

Resources

Presto ODBC

There are drivers available from: [Prestogres](#), [Starburst Data](#)

Prestogres

Prestogres

Project

[GitHub](#)

Maintained by

[Furuhashi Sadayuki](#)

Description

Prestogres is a gateway server that allows clients to use PostgreSQL protocol and thus the PostgreSQL ODBC driver to run queries on Presto.

Presto Resources

- [ODBC](#)
- [Libraries](#)
- [GUIs](#)
- [Management Tools](#)
- [Enterprise Support](#)
- [Hosted Presto](#)
- [Cloud](#)

Starburst

Starburst ODBC Driver

Project

[Starburst Enterprise ODBC/JDBC Presto Drivers](#)

Maintained by

[Starburst, The Presto company](#)

Description

The Starburst Enterprise Client Drivers for Presto include enterprise grade ODBC and JDBC drivers enabling you to use your preferred Business Intelligence tools with Presto. The drivers fully implement the ODBC and JDBC specifications and are compatible with the [Enterprise Starburst Presto release](#) which is available for download. There are ODBC drivers for Windows, Mac, and Linux platforms. Starburst additionally provides [enterprise Presto support and services](#)

Presto Libraries

The following client libraries can be used to run queries from several programming languages: [C](#), [Go](#), [Java](#), [Node.js](#), [Python](#), [R](#), [Ruby](#).

Language: C

PrestoClient C

Project

[GitHub](#)

Maintained by

[Ivo Herweijer](#)

Description

C client for Presto.

Language: Go

[presto-go-client](#)

presto-go-client

Project

[GitHub](#)

Maintained by

[Presto Go Client Team](#)

Description

Go client for Presto.

Language: Java

Presto JDBC Driver

Project

[Presto](#)

Maintained by

[Presto Team](#)

Description

JDBC driver for Presto.

Example

```
String sql = "SELECT * FROM example";
String url = "jdbc:presto://localhost:8080/catalog/schema";
try (Connection connection =
    DriverManager.getConnection(url, "test", null)) {
    try (Statement statement = connection.createStatement()) {
        try (ResultSet rs = statement.executeQuery(sql)) {
            while (rs.next()) {
                System.out.println(rs.getString("node_id"));
            }
        }
    }
}
```

Language: Node.js

presto-client-node

Project

[GitHub](#)

Maintained by

[Satoshi Tagomori](#)

Description

Node.js client for Presto.

Language: Python

presto-python-client

Project

[GitHub](#)

Maintained by

[Presto Python Client Team](#)

Description

Python client for Presto.

PyHive

[Project](#)

Project

[GitHub](#)

Maintained by

[Dropbox](#)

Description

PyHive is a collection of Python DB-API and SQLAlchemy interfaces for Presto and Hive.

PrestoClient Python

Project

[GitHub](#)

Maintained by

[Ivo Herweijer](#)

Description

Python client for Presto.

Language: R

RPresto

Project

[GitHub](#)

Maintained by

[RPresto Team](#)

Description

DBI-based adapter for Presto for R.

Language: Ruby

presto-client-ruby

Project

[GitHub](#)

Maintained by

[Furuhashi Sadayuki](#)

Description

Ruby client for Presto.

Presto GUIs

Airpal

Project

[Airpal](#)

Maintained by

[Airbnb](#)

Description

Airpal is a web-based, query execution tool which leverages Presto to make authoring queries and retrieving results simple for users. Airpal provides the ability to find tables, see metadata, browse sample rows, write and edit queries, then submit queries all in a web interface.

Redash

Project

[Redash](#)

Maintained by

[Arik Fraimovich](#)

Description

Redash is a take on freeing the data within our company in a way that will better fit our culture and usage patterns. It has Presto support as well as other backends, and offers a query editor with syntax highlighting and completion, and creating visualizations and dashboards from query results.

Shib

Project

[GitHub](#)

Maintained by

[Tagomori Satoshi](#)

Description

Shib is a web-client written in Node.js designed to query Presto and Hive. To run Shib install node.js, alter your config.js, and follow the instructions on the shib project page. Shib can also be used as an proxy server for query engines.

Superset

Project

[Apache Superset](#)

Maintained by

[Apache Superset](#)

Description

Superset enables users to consume data in many different ways: writing SQL queries, creating new tables, creating a visualization (slice), adding that visualization to one or many dashboards and downloading a CSV. SQL Lab is a part of Superset and provides rich SQL editor that enables users to both query and visualize data. You can explore and preview tables in Presto, effortlessly compose SQL queries to access data. From there, you can either export a CSV file or immediately visualize your data in the Superset "Explore" view.

yanagishima

Project

[GitHub](#)

Maintained by

[wyukawa, okazou](#)

Description

yanagishima is a web application for Presto. yanagishima provides the ability to execute query, show query, kill query, bookmark query, search table, share query/query result, format query, download as CSV/TSV file, insert chart, substitute query parameter, and so on.

Presto Management Tools

Presto-Admin

Project

[Presto Admin](#)

Maintained by[Starburst, The Presto company](#)**Description**

Presto-Admin is a tool for installing and managing the Presto query engine on a cluster. It provides easy-to-use commands:

- Install and uninstall Presto across your cluster
- Configure your Presto cluster
- Start and stop the Presto servers
- Gather status and log information from your Presto cluster

If you need any assistance with your Presto cluster management reach out to [Starburst](#) for their [Presto enterprise support](#) offering and other Presto related services.

Examples

```
presto-admin server start|stop|restart|status
presto-admin server install path-to-presto-rpm
presto-admin connector add connector-name
```

Enterprise Support for Presto

Starburst Data

Website[Starburst, The Presto company](#)**Scope**

Enterprise 24/7 Support, Installation, Configuration, Training, Custom Development, Tuning

Description

At Starburst, our team is a major contributor to the open source Presto project. We consist of many of the **experts and committers** who have been contributing to and advancing the Presto product over the last few years. Starburst provides an [enterprise ready Presto distribution](#) and [Presto support](#). Starburst's distribution of Presto consists of additional tooling and configurations to make it work well in the enterprise. Further, it is rigorously tested at scale and patched as needed. And with our wide range of services, support, and training we will help you be successful in this new world of open source technologies in the enterprise.

In addition to implementing Presto internals to make it the most reliable, robust, and performant open source distributed query engine, our team at Starburst is dedicated to continually improve and add the needed features for the enterprise. Our major enterprise focus areas are ease-of-use, security, robustness, wide range of integrations, and reliable support SLAs.

Hosted Presto

Qubole

Product[Managed Presto Service](#)**Description**

Qubole integrates an enhanced and cloud-optimized version of Presto. Qubole's Presto implementation is an enterprise-ready and secure distributed SQL query engine, which allows analysts to quickly derive business insights from data.

Qubole has optimized Presto for the cloud. Qubole's enhancements allow for dynamic cluster sizing based on workload and termination of idle clusters – ensuring high reliability while reducing compute costs. Qubole's Presto clusters support multi-tenancy and provide logs and metrics to track performance of queries.

Amazon Athena

Website

[Amazon Athena](#)

Maintainer

[Amazon Web Services](#)

Amazon Athena is an interactive query service based on Presto that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

Amazon Athena uses Presto with full standard SQL support and works with a variety of standard data formats. Athena is out-of-the-box integrated with AWS Glue Data Catalog, allowing you to create a unified metadata repository across various services, crawl data sources to discover schemas and populate your Catalog with new and modified table and partition definitions, and maintain schema versioning.

Presto Cloud

Presto is also readily available in [AWS](#) and [Azure](#) cloud environments.

AWS

Starburst Presto

Website

[Presto on AWS](#)

Maintainer

[Starburst, The Presto company](#)

Starburst Presto on AWS combines the reliable, scalable, and cost-effective cloud computing services provided by Amazon Web Services (AWS) with the power of Presto, the fastest growing distributed SQL query engine within the industry.

Through the use of Starburst's CloudFormation template and Starburst Presto AMI, Starburst Presto on AWS enables you to run analytic SQL queries across a wide variety of data sources with elastic scaling and usage-based pricing. Read more how to use Presto on AWS on our [AWS Documentation site](#).

Amazon EMR

Website

[Amazon EMR](#)

Maintainer

[Amazon Web Services](#)

Amazon EMR provides a managed Hadoop framework that makes it easy, fast, and cost-effective to process vast amounts of data across dynamically scalable Amazon EC2 instances. With EMR, you can launch a large Presto cluster in minutes. You don't need to worry about node provisioning, cluster setup or tuning.

Using Presto on EMR provides these benefits to customers:

- Elasticity: With Amazon EMR, you can provision one, hundreds, or thousands of compute instances to process data at any scale. You can easily increase or decrease the number of instances manually or with Auto Scaling, and you only pay for what you use.
- Simple and predictable pricing: You pay a per-second rate for every second used, with a one-minute minimum charge.

Azure

Starburst Presto

Website

[Presto on Azure](#)

Maintainer

[Starburst, The Presto company](#)

Starburst Presto for HDInsight is a distributed SQL query engine that is integrated into Azure HDInsight Platform and available via the Azure Marketplace.

Azure HDInsight is a fully-managed cloud service that makes it easy, fast, and cost-effective to process massive amounts of data. Presto and HDInsight were both designed for the separation of storage and compute which allows for flexible performance and cost. You pay only for what you use by creating clusters on demand scaling them up and down. Read more how to use Presto on Azure on our [Azure Documentation site](#).

Integrating Presto with HDInsight provides Azure users with two new capabilities:

- A fast, scalable, interactive SQL interface to data in Azure Blob and Azure Data Lake Storage
- An easy way to integrate data with HDInsight by leveraging Presto's vast portfolio of data connectors.

[Documentation](#) · [Community](#) · [Resources](#) · [Development](#) · [Download](#)



Presto is open source software licensed under the Apache License and supported by the Presto Software Foundation