CERN openlab

Evaluation of the Presto Query Engine for integrating relational databases with big data platforms at scale

Lightning Talk

Andrew Waldman

14/08/2018

WHAT IS PRESTO?

A distributed SQL query engine
- Typically run on top of a Hadoop cluster
- Open source
- Already well established in the cloud sector

SQL-on-Anything
- HDFS (Parquet, Avro etc.), S3
- Relational DBS (Oracle, MySQL, PostgreSQL, SQLServer…)
- NoSQL (Cassandra, Kudu)
- Apache Kafka and more (data sources are pluggable)
- Local File System

You can Query different data sources from one query!!!

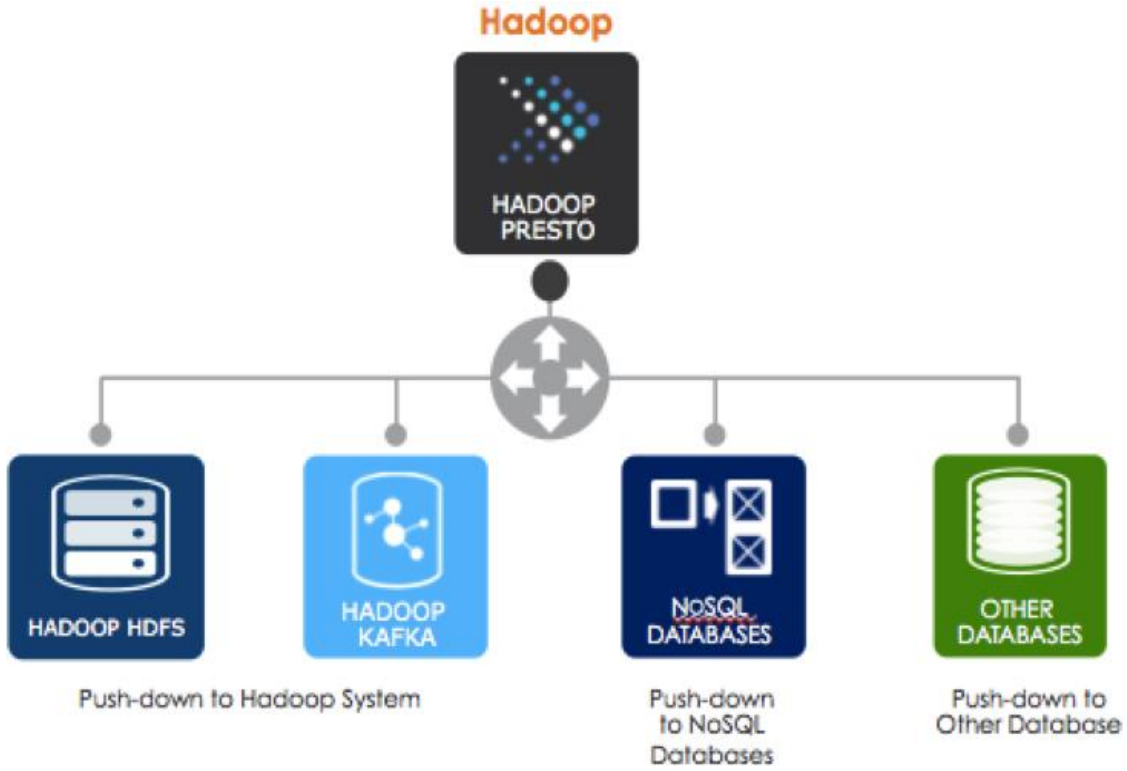

CERN openlab

Objectives of the Project

Is it worth adding to the current Hadoop service portfolio?

- Get Presto running (installation and configuration)
- Performance comparison on different data sets (with current frameworks)
- Evaluation of native and open source connectors
- Usability of PaaS (Presto as a Service)

Control System Dataset Tests (WinCC*)

- Begin evaluating Presto
  - Using Hive connector to connect to Hive metastore
  - Run 5 different types of queries varying in complexity
  - Using numerous sets of different resource configurations

- 4 configurations
  - Single core, CPU and Thread
  - 4 Machines, single CPU and single Thread
  - 8 Machines, single CPU and single Thread
  - 8 Machines, all CPU and Multi Threading

```
with daily as (select day,stddev(value_number)
dev,element_id from psen.eventhistory_00000008 group by
day,element_id) select element_id,stddev(dev) from
daily group by element_id having stddev(dev)>100000
order by 2;
```

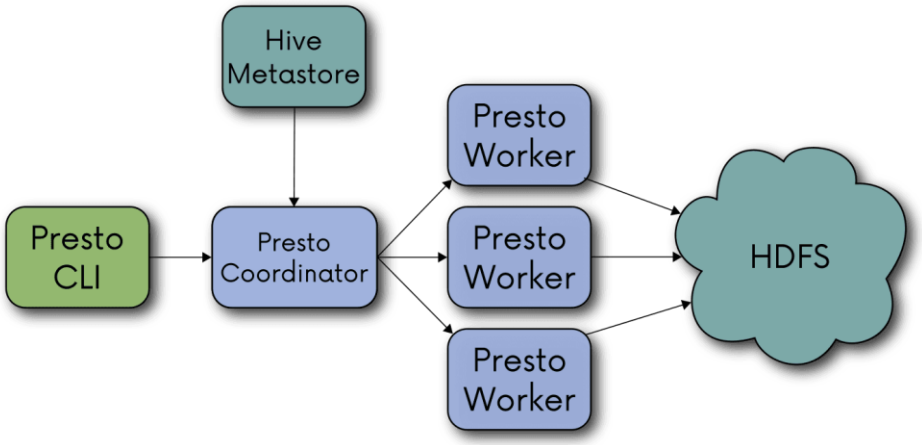

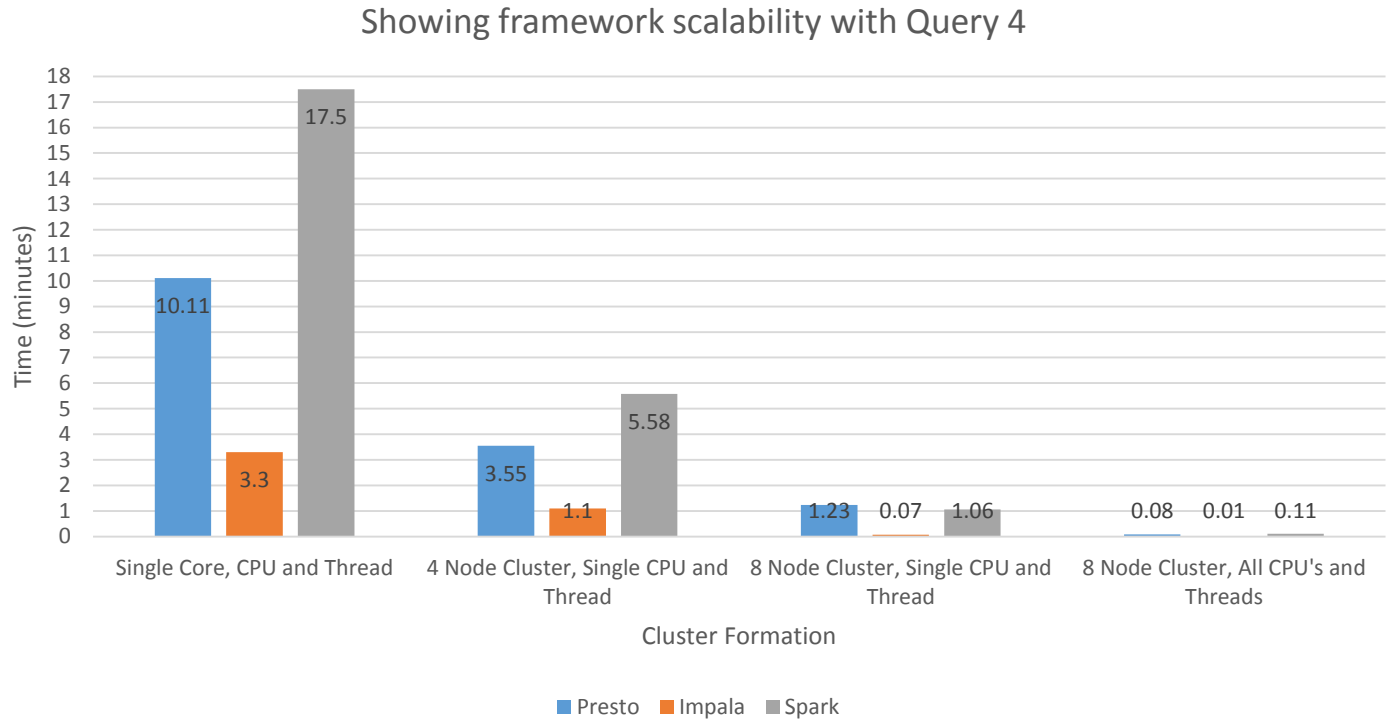

*Time is x/100 after decimal not seconds

*CERN production data set

TPC-DS Test Set 2

100GB (Presto, Impala, Spark)

- Same set of queries as for the previous test set
- Automated result collection
- Ran on the Hadalytic cluster

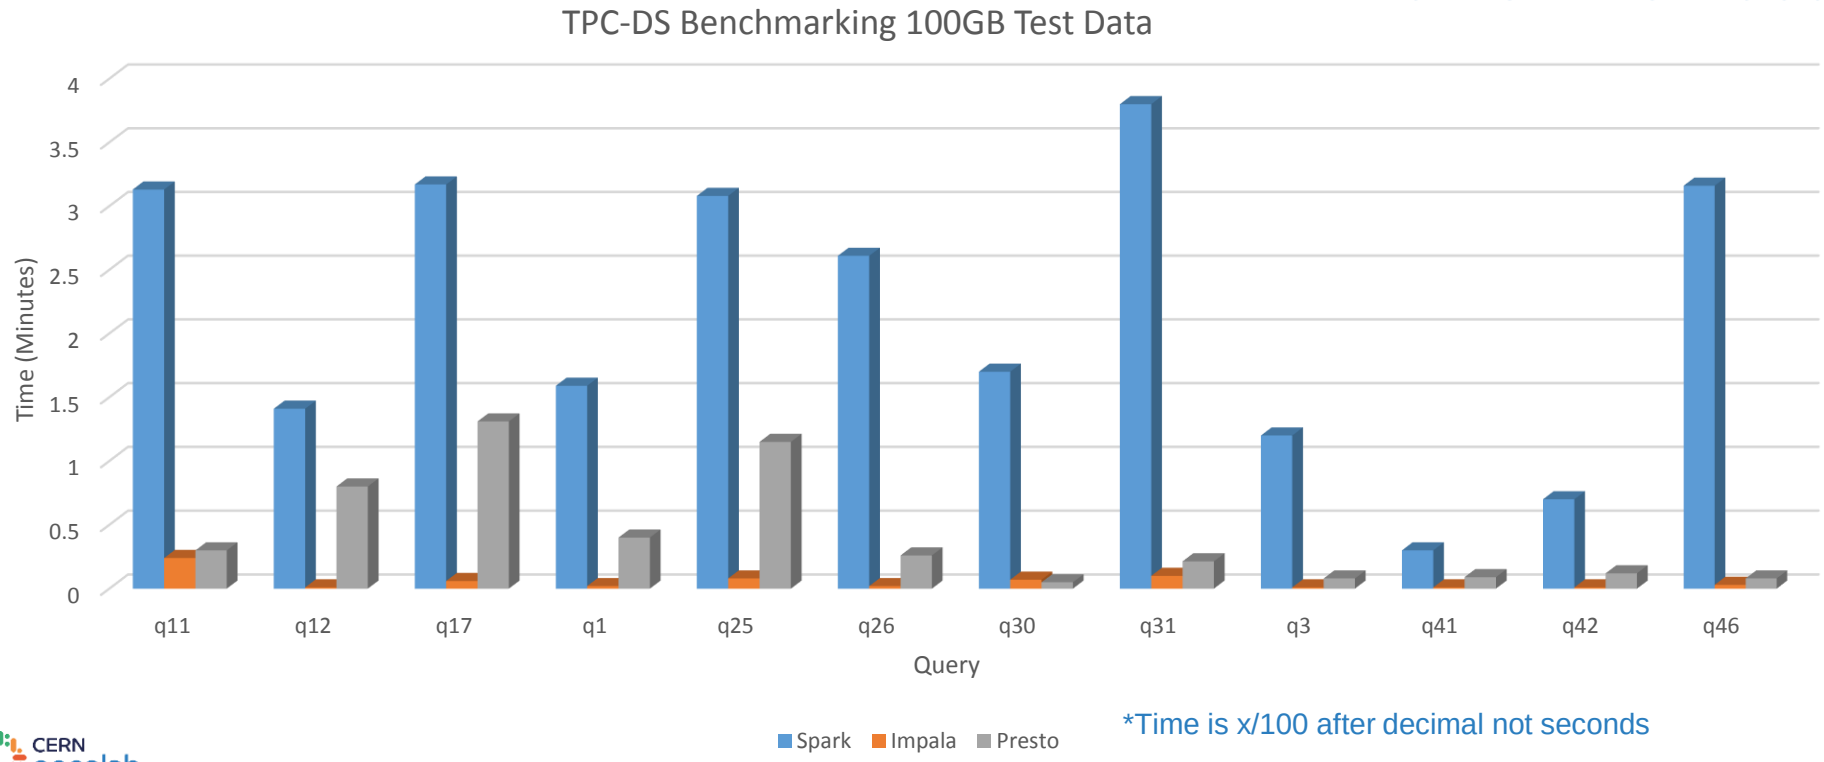

*Time is x/100 after decimal not seconds

Connectors and Catalogs

- A connector is software to make a connection between Presto and endpoint
- A catalog is an instance of a connector (a configured connector to access particular endpoint).

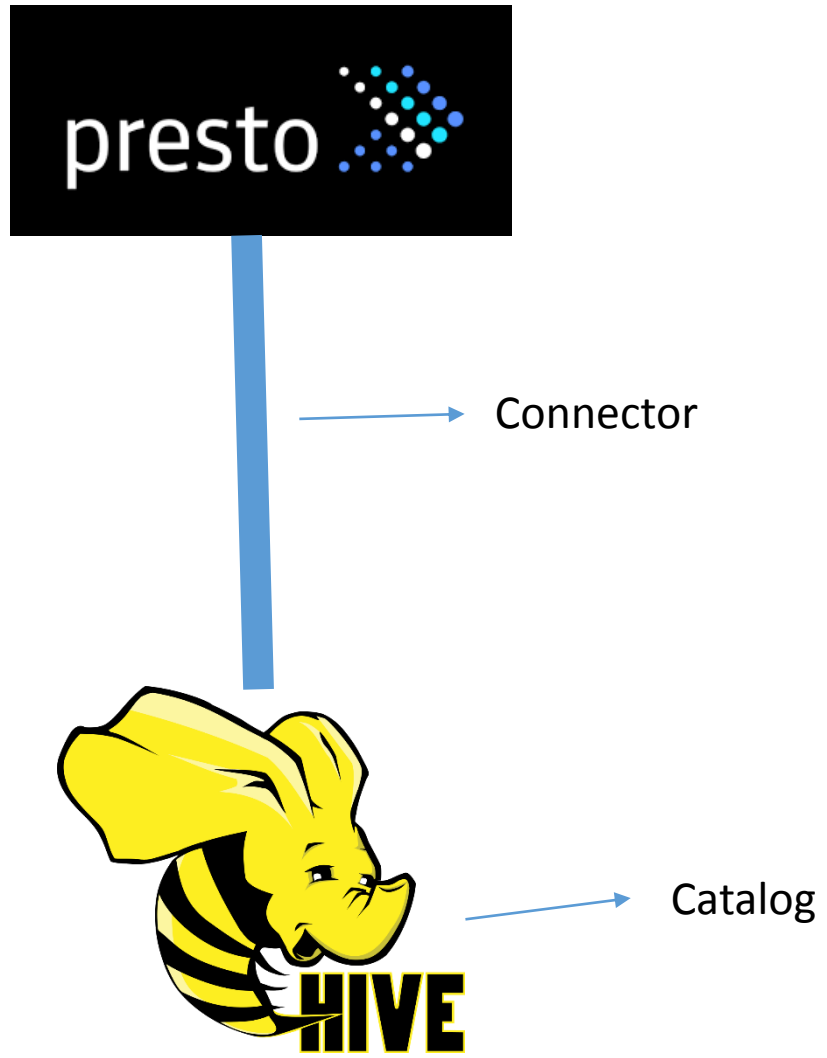

Summary of Connector Testing

- Native connectors where very easy to configure and get working
- The Native connectors could perform aggregations and predicate queries with no problems
- The open source connectors where less consistent in working than the Native ones. With the Oracle open source connectors not working and so preventing full evaluation

Connector	Successful connection to endpoint	Querying Tables	Filter predicate pushdown	Scale out (Input Splits)
Marcelopaesrech Oracle				
ITGrocery Oracle	✔	✔		
MartinWeindel Kudu	✔	✔	✔	✔
MySQL	✔	✔	✔	✔
Kafka	✔	✔		✔
Cassandra	✔	✔	✔	✔

Conclusion

- The presto query engine when running against TPC-DS benchmark data and WinCC data, performance was slightly behind the current framework, Impala
- However, when it comes to compatibility and being able to connect to multiple data sources and query data even at once
- Presto is more future-proof.
  Also eases building of hybrid systems (OLTP + archive)
- This leads to the conclusion that it would be worth considering as an additional framework to be added to production alongside Spark and Impala

QUESTIONS?

Andrew.waldman@cern.ch

CERN openlab

CONTACTS

ANDREW WALDMAN
CERN openlab Student
andrew.waldman@cern.ch

ZBIGNIEW BARANOWSKI
CERN openlab Supervisor
zbigniew.baranowski@cern.ch

EMIL KLESZCZ
CERN IT-DB
emil.kleszcz@cern.ch

ALBERTO DI MEGLIO
CERN openlab Head
alberto.di.meglio@cern.ch

ANDREW PURCELL
CERN openlab Communications Officer
andrew.purcell@cern.ch

KRISTINA GUNNE
CERN openlab Administration/Finance Officer
kristina.gunne@cern.ch

www.cern.ch/openlab