Presto is a high performance, distributed SQL query engine for big data.

DOWNLOAD GITHUB

High performance

Presto is a highly parallel and distributed query engine, that is built from the ground up for efficient, low latency analytics.

Versatile

Supports diverse use cases: ad-hoc analytics at interactive speeds, massive multi-hour batch queries, and high volume apps that perform sub-second queries.

In-place analysis

Can natively query data in Hadoop, S3, Cassandra, MySQL, and many others, without the need for complex and error-prone processes for copying the data to a proprietary storage system.

Query federation

Access data from multiple systems within a single query. For example, join historic log data stored in S3 with real-time customer data stored in MySQL.

Works with existing BI tools

Presto is an ANSI SQL compliant query engine and works with BI tools such as R, Tableau, MicroStrategy, Power BI, and Superset.

Trusted

Presto is used for for critical business operations, including financial results for public markets, by some of the largest organizations in the world.

Runs everywhere

Presto is optimized for both on-premise and cloud environments such as Amazon, Azure, and Google Cloud.

Scalable

The largest organizations in the world use Presto to run exabyte scale warehouses.

Open

The Presto project is community driven project under the non-profit Presto Software Foundation.