


Integration of Oracle and Hadoop: hybrid databases affordable at scale

Luca Canali, Zbigniew Baranowski, Prasanth Kothuri
CERN IT, Geneva (CH)



22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016


Advantages of integrating Oracle and Hadoop

- Best of two worlds:
 - Oracle, optimized for Online Transactional System
 - Hadoop, scalable distributed data processing platform
- Hybrid systems:
 - Move (read-only) data from Oracle to Hadoop
 - Query Hadoop data from Oracle (using Oracle APIs)
 - Also possible: query Oracle from Hadoop
- Increase **scalability** and lower ratio **cost**/performance
 - Hadoop data formats and engines for high performance analytics
 -without need of changing the end-user apps connecting to Oracle

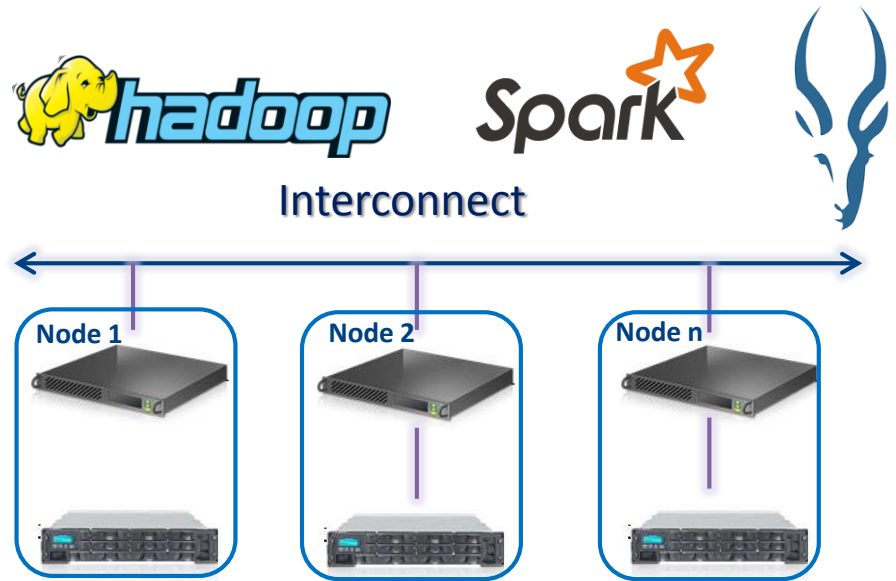
ORACLE®



Oracle optimized for OLTP, Hadoop affordable at scale



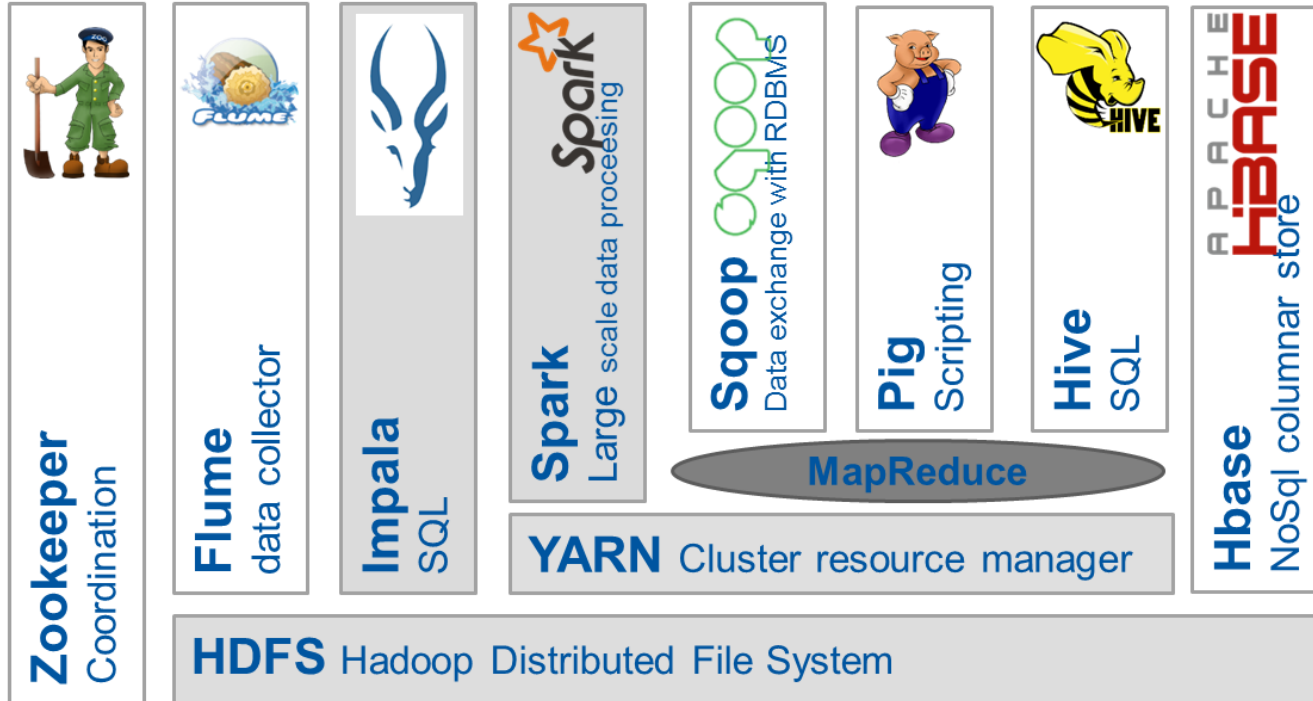
Example of Oracle RAC deployed with **shared storage**



The **shared nothing** architecture allows to **scale** for high capacity and throughput on commodity HW

The Hadoop ecosystem is heterogeneous and evolving

Main Hadoop components at the CERN-IT Hadoop service (2016):



Hadoop data formats and data ingestion

- Hadoop data formats
 - Are an important dimension to the architecture
 - Compression and encoding for analytic workloads
 - Columnar formats (Parquet)
- Data ingestion also very important
 - Sqoop to transfer from databases
 - Kafka, very successful for message-oriented ingestion



Examples from CERN and HEP

- ATLAS DDM for reports
- FCC reliability studies
- Accelerator and industrial controls at CERN
- Archive and reporting use cases to be explored

- Example:
 - Speedup of a reporting query from CERN Network experts
 - Running in Oracle in 12 hours (no parallel query allowed in prod)
 - Moved to Hadoop and run in parallel in minutes (throw HW to the problem)

The biggest (RDBMS) database at CERN


- CERN Accelerator Logging System
 - An archive system for metrics from most of devices and systems installed and used by LHC
 - **500 TB** in Oracle
 - 500 GB produced per day (15 billion data points)
- Offloaded to Hadoop with Apache Sqoop
 - Daily export with 10 parallel streams, duration 3 hours (40 MB/s)
 - **Parquet** format with Snappy compression: factor 3.3
 - Size on HDFS: 140 TB

Techniques for integrating Oracle and Hadoop

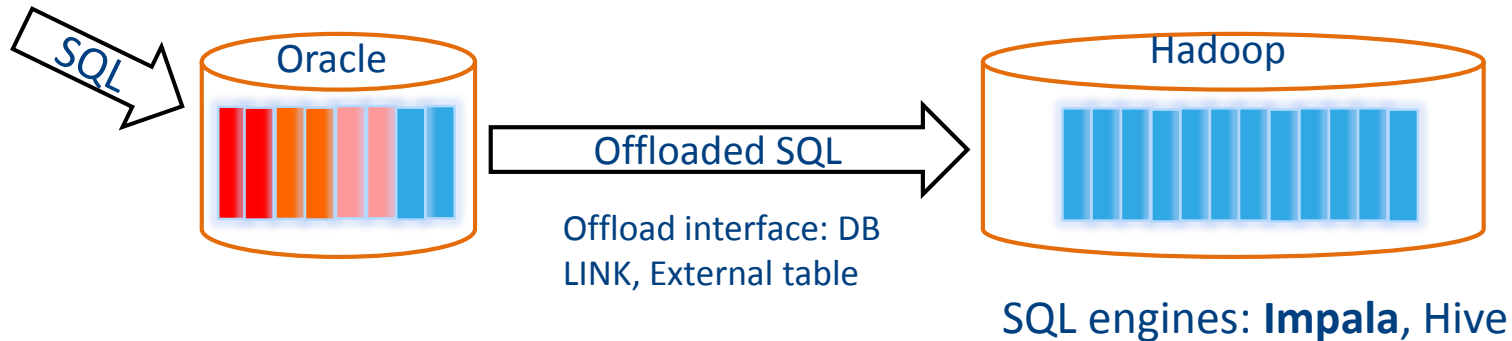
- Export data from Oracle to HDFS
 - **Sqoop** good enough for most cases
 - Other options possible (custom ingestion, Oracle DataPump, streaming, ..)
- Query Hadoop from Oracle
 - Access tables in Hadoop engines using **DB links** in Oracle
 - Build **hybrid views**: transparently combine data in Oracle and Hadoop
- Use Hadoop frameworks to process data in Oracle DBs
 - Use Hadoop engines (Impala, Spark) to process **data exported** from Oracle
 - Read data in a RDBMS **directly** from Spark SQL with JDBC

Offloading from Oracle to Hadoop

- Step1: Offload **data** to Hadoop



- Step2: Offload **queries** to Hadoop

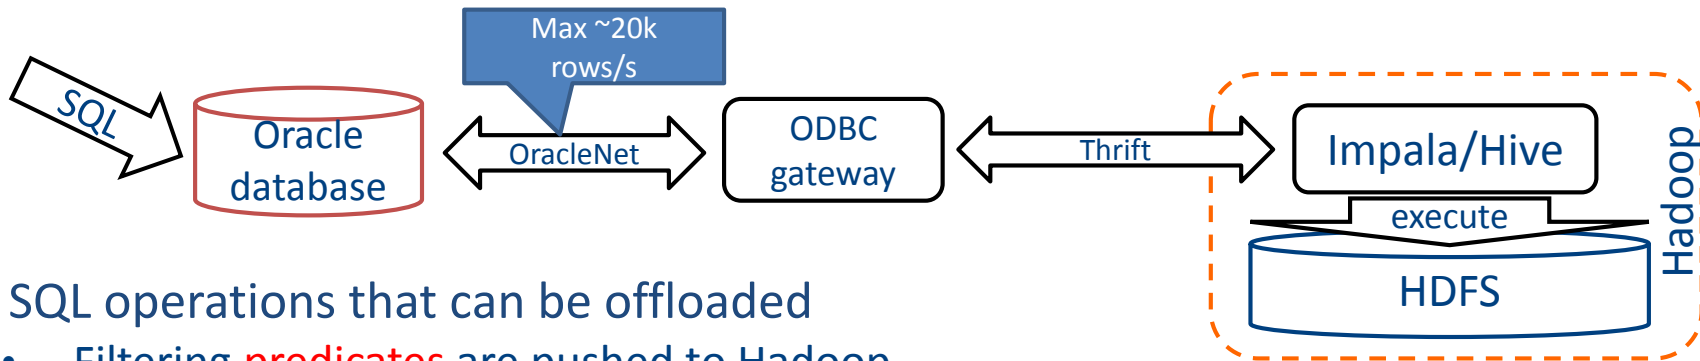


How to access Hadoop from an Oracle query

- Query Apache Hive/Impala tables from Oracle using a **database link**

```
create database link my_hadoop using 'impala-gateway';  
select * from big_table@my_hadoop where coll = :vall;
```

- Query offloaded via ODBC gateway to Impala (or Hive)



- SQL operations that can be offloaded
 - Filtering **predicates** are pushed to Hadoop
 - Problem: grouping aggregates are not pushed
- There are techniques to work around this problem
 - Create aggregation with views in Hive/Impala
 - DBMS_HS_PASSTHROUGH – to push exact SQL statement to Hadoop

Making data sources transparent to end-user

- Hybrid views on Oracle
 - **Recent** (read-write) data in **Oracle**
 - **Archive** data in **Hadoop**
 - **Advantages:** Hadoop performance with unchanged applications (Oracle APIs)

Split point has to be updated after each successful data offload

```
create view hybrid_view as
  select * from online_table where date > '2016-10-01'
union all
  select * from archival_table@hadoop where date <= '2016-10-01'
```

Specialized products for hybrid DB and offloading

- Oracle BigData SQL
 - Custom engine to query Hadoop from Oracle
 - Based on Oracle external tables (predicate pushing)
- Gluent Inc
 - Similar functionality to BigData SQL
 - Uses Apache Impala to process data on Hadoop
 - Leverages **hybrid** views on Oracle for data integrity
 - Predicate pushing and partition pruning
 - Data retrieval >10x **faster** than ODBC gateway

The Oracle logo, consisting of the word "ORACLE" in a bold, red, sans-serif font with a registered trademark symbol.The Gluent logo, consisting of the word "gluent." in a bold, green, sans-serif font with a period at the end.

Query Oracle from Apache Spark

- Spark SQL using JDBC to query Oracle directly
 - To access metadata and/or lookup tables
 - Use to retrieve metadata that can quickly become stale
- Example code to query Oracle into a Spark DataFrame (Python):

```
df = sqlContext.read.format('jdbc').options(  
    url="jdbc:oracle:thin:@ORACLE_DB/orcl.cern.ch",  
    user="myuser",  
    password="mypass",  
    fetchSize=1000,  
    dbtable="(select id, payload from my_oracle_table) df",  
    driver="oracle.jdbc.driver.OracleDriver"  
).load()
```

Conclusions

- **Hadoop** performs at scale, excellent for data **analytics**
- **Oracle** proven for concurrent **transactional** workloads
- Solutions are available to **integrate** Oracle and Hadoop

- There is value in hybrid systems (Oracle + Hadoop):
 - Oracle APIs for legacy applications and **OLTP** workloads
 - **Scalability** on commodity HW for analytic workloads

Hadoop Service at CERN

- **Service** provided by CERN-IT for Experiments and CERN users
- Projects ongoing with Experiments, Accelerators sector and IT
- Hadoop Users Forum for open discussions: subscribe to egroup **it-analytics-wg**
- Getting started material: Hadoop **tutorials** <https://indico.cern.ch/event/546000/>

- Related talks/posters at CHEP 2016:
 - First results from a combined analysis of CERN computing infrastructure metrics, talk on Tuesday at 12:00
 - A study of data representations in Hadoop to optimize data storage and search performance of the ATLAS EventIndex, poster on Tuesday 16:30
 - Big Data Analytics for the Future Circular Collider Reliability and Availability Studies, talk on Thursday at 14:15
 - Hadoop and friends - first experience at CERN with a new platform for high throughput analysis steps, talk on Thursday at 14:45
 - Developing and optimizing applications for the Hadoop environment, talk on Thursday at 15:15