




presto 

Next Chapter for CBO

Kamil Bajda-Pawlikowski
Co-founder and CTO
www.starburstdata.com

Starburst



Starburst Data

The Presto  Experts.

Founded by Presto committers:


- Over 4 years of contributions to Presto
- Presto distro for on-prem and cloud env
- Supporting large customers in production
- Enterprise subscription add-ons

Notable features contributed:

- ANSI SQL syntax enhancements
- Execution engine improvements
- Security integrations
- Spill to disk
- Cost-Based Optimizer



Presto Cluster



Object Storage



NoSQL Sources



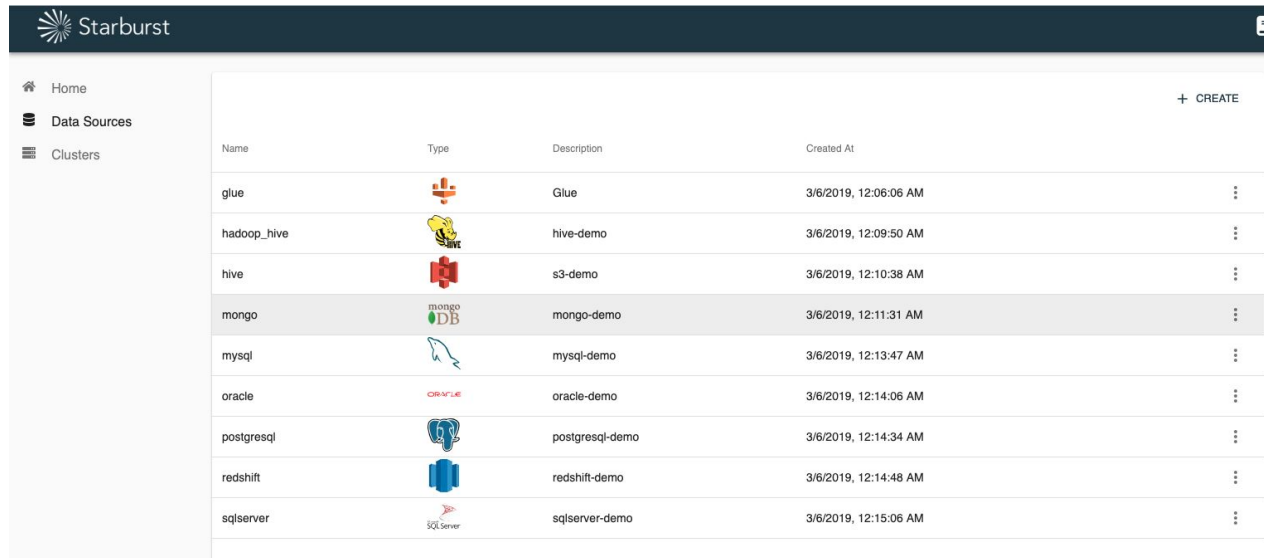
RDBMS Storage












Hadoop



Starburst Presto & Cloud



The screenshot shows the Starburst dashboard interface. On the left, there is a navigation menu with 'Home', 'Data Sources', and 'Clusters'. The main area displays a table of data sources with columns for Name, Type, Description, and Created At. A '+ CREATE' button is visible in the top right of the table area.

Name	Type	Description	Created At
glue		Glue	3/6/2019, 12:06:06 AM
hadoop_hive		hive-demo	3/6/2019, 12:09:50 AM
hive		s3-demo	3/6/2019, 12:10:38 AM
mongo		mongo-demo	3/6/2019, 12:11:31 AM
mysql		mysql-demo	3/6/2019, 12:13:47 AM
oracle		oracle-demo	3/6/2019, 12:14:06 AM
postgresql		postgresql-demo	3/6/2019, 12:14:34 AM
redshift		redshift-demo	3/6/2019, 12:14:48 AM
sqlserver		sqlserver-demo	3/6/2019, 12:15:06 AM




Azure HDInsight

Starburst Presto

<https://www.starburstdata.com/technical-blog/announcing-starburst-enterprise-302e-with-mission-control/>

Presto CBO



Built for Performance

Query Execution Engine:

- MPP-style **pipelined** in-memory execution
- **Columnar** and **vectorized** data processing
- Runtime query **bytecode compilation**
- Memory **efficient** data structures
- Multi-threaded **multi-core** execution
- Optimized readers for **columnar formats** (ORC and Parquet)
- Predicate and column projection **pushdown**
- Now also **Cost-Based Optimizer**

CBO in a nutshell

Cost-Based Optimizer v1 includes:

- support for **statistics** stored in Hive Metastore
- **join reordering** based on selectivity estimates and cost
- automatic **join type** selection (repartitioned vs broadcast)
- automatic left/right **side selection** for joined tables

<https://www.starburstdata.com/technical-blog/>

Statistics and Cost

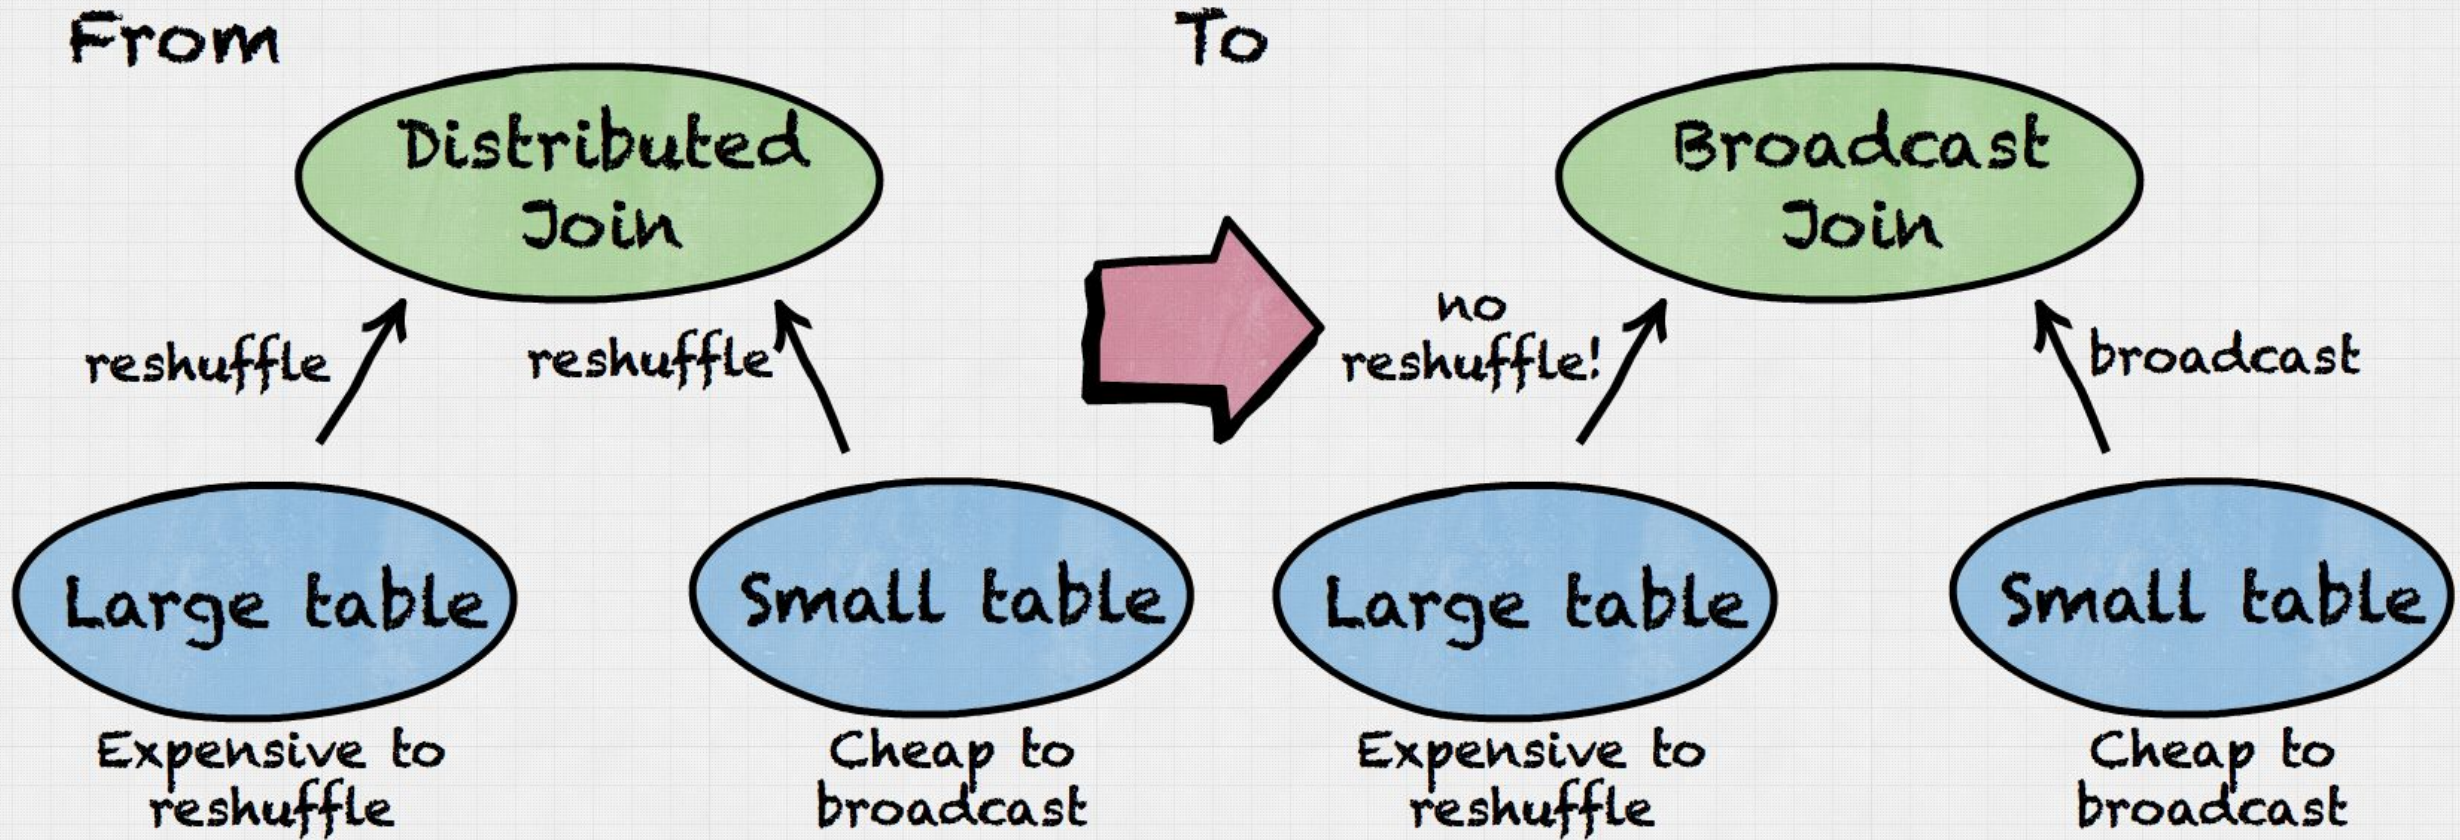
Hive Metastore statistics:

- number of rows in a table
- number of distinct values in a column
- fraction of NULL values in a column
- minimum/maximum value in a column
- average data size for a column

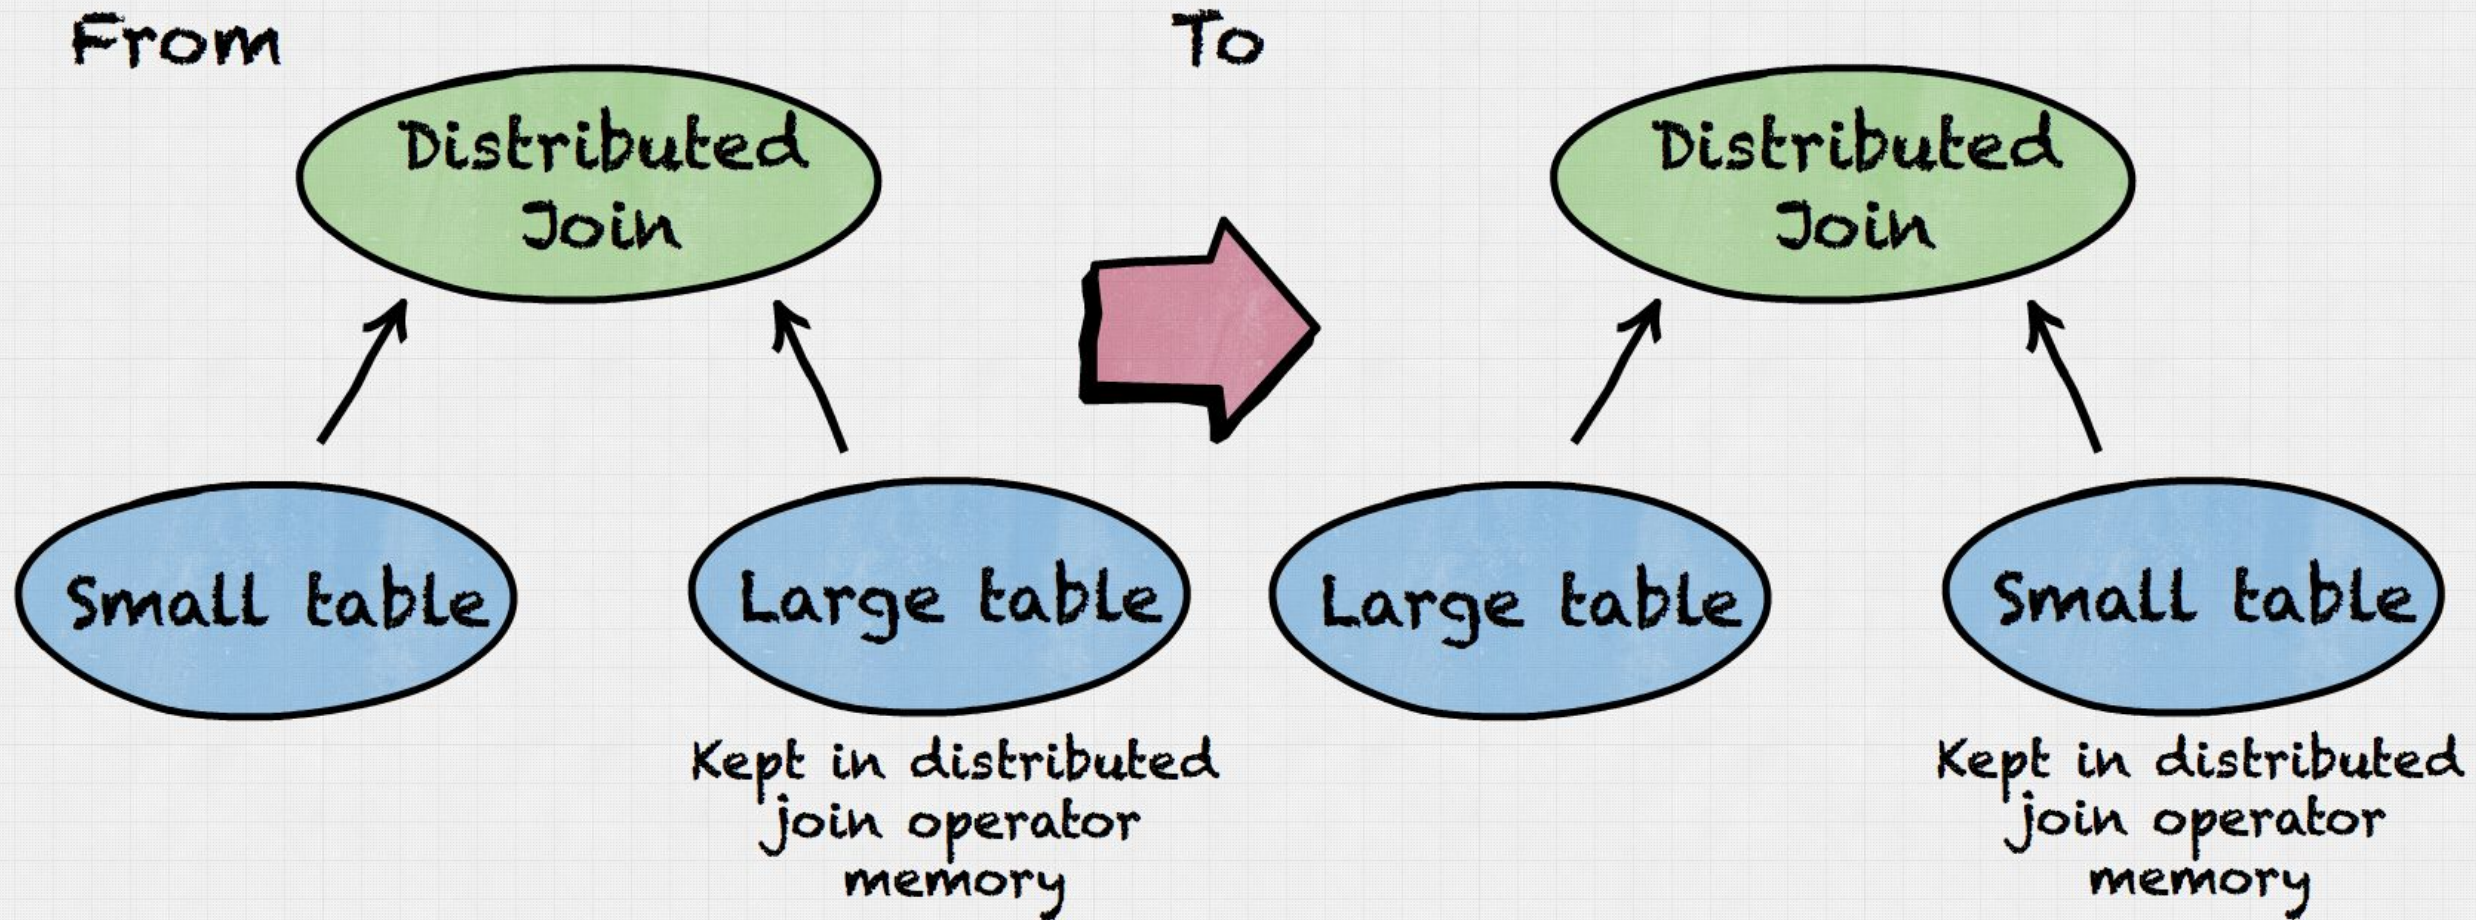
Cost calculation includes:

- CPU
- Memory
- Network I/O

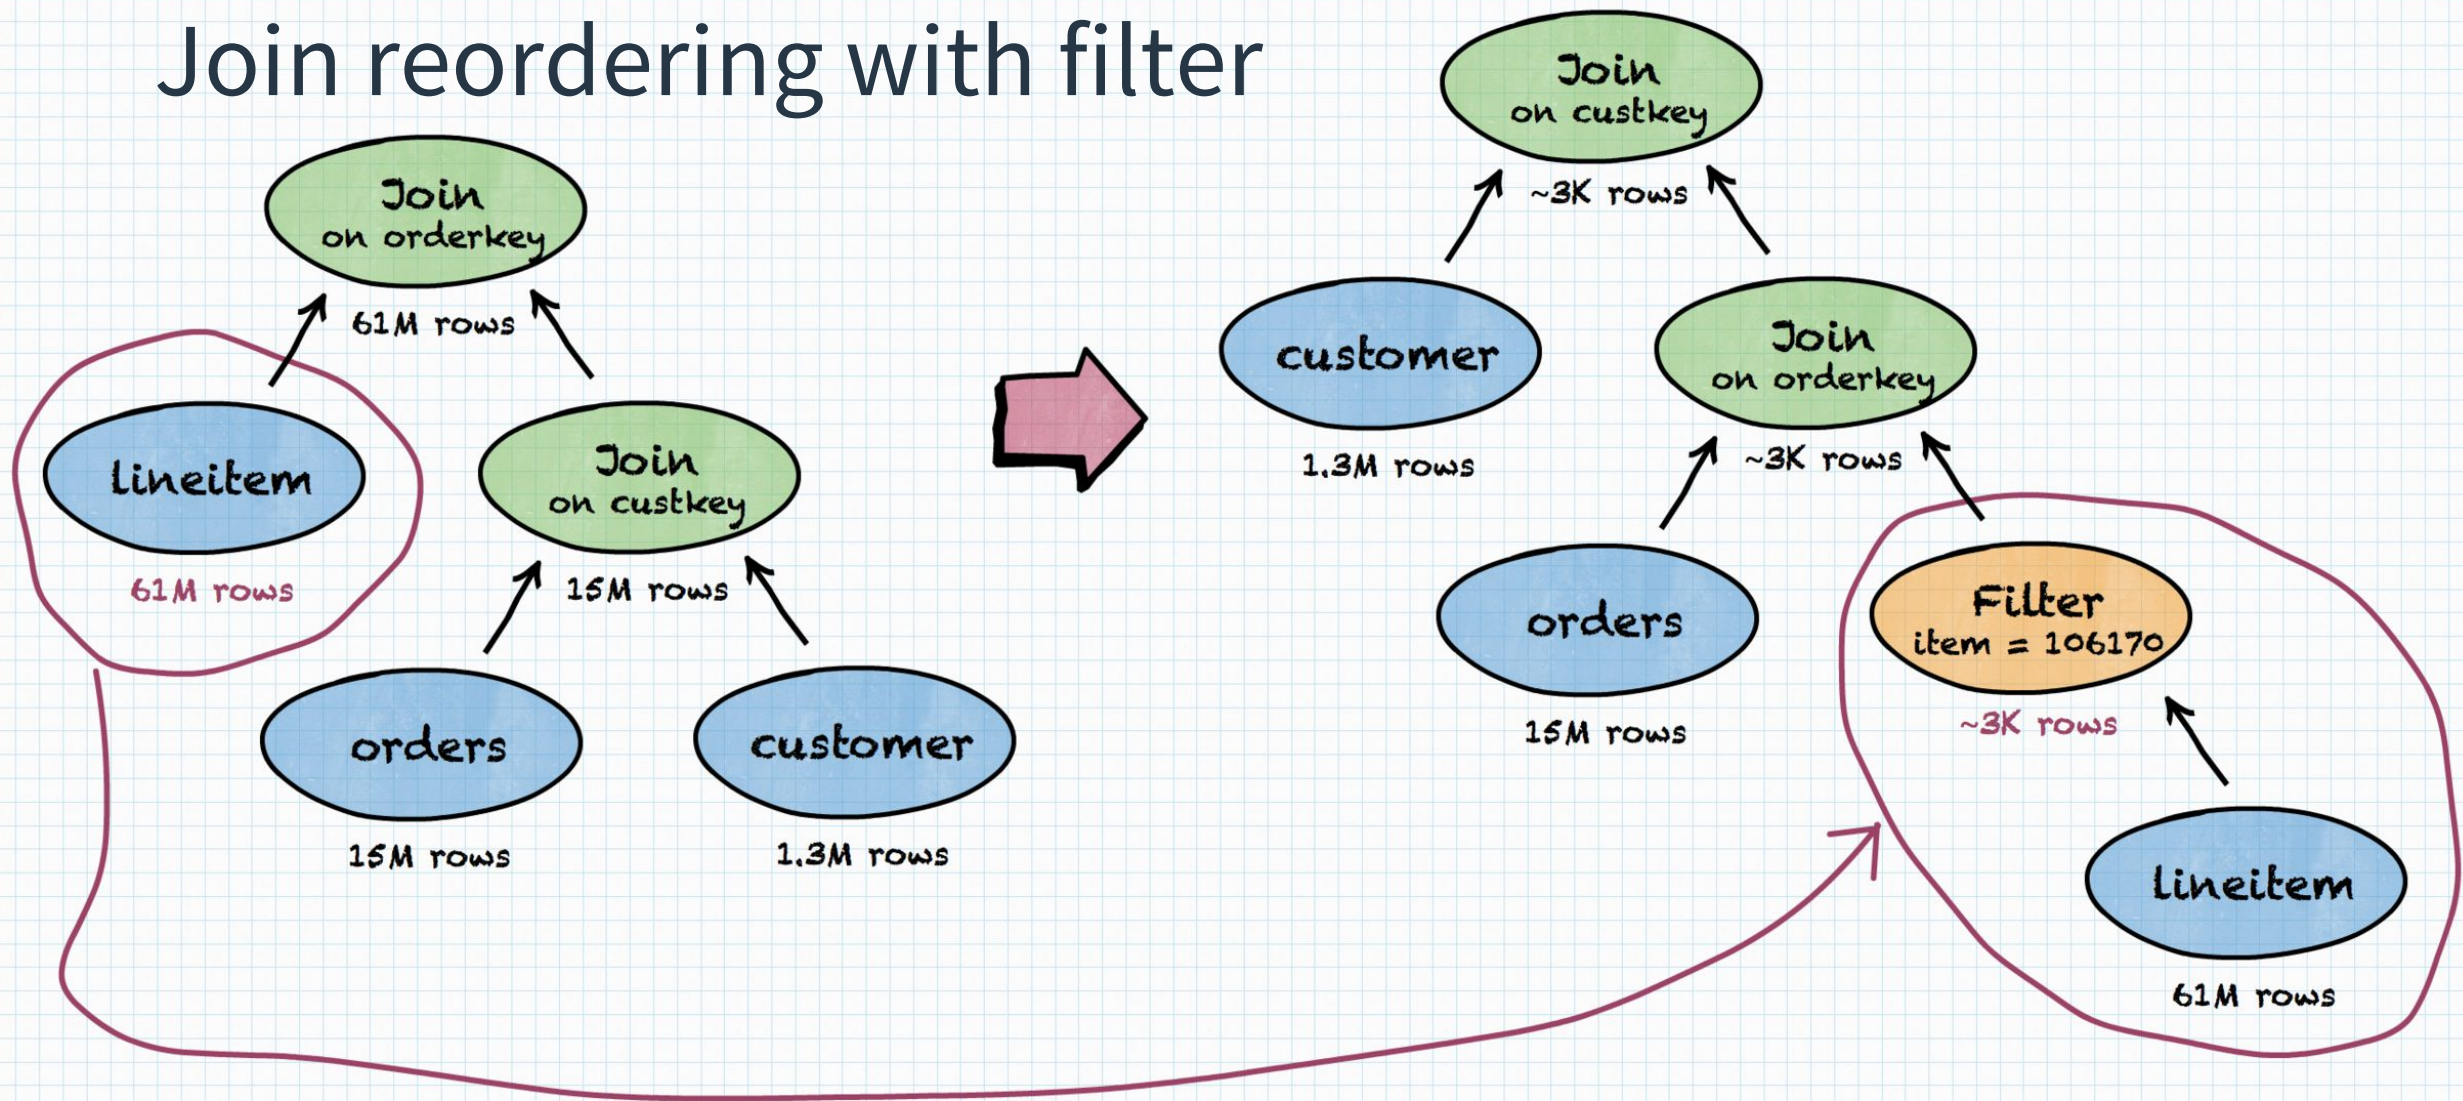
Join type selection



Join left/right side decision

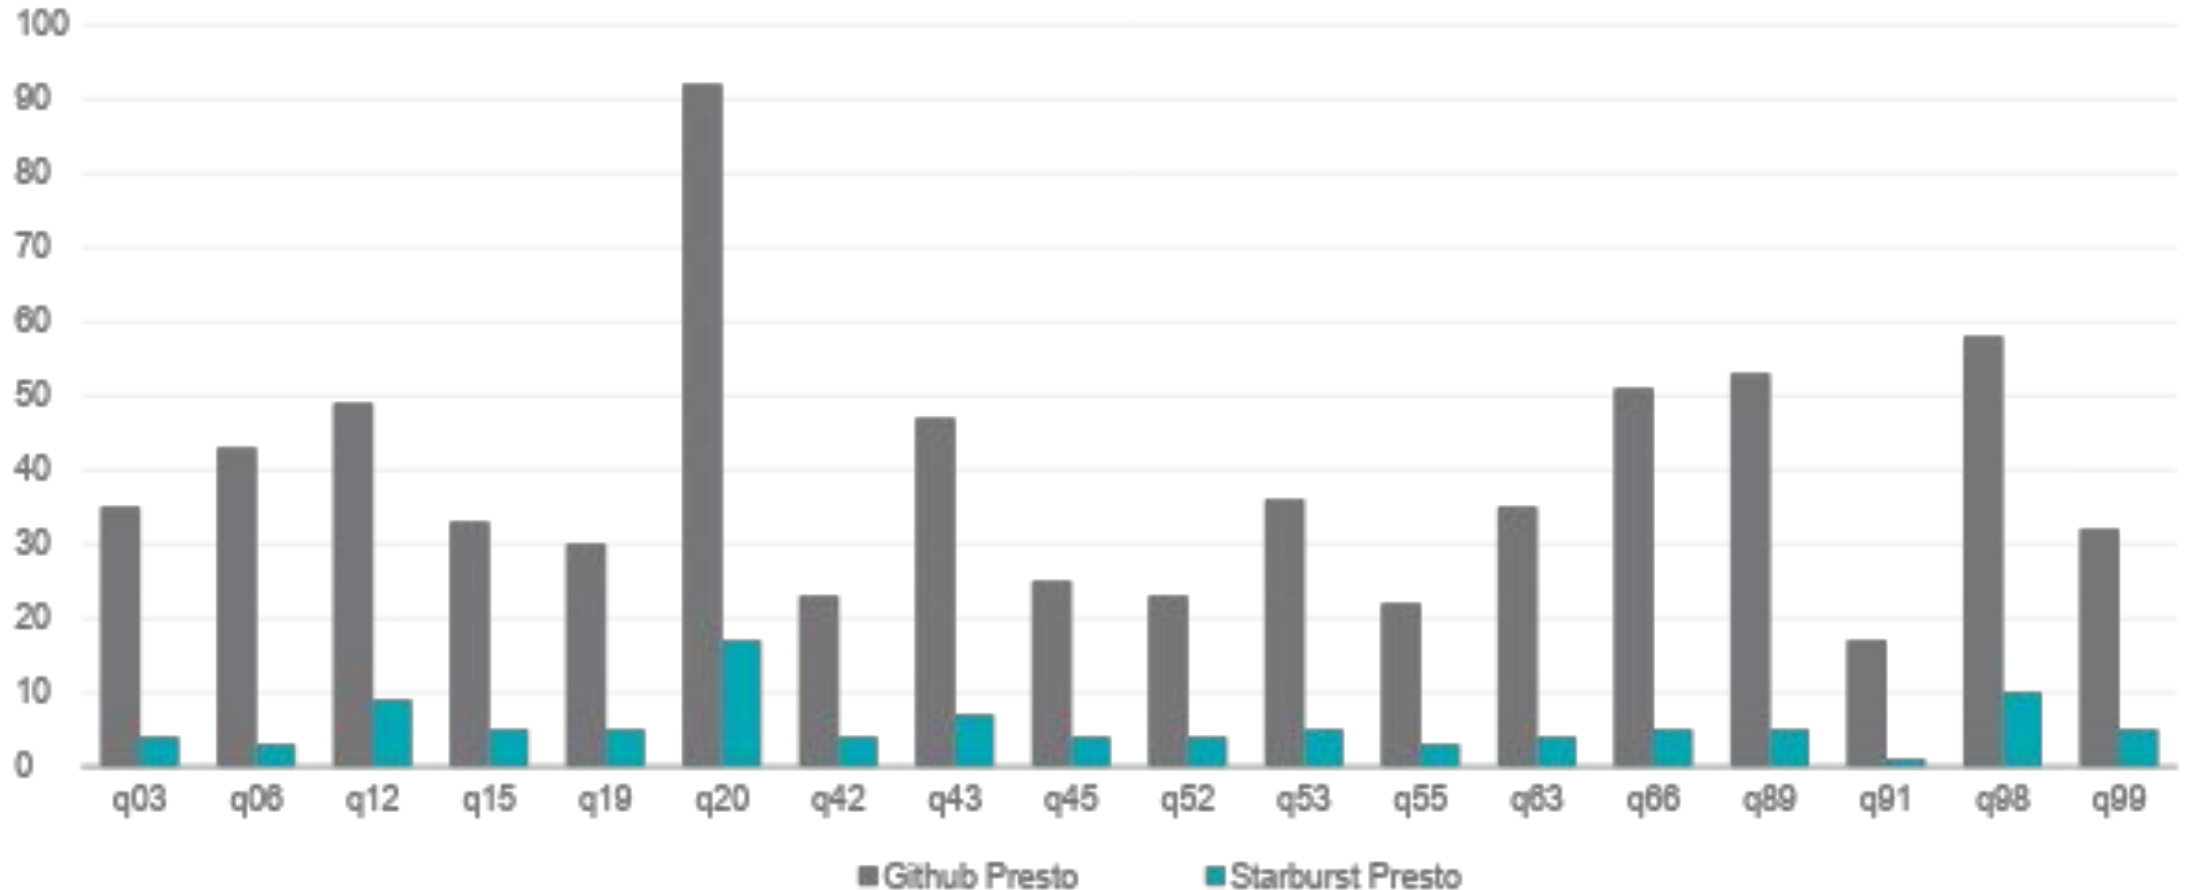


Join reordering with filter



Presto CBO Speedup

Duration of TPC-DS queries (lower is better)




<https://www.starburstdata.com/presto-benchmarks/>


Cloud cost reduction

- on average 7x improvement vs EMR Presto
- EMR Presto cannot execute many TPC-DS queries
- All TPC-DS queries pass on Starburst Presto

Starburst Presto (CBO) vs EMR Presto speedup




EMR Presto TPC-DS passed queries %



<https://www.starburstdata.com/presto-aws/>


Feedback and progress



Enhancements

- Deciding on semi-join distribution type based on cost
- Support for outer joins
- Capping a broadcasted table size
- various minor fixes in cardinality estimation
- ANALYZE table (native in Presto)
- Stats for AWS Glue Catalog (exclusive from Starburst)

Roadmap



What's next

- Stats support
 - Improved stats for Hive
 - Stats for DBMS connectors
 - Stats for NoSQL connectors
- Core CBO enhancements
 - Involve connectors in optimizations
 - Adjust cost model weights based on the hardware
 - Cost more operators
 - Introduce Traits
 - Adaptive optimizations

Further reading

<https://prestosql.io/>

<https://www.starburstdata.com/technical-blog/>

<https://fivetran.com/blog/warehouse-benchmark>

<https://www.concurrencylabs.com/blog/starburst-presto-vs-aws-emr-sql/>

<http://bytes.schibsted.com/bigdata-sql-query-engine-benchmark/>

<https://virtuslab.com/blog/benchmarking-spark-sql-presto-hive-bi-processing-goo-gles-cloud-dataproc/>

presto



Thank You!

Twitter: @starburstdata @prestosql

Blog: www.starburstdata.com/technical-blog/

Newsletter: www.starburstdata.com/newsletter