
Delft University of Technology
Parallel and Distributed Systems Report Series

An Empirical Performance Evaluation of
Distributed SQL Query Engines: Extended Report

Stefan van Wouw, José Viña,
Alexandru Iosup, and Dick Epema
stefanvanwouw@gmail.com, jose@azavista.com,
{a.iosup,d.h.j.epema}@tudelft.nl

Completed September 2014.

Report number PDS-2014-002



ISSN 1387-2109

Published and produced by:
Parallel and Distributed Systems Group
Department of Software and Computer Technology
Faculty Electrical Engineering, Mathematics, and Computer Science
Delft University of Technology
Mekelweg 4
2628 CD Delft
The Netherlands

Information about Parallel and Distributed Systems Report Series:
reports@pds.ewi.tudelft.nl

Information about Parallel and Distributed Systems Section:
<http://www.pds.ewi.tudelft.nl/>

© 2014 Parallel and Distributed Systems Group, Department of Software and Computer Technology, Faculty Electrical Engineering, Mathematics, and Computer Science, Delft University of Technology. All rights reserved. No part of this series may be reproduced in any form or by any means without prior written permission of the publisher.



Abstract

Distributed SQL Query Engines (DSQEs) are increasingly used in a variety of domains, but especially users in small companies with little expertise may face the challenge of selecting an appropriate engine for their specific applications. Although both industry and academia are attempting to come up with high level benchmarks, the performance of DSQEs has never been explored or compared in-depth. We propose an empirical method for evaluating the performance of DSQEs with representative metrics, datasets, and system configurations. We implement a micro-benchmarking suite of three classes of SQL queries for both a synthetic and a real world dataset and we report response time, resource utilization, and scalability. We use our micro-benchmarking suite to analyze and compare three state-of-the-art engines, viz. Shark, Impala, and Hive. We gain valuable insights for each engine and we present a comprehensive comparison of these DSQEs. We find that different query engines have widely varying performance: Hive is always being outperformed by the other engines, but whether Impala or Shark is the best performer highly depends on the query type.



Contents

1	Introduction	3
2	Related Work	3
3	Query Engine Selection	4
4	Experimental Method	5
4.1	Workload	5
4.2	Performance Aspects and Metrics	6
4.3	Evaluation Procedure	7
5	Experimental Setup	7
6	Experimental Results	9
6.1	Processing Power	9
6.2	Resource Consumption	11
6.3	Resource Utilization over Time	12
6.4	Scalability	13
7	Conclusions and Future Work	15
A	Detailed Distributed SQL Query Engine Performance Metrics	18
B	Detailed Distributed SQL Query Engine Resource Utilization	20

1 Introduction

With the decrease in cost of storage and computation of public clouds, even small and medium enterprises (SMEs) are able to process large amounts of data. This causes businesses to increase the amounts of data they collect, to sizes that are difficult for traditional database management systems to handle. This has led to Hadoop-oriented distributed query engines such as Hive [17], Impala [5], Shark [19], and more recently, Presto [7], Drill [10], and Hive-on-Tez [3]. Selecting the most suitable of these systems for a particular SME is a big challenge, because SMEs are not likely to have the expertise and the resources available to perform an in-depth study. We remove this burden from SMEs by addressing the following research question: *How well do Distributed SQL Query Engines (DSQEs) perform on SME workloads?*

Although performance studies do exist for Distributed SQL Query Engines [1, 6, 8, 9, 14, 19], many of them only use synthetic workloads or very high-level comparisons that are only based on query response time. Our work evaluates performance much more in-depth by reporting more metrics and evaluating more performance aspects. In addition to reporting query response times, we also show scalability and detailed resource utilization. The latter performance aspects are particularly important for an SME in order to choose a query engine.

In order to answer the research question we define a comprehensive performance evaluation method to assess different aspects of query engines. We compare Hive, a somewhat older but still widely used query engine, with Impala and Shark, both state-of-the-art distributed query engines. This method can be used to compare current and future query engines, despite not covering all the methodological and practical aspects of a true benchmark. The method focuses on three performance aspects: processing power, resource utilization and scalability. With the results from this study, system developers and data analysts can make informed choices related to both cluster infrastructure and query tuning.

Using both a real world and a synthetic dataset with queries representative of SME workloads, we evaluate the query engines' performance. We find that different query engines have widely varying performance, with Hive always being outperformed by the other engines. Whether Impala or Shark is the best performer highly depends on the query type and input size.

Our main contributions are:

- We propose a method for performance evaluation of DSQEs (Section 4), which includes defining a workload representative for SMEs as well as defining the performance aspects of the query engines: processing power, resource utilization and scalability.
- We define a micro-benchmark setup for three major query engines, namely Shark, Impala and Hive (Section 5).
- We provide an in-depth performance comparison between Shark, Impala and Hive using our micro-benchmark suite (Section 6).

2 Related Work

We wanted to evaluate the major Distributed SQL Query Engines currently on the market using a cluster size and dataset size that is representative for SMEs, but still comparable to similar studies. Table 1 summarizes the related previous works. Some of them run a subset or enhanced version of the TPC-DS benchmark [15] which has only recently been adopted for Big Data analytics in the form of BigBench [9]. Other studies run a variant of the Pavlo et al. micro-benchmark [14] which is widely accepted in the field.

Table 1: Overview of Related Work. Legend: Real World (R), Synthetic (S), Modified Workload (+)

Query Engines	Workload	Dataset Type	Largest Dataset	Cluster Size
Hive, Shark [19]	Pavlo+, other	R, S	1.55 TiB	100
Redshift, Hive, Shark, Impala, Tez [1]	Pavlo+	S	127.5 GiB	5
Impala, Tez, Shark, Presto [6]	TPC-DS+	S	13.64 TiB	20
Teradata DBMS [9]	TPC-DS+	S	186.24 GiB	8
Hive, Impala, Tez [8]	TPC-DS/H+	S	220.72 GiB	20
DBMS-X, Vertica [14]	Pavlo	S	931.32 GiB	100
Our Work	Pavlo+, other	R, S	523.66 GiB	5

Overall, most studies use synthetic workloads, of which some are very large. Synthetic workloads do not necessarily characterise real world datasets very well. For our work we have also taken a real world dataset in use by an SME. Besides our work, only one other study uses real world datasets [19]. However, like most of the other studies, it only reports on query response times. Our work evaluates performance much more in-depth by reporting more metrics and evaluating more performance aspects including scalability and detailed resource utilization. We argue that scalability and resource utilization are also very important when deciding which query engine will be used by an SME.

3 Query Engine Selection

In this study we initially attempted to evaluate 5 state-of-the-art Distributed SQL Query engines: Drill, Presto, Shark, Impala and Hive. However, we ended discarding Drill and Presto because these systems lacked required functionality at the time of testing. Drill only had a proof of concept one node version, and Presto did not have the functionality needed to write output to disk (which is required for the kind of workloads we wanted to evaluate).

Shark [19] is a DSQE built on top of the Spark [21] execution engine, which in turn heavily relies on the concept of Resilient Distributed Datasets (RDDs) [20]. In short this means that whenever Shark receives an SQL query, it will convert it to a Spark job, execute it in Spark, and then return the results. Spark keeps all intermediate results in memory using RDDs, and only spills them to disk if no sufficient memory is available. Mid-query fault tolerance is provided by Spark. It is also possible to have the input and output dataset cached entirely in memory.

Impala [5] is a DSQE being developed by Cloudera and is heavily inspired by Google’s Dremel [13]. It employs its own massively parallel processing (MPP) architecture on top of HDFS instead of using Hadoop MapReduce as execution engine (like Hive below). One large downside to this engine is that it does not provide fault tolerance. Whenever a node dies in the middle of query execution, the whole query is aborted.

Hive [17] was one of the first DSQEs, introduced by Facebook and built on top of the Hadoop platform [2]. It provides a Hive Meta Store service to put a relational database-like structure on top of the raw data stored in HDFS. Whenever a HiveQL (SQL dialect) query is submitted to Hive, Hive will convert it to a job to be run on Hadoop MapReduce. Although Hive provides mid-query fault tolerance, it relies on Hadoop MapReduce and is slowed down whenever this system stores intermediate results on disk.

Table 2: Summary of Datasets.

Table	# Columns	Description
<code>uservisits</code>	9	Structured server logs per page.
<code>rankings</code>	3	Page rank score per page.
<code>hotel_prices</code>	8	Daily hotel prices.

4 Experimental Method

In this section we present the method of evaluating the performance of Distributed SQL Query Engines. First we define the workload as well as the aspects of the engines used for assessing this performance. Then we describe the evaluation procedure.

4.1 Workload

During the performance evaluation we use both synthetic and real world datasets with three SQL queries per dataset. We carefully selected the different types of queries and datasets to match the scale and diversity of the workloads SMEs deal with.

1) *Synthetic Dataset*: Based on the benchmark from Pavlo et al. [14], the UC Berkeley AMPLab introduced a general benchmark for DSQEs [1]. We have used an adapted version of AMPLab’s Big Data benchmark where we leave out the query testing User Defined Functions (UDFs), since not all query engines support UDF in similar form. The synthetic dataset used by these 3 queries consists of 118.29 GiB of structured server logs per URL (the `uservisits` table), and 5.19 GiB of page ranks (the `rankings` table) per website, as seen in Table 2.

Is this dataset representative for SME data? The structure of the data closely resembles the structure of click data being collected in all kinds of SMEs. The dataset size might even be slightly large for SMEs, because as pointed out by Rowstron et al. [16] analytics production clusters at *large* companies such as Microsoft and Yahoo have median job input sizes under 13.03 GiB and 90% of jobs on Facebook clusters have input sizes under 93.13 GiB.

On this dataset, we run queries 1 to 3 to test raw data processing power, aggregation and *JOIN* performance respectively. We describe each of these queries below in addition to providing query statistics in Table 3.

Query 1 performs a data scan on a relatively small dataset. It simply scans the whole `rankings` table and filters out certain records.

Query 2 computes the sum of ad revenues generated per visitor from the `uservisits` table in order to test aggregation performance.

Query 3 joins the `rankings` table with the `uservisits` table in order to test *JOIN* performance.

2) *Real World Dataset*: We collected price data of hotel rooms on a daily basis during a period of twelve months between November 2012 and November 2013. More than 21 million hotel room prices for more than 4 million hotels were collected on average every day. This uncompressed dataset (the `hotel_prices` table) is 523.66 GiB on disk as seen in Table 3. Since the price data was collected every day, we decided to partition the dataset in daily chunks as to be able to only use data of certain collection days, rather than having to load the entire dataset all the time.

Is this dataset representative for SME data? The queries we selected for the dataset are in use by Azavista, an SME specialized in meeting and event planning software. The real world scenarios for these queries relate to reporting price statistics per city and country.

Table 3: Summary of SQL Queries.

Query	Input Size		Output Size		Tables
	<i>GiB</i>	<i>Records</i>	<i>GiB</i>	<i>Records</i>	
1	5.19	90M	5.19	90M	rankings
2	118.29	752M	40	254M	uservisits
3	123.48	842M	$< 10^{-7}$	1	uservisits, rankings
4	523.66	7900M	$< 10^{-2}$	113K	hotel_prices
5	20	228M	4.3	49M	hotel_prices subsets
6	8	94.7M	4	48M	hotel_prices subsets

On this dataset, we run queries 4 to 6 to also (like queries 1 to 3) test raw data processing power, aggregation and *JOIN* performance respectively. However, these queries are not interchangeable with queries 1 to 3 because they are tailored to the exact structure of the hotel price dataset, and by using different input and output sizes we test different aspects of the query engines. We describe each of the queries 4 to 6 below in addition to providing query statistics in Table 3.

Query 4 computes average prices of hotel rooms grouped by certain months.

Query 5 computes linear regression pricing curves over a timespan of data collection dates.

Query 6 computes changes in hotel room prices between two collection dates.

3) *Total Workload*: Combining the results from the experiments with the two datasets gives us insights in performance of the query engines on both synthetic and real world data. In particular we look at how the engines deal with data scans (Query 1 and 4), heavy aggregation (Query 2 and 5), and the JOINS (Query 3 and 6).

4.2 Performance Aspects and Metrics

In order to be able to reason about the performance differences between different query engines, the different aspects contributing to this performance need to be defined. In this study we focus on three performance aspects:

1. *Processing Power*: the ability of a query engine to process a large number of SQL queries in a set amount of time. The more SQL queries a query engine can handle in a set amount of time, the better. We measure the processing power in terms of response time, that is, the time between submitting an SQL query to the system and getting a response. In addition, we also calculate the throughput per SQL query: the number of input records divided by response time.
2. *Resource Utilization*: the ability of a query engine to efficiently use the system resources available. This is important, because especially SMEs cannot afford to waste precious system resources. We measure the resource utilization in terms of mean, maximum and total CPU, memory, disk and network usage.

3. *Scalability*: the ability of a query engine to maintain predictable performance behaviour when system resources are added or removed from the system, or when input datasets grow or shrink. We perform both horizontal scalability and data input size scalability tests to measure the scalability of the query engines. Ideally, the performance should improve (at least) linearly with the amount of resources added, and should only degrade linearly with every unit of input data added. In practice this highly depends on the type of resources added as well as the complexity of the queries and the overhead of parallelism introduced.

4.3 Evaluation Procedure

Our procedure for evaluating the DSQEs is as follows: we run each query 10 times on its corresponding dataset while taking snapshots of the resource utilization using the monitoring tool `collectl` [4]. After the query completes, we also store its response time. When averaging over all the experiment iterations, we report the standard deviation as indicated with error bars in the experimental result figures. Like that, we take into account the varying performance of our cluster at different times of the day, intrinsic to the cloud [12].

The queries are submitted on the master node using the command line tools each query engine provides, and we write the output to a dedicated table which is cleared after every experiment iteration. We restart the query engine under test at the start of every experiment iteration in order to keep it comparable with other iterations.

5 Experimental Setup

We define a full micro-benchmarking setup by configuring the query engines as well as tuning their data caching policies for optimal performance. We evaluate the most recent stable versions of Shark (v0.9.0), Impala (v1.2.3) and Hive (v0.12). Many different parameters can influence the query engine's performance. In the following we define the hardware and software configuration parameters used in our experiments.

Hardware: To make a fair performance comparison between the query engines, we use the same cluster setup for each when running the experiments. The cluster consists of 5 `m2.4xlarge` worker Amazon EC2 VMs and 1 `m2.4xlarge` master VM, each having 68.4 GiB of memory, 8 virtual cores and 1.5 TiB instance storage. This cluster has sufficient storage for the real-world and synthetic data, and also has the memory required to allow query engines to benefit from in-memory caching of query inputs or outputs. Contrary to other Big Data processing systems, DSQEs (especially Impala and Shark) are tuned for nodes with large amounts of memory, which allows us to use fewer nodes than in comparable studies for batch processing systems to still get comparable (or better!) performance. An additional benefit of this specific cluster setup is the fact it is the same cluster setup used in the AMPLab benchmarks previously performed on older versions of Shark (v0.8.1), Impala (v1.2.3) and Hive (v0.12) [1]. By using the same setup, we can also compare current versions of these query engines with these older versions and see if significant performance improvements have been made.

Software: Hive uses YARN [18] as resource manager while we have used Impala's and Shark's standalone resource managers respectively, because at the time of testing the YARN compatible versions were not mature yet. All query engines under test run on top of a 64-bit Ubuntu 12.04 operating system. Since the queries we run compute results over large amounts of data, the configuration parameters of the distributed file system this data is stored on (HDFS) are crucial. It is therefore imperative that we keep these parameters fixed across all query engines under test. One of these parameters includes the HDFS block size, which we keep to the default of 64 MB. The number of HDFS files used per dataset, and how these files are structured and compressed

Table 4: Different ways to configure Shark with caching.

Abbreviation	OS Disk Cache	Input Cache	Output Cache
<i>Cold</i>	No	No	No
<i>OC</i>	No	No	Yes
<i>OSC</i>	Yes	No	No
<i>IC</i>	N/A	Yes	No
<i>OSC+OC</i>	Yes	No	Yes
<i>IC+OC</i>	N/A	Yes	Yes

is also kept fixed. While more sophisticated file formats are available (such as RCFile [11]) we selected the Sequence file key-value pair format because unlike the more sophisticated formats this is supported by all query engines, and this format uses less disk space than the plain text format. The datasets are compressed on disk using the Snappy compression type, which aims for reasonable compression size while being very fast at decompression.

Each worker has 68.4 GiB of memory available of which we allow a maximum of 60GiB for the query engines under test. This leaves a minimum of 8 GiB of free memory for other processes running on the same system. By doing this we ensure that all query engines under test have an equal amount of maximum memory reserved for them while still allowing the OS disk buffer cache to use more than 8 GiB when the query engine is not using a lot of memory.

Dataset Caching: Another important factor that influences query engine performance is whether the input data is cached or not. By default the operating system will cache files that were loaded from disk in an OS disk buffer cache. Because both Hive and Impala do not have any configurable caching policies available, we will simply run the queries on these two query engines both with and without the input dataset loaded into the OS disk buffer cache. To accomplish this, we perform a `SELECT` query over the relevant tables, so all the relevant data is loaded into the OS disk buffer cache. The query engines under test are restarted after every query as to prevent any other kind of caching to happen that might be unknown to us (e.g., Impala has a non-configurable internal caching system).

In contrast, Shark has more options available regarding caching. In addition to just using the OS disk buffer caching method, Shark also has the option to use an in-memory cached table as input and an in-memory cached table as output. This completely removes the (need for) disk I/O once the system has warmed up. To establish a representative configuration for Shark, we first evaluate the configurations as depicted in Table 4. OS Disk Cache means the entire input tables are first loaded through the OS disk cache by means of a `SELECT`. Input Cache means the input is first cached into in-memory Spark RDDs. Lastly, Output Cache means the result is kept in memory rather than written back to disk.

Figure 1 shows the resulting average response times for running a simple `SELECT *` query using the different possible Shark configurations. Note that no distinction is made between OS disk buffer cache being cleared or not when a cached input table is used, since in this case Shark does not read from disk at all.

The configuration with both input and output cached tables enabled (IC+OC) is the fastest setup for both the small and large data set. But the IC+OC and the IC configuration can only be used when the entire input data set fits in memory, which is often not the case with data sets of multiple TBs in size. The second fastest configuration (OSC+OC) only keeps the output (which is often much smaller) in memory and still reads the input from disk. The configuration which yields the worst results is using no caching at all (as expected).

In the experiments in Section 6 we use the optimistic IC+OC configuration when the input

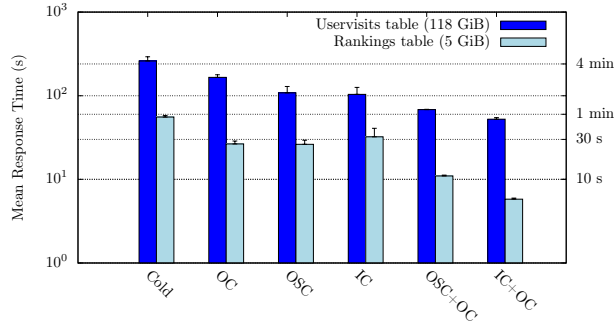


Figure 1: Response time for different Shark caching configurations (vertical axis is in log-scale).

data set fits in memory and the OSC+OC configuration when it does not, representing the best-case scenario. In addition the Cold configuration will be used to represent worst-case scenarios.

6 Experimental Results

We evaluate the three query engines selected in Section 3 on the performance aspects described in Section 4.2 using the workloads described in Section 4.1. We evaluate processing power in Section 6.1, resource consumption in Section 6.2, and scalability in Section 6.4.

6.1 Processing Power

We have used the fixed cluster setup with a total of 5 worker nodes and 1 master node as described in Section 5 to evaluate the response time and throughput (defined as the number of input records divided by the response time) of Hive, Impala and Shark on the 6 queries in the workloads. The results of the experiments are depicted in Figure 2. All experiments have been performed 10 times except for Query 4 with Impala since it took simply too long to complete. Only 2 iterations have been performed for this particular query. We used the dataset caching configurations explained in Section 5. For Impala and Hive we used disk buffer caching and no disk buffer caching for the warm and cold situations, respectively. For Shark we used the *Cold* configuration for the cold situation. In addition we used input and output dataset caching (*IC+OC*) for the warm situation of queries 1 to 3, and disk buffer caching and output caching (*OSC+OC*) for the warm situation of queries 4 to 6, since the price input dataset does not entirely fit in memory.

Key Findings:

- Input data caching generally does not cause a significant difference in response times.
- Performance is relatively stable over different iterations.
- Impala and Shark have similar performance and Hive is the worst performer in most cases. There is no overall winner.
- Impala does not handle large input sizes very well (Query 4).

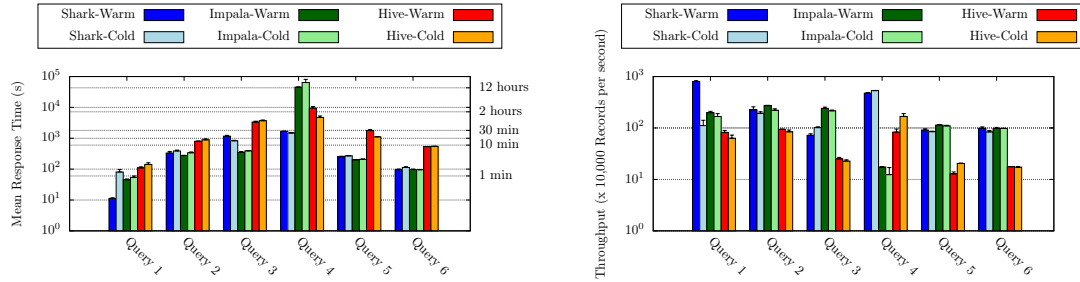


Figure 2: Query Response Time (left) and Throughput (right). Vertical axis is in log-scale.

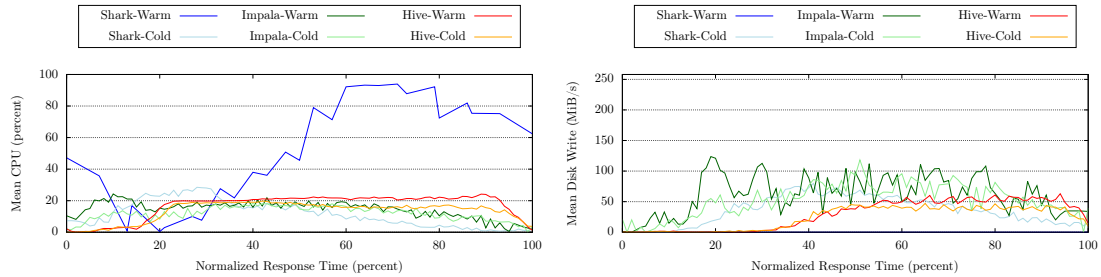


Figure 3: CPU utilization (left) and Disk Write (right) for query 1 over normalized response time.

The main reason why Hive is much slower than Impala and Shark is because of the high intermediate disk I/O. Because most queries are not disk I/O bound, data input caching makes little difference in performance. We elaborate on these two findings in more detail in Section 6.3.

In the following we discuss the response times from the 6 queries in a pair-wise manner. We evaluate the data scan queries 1 and 4, the aggregation queries 2 and 5, and the JOIN performance queries 4 and 6 depicted in Figure 2.

1) *Scan performance*: Shark’s response time for query 1 with data input and output caching enabled is significantly better than that of the other query engines. This is explained by the fact that query 1 is CPU-bound for the *Shark-Warm* configuration, but disk I/O bound for all other configurations as depicted in Figure 3. Since *Shark-Warm* caches both the input and output, and the intermediate data is so small that no spilling is required, no disk I/O is performed at all for *Shark-Warm*.

Results for query 4 for Impala are particularly interesting. Impala’s response time is 6 times as high as Hive’s, while resource utilization is much lower, as explained in Section 6.3. No bottleneck can be detected in the resource utilization logs and no errors are reported by Impala. After re-running the experiments for Impala on query 4 on a different set of Amazon EC2 instances, similar results are obtained, which makes it highly unlikely an error occurred during experiment execution. A more in-depth inspection is needed to get to the cause of this problem, which is out of the scope of our work.

2) *Aggregation performance*: Both the aggregation query 2 and 5 are handled quite well by all the engines. The main reason why even though query 5 has a much smaller input dataset, the response times are close to the ones of query 2 is that this query is relatively much more compute intensive (see Figure 4).

3) *JOIN performance*: The query engines perform quite similar on the JOIN queries 3 and 6.




Figure 4: CPU Core Seconds. Vertical axis is in log-scale.

A remarkable result is that the fully input and output cached configuration *Shark-Warm* starts to perform worse than its cold equivalent when dataset sizes grow. This is explained in more detail in Section 6.4.

6.2 Resource Consumption

Although the cluster consists of both a master and 5 worker nodes, we only evaluate the resource consumption on the workers, since the master is only used for coordination and remains idle the most of the time. For any of the queries the master never used more than 6 GiB of memory (<10% of total available), never exceeded more than 82 CPU Core Seconds (<0.0005% of the workers' maximum), has negligible disk I/O, and never exceeded total network I/O of 4 GiB (<0.08% of the workers' maximum).

Key Findings:

- Impala is a winner in total CPU consumption. Even when Shark outperforms Impala in terms of response time, Impala is still more CPU efficient (Figure 4).
- All query engines tend to use the full memory assigned to them.
- Disk I/O is as expected significantly higher for the queries without data caching vs. the queries with data caching. Impala has slightly less disk I/O than Shark. Hive ends last (Figure 5).
- Network I/O is comparable in all query engines, with the exception of Hive, which again ends last (Figure 5).

In the following we discuss the resource consumption per query averaged over 5 workers and 10 iterations (50 datapoints per average). We show the CPU Core Seconds per query in Figure 4. This shows how much total CPU a query uses during the query execution. The CPU Core Seconds metric is calculated by taking the area under the CPU utilization graphs of all workers, and then multiplying this number by the number of cores per worker (8 in our setup). For example, a hypothetical query running 2 seconds using 50% of CPU uses 1 CPU Second which is equal to 8 CPU Core Seconds in our case. A query running on Impala is the most efficient in terms of total CPU utilization, followed by Shark. This is as expected since although Shark and Impala are quite close in terms of response time, Impala is written in C/C++ instead of Scala, which is slightly more efficient.

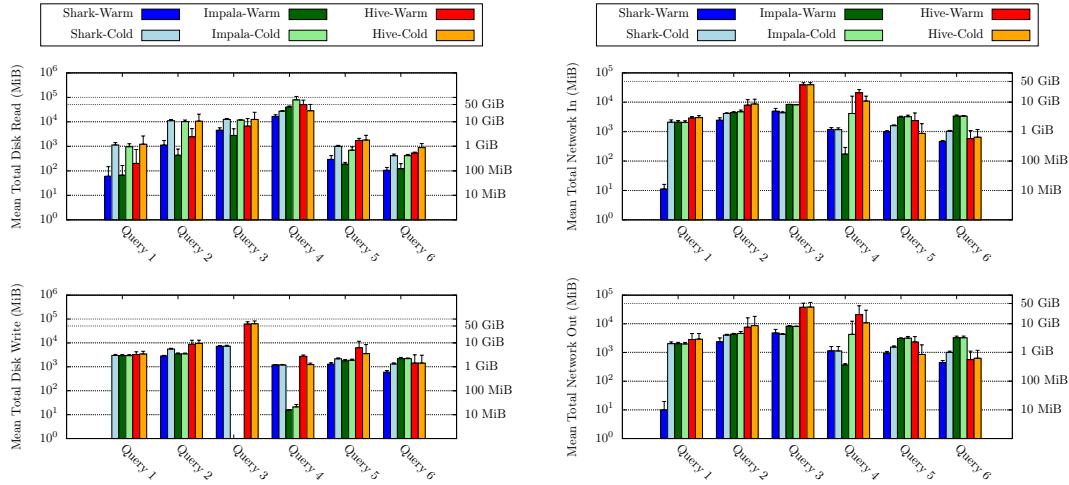


Figure 5: Total Disk Read (left-top), Disk Write (left-bottom), Network In (right-top) and Network Out (right-bottom). Vertical axis is in log-scale.

The total disk I/O per query is depicted in Figure 5. It shows that data caching does make a significant difference in the number of disk reads performed, but as shown in Section 6.3, this is hardly ever the bottleneck. Although the input datasets are entirely cached in the *Shark-Warm* configuration for queries 1 to 3, disk reading still occurs. This can be explained by the fact that Shark writes shuffle output of intermediate stages to the disk buffer cache (which eventually spills to disk). For queries 4 to 6 less significant differences occur since *Shark-Warm* only uses the OS disk buffer cache mechanism, like Impala and Hive. Note that because the input and output are compressed (compression ratio around 10), generally no more than 10% of the datasets is read or written to disk. Query 1 and 3 have very small output datasets, which makes *Shark-Warm*'s output not visible in the figure for query 1. Similarly for query 3 Impala does not show up at all because Impala does not write intermediate shuffle data to disk.

Figure 5 (right) shows the network I/O per query. Since most network I/O occurs between the workers, the network in and out look similar. Hive has a very variable network output total. Exact resource consumption numbers can be found in the tables of Appendix A

6.3 Resource Utilization over Time

In addition to reporting the average resources used per query, we show a more in depth analysis on the resource utilization over time for each query. Since queries have different response times each iteration, we normalize the response time to a range between 0 and 100% of query completion, and we show the mean values of each resource for all worker nodes averaged over all iterations (5 worker nodes, 10 iterations equals 50 data points per average total). We do not show network in performance in the figures since this is very similar to the network out, and we omit memory utilizations since for the majority of query engines the memory does not get deallocated until after the query execution has been completed. All the resulting figures are in Appendix B for brevity.

Key Findings:

- While the scan queries were expected to be disk I/O bound, query 1 is generally I/O bound, and query 4 is CPU bound because it also performs complex aggregation (Figure 7 and 10).
- The aggregation queries are both CPU bound (Figure 8 and 11).
- The JOIN queries are network I/O bound as expected (Figure 9 and 12).

In the following we go over each of the queries in detail in pair wise manner.

1) *Scan performance*: As explained in Section 6.1 query 1 is bound by disk and network I/O with the exception of the *Shark-Warm* configuration because the dataset is so small and entirely cached in memory (see Figure 7). As seen in Figure 10, query 4 is so computation heavy due to the additional aggregations performed, that it can suffice with a very low Disk read per second to keep up with the calculations. Because of the same reason the network out is low. And since the output size is so small, little to no disk writes are performed.

2) *Aggregation performance*: Both query 2 and 5 are CPU bound as shown in Figure 8 and 11 respectively. Both the queries show typical behaviour with higher CPU load at the start of query execution and higher disk I/O at the end.

3) *JOIN performance*: As expected both query 3 and 6 are network I/O heavy (see Figure 9 and 12 respectively). This is because one of the tables being joined is distributed over the network of workers.

6.4 Scalability

In this section we analyze both the horizontal and the data size scalability of the query engines. We decided to scale horizontally down instead of up because a cluster of 5 nodes of this caliber is already quite expensive for SMEs (more than \$7000 a month), and from our experimental results it shows that some queries already do no longer scale well from 4 to 5 worker nodes. We used queries 1 to 3 for data size scaling (since this dataset was already synthetic in the first place) and queries 4 to 6 for horizontal scaling.

Key Findings:

- Both Impala and Shark have near linear data size scalability on the scan and aggregation queries, whereas scaling on the JOIN queries is sub-linear.
- Hive has sub-linear data size scalability in all cases.
- If *Shark-Warm*'s input dataset is too large for its data storage memory, the response time will increase beyond *Shark-Cold* due to swapping (Figure 6; left).
- Shark and Impala horizontally scale reasonably well on all types of queries up to 3 nodes, whereas Hive only scales well on queries with large input sizes (Figure 6; right).
- The query engines do not benefit from having more than 3 or 4 nodes in the cluster. Impala even performs worse for query 6 at bigger cluster sizes (Figure 6; right).

1) *Data Size Scalability*: The data size scalability of the query engines is depicted in Figure 6 (left). We have sampled subsets from the original dataset and display these along the horizontal axis of the figure. For small datasets the query engines do not scale well since in that case there




Figure 6: Data size scalability: Response time Query 1, 2 and 3 (left; from top to bottom). Horizontal scalability: Response time Query 4, 5 and 6 (right; from top to bottom) vertical axis in log-scale.

is more engine overhead than data to be processed. For example when looking at the difference in response time between 5% and 10% dataset size for query 1, it takes almost the same time to process 5% as 10% of the dataset. However, when the dataset grows beyond a certain threshold the query engines all scale about linearly on query 1 (25% to 100%) except for Hive. Similar phenomena occur for query 2 and 3.

An interesting phenomenon occurs with Shark for query 3. When the data input size grows and passes 75%, *Shark-Warm* actually starts becoming slower than *Shark-Cold*. This is caused by the fact that Shark allocates only around 34 GiB of the 60 GiB it was assigned for data storage and uses the remaining amount as JVM overhead. This means that the total cluster can only store about 170 GiB of data instead of the 300 GiB it was assigned. The input dataset is 123.48 GiB for query 3, which has some of the worker nodes' memory fully filled at the start of query execution. When JOINS are performed even more memory is needed for the shuffled data received from the other worker nodes, causing the node to spill some of its input data back to disk.

2) *Horizontal Scalability*: The horizontal scalability of the query engines is depicted in Figure 6 (right) (note that we are scaling down instead of up). For Impala we only ran 4 and 5 nodes since it already took 12 hours to complete. Both Shark and Hive scale near linearly on the number of nodes. Hive only scales well on query 4 since Hive's Hadoop MapReduce overhead likely outweighs the computation time for query 5 and 6. This is because they have a relatively small input size. Impala actually starts to perform worse on query 6 if more than 3 nodes are added to the cluster. Similarly, both Shark and Impala no longer improve performance after more than 4 nodes are added to the cluster for query 4 and 5.

This remarkable result for horizontal scaling shows that whenever a query is not CPU-bound on a cluster with some number of nodes, performance will not improve any further when adding even more nodes. In the case of network I/O bound queries like query 6, it might even be more beneficial to bind these to a smaller number of nodes so less network overhead occurs.

7 Conclusions and Future Work

In recent years an increasing number of Distributed SQL Query Engines have become available. They all allow for large scale Big Data processing using SQL as interface language. In this work we compare three major query engines (Hive, Impala and Shark) with the requirements of SMEs in mind. SMEs have only little resources available to run their big data analytics on, and cannot afford running a query engine with large overhead.

In this work we have defined an empirical evaluation method to assess the performance of different query engines. Despite not covering all the methodological aspects of a scientific benchmark, this micro-benchmark gives practical insights for SMEs to take informed decisions when selecting a specific tool. Moreover, it can be used to compare current and future engines. The method focuses on three performance aspects: processing power, resource utilization, and scalability.

Using both a real world and a synthetic dataset with representative queries, we evaluate the query engines' performance. We find that different query engines have widely varying performance. Although Hive is almost always outperformed by the other engines, it highly depends on the query type and input size whether Impala or whether Shark is the best performer. Shark can also perform well on queries with over 500 GiB in input size in our cluster setup, while Impala starts to perform worse for these queries. Overall Impala is the most CPU efficient, and all query engines have comparable resource consumption for memory, disk and network. A remarkable result found is that query response time does not always improve when adding more nodes to the

cluster. Remaining key findings can be found at the top of every experiment in Section 6.

This work has been an attempt to get insights in DSQE performance in order to make life easier for SMEs picking the query engine that best suits their needs. Query engine performance in a multi-tenant environment has not been evaluated, and is part of future work.

References

- [1] AMPLab’s Big Data Benchmark. <https://amplab.cs.berkeley.edu/benchmark/>. [Online; Last accessed 1st of September 2014]. 3, 4, 5, 7
- [2] Apache Hadoop. <http://hadoop.apache.org>. [Online; Last accessed 1st of September 2014]. 4
- [3] Apache Tez. <http://tez.apache.org>. [Online; Last accessed 1st of September 2014]. 3
- [4] Collectl Resource Monitoring. <http://collectl.sourceforge.net>. [Online; Last accessed 1st of September 2014]. 7
- [5] Impala. <http://blog.cloudera.com/blog/2012/10/cloudera-impala-real-time-queries-in-apache-hadoop-for-real/>. [Online; Last accessed 1st of September 2014]. 3, 4
- [6] Impala Benchmark. <http://blog.cloudera.com/blog/2014/05/new-sql-choices-in-the-apache-hadoop-ecosystem-why-impala-continues-to-lead/>. [Online; Last accessed 1st of September 2014]. 3, 4
- [7] Presto. <http://www.prestodb.io>. [Online; Last accessed 1st of September 2014]. 3
- [8] A. Floratou, U. F. Minhas, and F. Ozcan. Sql-on-hadoop: Full circle back to shared-nothing database architectures. *Proceedings of the VLDB Endowment*, 7(12), 2014. 3, 4
- [9] A. Ghazal, T. Rabl, M. Hu, F. Raab, M. Poess, A. Crolotte, and H.-A. Jacobsen. Bigbench: Towards an industry standard benchmark for big data analytics. In *Proceedings of the 2013 international conference on Management of data*, pages 1197–1208, 2013. 3, 4
- [10] M. Hausenblas and J. Nadeau. Apache Drill: Interactive Ad-Hoc Analysis at Scale. *Big Data*, 2013. 3
- [11] Y. He, R. Lee, Y. Huai, Z. Shao, N. Jain, X. Zhang, and Z. Xu. RCFile: A fast and space-efficient data placement structure in MapReduce-based warehouse systems. In *Data Engineering (ICDE), 2011 IEEE 27th International Conference on*, pages 1199–1208, 2011. 8
- [12] A. Iosup, N. Yigitbasi, and D. Epema. On the performance variability of production cloud services. In *Cluster, Cloud and Grid Computing (CCGrid), 2011 11th IEEE/ACM International Symposium on*, pages 104–113, 2011. 7
- [13] S. Melnik, A. Gubarev, J. J. Long, G. Romer, S. Shivakumar, M. Tolton, and T. Vassilakis. Dremel: interactive analysis of web-scale datasets. *Proceedings of the VLDB Endowment*, 3(1-2):330–339, 2010. 4

- [14] A. Pavlo, E. Paulson, A. Rasin, D. J. Abadi, D. J. DeWitt, S. Madden, and M. Stonebraker. A comparison of approaches to large-scale data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, pages 165–178, 2009. [3](#), [4](#), [5](#)
- [15] M. Poess, R. O. Nambiar, and D. Walrath. Why you should run TPC-DS: a workload analysis. In *Proceedings of the 33rd international conference on Very large data bases*, pages 1138–1149. VLDB Endowment, 2007. [3](#)
- [16] A. Rowstron, D. Narayanan, A. Donnelly, G. O’Shea, and A. Douglas. Nobody ever got fired for using Hadoop on a cluster. In *Proceedings of the 1st International Workshop on Hot Topics in Cloud Data Processing*, page 2. ACM, 2012. [5](#)
- [17] A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, S. Anthony, H. Liu, P. Wyckoff, and R. Murthy. Hive: a warehousing solution over a map-reduce framework. *Proceedings of the VLDB Endowment*, 2(2):1626–1629, 2009. [3](#), [4](#)
- [18] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, et al. Apache Hadoop YARN: Yet another resource negotiator. In *Proceedings of the 4th annual Symposium on Cloud Computing*, page 5. ACM, 2013. [7](#)
- [19] R. S. Xin, J. Rosen, M. Zaharia, M. J. Franklin, S. Shenker, and I. Stoica. Shark: Sql and rich analytics at scale. In *Proceedings of the 2013 ACM International Conference on Management of Data*, pages 13–24, 2013. [3](#), [4](#)
- [20] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2, 2012. [4](#)
- [21] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, pages 10–10, 2010. [4](#)

A Detailed Distributed SQL Query Engine Performance Metrics

In this appendix detailed performance metrics of the experiments are reported as-is, extending on the experimental results in Section 6. This information can be used to get more precise insights into performance differences across query engines. For all tables holds that the green coloured cells are marking the system that has the best performance for a query for a certain metric, and that cells indicating a 0 are actually presenting values *close* to 0 (because of rounding).

Query 1 (cold)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	154	1.58	666	53	0.08	60	78	0.16	104
Disk Read (MiB/s)	8	1.82	110	17	2.05	204	13	1.40	119
Disk Read Total (MiB)	1,229	1.17	5,211	970	0.32	1,202	1,138	0.25	1,871
Disk Write (MiB/s)	24	1.35	172	53	1.05	245	35	1.11	204
Disk Write Total (MiB)	3,455	0.30	5,721	2,954	0.10	3,653	2,989	0.11	3,682
Memory (MiB)	5,320	0.51	16,171	11,346	0.27	17,059	9,696	0.37	16,494
Network In (MiB/s)	21	0.97	97	37	0.88	141	25	1.22	152
Network In Total (MiB)	3,015	0.16	3,637	2,070	0.14	2,863	2,105	0.20	2,914
Network Out (MiB/s)	20	1.21	117	36	0.84	141	24	1.20	134
Network Out Total (MiB)	2,919	0.57	5,452	1,996	0.13	2,681	2,047	0.15	2,842
Response Time (s)	142	0.13	191	54	0.12	69	81	0.21	108

Table 5: Statistics for Query 1 (cold), Data Scale: 100%, Number of Nodes: 5.

Query 1 (warm)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	149	1.69	674	54	0.08	63	59	0.07	65
Disk Read (MiB/s)	2	3.51	78	1	7.09	146	5	2.98	117
Disk Read Total (MiB)	202	2.68	2,674	67	1.47	384	60	1.46	341
Disk Write (MiB/s)	29	1.27	180	63	1.04	248	0	2.97	0
Disk Write Total (MiB)	3,272	0.31	5,609	2,971	0.10	3,809	0	0.56	0
Memory (MiB)	17,041	0.27	27,134	66,865	0.02	68,249	53,384	0.13	67,168
Network In (MiB/s)	26	0.91	93	44	0.73	142	1	2.15	11
Network In Total (MiB)	2,929	0.11	3,642	2,102	0.14	2,911	11	0.44	23
Network Out (MiB/s)	25	1.20	113	43	0.70	132	1	2.91	14
Network Out Total (MiB)	2,837	0.62	5,597	2,029	0.12	2,603	10	0.97	33
Response Time (s)	110	0.08	129	45	0.06	50	11	0.04	12

Table 6: Statistics for Query 1 (warm), Data Scale: 100%, Number of Nodes: 5.

Query 2 (cold)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	2,852	0.41	4,807	417	0.14	638	1,346	0.08	1,563
Disk Read (MiB/s)	12	1.67	116	30	1.42	193	28	0.87	101
Disk Read Total (MiB)	10,680	0.94	38,117	10,237	0.16	11,491	11,145	0.12	13,330
Disk Write (MiB/s)	11	1.87	184	10	2.97	235	14	1.46	148
Disk Write Total (MiB)	9,502	0.38	17,079	3,499	0.08	4,209	5,554	0.08	6,488
Memory (MiB)	20,100	0.65	63,388	28,471	0.35	55,062	31,317	0.36	62,393
Network In (MiB/s)	10	1.59	185	14	1.10	85	11	2.52	171
Network In Total (MiB)	8,747	0.46	14,998	4,758	0.14	6,449	4,163	0.05	4,580
Network Out (MiB/s)	10	1.88	151	14	1.18	102	10	2.49	148
Network Out Total (MiB)	8,687	1.10	35,624	4,648	0.16	6,508	4,035	0.07	4,493
Response Time (s)	891	0.08	1,004	340	0.06	379	389	0.07	438

Table 7: Statistics for Query 2 (cold), Data Scale: 100%, Number of Nodes: 5.

Query 2 (warm)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	2,648	0.49	5,106	399	0.02	425	938	0.08	1,102
Disk Read (MiB/s)	3	2.49	96	2	4.11	116	3	1.47	35
Disk Read Total (MiB)	2,439	1.17	11,615	433	0.80	1,373	1,129	0.53	2,250
Disk Write (MiB/s)	12	1.84	187	13	2.65	250	9	1.65	111
Disk Write Total (MiB)	8,779	0.48	17,612	3,512	0.09	4,176	2,861	0.06	3,178
Memory (MiB)	47,685	0.25	68,255	62,154	0.07	68,249	68,206	0.00	68,252
Network In (MiB/s)	11	1.87	203	16	1.09	86	7	2.30	209
Network In Total (MiB)	8,013	0.53	16,480	4,457	0.07	5,178	2,458	0.21	3,857
Network Out (MiB/s)	10	2.02	129	16	1.05	93	7	2.50	155
Network Out Total (MiB)	7,620	1.12	30,953	4,358	0.06	5,083	2,381	0.36	6,491
Response Time (s)	800	0.01	811	276	0.02	283	333	0.13	410

Table 8: Statistics for Query 2 (warm), Data Scale: 100%, Number of Nodes: 5.

Query 3 (cold)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	6,647	0.39	10,853	551	0.01	569	1,726	0.08	1,924
Disk Read (MiB/s)	3	3.56	143	30	1.47	192	15	1.29	125
Disk Read Total (MiB)	12,624	0.95	45,355	11,959	0.03	12,449	12,645	0.09	14,605
Disk Write (MiB/s)	18	1.74	250	0	20.41	0	9	0.72	46
Disk Write Total (MiB)	64,007	0.30	106,898	0	1.45	1	7,353	0.09	8,384
Memory (MiB)	38,205	0.38	68,250	34,437	0.42	55,775	33,857	0.48	68,250
Network In (MiB/s)	11	1.75	198	21	1.30	102	5	4.40	216
Network In Total (MiB)	39,728	0.17	53,539	8,260	0.01	8,391	4,414	0.08	5,157
Network Out (MiB/s)	11	1.92	160	21	1.37	152	5	3.48	152
Network Out Total (MiB)	38,552	0.44	95,724	8,080	0.04	8,466	4,293	0.08	4,991
Response Time (s)	3,699	0.06	3,947	390	0.03	422	829	0.04	882

Table 9: Statistics for Query 3 (cold), Data Scale: 100%, Number of Nodes: 5.

Query 3 (warm)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	6,347	0.36	13,671	561	0.02	606	1,413	0.10	1,664
Disk Read (MiB/s)	2	3.87	118	8	2.63	170	4	1.31	27
Disk Read Total (MiB)	6,704	1.03	24,295	2,786	0.88	9,649	4,554	0.26	5,906
Disk Write (MiB/s)	19	1.70	249	0	6.42	0	6	1.25	53
Disk Write Total (MiB)	62,028	0.24	100,819	0	0.14	0	7,194	0.08	8,074
Memory (MiB)	53,394	0.20	68,253	59,443	0.16	68,249	67,823	0.03	68,257
Network In (MiB/s)	12	1.79	197	24	1.29	112	4	2.52	202
Network In Total (MiB)	39,201	0.18	57,428	8,460	0.00	8,575	4,974	0.25	7,825
Network Out (MiB/s)	12	1.90	161	23	1.32	132	4	2.79	154
Network Out Total (MiB)	37,796	0.40	92,610	8,268	0.03	8,578	4,812	0.33	10,326
Response Time (s)	3,359	0.06	3,725	350	0.05	384	1,168	0.06	1,253

Table 10: Statistics for Query 3 (warm), Data Scale: 100%, Number of Nodes: 5.

Query 4 (cold)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	29,615	0.05	32,528	15,022	1.48	78,620	10,741	0.03	11,321
Disk Read (MiB/s)	6	1.11	79	1	2.96	105	18	0.26	78
Disk Read Total (MiB)	28,445	0.82	77,034	79,334	0.38	159,037	27,229	0.06	29,998
Disk Write (MiB/s)	0	6.29	95	0	6.55	0	1	2.32	20
Disk Write Total (MiB)	1,245	0.15	1,618	21	0.26	29	1,205	0.04	1,297
Memory (MiB)	34,162	0.41	68,251	41,967	0.50	68,250	45,285	0.32	68,250
Network In (MiB/s)	2	0.78	67	0	23.04	112	1	6.79	143
Network In Total (MiB)	11,011	0.47	16,226	4,158	2.90	40,329	1,175	0.15	1,510
Network Out (MiB/s)	2	2.04	85	0	21.58	65	1	6.46	125
Network Out Total (MiB)	10,843	1.76	50,992	4,287	1.88	25,404	1,133	0.43	2,212
Response Time (s)	4,727	0.12	5,809	63,936	0.27	81,178	1,489	0.01	1,521

Table 11: Statistics for Query 4 (cold), Data Scale: 100%, Number of Nodes: 5.

Query 4 (warm)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	58,536	0.05	66,008	13,182	0.02	13,579	11,276	0.02	11,733
Disk Read (MiB/s)	6	0.86	60	1	3.59	103	10	0.49	48
Disk Read Total (MiB)	50,883	0.51	87,299	40,925	0.09	45,915	16,687	0.17	24,282
Disk Write (MiB/s)	0	8.95	205	0	6.26	0	1	2.50	21
Disk Write Total (MiB)	2,753	0.14	3,796	16	0.01	16	1,205	0.04	1,322
Memory (MiB)	67,780	0.02	68,251	68,239	0.00	68,250	67,998	0.02	68,251
Network In (MiB/s)	2	0.86	127	0	17.10	18	1	3.64	48
Network In Total (MiB)	21,062	0.28	29,399	173	0.68	407	1,188	0.16	1,680
Network Out (MiB/s)	2	1.38	131	0	8.29	21	1	4.12	74
Network Out Total (MiB)	21,095	1.02	51,920	368	0.11	428	1,148	0.41	2,022
Response Time (s)	9,507	0.13	10,957	45,396	0.00	45,396	1,668	0.02	1,738

Table 12: Statistics for Query 4 (warm), Data Scale: 100%, Number of Nodes: 5.

Query 5 (cold)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	1,470	1.76	7,351	269	0.16	445	1,355	0.11	1,664
Disk Read (MiB/s)	2	1.16	40	3	4.37	189	4	0.77	38
Disk Read Total (MiB)	1,808	0.57	4,628	703	0.36	1,064	999	0.10	1,245
Disk Write (MiB/s)	3	4.91	248	9	3.08	234	8	1.90	144
Disk Write Total (MiB)	3,577	1.39	16,997	1,900	0.11	2,293	2,164	0.09	2,662
Memory (MiB)	4,510	0.70	16,746	13,045	0.30	26,812	23,039	0.25	30,706
Network In (MiB/s)	1	7.80	168	15	1.04	78	6	3.30	177
Network In Total (MiB)	872	1.14	2,902	3,242	0.12	4,825	1,601	0.06	1,791
Network Out (MiB/s)	1	7.80	156	15	1.06	79	6	3.02	129
Network Out Total (MiB)	854	1.17	2,923	3,173	0.14	4,840	1,551	0.10	1,882
Response Time (s)	1,116	0.01	1,134	208	0.03	221	269	0.02	277

Table 13: Statistics for Query 5 (cold), Data Scale: 100%, Number of Nodes: 5.

Query 5 (warm)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	4,452	0.94	15,631	263	0.03	278	1,338	0.11	1,656
Disk Read (MiB/s)	1	1.10	9	1	1.53	53	1	1.04	8
Disk Read Total (MiB)	1,743	0.22	3,611	188	0.18	281	295	0.42	772
Disk Write (MiB/s)	4	3.52	186	9	3.46	251	5	1.79	81
Disk Write Total (MiB)	6,324	0.86	22,651	1,793	0.12	2,466	1,302	0.15	1,795
Memory (MiB)	62,400	0.07	68,252	27,611	0.16	38,695	63,480	0.05	68,252
Network In (MiB/s)	1	5.37	195	15	1.08	78	4	4.55	194
Network In Total (MiB)	2,389	0.82	6,272	3,136	0.07	3,692	982	0.10	1,108
Network Out (MiB/s)	1	5.20	150	15	1.10	92	4	4.07	160
Network Out Total (MiB)	2,339	0.54	4,718	3,068	0.06	3,591	953	0.15	1,335
Response Time (s)	1,775	0.08	2,045	200	0.02	208	250	0.05	281

Table 14: Statistics for Query 5 (warm), Data Scale: 100%, Number of Nodes: 5.

Query 6 (cold)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	426	1.29	1,908	122	0.07	132	332	0.10	385
Disk Read (MiB/s)	2	1.39	27	4	3.82	132	4	1.09	22
Disk Read Total (MiB)	890	0.46	2,012	418	0.12	468	415	0.17	551
Disk Write (MiB/s)	3	3.84	172	22	1.68	231	11	1.84	148
Disk Write Total (MiB)	1,441	1.14	5,886	2,205	0.07	2,384	1,318	0.11	1,647
Memory (MiB)	3,030	0.39	7,151	12,716	0.28	17,151	12,758	0.49	23,980
Network In (MiB/s)	1	6.08	165	34	0.69	89	9	2.24	125
Network In Total (MiB)	650	0.83	1,859	3,378	0.03	3,523	1,035	0.07	1,186
Network Out (MiB/s)	1	5.98	162	33	0.77	108	9	2.10	146
Network Out Total (MiB)	631	0.89	1,733	3,276	0.14	4,085	1,008	0.11	1,234
Response Time (s)	547	0.01	561	96	0.00	96	113	0.06	124

Table 15: Statistics for Query 6 (cold), Data Scale: 100%, Number of Nodes: 5.

Query 6 (warm)	Hive			Impala			Shark		
	Mean	CV.	Max	Mean	CV.	Max	Mean	CV.	Max
CPU Seconds	427	1.38	1,944	121	0.07	139	333	0.08	381
Disk Read (MiB/s)	1	1.09	7	1	4.22	121	1	1.08	7
Disk Read Total (MiB)	536	0.09	827	123	0.61	340	106	0.31	231
Disk Write (MiB/s)	3	4.07	167	22	1.70	234	6	1.59	48
Disk Write Total (MiB)	1,437	1.23	5,968	2,218	0.12	2,700	598	0.18	787
Memory (MiB)	6,974	0.29	12,306	15,839	0.23	22,259	19,358	0.23	29,049
Network In (MiB/s)	1	5.83	162	34	0.72	90	5	3.32	103
Network In Total (MiB)	582	0.88	1,837	3,381	0.10	4,045	465	0.08	534
Network Out (MiB/s)	1	5.84	150	33	0.79	117	5	3.07	124
Network Out Total (MiB)	568	0.97	1,809	3,289	0.13	4,159	449	0.18	612
Response Time (s)	536	0.01	543	97	0.04	103	96	0.06	106

Table 16: Statistics for Query 6 (warm), Data Scale: 100%, Number of Nodes: 5.

B Detailed Distributed SQL Query Engine Resource Utilization

This appendix gives more insights in how resources are utilized over time. See Section 6.3 for accompanying explanations.

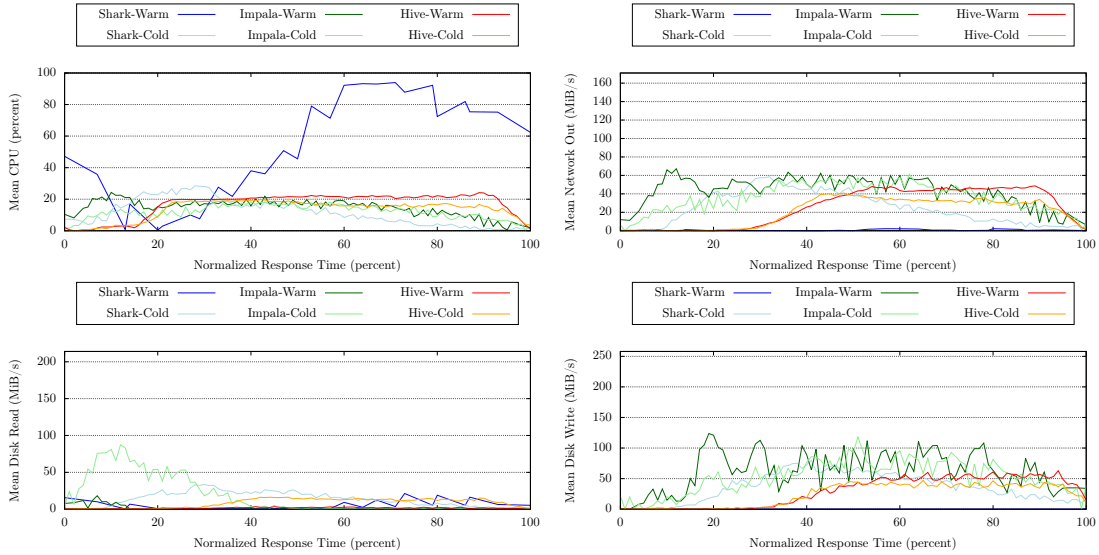


Figure 7: CPU utilization (top-left), Network Out (top-right), Disk Read (bottom-left), Disk Write (bottom-right) for query 1 over normalized response time.

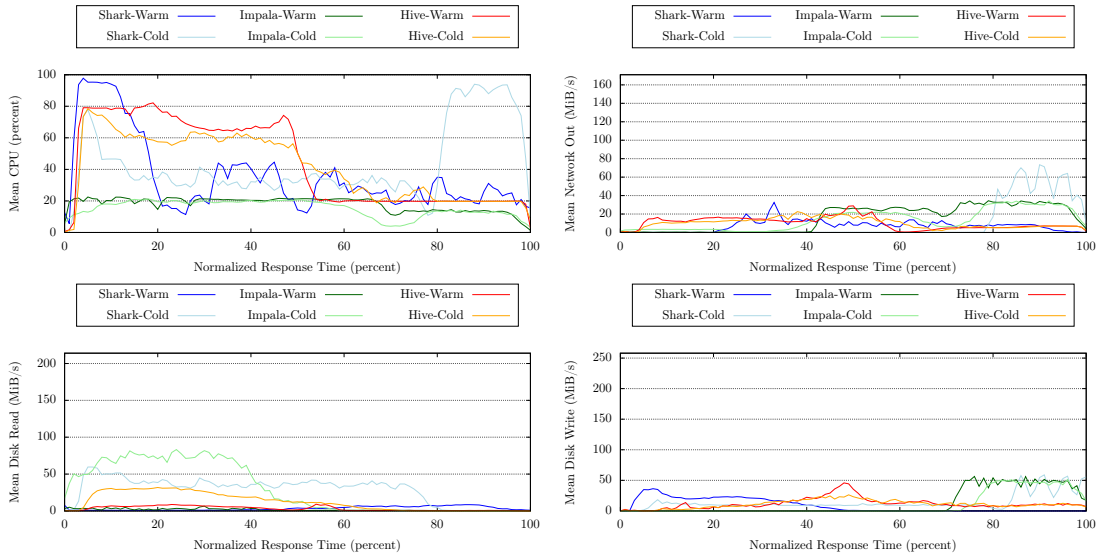


Figure 8: CPU utilization (top-left), Network Out (top-right), Disk Read (bottom-left), Disk Write (bottom-right) for query 2 over normalized response time.

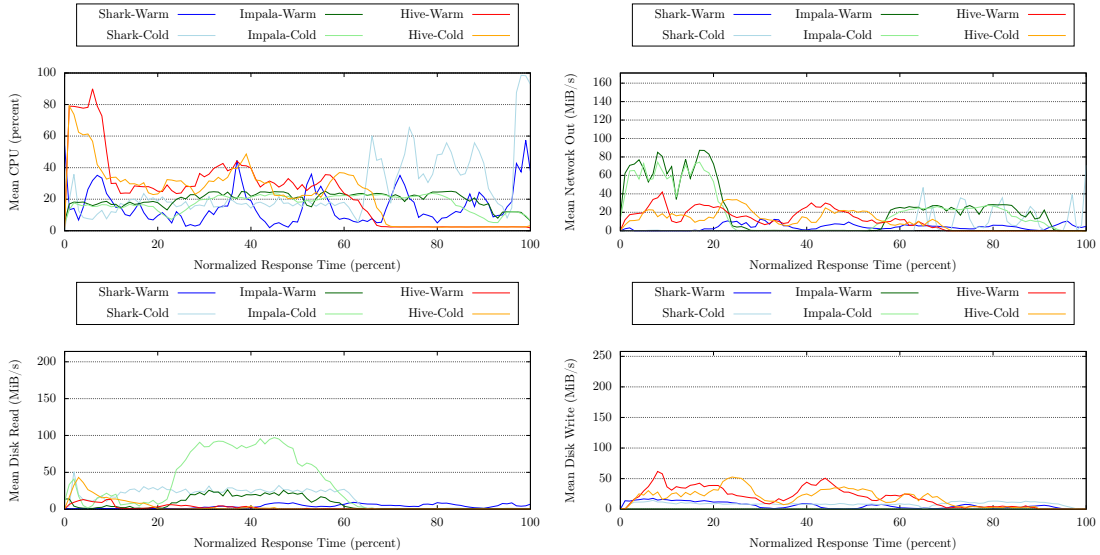


Figure 9: CPU utilization (top-left), Network Out (top-right), Disk Read (bottom-left), Disk Write (bottom-right) for query 3 over normalized response time.

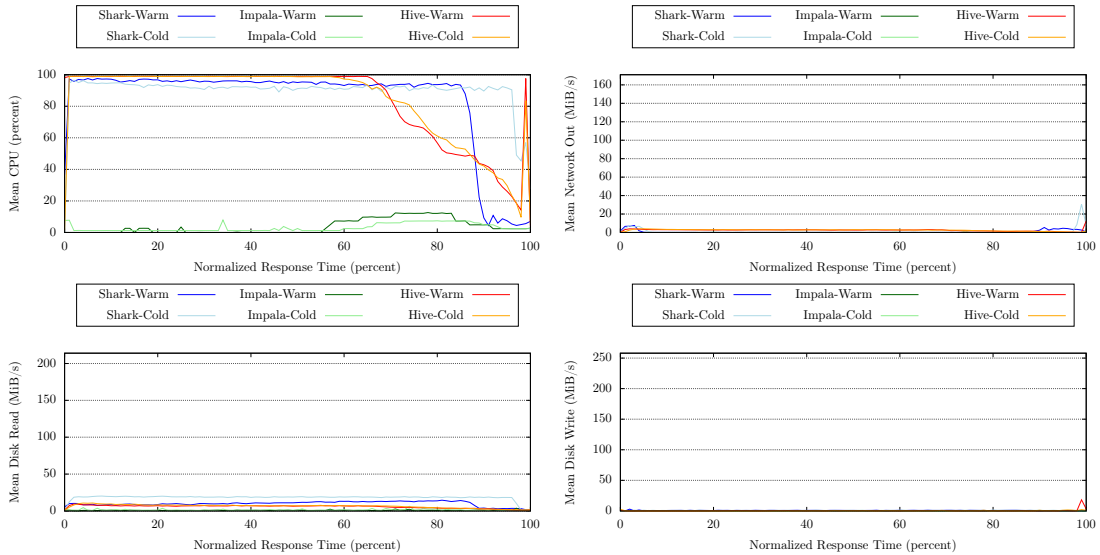


Figure 10: CPU utilization (top-left), Network Out (top-right), Disk Read (bottom-left), Disk Write (bottom-right) for query 4 over normalized response time.

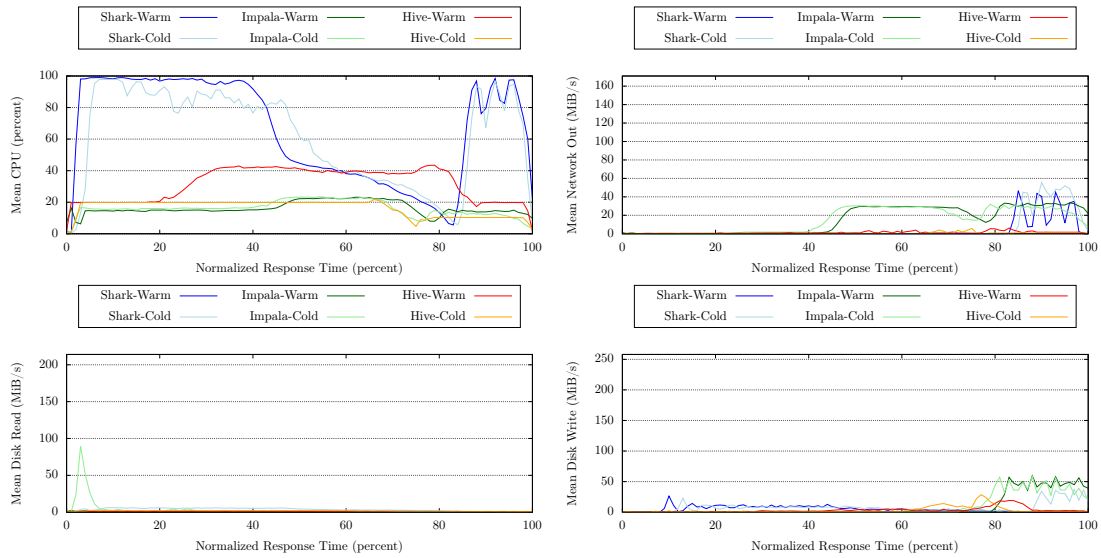


Figure 11: CPU utilization (top-left), Network Out (top-right), Disk Read (bottom-left), Disk Write (bottom-right) for query 5 over normalized response time.

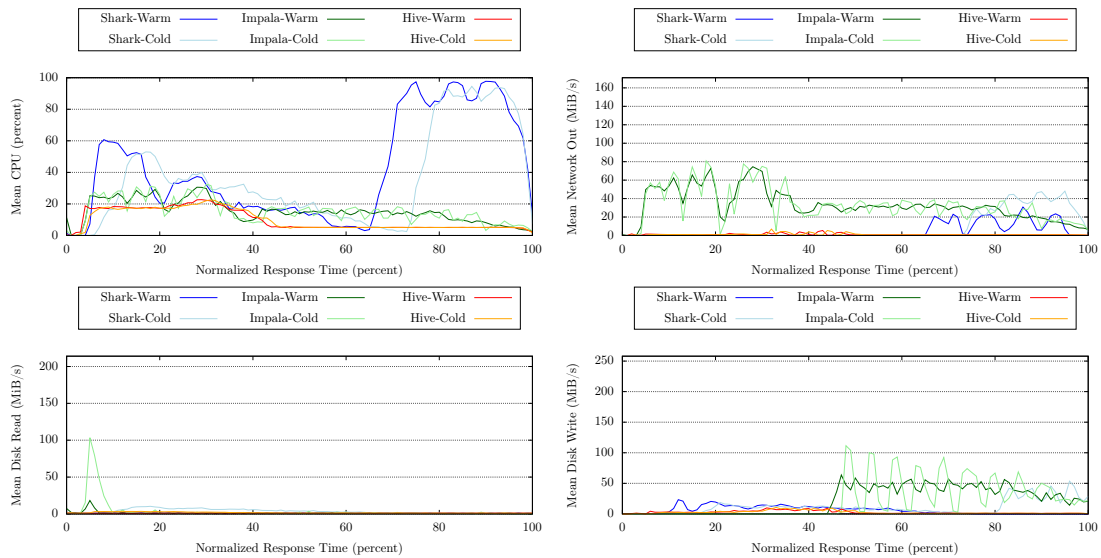


Figure 12: CPU utilization (top-left), Network Out (top-right), Disk Read (bottom-left), Disk Write (bottom-right) for query 6 over normalized response time.