

# Starburst Enterprise

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, Trino, formerly known as Presto SQL, 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 is purpose-built to ensure Trino 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 Trino in the public cloud or on-premises.
- Separate SQL compute from data storage and scale each independently.
- Trino 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 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:** 40+ 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 Trino experts; fully-tested, stable releases, etc.

## Starburst Enterprise

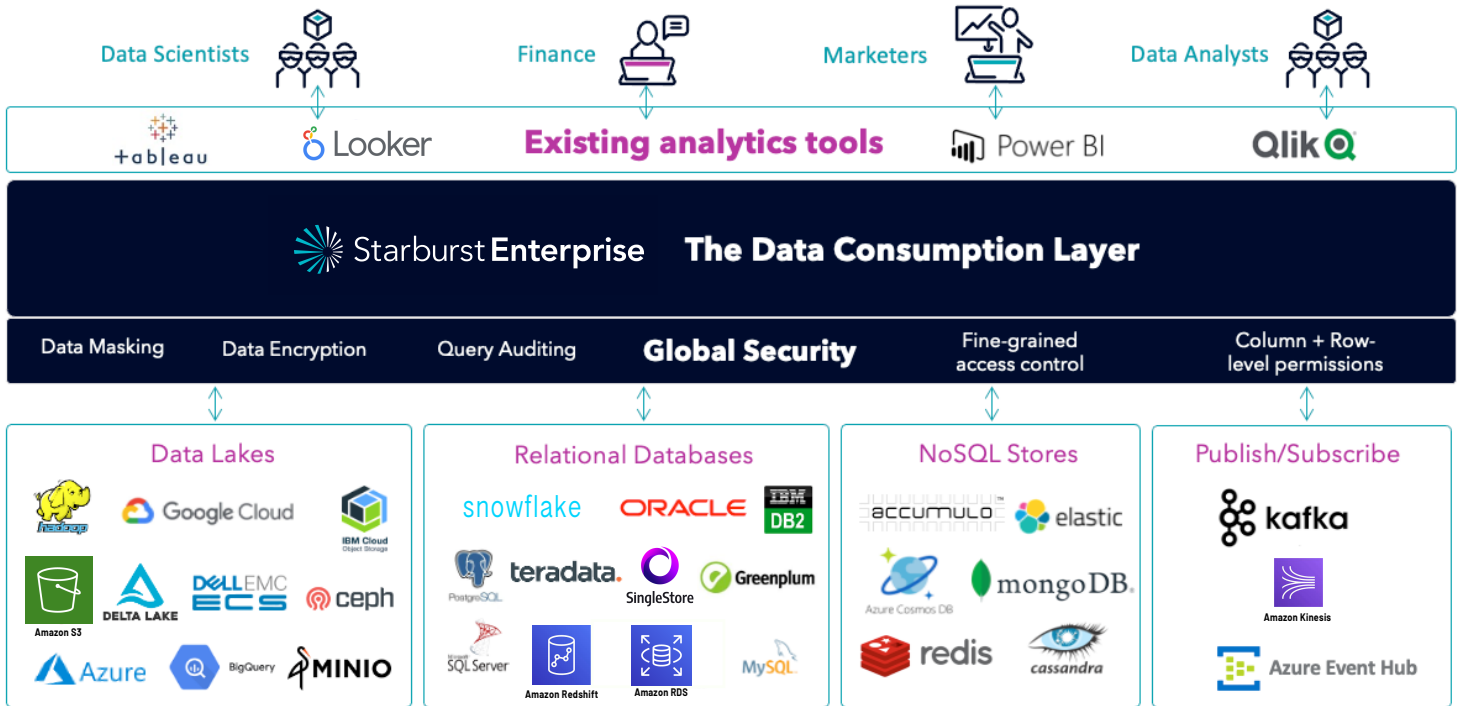
- Single point of access to all your data, everywhere
- Join across all your data sources
- Directly query any distributed object storage (S3, ADLS, etc.)
- Fine-grained access control and end-to-end encryption
- 24x7 support from the Presto experts
- Fully-tested, stable releases
- Easy access from 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 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.

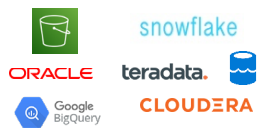




# SINGLE-POINT-OF-ACCESS



## Starburst Enterprise Explained

# Starburst Enterprise

Performance	Connectivity	Security	Management	Support
<p>From petabytes to exabytes – query data from disparate sources using SQL – with high concurrency</p> <p>Control your price/performance with the latest cost-based optimizer</p> <p>Caching available for frequently accessed data</p>	<p>40+ supported enterprise connectors</p> <p>High performance parallel connectors for Oracle, Teradata, Snowflake, Cloudera, and more</p> 	<p>Kerberos &amp; LDAP integration</p> <p>Global Security for fine-grained Access Control</p> <p>Data encryption</p> <p>Data masking</p> <p>Query auditing</p> 	<p>Configuration</p> <p>Autoscaling</p> <p>High availability</p> <p>Monitoring</p> <p>Deploy anywhere</p> 	<p>The largest team of Trino experts in the world</p> <p>Fully-tested, stable releases, curated by the Trino creators</p> <p>Hot fixes &amp; security patches</p> <p>24x7 support, 365 – we've got your back</p>

Starburst optimizes Trino for the enterprise. High-performance Starburst Enterprise 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

## Starburst Insights

Starburst Insights provides a visual overview of important metrics about your Starburst Enterprise cluster for all types of users, from platform administrators to data consumers. From the Insights interface, you can access detailed query history, including single-query statistics and query plans, and cluster performance information from a selected date range.

## Cached Views

Starburst Cached Views is a collection of performance features that provides customers with options on accessing frequently used data. Starburst customers benefit from materialized views, table scan redirections, and the cache service/CLI for easy management, configuration, and most importantly, speed for accessing the same, or similar, data sets repeatedly. Starburst Cached Views can now improve read performance up to 20x and update 100's of times faster. More efficiency, less work.



## The Biggest Brands in the World Rely on Trino

- Comcast runs over 350,000 Trino queries daily
- Thousands of users globally use Trino to run queries daily
- Netflix, Lyft, Airbnb Twitter, LinkedIn, Amazon, and Alibaba all depend on Trino
- More and more of the largest, best-known enterprises in the world are upgrading to Starburst Enterprise 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 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 Starburst?

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!