



Hadoop




- Oracle -> Hadoop @CERN / Atlas
- Hadoop = HDFS + Map/Reduce
- API:
 - Native
 - Pig, Grunt, Hive
 - HBase
 - Web Service

Oracle->Hadoop @CERN



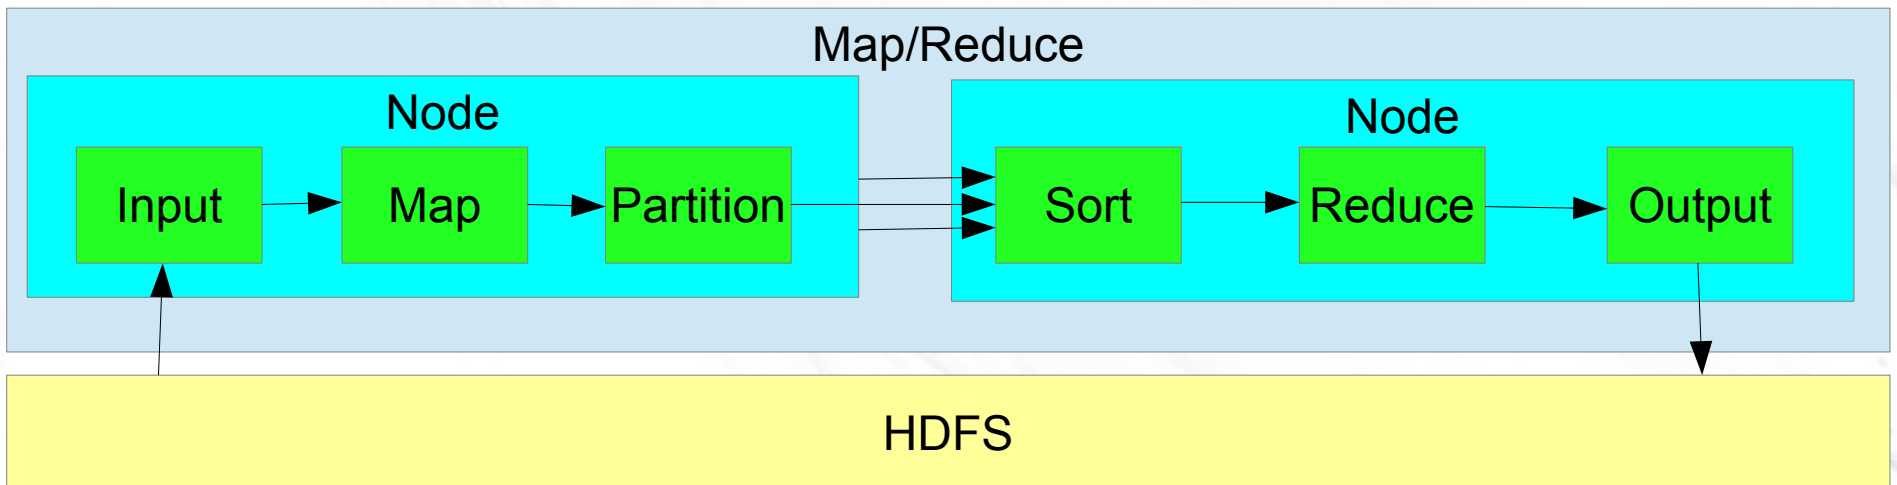
- Oracle:
 - Oracle doesn't seem to satisfy our performance requirements, needs a lot of tuning by highly experienced staff
 - Oracle is expensive (CERN+Tier1)
 - Our data are not table-like (SQL), but column-wise or unstructured (NoSQL)
- Some preliminary tests with Hadoop done last year in CERN & Atlas – with very positive results
- Hadoop training (5 days) organized
- Atlas seems to have decided to migrate its large SQL databases from Oracle to Hadoop
 - starts with TAG DB, may be followed by Conditions, Geometry, ...
 - during machine upgrade (2013)

Hadoop = HDFS + Map/Reduce




➤ HDFS:

- Distributed file storage (each data component replicated, failover capability)
 - Transparent access from client
 - Many file formats supported, others can be added
- ## ➤ Map/Reduce:
- Tasks are executed on servers carrying data
 - Results are recombined and consolidated
 - Cheap hw
 - Small tasks slow, big tasks fast
 - Data represented by key+value pairs




API




- Written in **Java + Ruby** scripting
 - native API Java
 - all **JVM** languages transparently supported (Python, Ruby, Groovy, Scala,...)
 - more obscure language supported via opaque API or streaming API
- Special purpose languages (client translates job into JAR file and sends it for execution to Hadoop):
 - **Pig**: Map/Reduce
 - **Sqoop**: interface to SQL db
 - **Hive**: SQL-like
- **Templeton**: HTTP REST Web Service

HBase



- NoSQL database
- Interactive
- Schema-free
- Three-dimensional: key-value-timestamp
- No transactions, not ACID
- Stores in HDFS

@ Atlas



➤ 2012:

- ServiceCatalog implementation to test Hbase – nice API
- Performance evaluation of HDFS (storing big Root files) – excellent performance

➤ 2013:

- New Hadoop cluster in CERN/IT (Linux)
 - (there is already a Sun cluster used by LHC machine)
- TAG DB being replicated into HDFS using different ways
- Evaluation of storage architecture
- Workshop in 2 weeks