

STARBURST ENTERPRISE PRESTO

DEPLOY ANYWHERE, QUERY ANYTHING

The Fastest Path from Big Data to Better Decisions

The traditional response to most new analytic needs is to move data from various sources into a data warehouse or data lake. Doing this can take months and cost a great deal of money.

Originally developed and open sourced by Facebook, Presto is a distributed SQL engine for running fast, interactive analytic queries against various data sources ranging in size from terabytes to exabytes. It enables immediate analysis of siloed data without expensive data warehouse appliances and significantly reduces the need to move or copy data.

Starburst Enterprise Presto is purpose-built to ensure Presto security and performance at scale, and includes multiple enterprise-grade features, global security enhancements, and 24x7 support from the Presto experts.

Run Fast SQL Queries at Scale

- Easily provision and scale Presto SQL in the public cloud or on-premises.
- Separate SQL compute from data storage and scale each independently.
- Presto users like Uber, Netflix, and Comcast run 100,000s of queries per day against petabytes of data.
- In-memory, columnar architecture for fast query response.
- Get big data projects running in days, not months.

Enterprise-grade Analytics Anywhere

The Starburst Enterprise Presto platform delivers additional benefits for the world's largest organizations, including:



Performance: Our distribution includes the updated cost-based query optimizer; and optional caching of frequently accessed data



Connectivity: 27+ supported connectors; high-performance parallel connectors



Security: fine-grained access control, end-to-end encryption and more



Management: configuration; autoscaling; high availability; monitoring



Support: 24x7 support from the Presto experts; fully-tested, stable releases, etc.

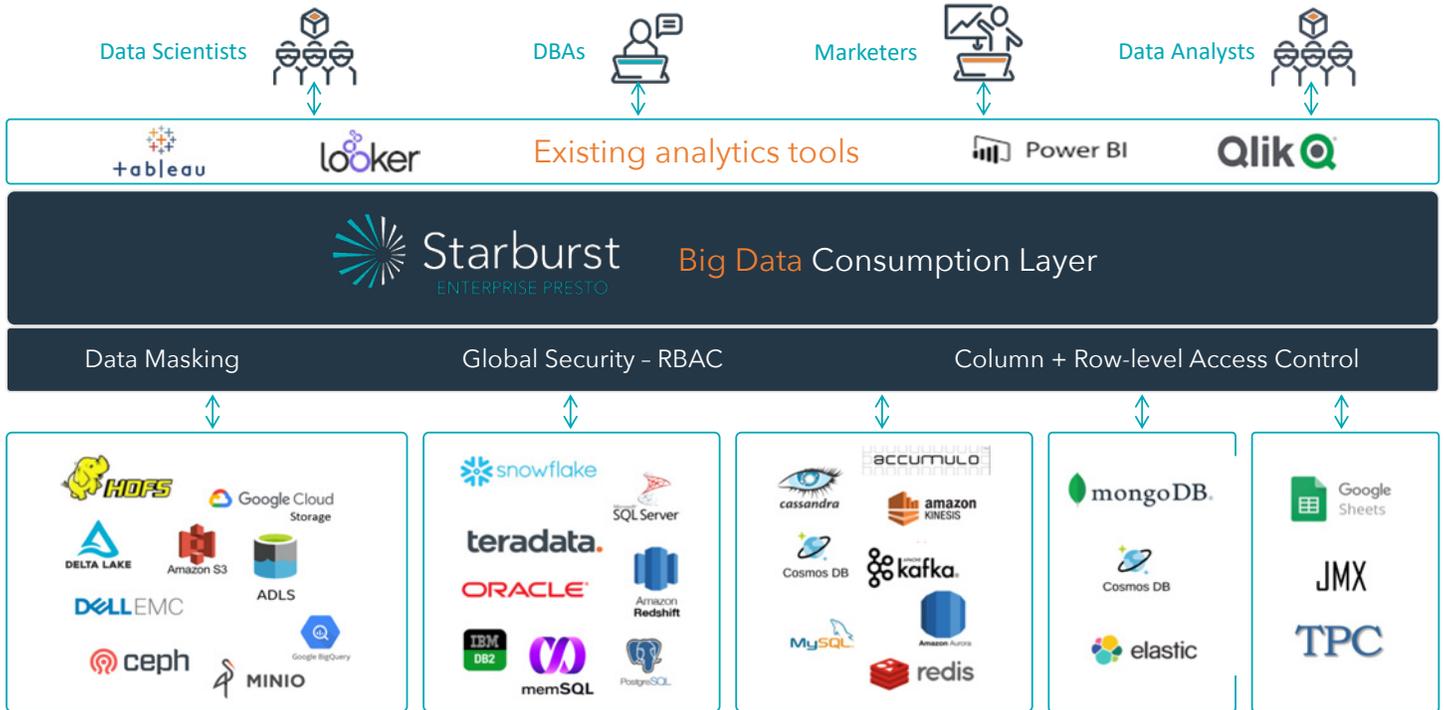
Starburst Enterprise Presto

- Single point of access to all your data, everywhere
- Join across all your data sources
- Directly query any distributed object storage (S3, ADLS, etc.)
- Fine-grained access control and end-to-end encryption
- 24x7 support from the Presto experts
- Fully-tested, stable releases
- Easy access from any popular BI tool
- Reduced big data costs by optimizing cloud spend
- Up to 10x performance improvement over alternatives like Hive

Introducing The Data Consumption Layer

Starburst solves for the intersection of ever-increasing data volumes and ever-increasing data consumers. Our perspective is that the historical approach of constantly moving data into a single-source-of-truth method is no longer possible - there's too much data, and too many people trying to ask new questions of the data. We help our customers de-prioritize data movement while they give their teams consistent, easy, and fast SQL based access to all their data. We call this the data consumption layer.

The Data Consumption Layer



Starburst Enterprise Presto is a data consumption layer that unlocks siloed data by providing fast access to any data source, via SQL - the tool language your team already knows and then can be utilized in the analytics tool you already use. With advanced security features – this data consumption layer provides a secure single-point-of-access for all your data consumers to all your data.

SINGLE-POINT-OF-ACCESS

Starburst Enterprise Presto Explained

Starburst Enterprise Presto

The graphic features the Starburst logo at the top center, with the text 'Starburst ENTERPRISE PRESTO' to its right. Below the logo are five teal-colored boxes, each representing a key feature area: Performance, Connectivity, Security, Management, and Support. Each box contains a list of specific capabilities and logos of associated partners or technologies.

- Performance:** Includes the latest cost-based query optimizer; Caching available for frequently accessed data.
- Connectivity:** 26+ supported enterprise connectors; High performance parallel connectors for Oracle, Teradata, Snowflake and more. Logos for Teradata, Snowflake, Oracle, and BigQuery are shown.
- Security:** Kerberos & LDAP integration; Global Security for fine-grained Access Control; Data encryption; Data masking; Query auditing. Logos for Apache Ranger, Okta, and Sentry are shown.
- Management:** Configuration; Autoscaling; High availability; Monitoring; Deploy anywhere. Logos for AWS, Google Cloud, Azure, and Snowflake are shown.
- Support:** The largest team of Presto experts in the world; Fully-tested, stable releases, curated by the Presto creators; Hot fixes & security patches; 24x7 support, 365 – we've got your back.

Starburst optimizes Presto for the enterprise. High-performance Starburst Enterprise Presto connectors enable queries from one or more sources, on-premises or in the cloud, and everything is governed by enterprise-grade global security features with fine-grained access control.

Deep Dive on Security

Starburst delivers global, enterprise-grade security features designed to meet the standards of the world's largest organizations. Security enhancements include:

- Fine-grained access control at the connector, table, column and row-level
- End-to-end data encryption, user authentication, and other enterprise security best practices further protect your Presto clusters and the data queried
- Kerberos & LDAP integration, data masking, and query auditing

The Biggest Brands in the World Rely on Presto

- Comcast runs over 350,000 Presto queries daily.
- 1,000+ Facebook users run Presto queries daily.
- Netflix, Lyft, Airbnb Twitter, LinkedIn, Amazon, and Alibaba all depend on Presto.
- More and more of the largest, best-known enterprises in the world are upgrading to Starburst Enterprise Presto each month, including Slack, Comcast, Zalando, FINRA, and Condé Nast
- Common use cases include interactive data investigation, BI dashboard and reports, data science, ETL, and as a high-performance data lake query engine.

FINRA: Protecting the Financial Markets from Fraud

Financial Industry Regulatory Authority analyzes billions of new trading events per day to detect securities fraud, insider trading, and other abuse.

“By separating SQL compute from storage, Starburst Presto makes it possible to economically analyze 25PB of trading data for fraud—100 billion rows of new data per day from 25+ sources.”

– John Hitchingham, FINRA

Why Presto?

“Presto, originally created at Facebook and commercially supported by Starburst Data, can query data not just in Hadoop and object storage, but also relational and non-relational databases (including NoSQL and Kafka), providing greater flexibility by enabling federated querying of multiple data sources, in multiple locations, at the same time.”

– Matt Aslett, 451 Research



Data is Everywhere!