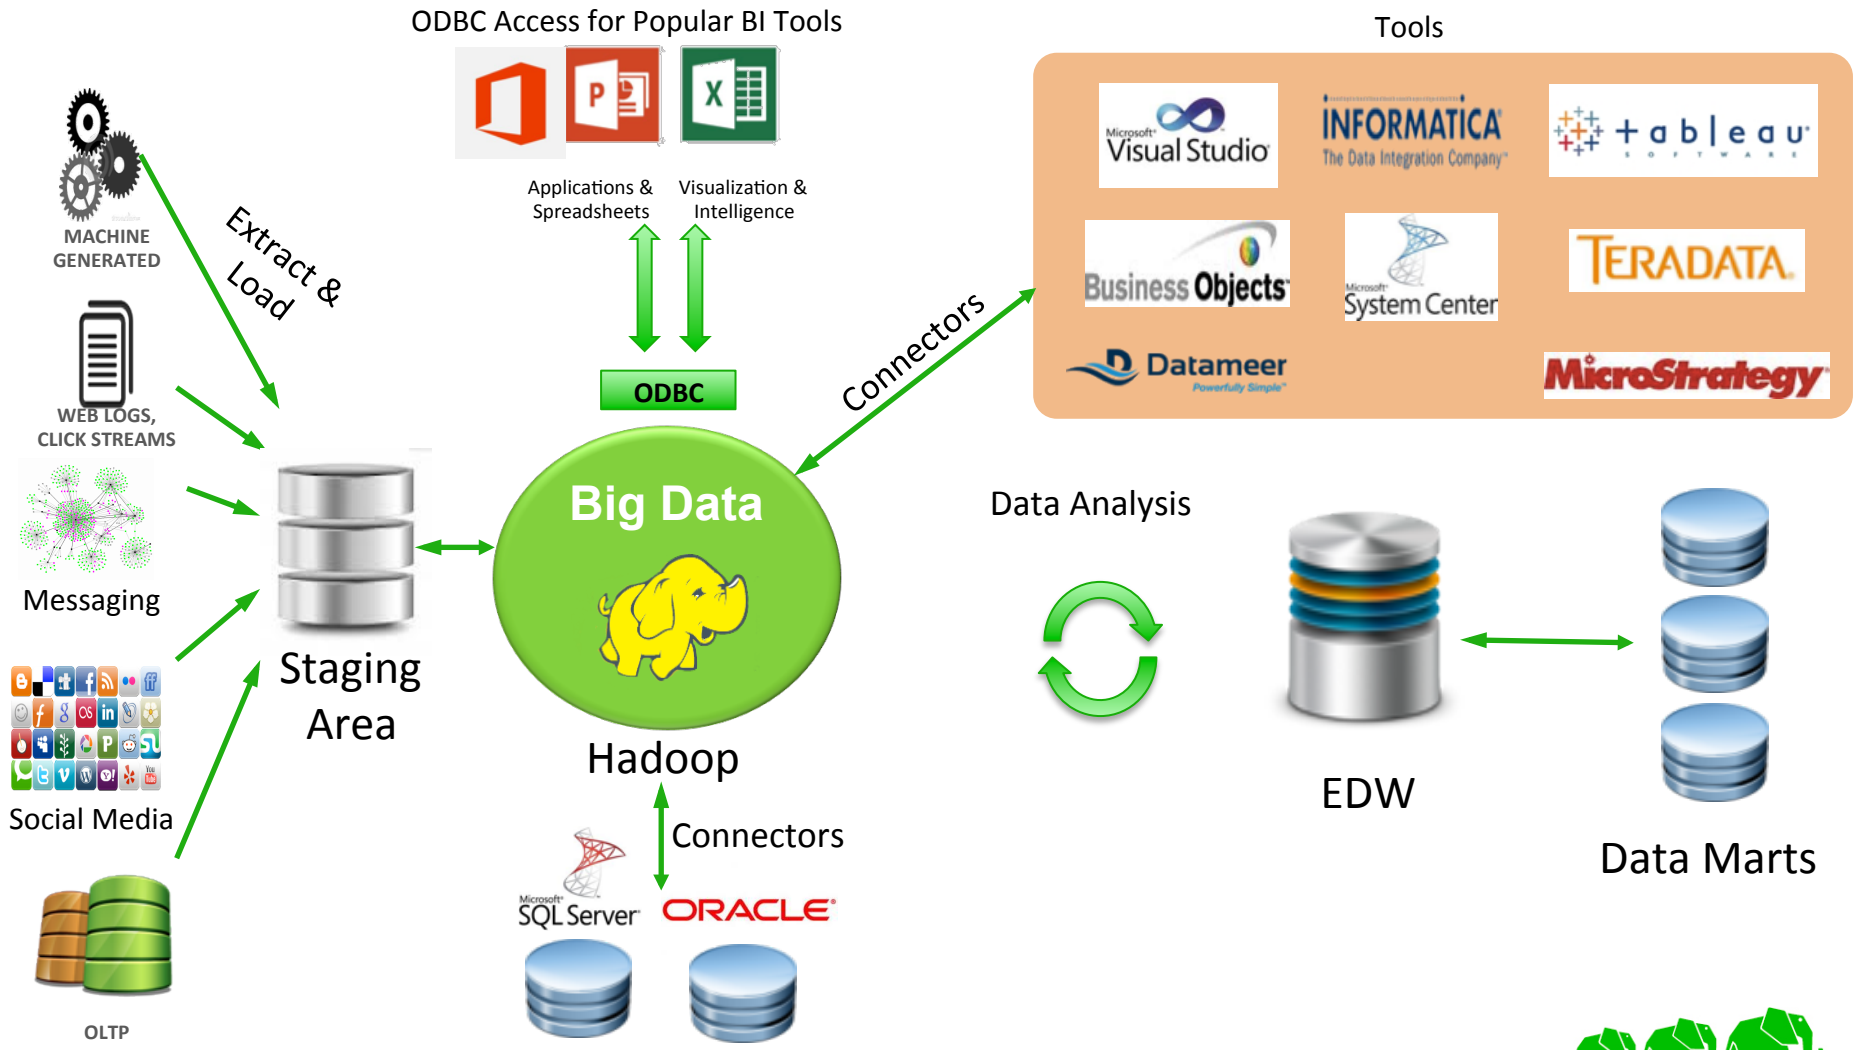

Hortonworks

Understanding Hadoop

What is Hadoop?

- Framework for solving data-intensive processes
- Designed to scale massively
- Processes all the contents of a file (instead of attempting to read portions of a file)
- Hadoop is very fast for very large jobs
- Hadoop is not fast for small jobs
- It does not provide caching or indexing (tools like HBase can provide these features if needed)
- Designed for hardware and software failures

Hortonworks

What is Hadoop 2.0?

- The Apache Hadoop 2.0 project consists of the following modules:
  - **Hadoop Common**: the utilities that provide support for the other Hadoop modules.
  - **HDFS**: the Hadoop Distributed File System
  - **YARN**: a framework for job scheduling and cluster resource management.
  - **MapReduce**: for processing large data sets in a scalable and parallel fashion.

Hortonworks

What is YARN? …

- YARN is a sub-project of Hadoop at the Apache Software Foundation that takes Hadoop beyond batch processing to enable broader data-processing
- It extends the Hadoop platform by supporting non-MapReduce workloads associated with other programming models
- The core concept of YARN was born out of a need to have Hadoop work for more real-time and streaming capabilities
  - As more and more data landed in Hadoop, enterprises have demanded that Hadoop extend its capabilities

What is YARN? ...

- As part of Hadoop 2.0, YARN takes the resource management capabilities that were in MapReduce and packages them so they can be used by new engines
- Streamlines MapReduce to do what it does best -- process data
- Run multiple applications in Hadoop, all sharing a common resource management
- Many organizations are already building applications on YARN in order to bring them IN to Hadoop
- With Hadoop 2.0 and YARN, organizations can use Hadoop for streaming, interactive and a world of other Hadoop-based applications

What is YARN?

- YARN is a re-architecture of Hadoop that allows multiple application to run on the same platform

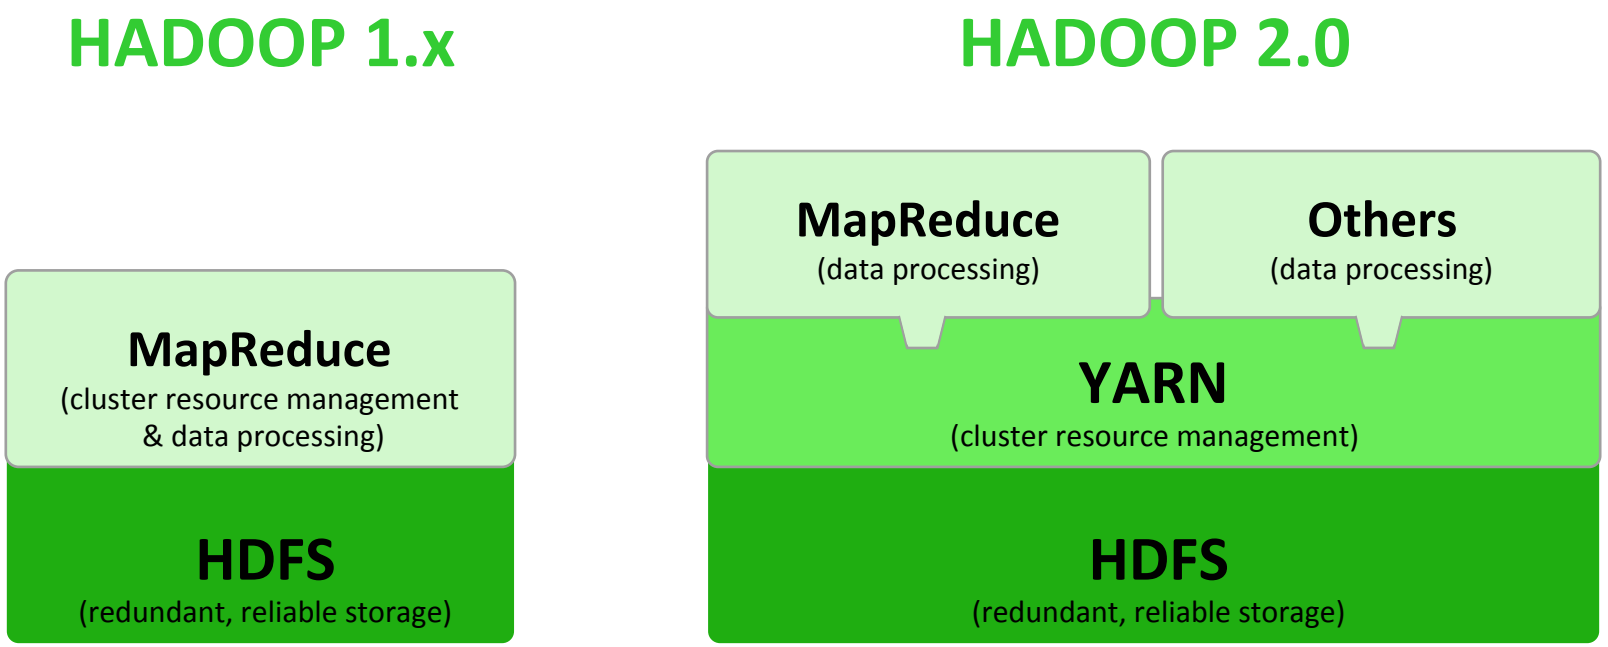

Hortonworks

YARN: Taking Hadoop Beyond Batch

- With YARN, applications run natively in Hadoop

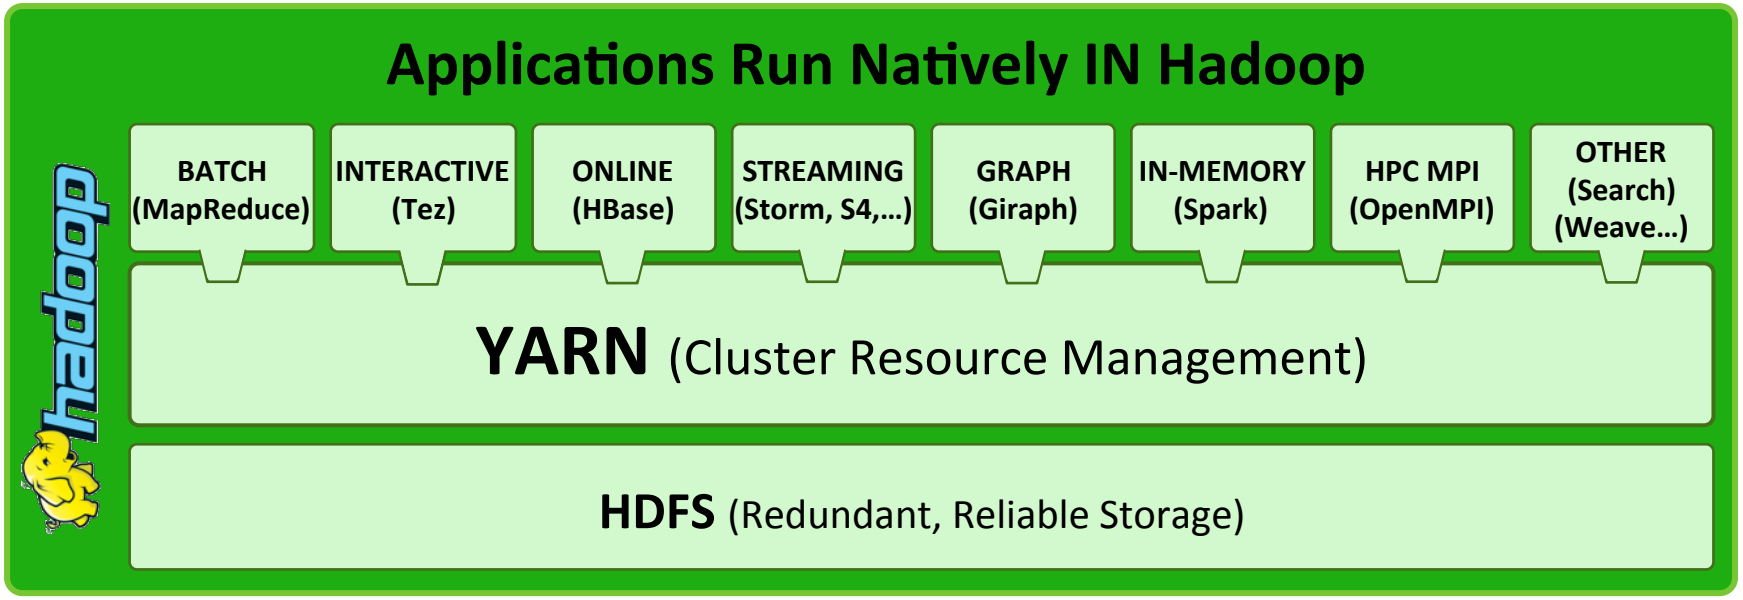

Hortonworks

The Hadoop Ecosystem

Apache Ambari
http://incubator.apache.org/ambari

sqoop

OOZIE

HIVE

Hadoop

APACHE HBASE

mahout

Hortonworks

The Path to ROI

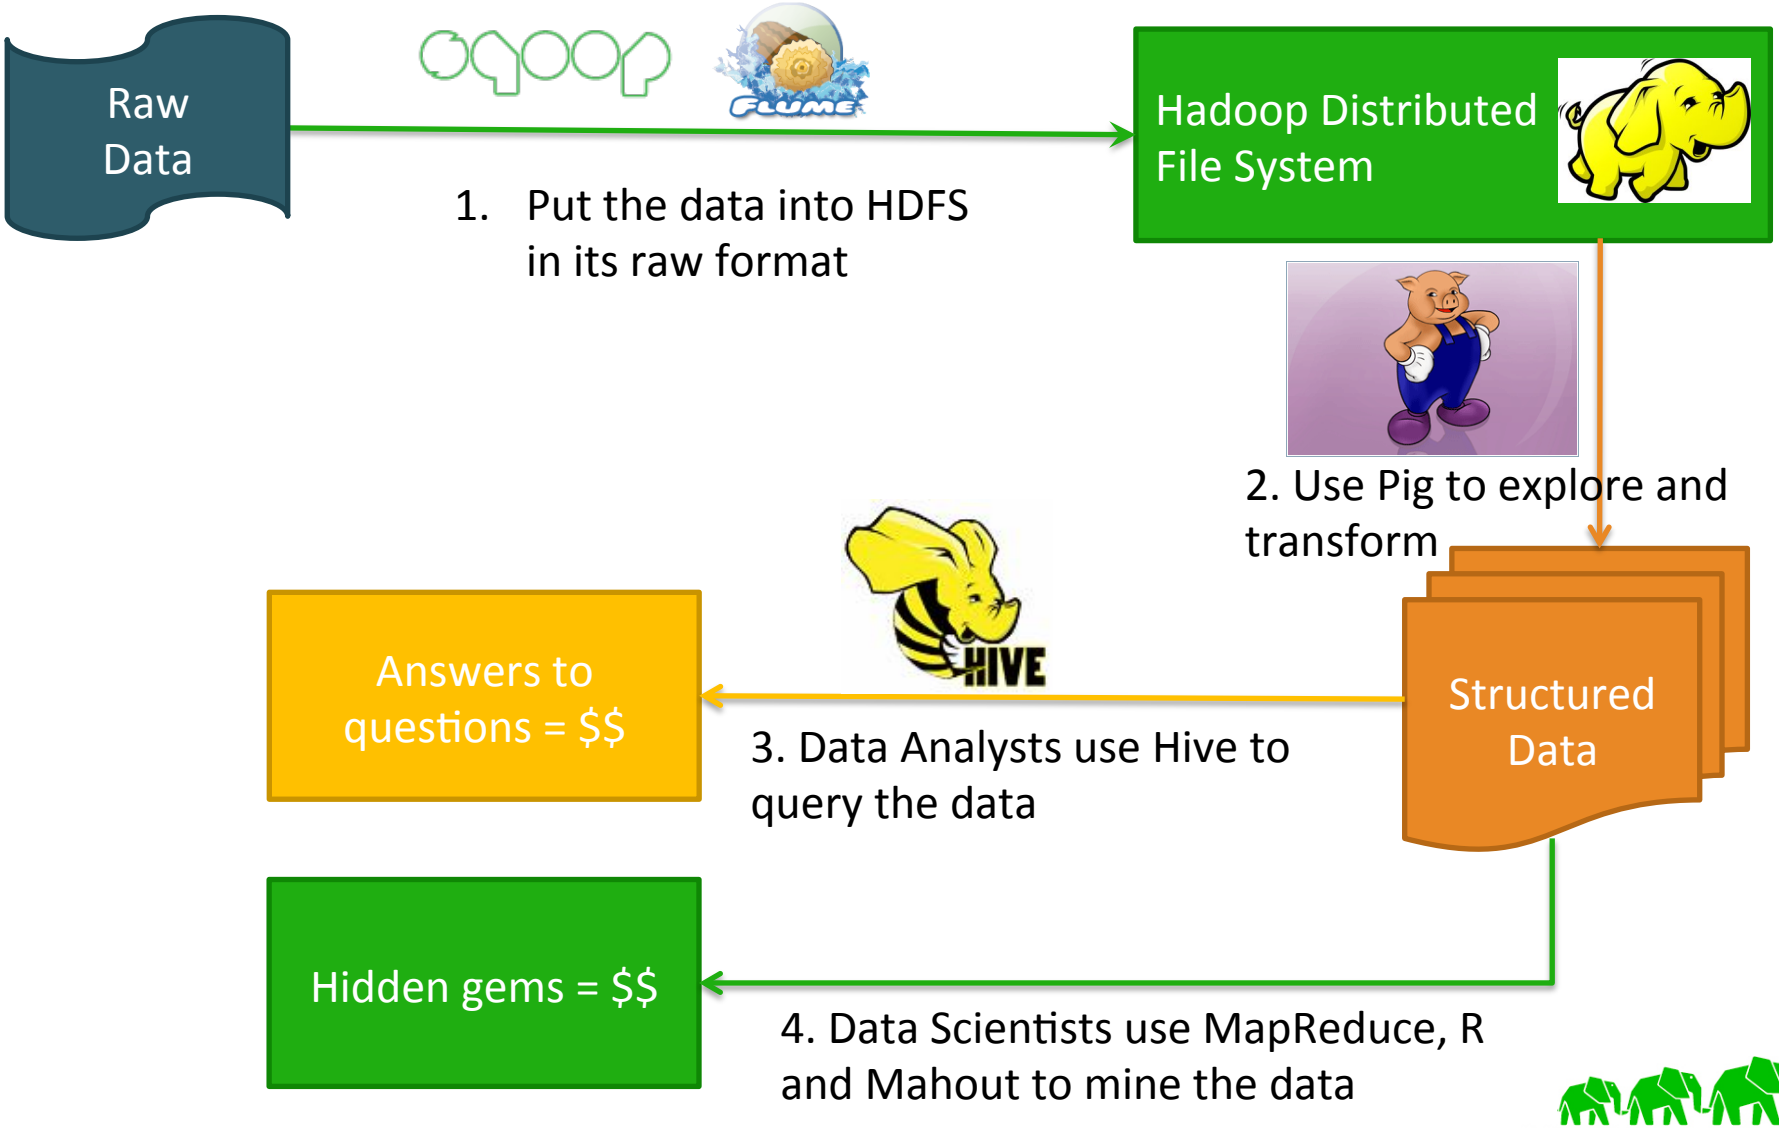

Hortonworks

Getting Relational data into Hadoop

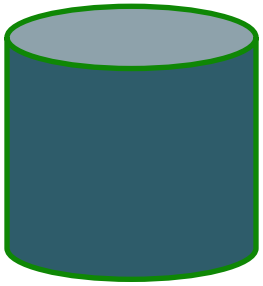

A table in an RDBMS

A Sqoop job

Sqoop is a tool for transferring data between an RDBMS and Hadoop.

Hadoop cluster

Hortonworks

Getting Sensor data into Hadoop

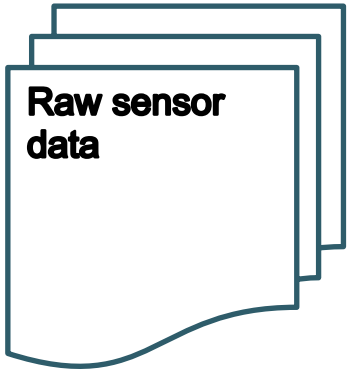

Flume Agent

Flume is a tool for streaming data into Hadoop.

Hadoop cluster

Hortonworks

What is Pig?

- Pig is an extension of Hadoop that simplifies the ability to query large HDFS datasets
- Pig is made up of two main components:
  - A data processing language called **Pig Latin**
  - A compiler that compiles and runs Pig Latin scripts
- Pig was created at Yahoo! to make it easier to analyze the data in HDFS without the complexities of writing a traditional MapReduce program
- With Pig, you can develop MapReduce jobs with a few lines of Pig Latin

Hortonworks

Running Pig

A Pig Latin script executes in three modes:

1. **MapReduce**: the code executes as a MapReduce application on a Hadoop cluster (the default mode)

```
$ pig myscript.pig
```

2. **Local**: the code executes locally in a single JVM using a local text file (for development purposes)

```
$ pig -x local myscript.pig
```

3. **Interactive**: Pig commands are entered manually at a command prompt known as the Grunt shell

```
$ pig
grunt>
```


Hortonworks

Pig example

- Suppose you have user data in one file, website data in another, and you need to find the top 5 most visited sites by users aged 18 - 25

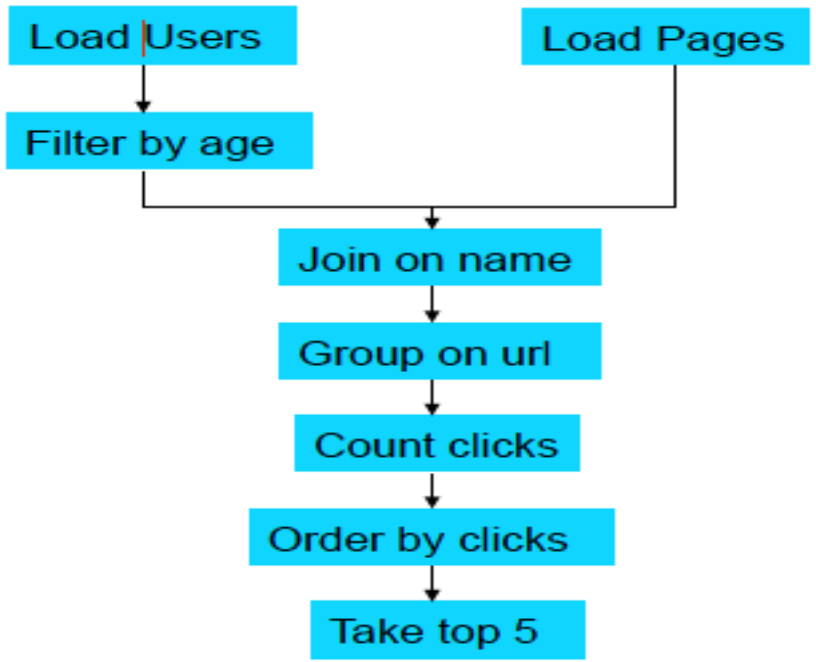

Hortonworks

In Java MapReduce ...

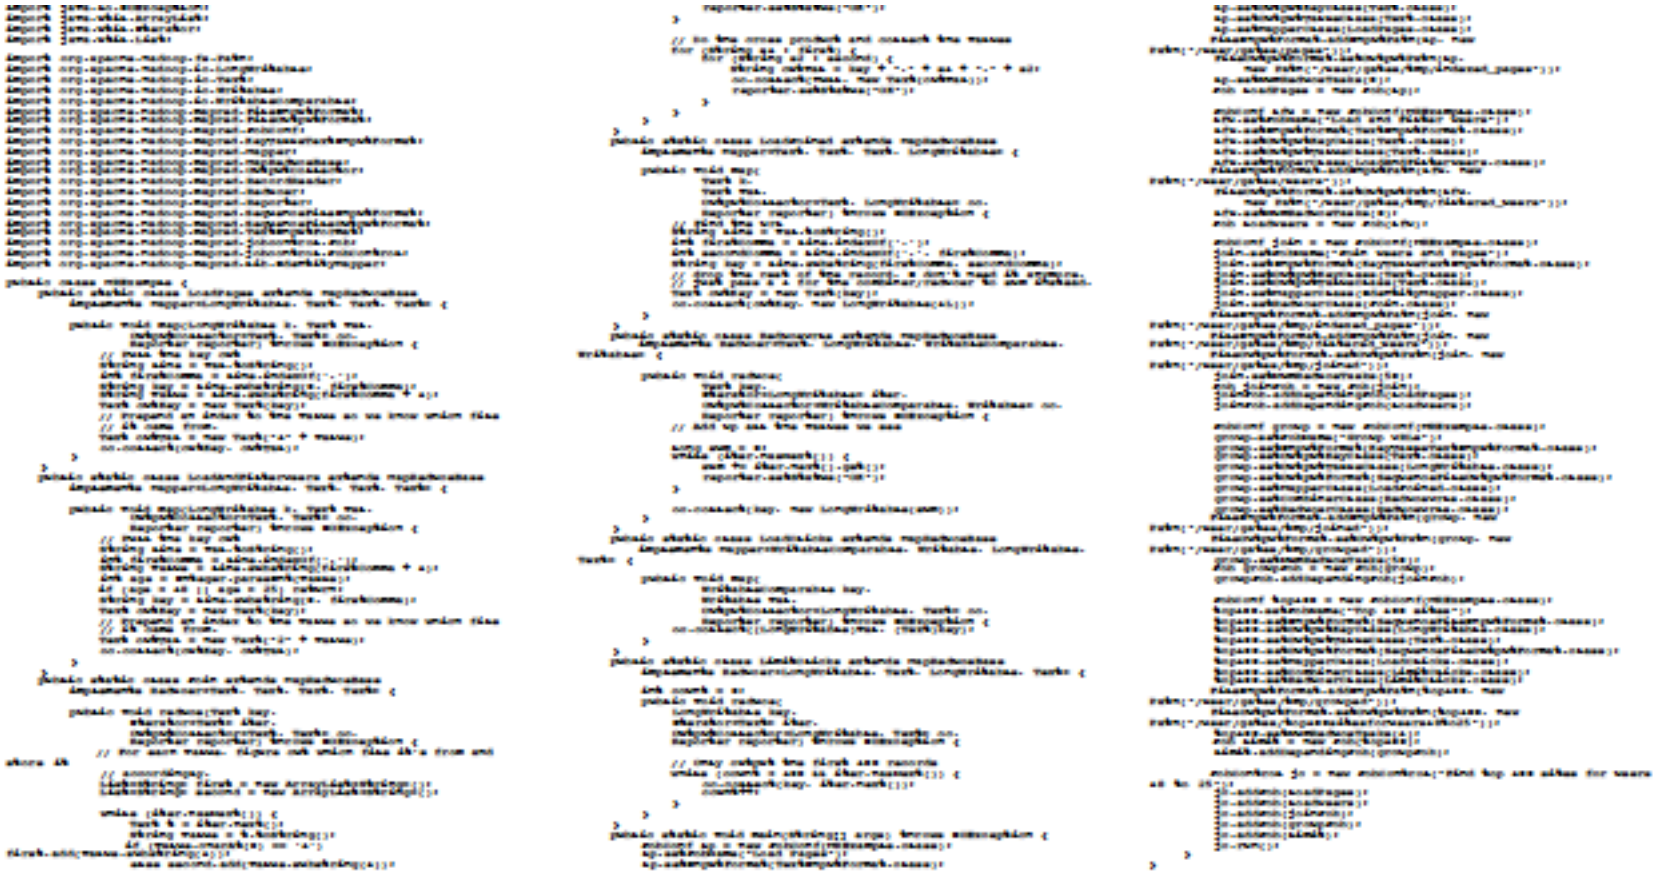

Hortonworks

In Pig Latin ...

```
Users = load 'input/users' using PigStorage(',')
        as (name:chararray, age:int);
Fltrd = filter Users by age >= 18 and age <= 25;
Pages = load 'input/pages' using PigStorage(',')
        as (user:chararray, url:chararray);
Jnd = join Fltrd by name, Pages by user;
Grpd = group Jnd by url;
Smmd = foreach Grpd generate group, COUNT(Jnd) as clicks;
Srtd = order Smmd by clicks desc;
Top5 = limit Srtd 5;
store Top5 into 'output/top5sites' using PigStorage(',');
```


Hortonworks

What is Hive?

- Hive is a subproject of the Apache Hadoop project that provides a data warehousing layer built on top of Hadoop
- Hive allows you to define a structure for your unstructured big data, simplifying the process of performing analysis and queries by introducing a familiar, SQL-like language called HiveQL
- Hive is for data analysts familiar with SQL who need to do ad-hoc queries, summarization and data analysis on their HDFS data

Hortonworks

Hive is not

- Hive is not a relational database
- Hive uses a database to store metadata, but the data that Hive processes is stored in HDFS
- Hive is not designed for on-line transaction processing and does not offer real-time queries and row level updates

Hortonworks

Hive Architecture

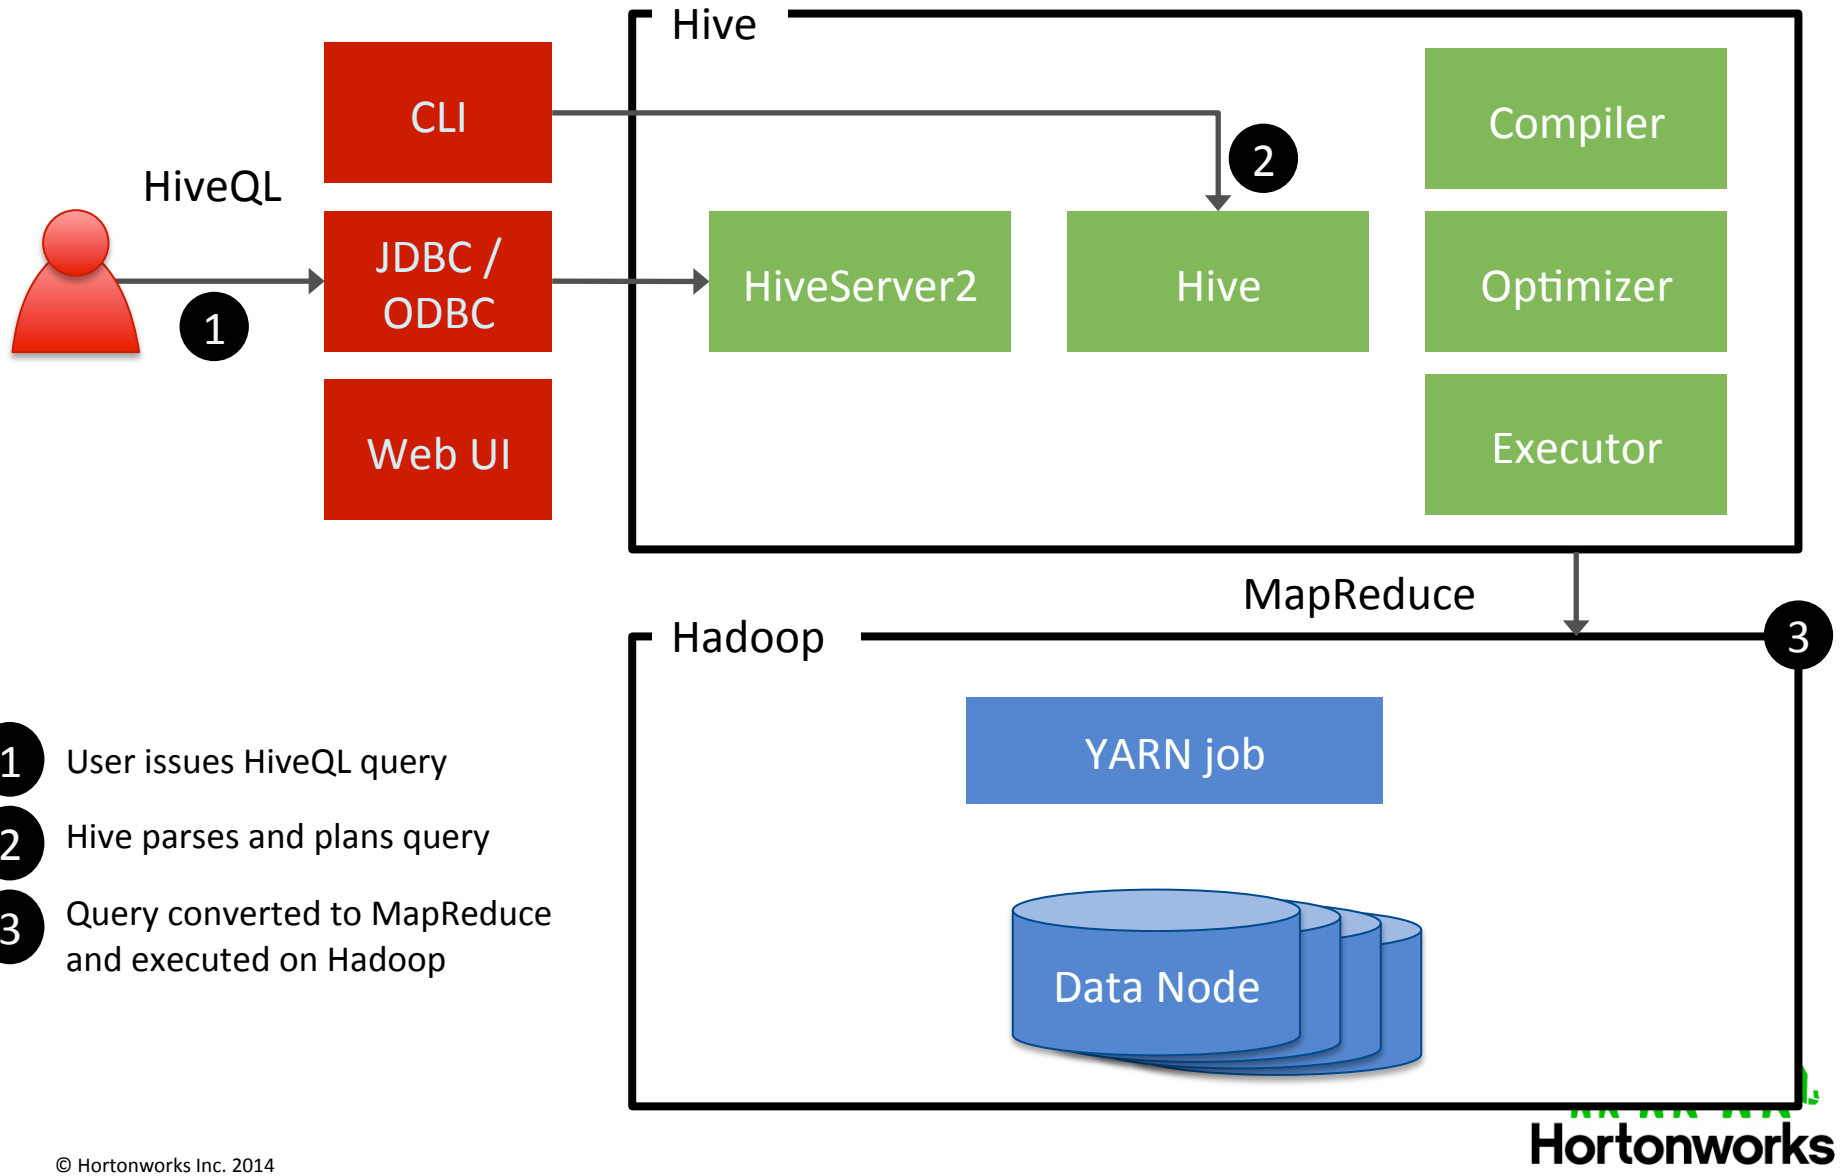

1. User issues HiveQL query
2. Hive parses and plans query
3. Query converted to MapReduce and executed on Hadoop

Hortonworks

Using HiveServer2

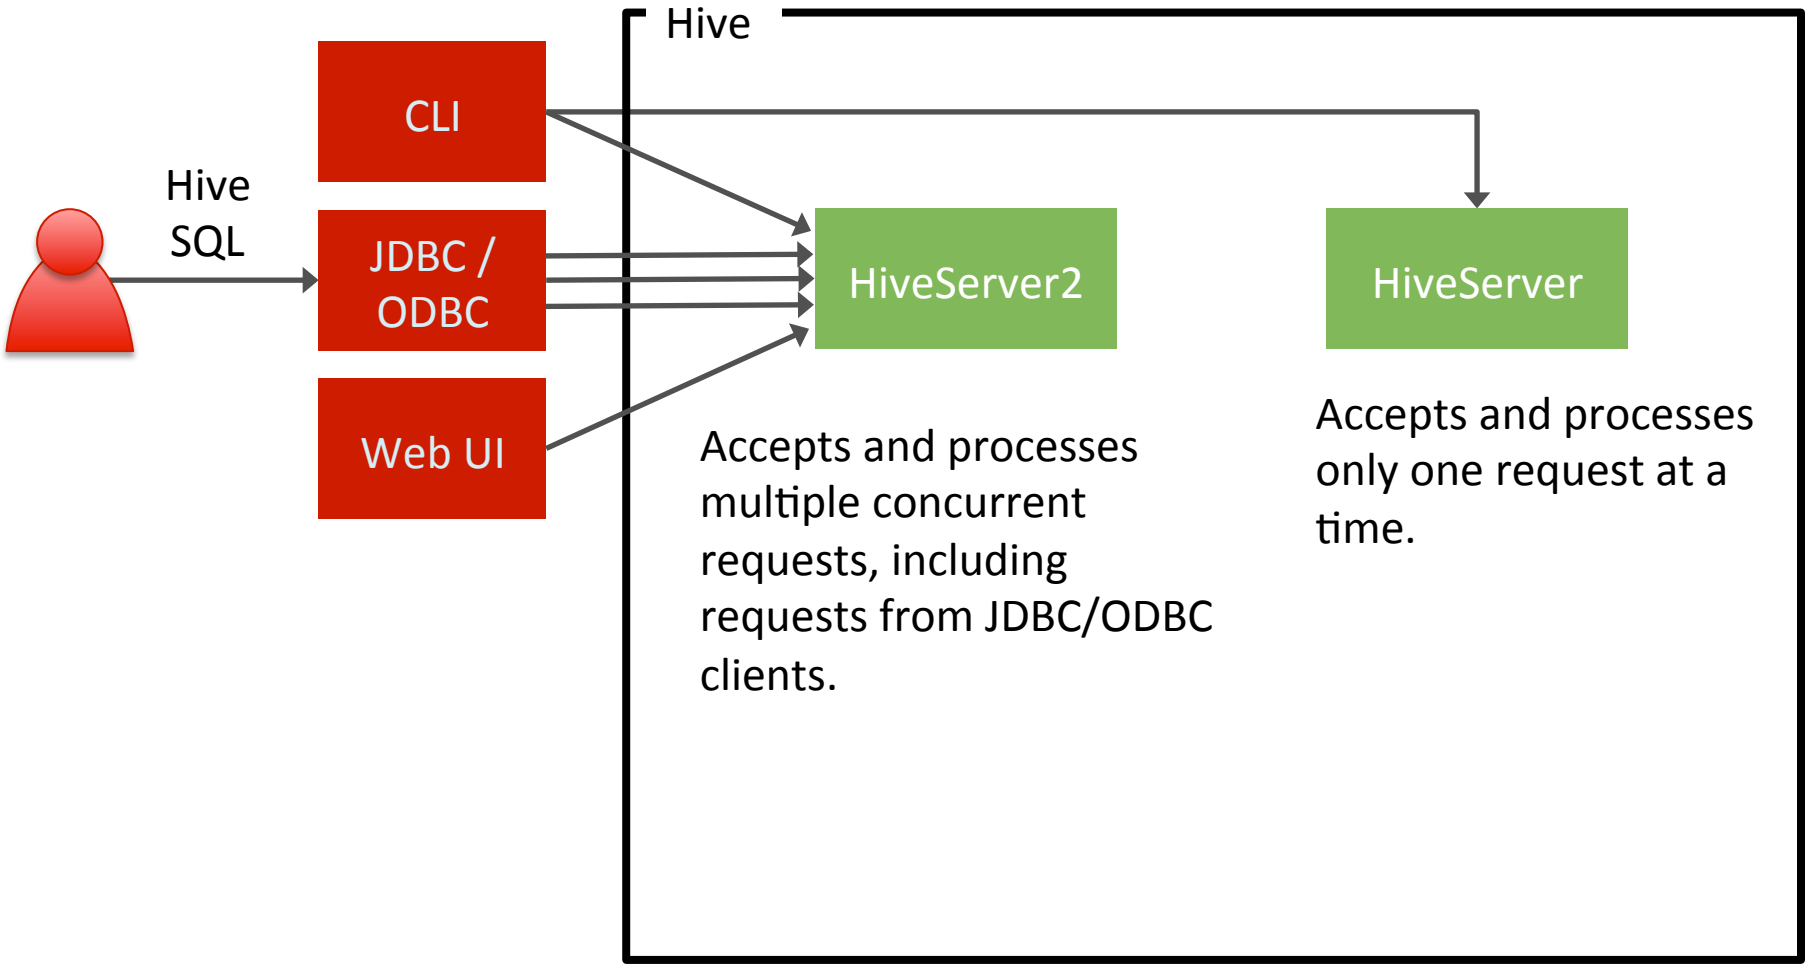

Hortonworks

HCatalog in the Ecosystem

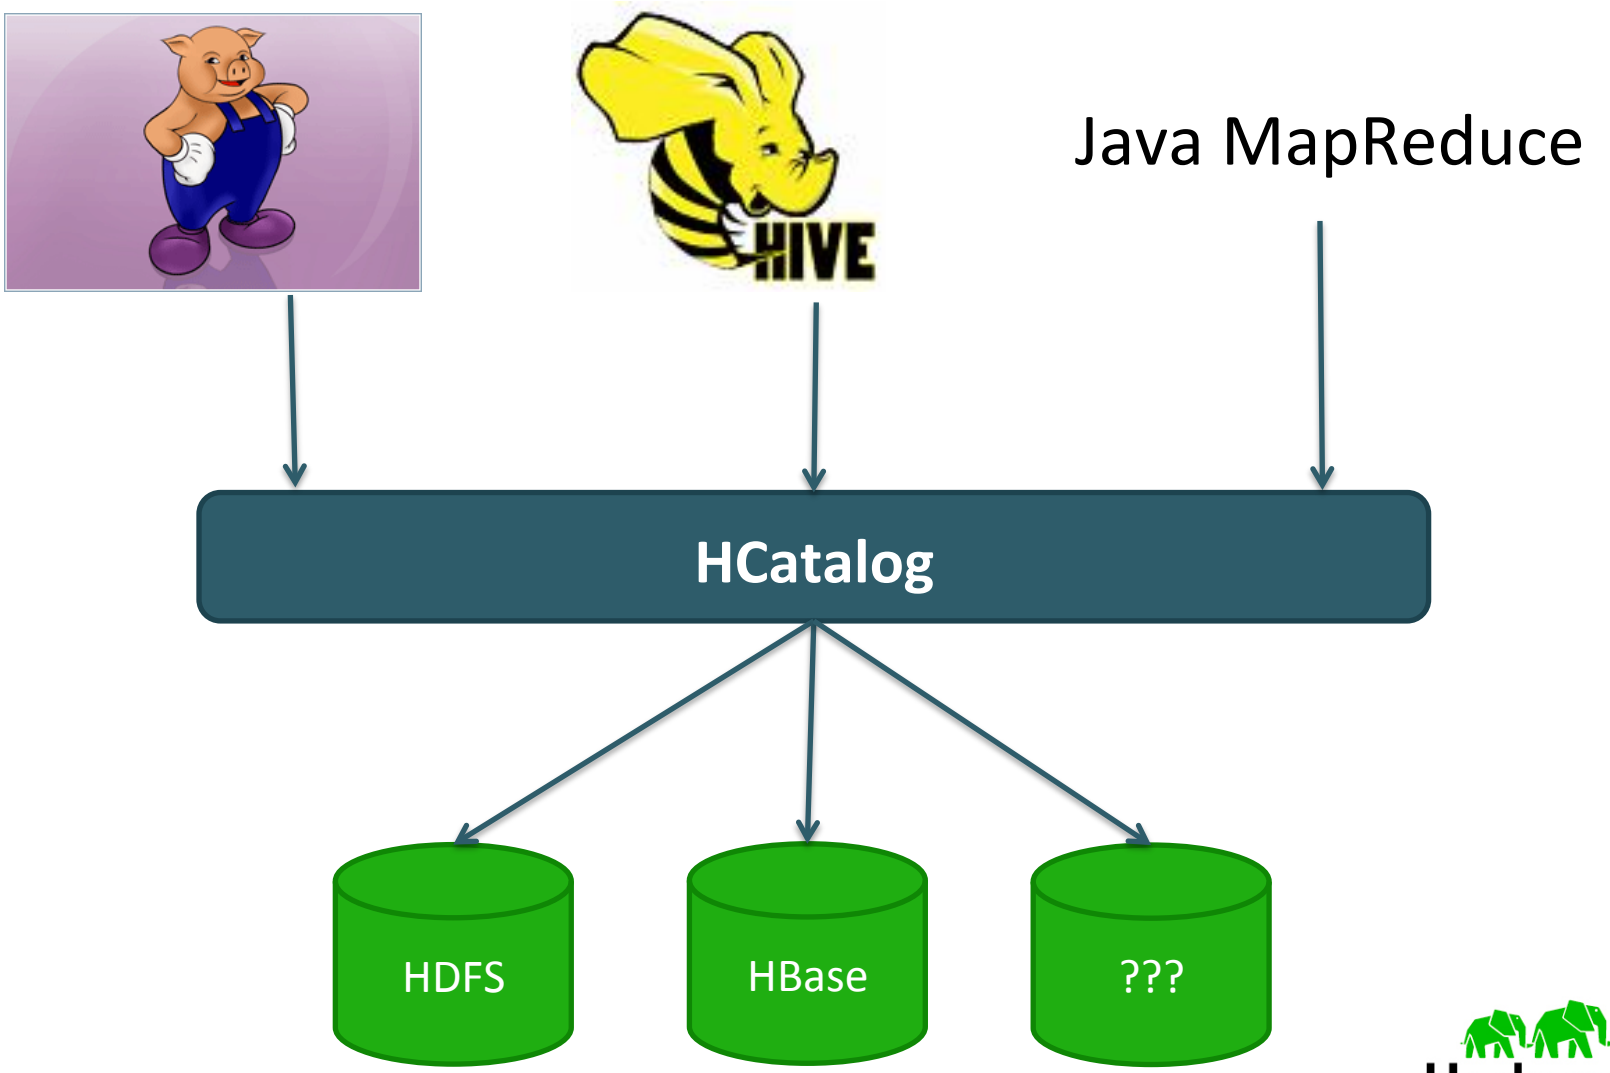

Hortonworks

Pig vs. Hive

- Pig and Hive work well together
- Hive is a good choice:
  - when you want to query the data
  - when you need an answer to a specific questions
  - if you are familiar with SQL
- Pig is a good choice:
  - for ETL (Extract -> Transform -> Load)
  - preparing your data so that it is easier to analyze
  - when you have a long series of steps to perform
- Many businesses use both Pig and Hive together

Hortonworks

Hadoop Architecture Fundamentals

What is HDFS?

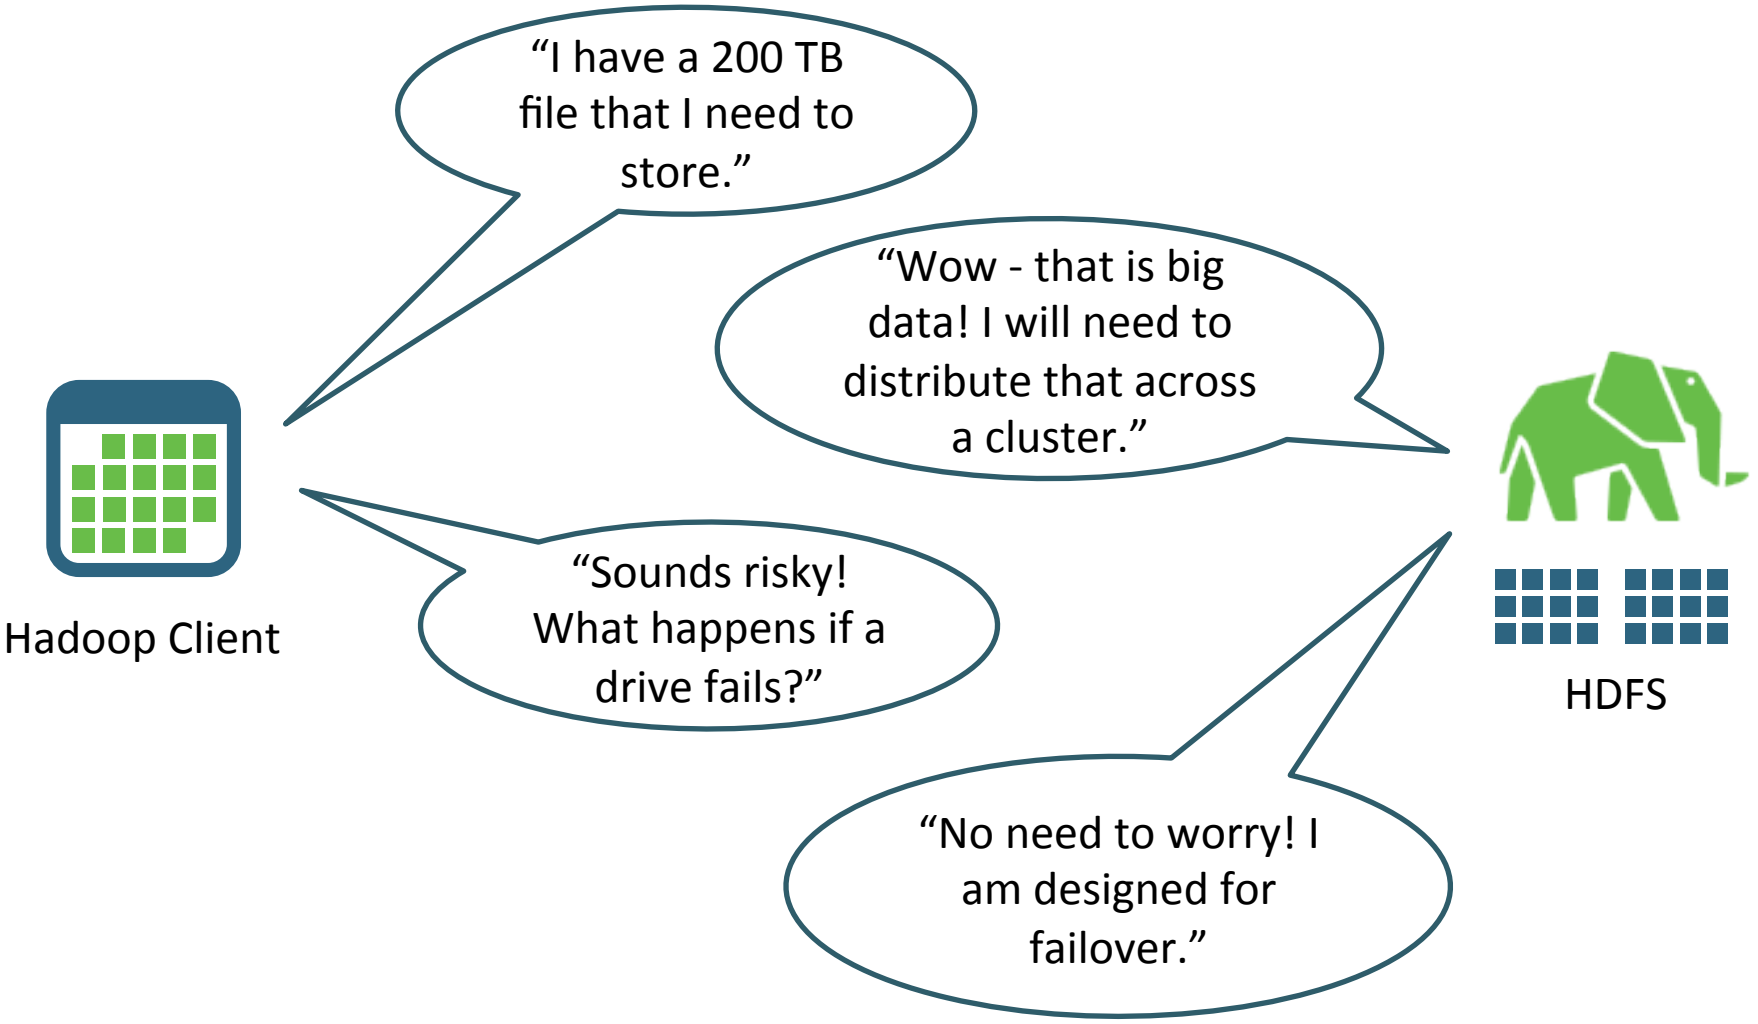

Hortonworks

The Components of HDFS

- NameNode
  - The “master” node of HDFS
  - Determines and maintains how the chunks of data are distributed across the DataNodes
- DataNode
  - Stores the chunks of data, and is responsible for replicating the chunks across other DataNodes

Hortonworks

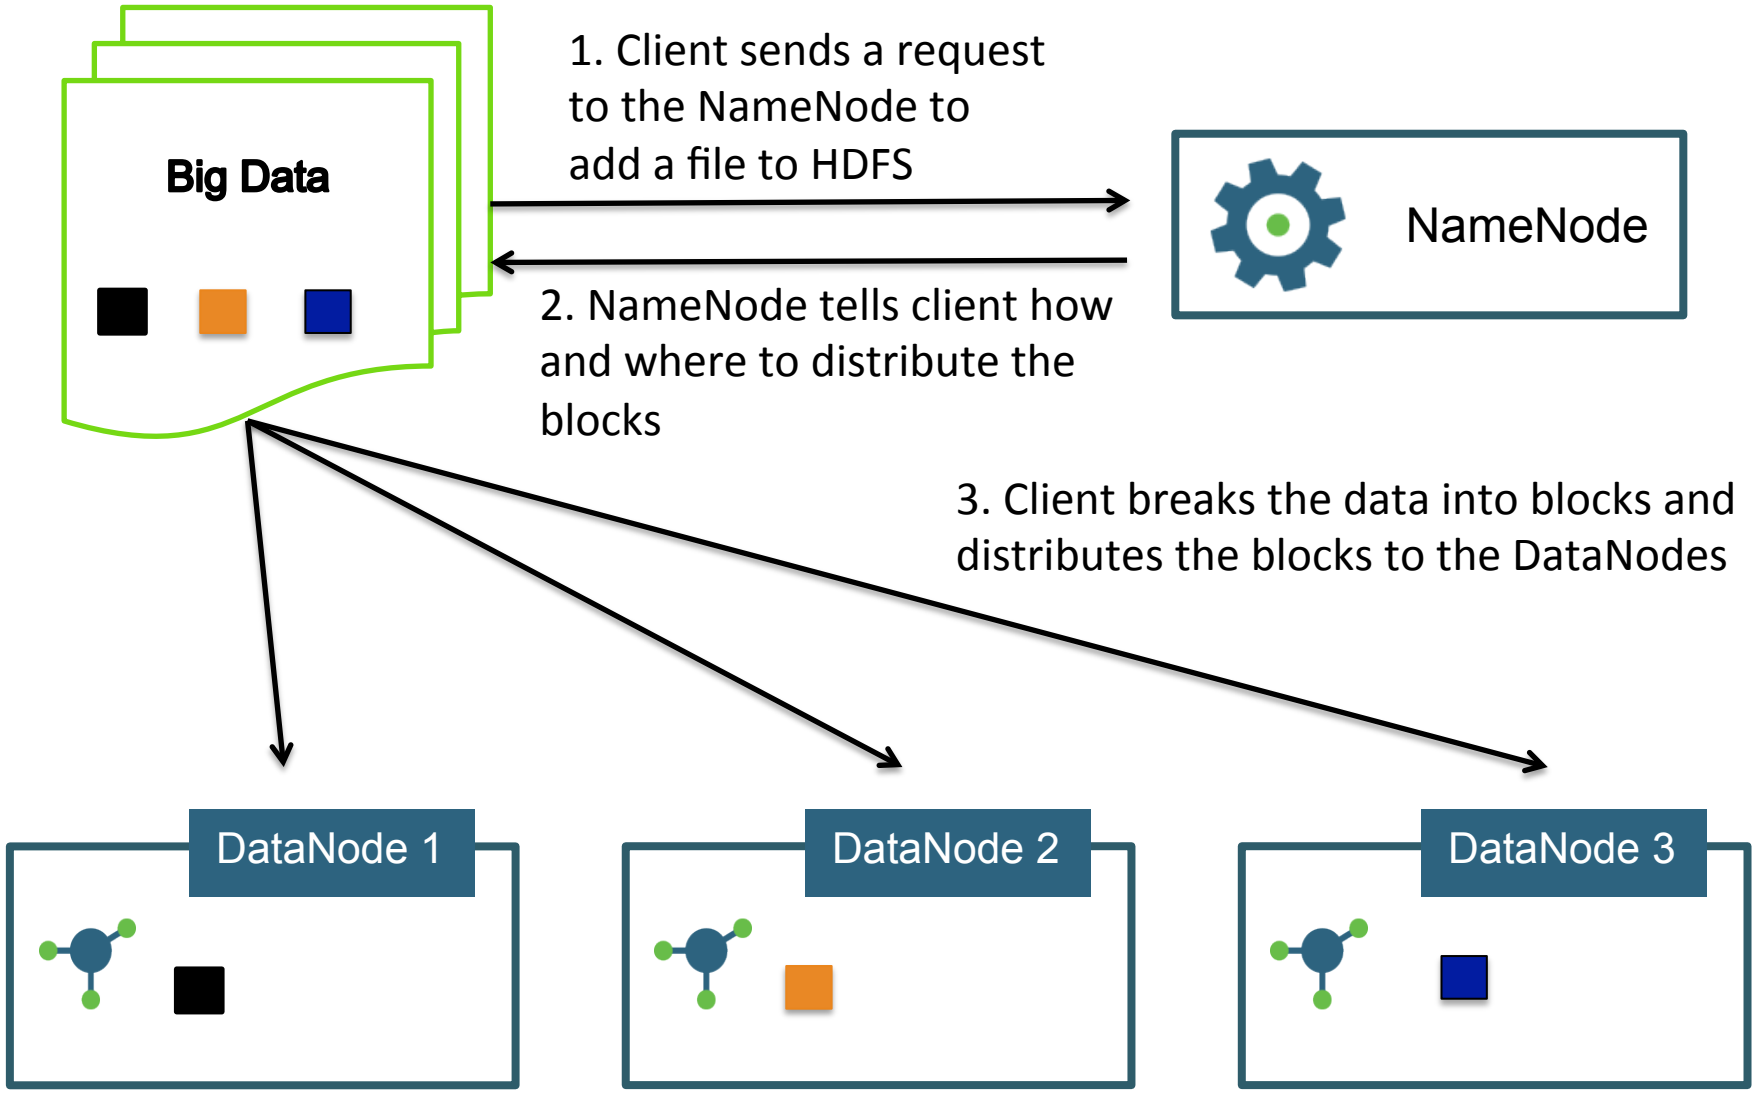

4. The DataNodes replicate the blocks (as instructed by the NameNode)

Hortonworks

The NameNode

1. When the NameNode starts, it reads the **fsimage** and **edits** files.

2. The transactions in **edits** are merged with **fsimage**, and **edits** is emptied.

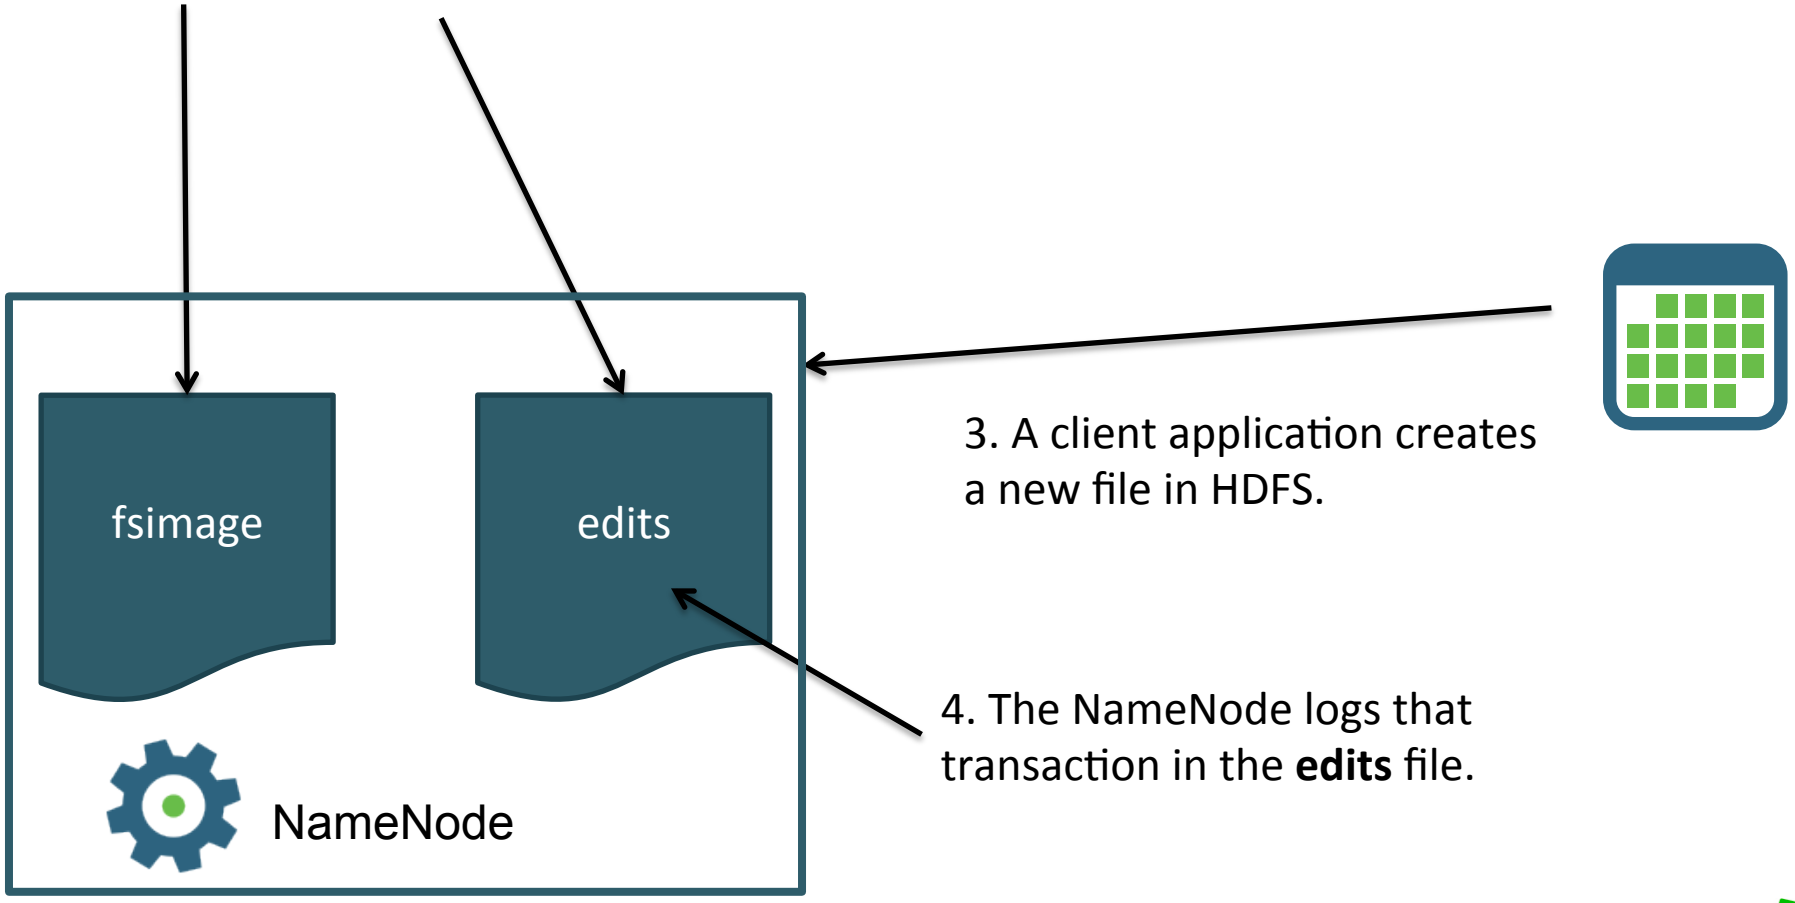

3. A client application creates a new file in HDFS.

4. The NameNode logs that transaction in the **edits** file.

Hortonworks

The DataNodes

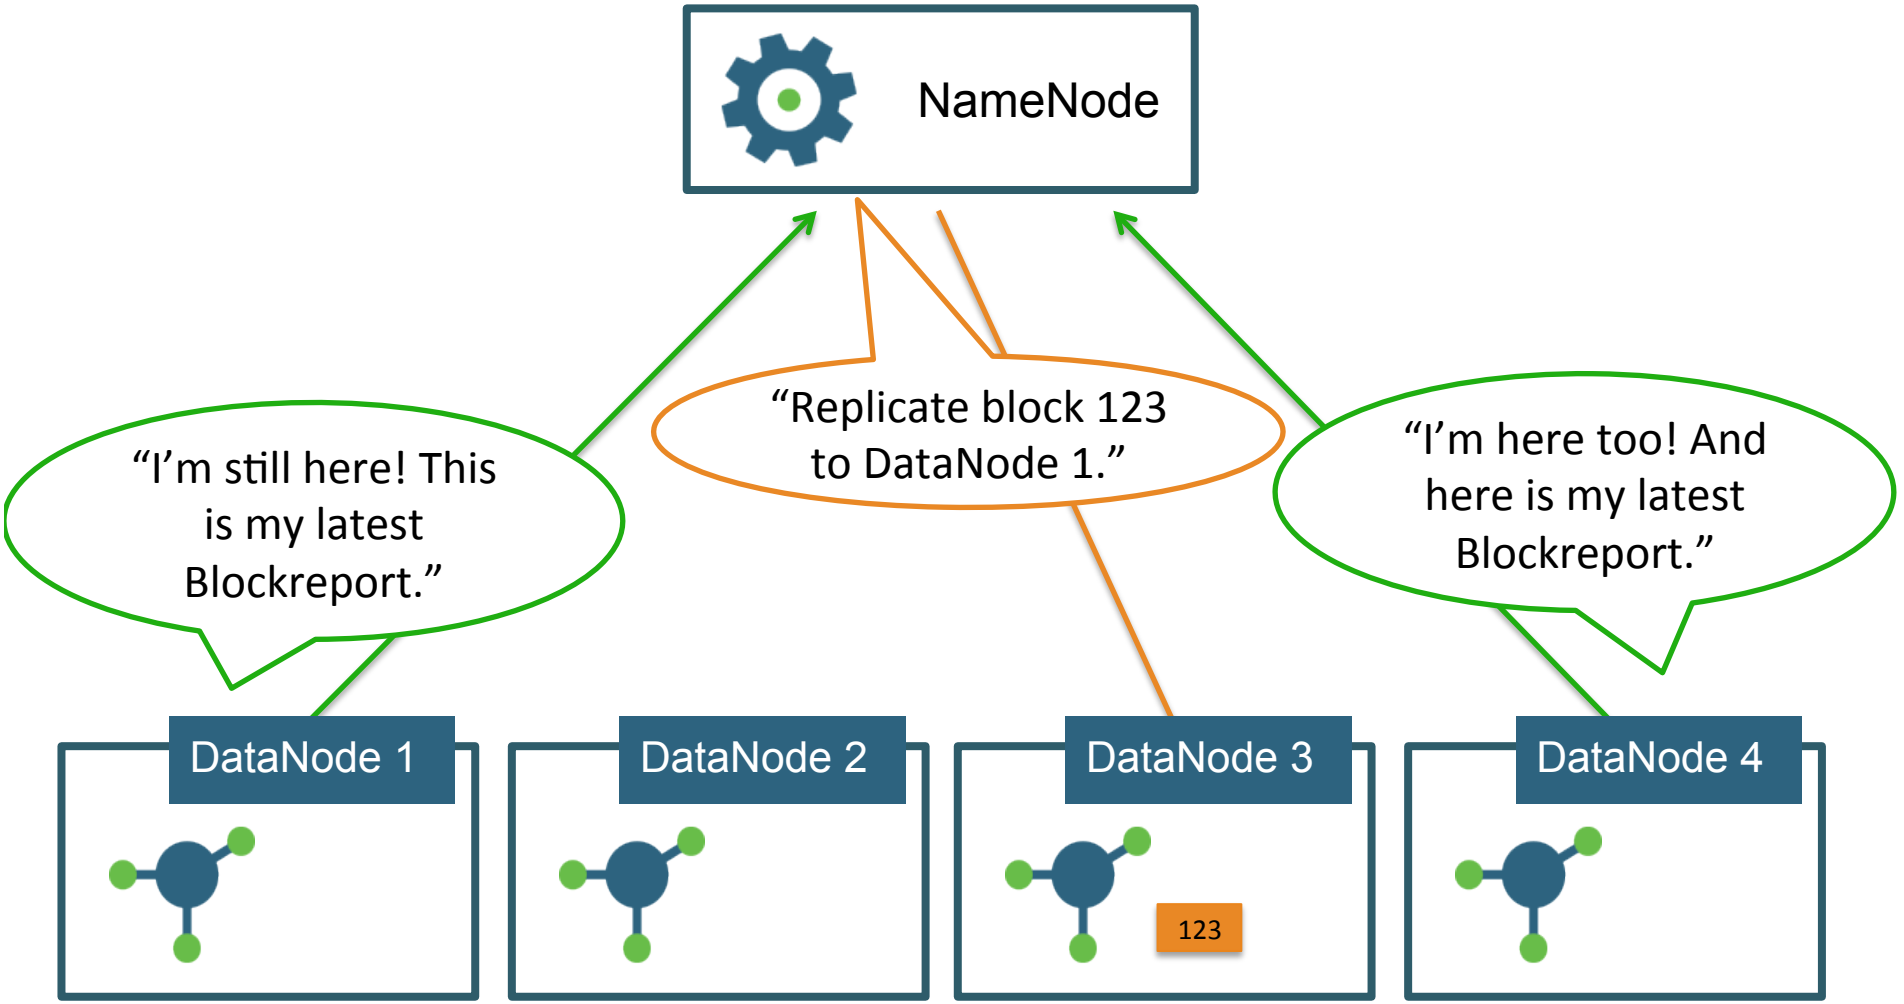

Hortonworks

Hadoop

- Assumes a task will require reading a significant amount of data off of a disk
- It does not maintain any data structure
- Simply reads the entire file
- Scales well (increase the cluster size to decrease the read time of each task)

RDBMS

- Uses indexes to avoid reading an entire file (very fast lookups)
- Maintains a data structure in order to provide a fast execution layer
- Works well as long as the index fits in RAM

- 2,000 blocks of size 256MB
- 1.9 seconds of disk read for each block
- On a 40 node cluster with 8 disks on each node, it would take about 14 seconds to read the entire 500 GB

61 minutes to read this data off of a disk (assuming a transfer rate of 1,030 Mbps)

Hortonworks

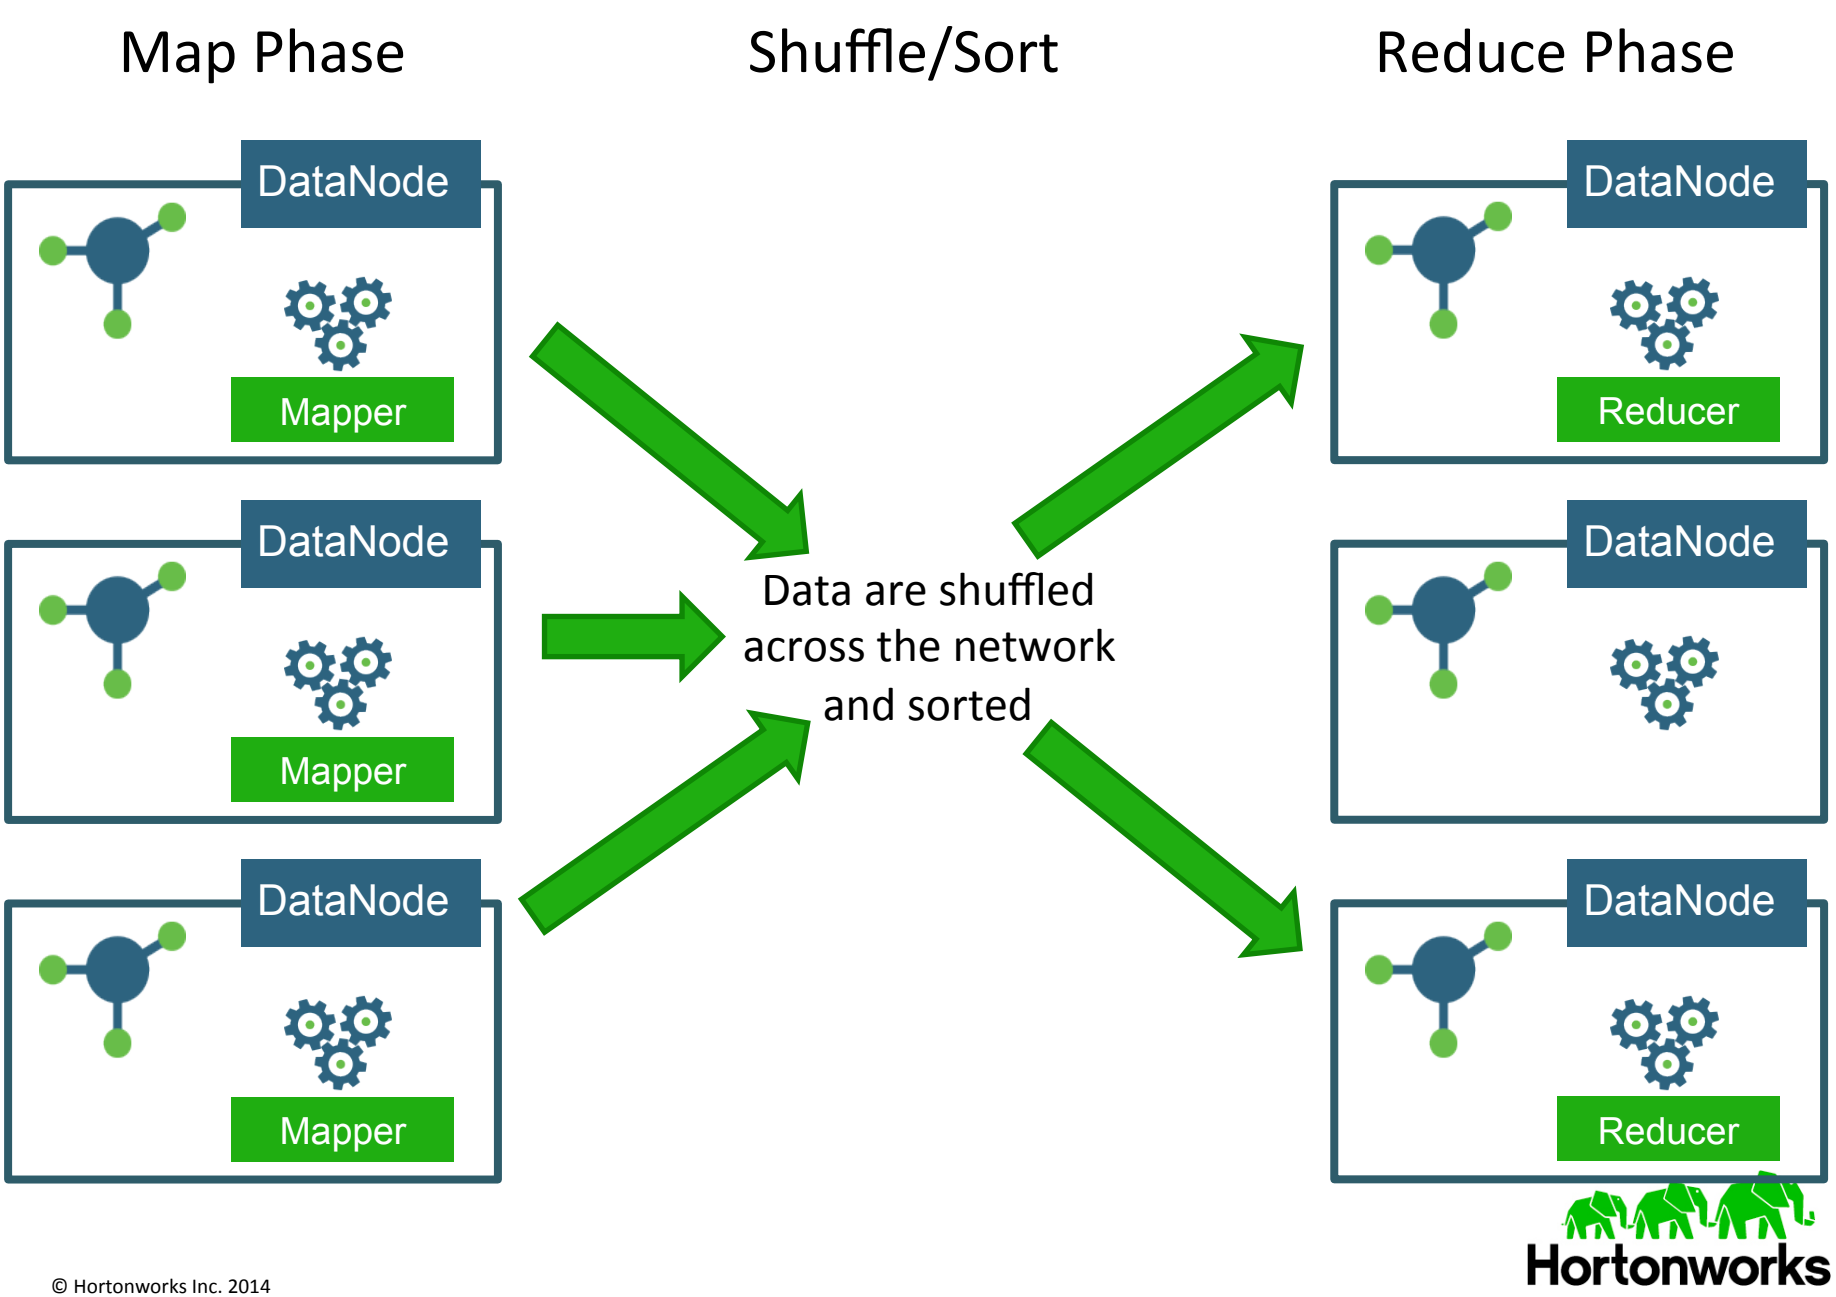

Hortonworks

WordCount in MapReduce

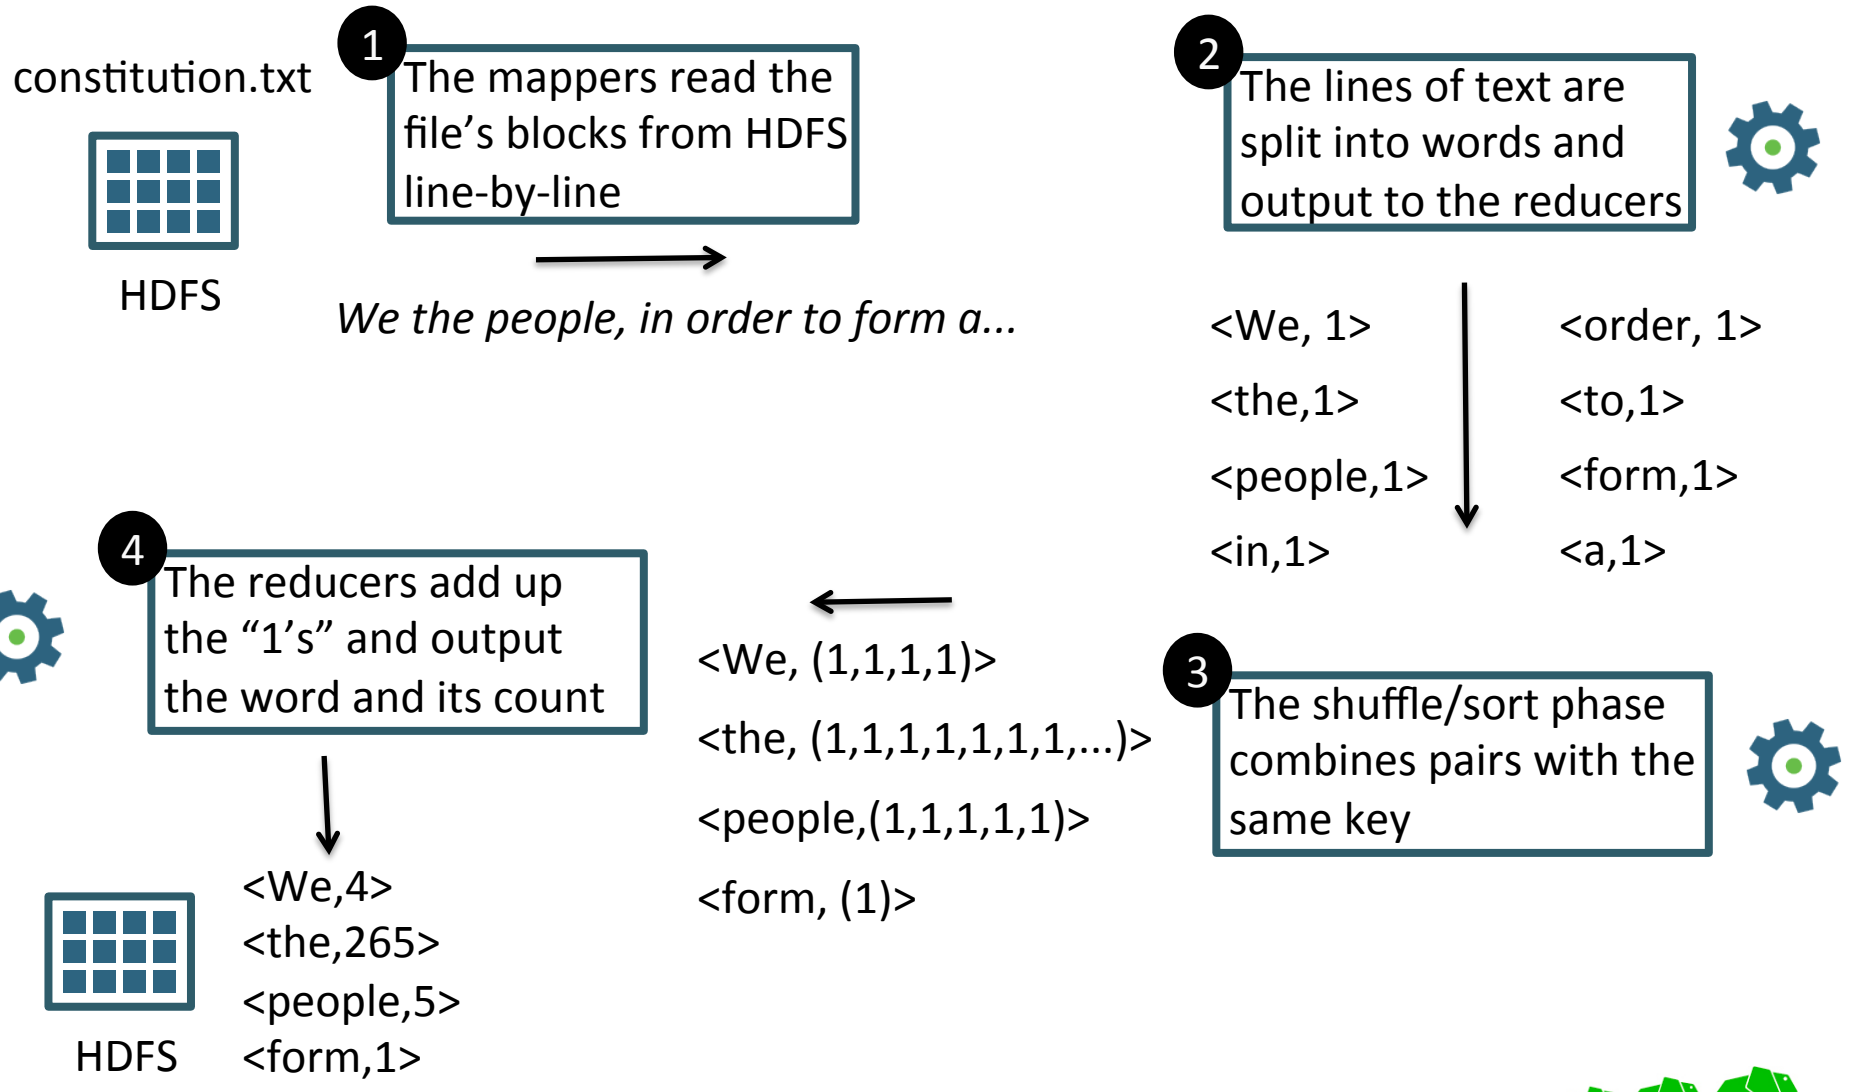

Hortonworks

The Future of Hadoop

Overview of Tez

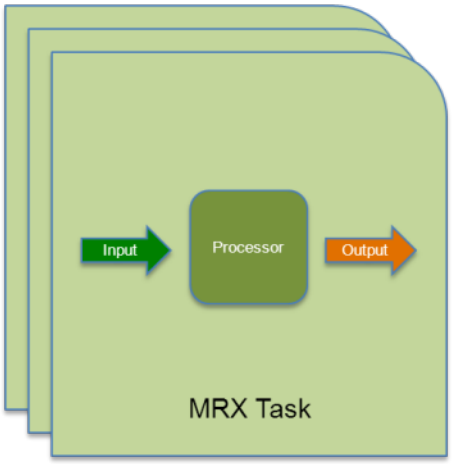

Tez Generalizes MapReduce

Simplified execution plans process data more efficiently

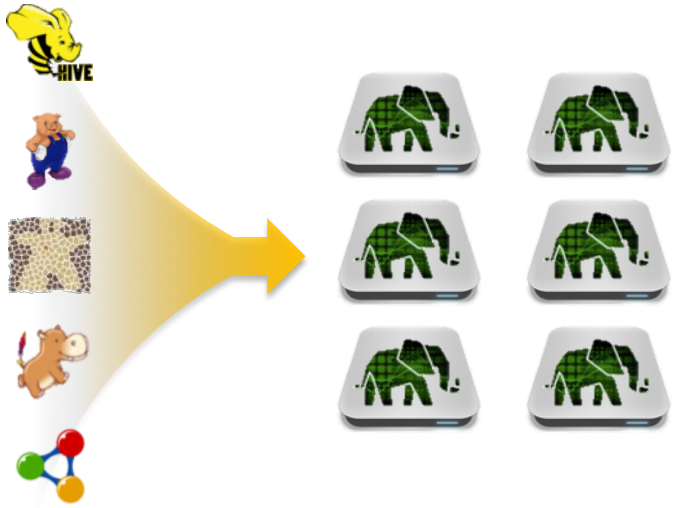

Always-On Tez Service

Low latency processing for all Hadoop data processing

Hortonworks

Tez example

```
SELECT a.state, COUNT(*), AVG(c.price)
FROM a
JOIN b ON (a.id = b.id)
JOIN c ON (a.itemId = c.itemId)
GROUP BY a.state
```

Tez avoids unneeded writes to HDFS

Hive – MapReduce

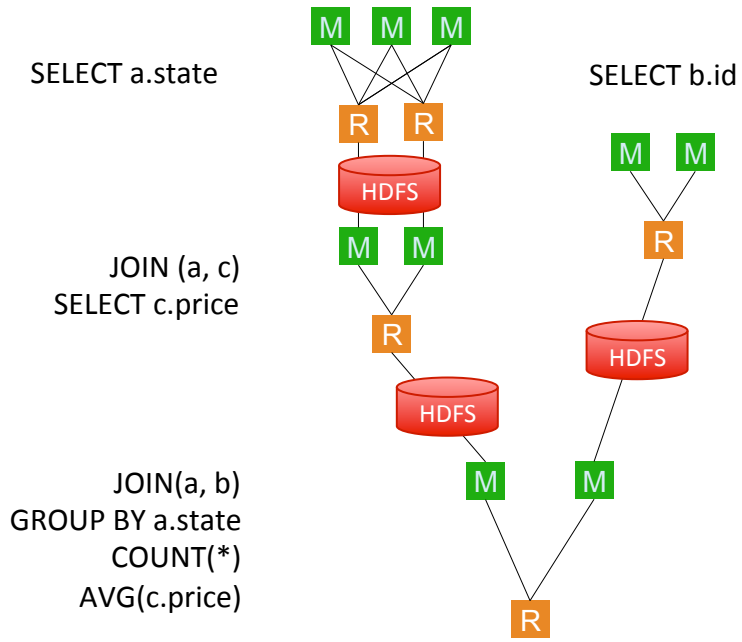

Hive – Tez

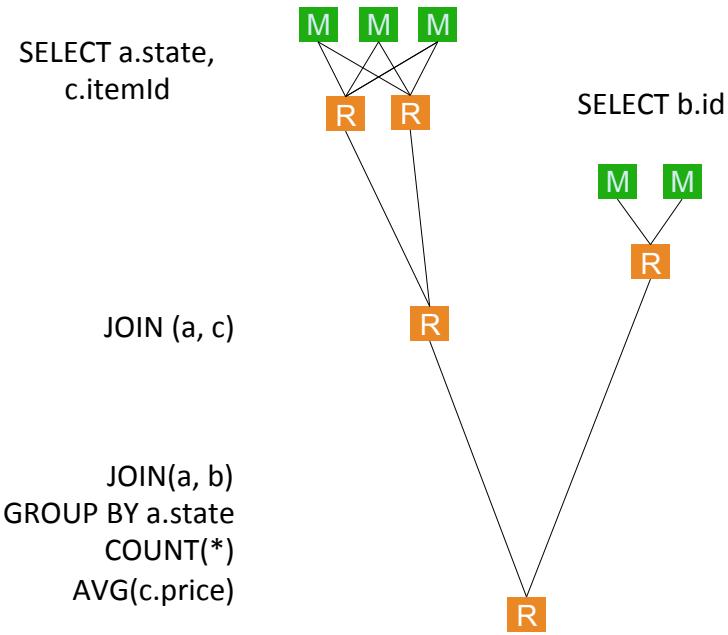

Hortonworks

Overview of Stinger

Performance Optimizations

100X+ Faster Time to Insight

Deeper Analytical Capabilities

Base Optimizations

Generate simplified DAGs
In-memory Hash Joins

YARN

Next-gen Hadoop data processing framework

+

Tez

Express tasks more simply
Eliminate disk writes
Pre-warmed Containers

ORCFile

Column Store
High Compression
Predicate / Filter Pushdowns

+

Vector Query Engine

Optimized for modern processor architectures

Query Planner

Intelligent Cost-Based Optimizer

Hortonworks

Knox: Gateway Hadoop Security

- **All users and services see one, secure end-point website or RESTful integration service**
- **Consistency across all interfaces and capabilities**
- **Safe, Firewalled cluster that no end users need to access**

Simplify Administration

Simplify Operations

Reduce Exposure & Risk

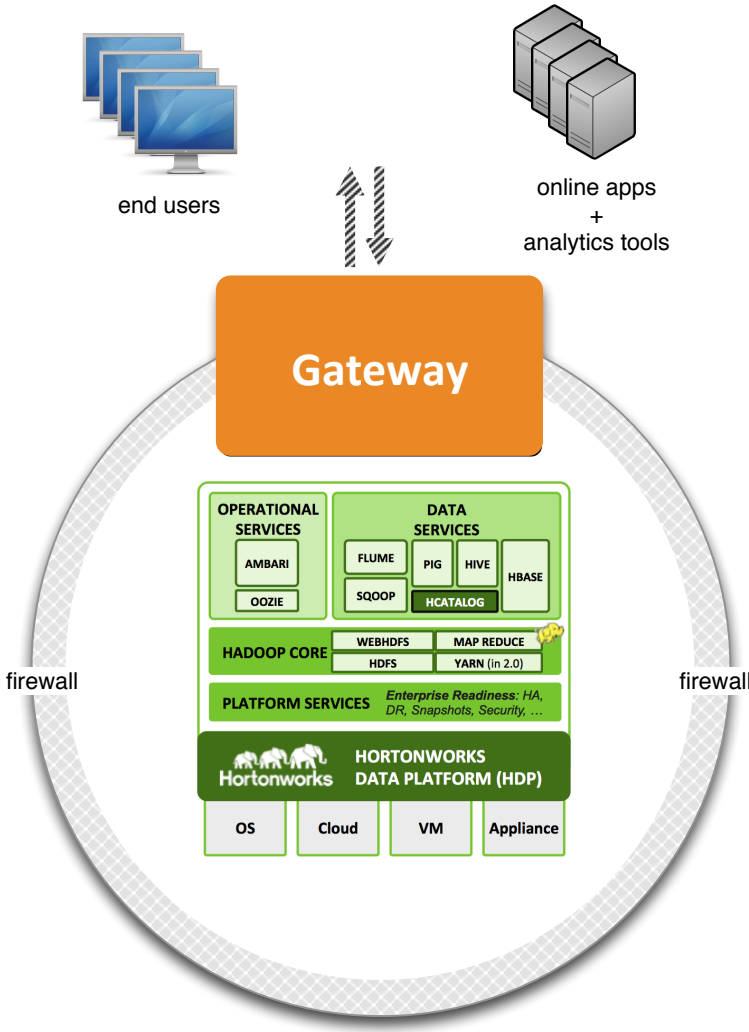

Hortonworks

Thank you!

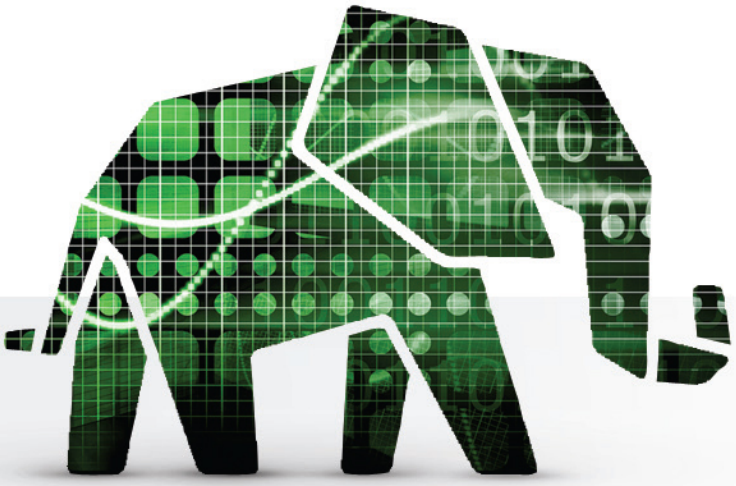

QCon

Please evaluate my talk via the mobile app!