

Presto, a Faster, More Scalable, More Flexible Open Source SQL on Hadoop Engine

APACHE HADOOP/OPEN SOURCE



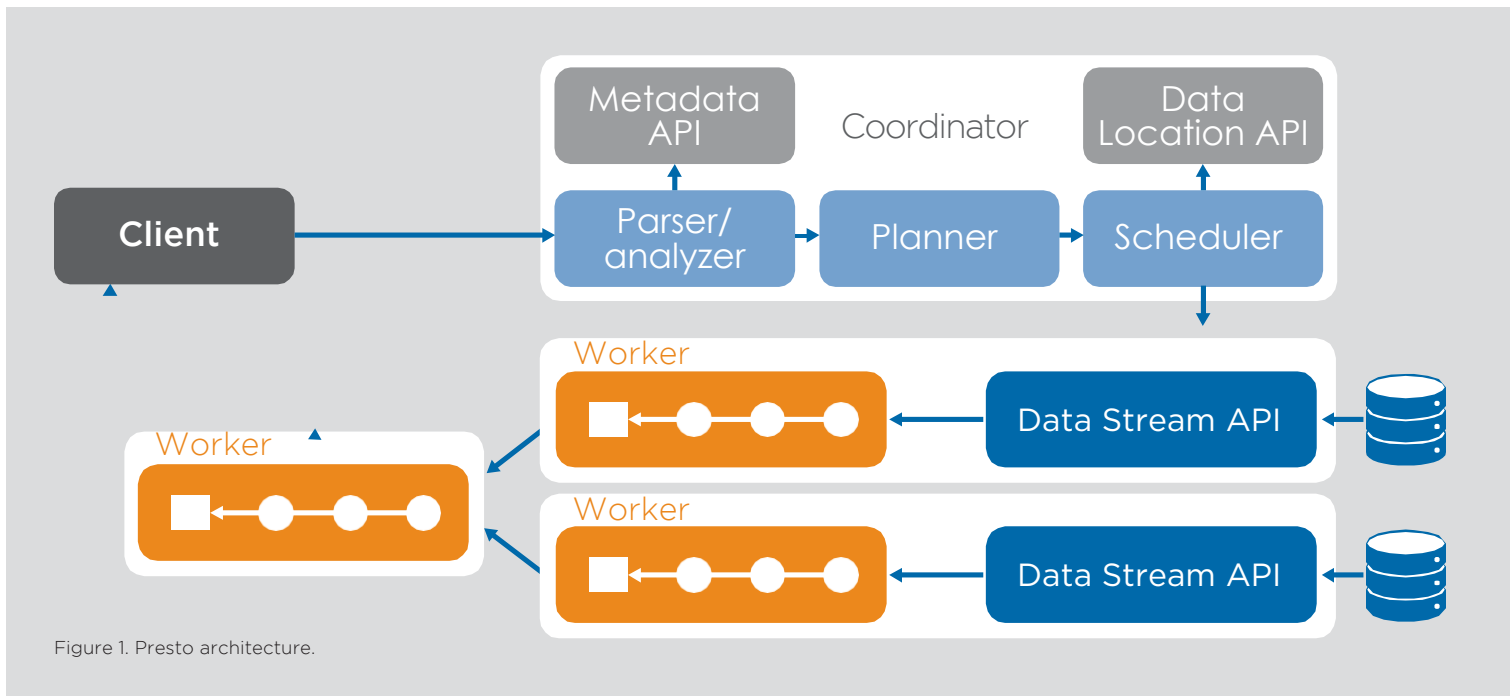
Presto

Originally developed by Facebook, Presto is an open source distributed SQL engine on Apache™ Hadoop®, for running interactive analytic queries against data sources of all sizes ranging from gigabytes to petabytes. Presto, unlike some other SQL on Hadoop engines, leverages standard ANSI SQL, and has been architected from the ground up for high-performance interactive query processing against Hadoop and other data sources. Many SQL on Hadoop engines are limited in performance by the restriction of writing processing steps to disk, while Presto's pure memory-based architecture is built for speed, allowing Presto to support large numbers of concurrent interactive queries against huge data sets. Presto's flexible architecture is unique in its ability to not only query Hadoop, but numerous other data sources such as Cassandra, MySQL, and PostgreSQL. Presto is not tied to a vendor, and runs on all the major Hadoop distributions, making it one of the most portable SQL-on-Hadoop tools in the market.

Figure 1 shows the simplified system architecture of Presto. The client sends SQL to the Presto coordinator for parsing, analysis, and planning of the query execution. The scheduler wires together the execution pipeline, assigns work to nodes closest to the data, and monitors progress. The pipelined execution model runs multiple stages at once, and streams data from one stage to the next as it becomes available. The Presto engine does not use MapReduce, rather it employs a modern query and execution engine with operators designed to support SQL semantics. In addition to improved scheduling, all processing is in memory and pipelined across the network between stages to avoid unnecessary I/O overhead and associated latency.

Teradata Commitment to Presto

Teradata has been a major supporter of the open source Presto project, and contributed numerous enterprise features and enhancements to the platform. Teradata was also the first commercial vendor to offer enterprise support for Presto. Our goal was to



increase adoption of this free, 100% open source software. Some of the major contributions that Teradata has made to the Project include new documentation, an administration tool, improved SQL Support, enterprise-class ODBC and JDBC drivers, Kerberos and LDAP support, Teradata™ QueryGrid™ support, BI tool certification, a SQL Server connector as well as numerous other fixes and performance improvements for Presto. Teradata is now focused on Presto's continued integration with QueryGrid and improvements as a crucial component of Teradata's overall QueryGrid and UDA strategy. All of Teradata's contributions to the core Presto engine are 100% open source under the Apache license and will advance Presto's modern code base, proven scalability, interactive querying, and cross-platform query capability.

Teradata worked closely with Facebook to outline these enhancements that have helped drive adoption of Presto within the enterprise. There continues to be a rapidly growing number of companies that are adopting and leveraging Presto. In addition, a vibrant and expanding number of users are contributing to the Presto project.

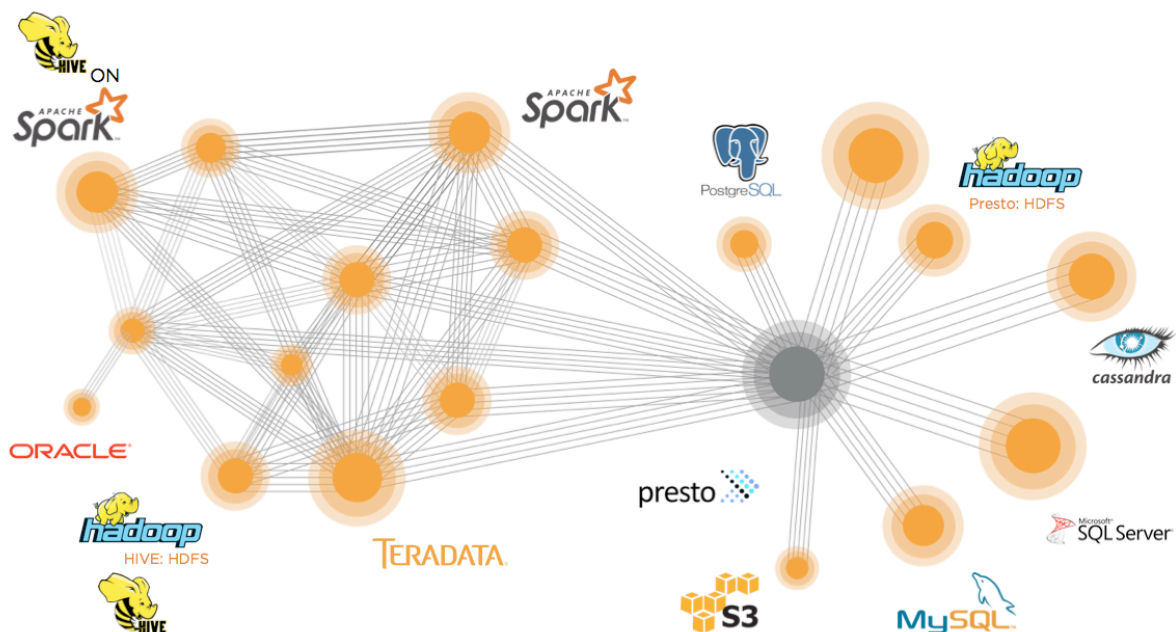
Teradata has partnered with Starburst Data (www.starburstdata.com) to deliver a stable enterprise ready distribution of Presto to its customers with continually improving enterprise features. These distributions are made available for download on Teradata.com/Presto.

SQL on Hadoop – Best Fit Engineering

Facebook originally created and open sourced the venerable Apache Hive™ in 2008. Recognizing that one tool cannot satisfy all the use cases at Facebook, Presto was developed and rolled out in 2013. However, Hive still plays a major role in the processing needs at Facebook, in conjunction with Presto, in what is commonly referred to as *best-fit-engineering*—whereby the use case is matched to the strengths of the technology.

- Hive is predicated on a declarative language known as Hive Query Language (HiveQL) which allows programmers familiar with the MapReduce framework to plug in custom mappers for greater data manipulation flexibility, and has made strides to include a subset of ANSI SQL syntax; whereas Presto has been built from the start to leverage the ubiquitous ANSI SQL protocol.
- Perhaps the strongest feature of Hive remains strong query fault tolerance designed to continue the execution of the query in event of the failure of some of its components.
- Third Party benchmarks from Radient Advisors has demonstrated order of magnitude increases in performance gains of Presto over Hive for many types of queries.
- While still in a nascent state, the optimizer represents an area where both tools have the opportunity to improve existing query performance and sophistication

Figure 2.



Comparing Apache Hive and Presto

	ANSI SQL Coverage	Query Fault Tolerance	Performance	Optimizer	Community	Platform Scope	Common Use Cases
Hive	HiveQL and Subset of ANSI SQL	Strong	Baseline, Improving	Basic Cost Based	Large, broad based adoption	Hadoop	ETL, Reporting, Deep Analysis
Presto	ANSI SQL	Weak	Highly Interactive	Basic Rules Based evolving to Cost Based	Early adopters from the world's largest internet companies	Hadoop, RDBMS, NoSQL	Interactive Analysis, Business Intelligence, Cross Platform Analysis, Self Service

through continued open source innovation and expertise borrowed from organizations with deep roots in database optimizers.

- Undoubtedly, Hive is extensively used across the Hadoop community given its heritage. Presto is currently being used by a community of well known, well respected Silicon Valley organizations with deep engineering roots - with growing interest from a variety of organizations as a result of Presto's existing capabilities and Teradata's commitment to enterprise features and support.
- Unlike most SQL-on-Hadoop tools like Hive, Presto goes beyond Hadoop and enables use cases that process data across the ecosystem.
- For the reasons mentioned, Hive is typically used for longer running ETL and reporting jobs, particularly where extreme sets of data are scanned and processed for deep analysis. Presto is typically leveraged for more interactive analysis, where ANSI SQL is needed for BI tools and the broadest self-service, and where the data required for analysis resides in multiple parts of the analytic ecosystem.

Presto and QueryGrid enhance the Teradata Unified Data Architecture™

Presto is a perfect fit with the Teradata Unified Data Architecture (UDA™) that allows organizations to capture, deploy, manage, and seamlessly access all of their data

across the analytic ecosystem (see Figure 2.) Advancing the UDA requires better SQL-on-Hadoop capabilities, whereby SQL is the interaction protocol across the UDA, and Teradata views Presto as a key engine to enable interactive querying against Hadoop with in the UDA. Within Teradata QueryGrid, Presto enables customers to more easily execute queries from Teradata with push down processing into Hadoop, as well as from Hadoop, with push down processing capabilities into Teradata Database. Presto enables QueryGrid to interface with numerous ecosystem components, such as Amazon S3, SQL Server, Cassandra, MySQL, PostgreSQL, and others to expand the scope and breath of the UDA.

Support and Professional Services

Similar to the early days of Hadoop, world-class customer support and professional services behind Presto will drive greater enterprise adoption. Teradata pioneered and was the first commercial vendor to support Presto. Now, in conjunction with Starburst Data, Teradata continues to provide the highest level of enterprise support for Presto. Teradata professional services also offers a variety of services around Presto.

Jumpstart allows companies to quickly pilot new functionality on site or in the cloud. Presto Training teaches engineers what they need to know about the *ins and outs* of Presto, understanding target use cases and the criteria for architecture decisions. Teradata professional services as well as Starburst Data can also help you design and build solutions based on Presto, from requirements analysis to solution design, to development and testing.

10000 Innovation Drive, Dayton, OH 45342 Teradata.com

QueryGrid, UDA, and Unified Data Architecture are trademarks, and Teradata, Aster, and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. Apache and Hadoop are registered trademarks of the Apache Software Foundation. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions, and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or Teradata.com for more information.

Copyright © 2016 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

10.16 EB8901

