11/16/2011, Stanford EE380 Computer Systems Colloquium

Introducing Apache Hadoop: The Modern Data Operating System

Dr. Amr Awadallah | Founder, CTO, VP of Engineering

aaa@cloudera.com, twitter: @awadallah

cloudera

Limitations of Existing Data Analytics Architecture

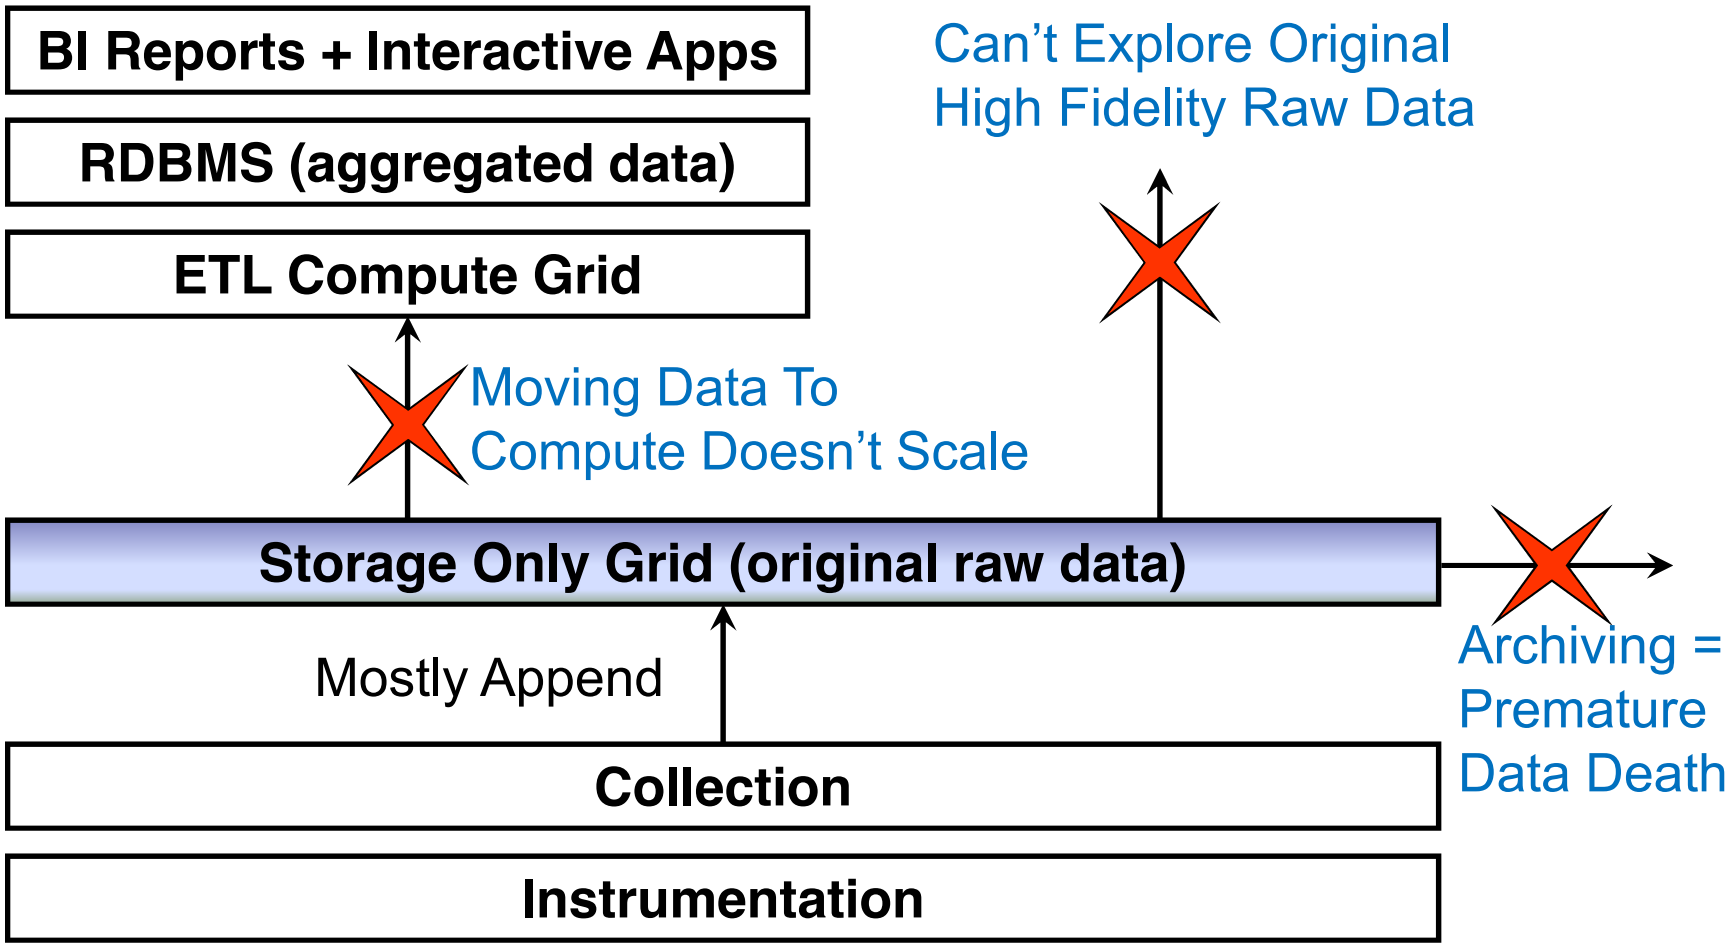

cloudera

So What is Apache  ?

- A scalable fault-tolerant distributed system for data storage and processing (open source under the Apache license).

- Core Hadoop has two main systems:
  - **Hadoop Distributed File System**: self-healing high-bandwidth clustered storage.
  - **MapReduce**: distributed fault-tolerant resource management and scheduling coupled with a scalable data programming abstraction.

HDFS: Hadoop Distributed File System

A given file is broken down into blocks (default=64MB), then blocks are replicated across cluster (default=3).

Optimized for:

- Throughput
- Put/Get/Delete
- Appends

Block Replication for:

- Durability
- Availability
- Throughput

Block Replicas are distributed across servers and racks.

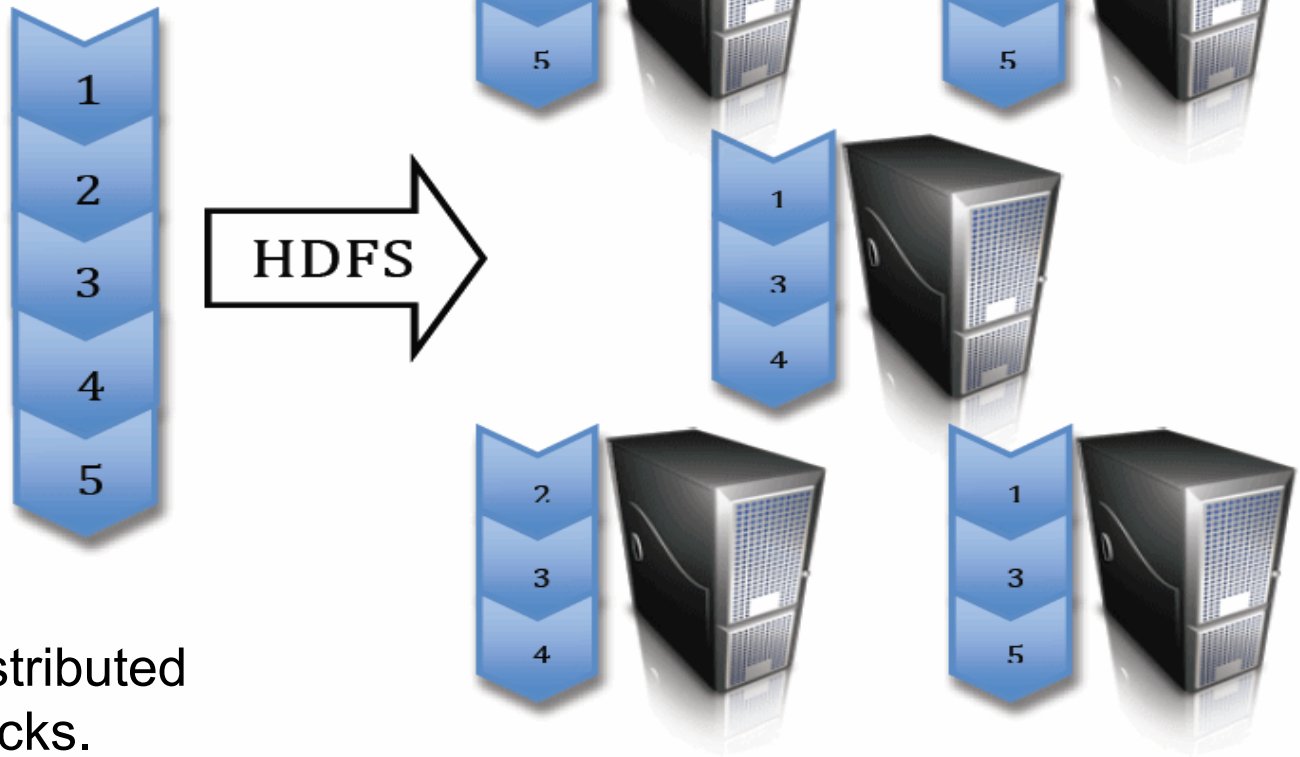

cloudera

MapReduce: Resource Manager / Scheduler

A given job is broken down into tasks, then tasks are scheduled to be as close to data as possible.

Three levels of data locality:

- Same server as data (local disk)
- Same rack as data (rack/leaf switch)
- Wherever there is a free slot (cross rack)

Optimized for:

- Batch Processing
- Failure Recovery

System detects *laggard* tasks and speculatively executes parallel tasks on the same slice of data.

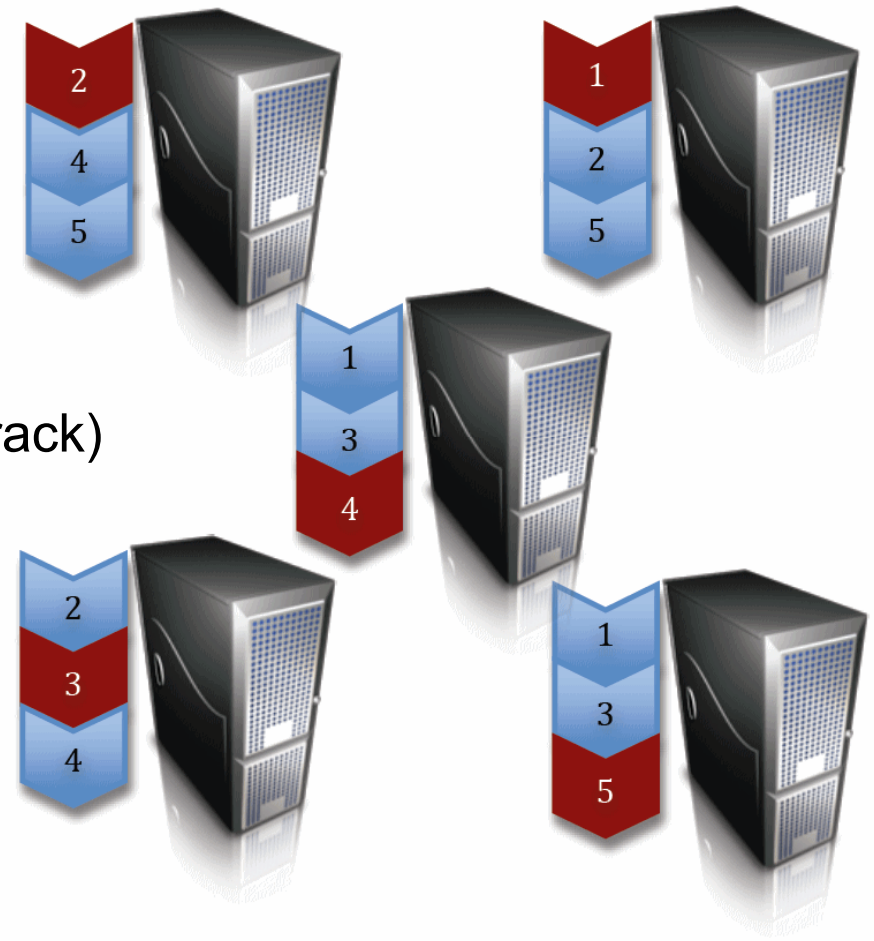

cloudera

Scalability: Scalable Software Development

Grows without requiring developers to re-architect their algorithms/application.

AUTO SCALE

cloudera

Use The Right Tool For The Right Job

Relational Databases:

Use when:

- **Interactive OLAP Analytics (<1sec)**
- **Multistep ACID Transactions**
- **100% SQL Compliance**

Hadoop:

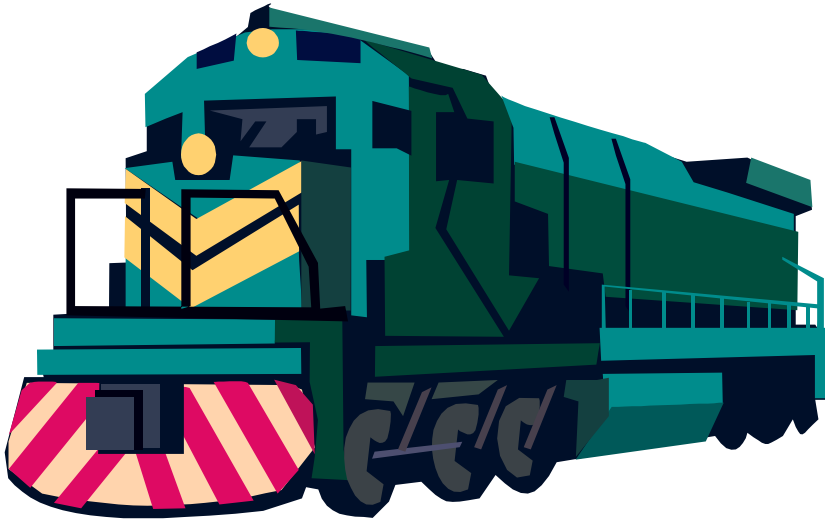

Use when:

- **Structured or Not (Flexibility)**
- **Scalability of Storage/Compute**
- **Complex Data Processing**

cloudera

Hadoop High-Level Architecture

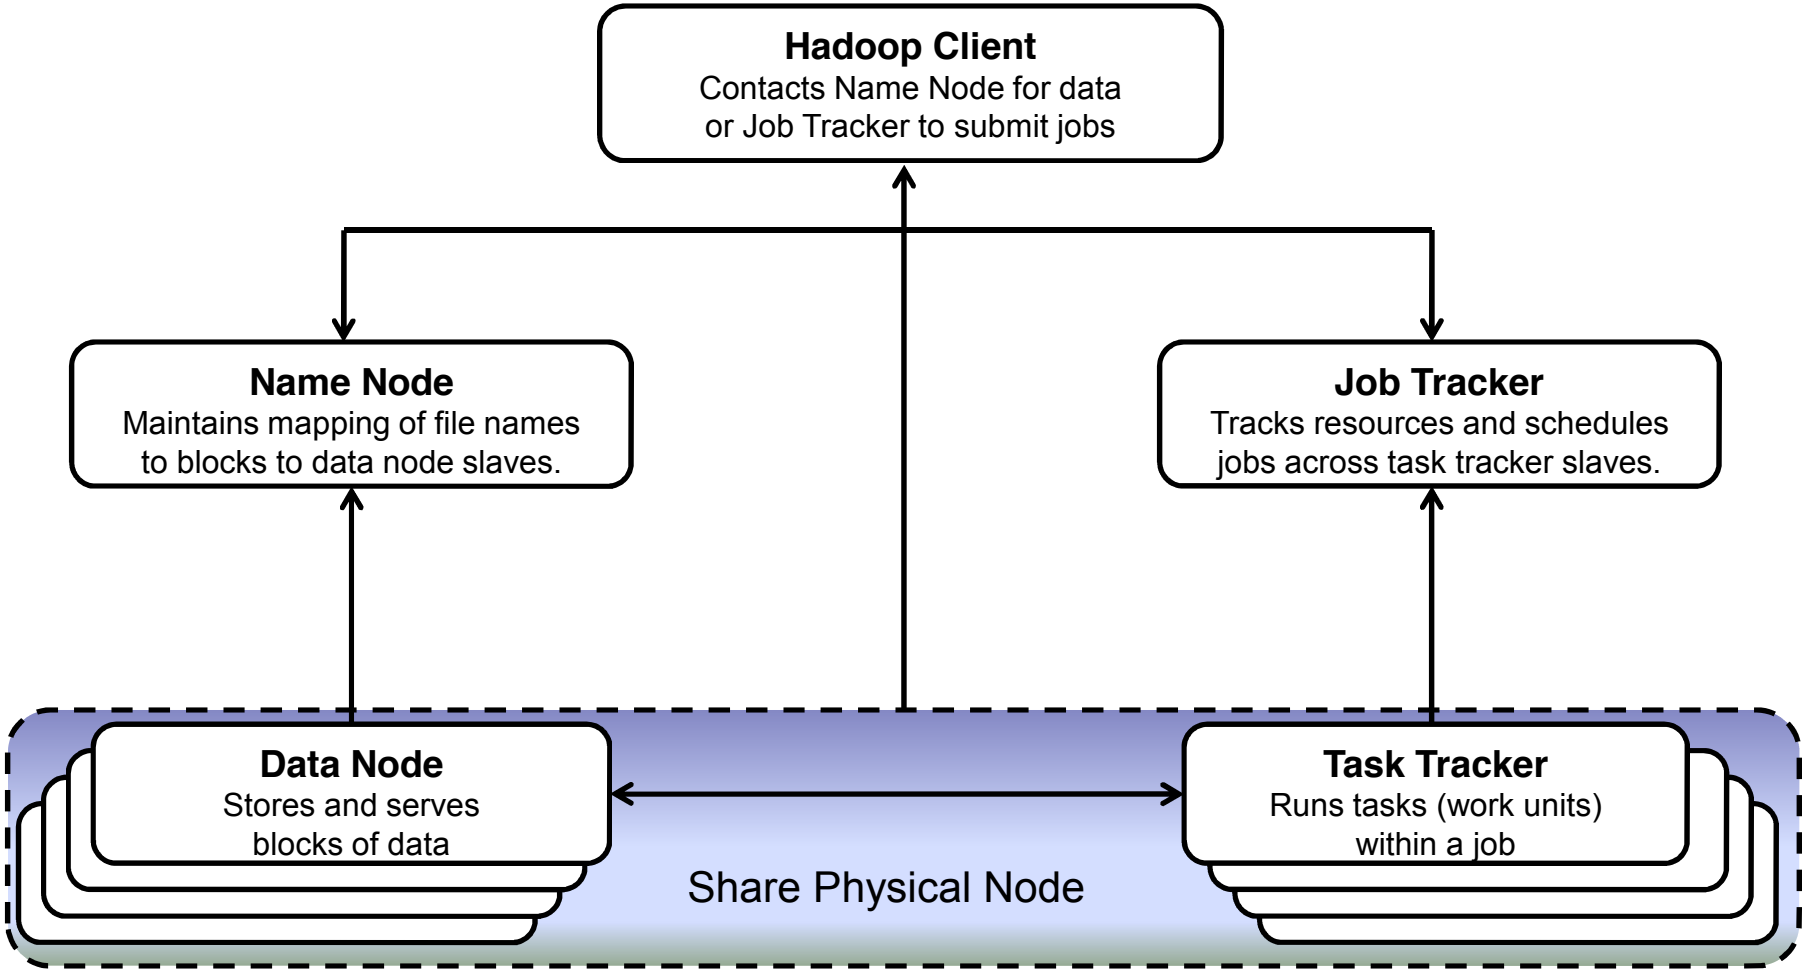

cloudera

Conclusion

- **The Key Benefits of Apache Hadoop:**
  - **Agility/Flexibility** (Quickest Time to Insight).
  - **Complex Data Processing** (Any Language, Any Problem).
  - **Scalability of Storage/Compute** (Freedom to Grow).
  - **Economical Storage** (Keep All Your Data Alive Forever).

- **The Key Systems for Apache Hadoop are:**
  - **Hadoop Distributed File System**: self-healing high-bandwidth clustered storage.
  - **MapReduce**: distributed fault-tolerant resource management coupled with scalable data processing.

cloudera

Appendix

BACKUP SLIDES

cloudera

But Networks Are Faster Than Disks!

Yes, however, core and disk density per server are going up very quickly:

- 1 Hard Disk = 100MB/sec (~1Gbps)
- Server = 12 Hard Disks = 1.2GB/sec (~12Gbps)
- Rack = 20 Servers = 24GB/sec (~240Gbps)
- Avg. Cluster = 6 Racks = 144GB/sec (~1.4Tbps)
- Large Cluster = 200 Racks = 4.8TB/sec (~48Tbps)
- Scanning 4.8TB at 100MB/sec takes 13 hours.

cloudera

Unstructured Data is Exploding

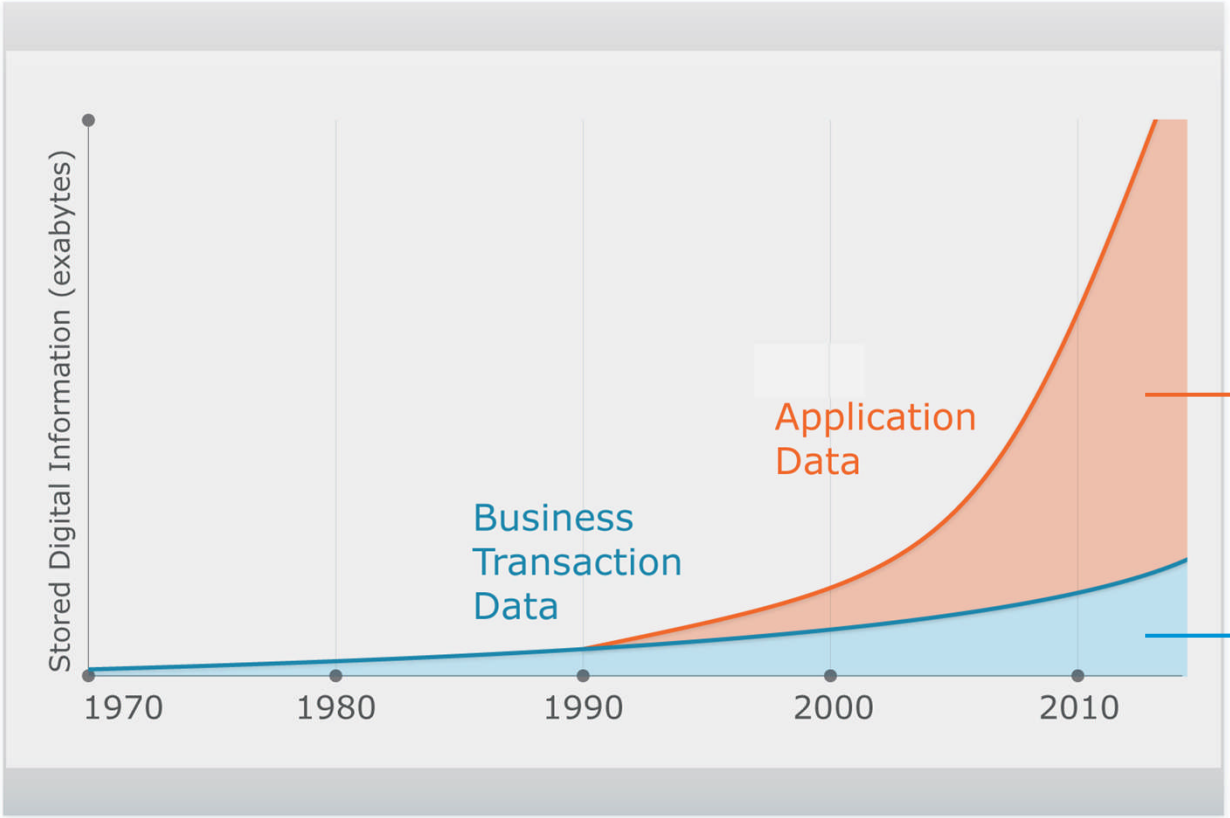

- 2,500 exabytes of new information in 2012 with Internet as primary driver
- Digital universe grew by 62% last year to 800K petabytes and will grow to 1.2 "zettabytes" this year

Source: IDC White Paper - sponsored by EMC.
As the Economy Contracts, the Digital Universe Expands. May 2009.

cloudera

Hadoop Creation History

Doug Cutting adds DFS & MapReduce support to Nutch

Doug Cutting & Mike Cafarella started working on Nutch

Google publishes GFS & MapReduce papers

Yahoo! hires Cutting, Hadoop spins out of Nutch

NY Times converts 4TB of image archives over 100 EC2s

Facebooks launches Hive: SQL Support for Hadoop

Fastest sort of a TB, 3.5mins over 910 nodes

cloudera Founded

• Fastest sort of a TB, 62secs over 1,460 nodes
• Sorted a PB in 16.25hours over 3,658 nodes

Doug Cutting joins Cloudera

Hadoop Summit 2009, 750 attendees

2002 2003 2004 2005 2006 2007 2008 2009

cloudera