

STARBURST ENTERPRISE PRESTO ON AZURE

Available on the Microsoft Azure marketplace, the Starburst Enterprise Presto platform is a fully supported, production-tested, enterprise-grade distribution of the open source Presto MPP SQL query engine. Starburst integrates the scalable cloud storage and computing services provided by Azure with a more stable, secure, efficient, and cost-effective way to query all your enterprise data, wherever it resides.

Leading organizations across multiple industries rely on Starburst Enterprise Presto and Azure.

Analytics Anywhere

Architected for the separation of storage and compute, Presto is ideal for querying data residing in multiple systems, from cloud data lakes to legacy data warehouses. Deployed via Azure Kubernetes Service (AKS), Starburst Enterprise Presto on Azure enables the user to run analytic queries across Azure data sources and on-prem systems such as Teradata, Oracle, and others via Presto clusters.

Within a single query, you can access multiple data stores, including Azure Blob Storage, ADLS, Azure SQL, SQL Server, plus Hadoop, Synapse, IBM DB2, Cassandra, Databricks Delta Lake, and many others. In minutes, users are able to provision from small to large clusters of compute instances and leverage the power of Presto's parallelism to analyze all enterprise data.



The Starburst Enterprise Presto Difference

Although they are built on the same basic query engine, there are several important differences between Starburst Enterprise Presto and the open source version of Presto. With Starburst, enterprises enjoy enhanced:

Performance:

Includes the latest optimizations; caching available for frequently accessed data; stable code that minimizes failed queries

Connectivity:

30+ supported enterprise connectors; high-performance connectors for Oracle, Teradata, Snowflake, IBM DB2, Databricks Delta Lake, and many more

Security:

Fine-grained, row- and column-level access control; Kerberos, OKTA, LDAP integration; data encryption & masking; query auditing to see who is doing what; plus standard role-based access control (via Apache Ranger)

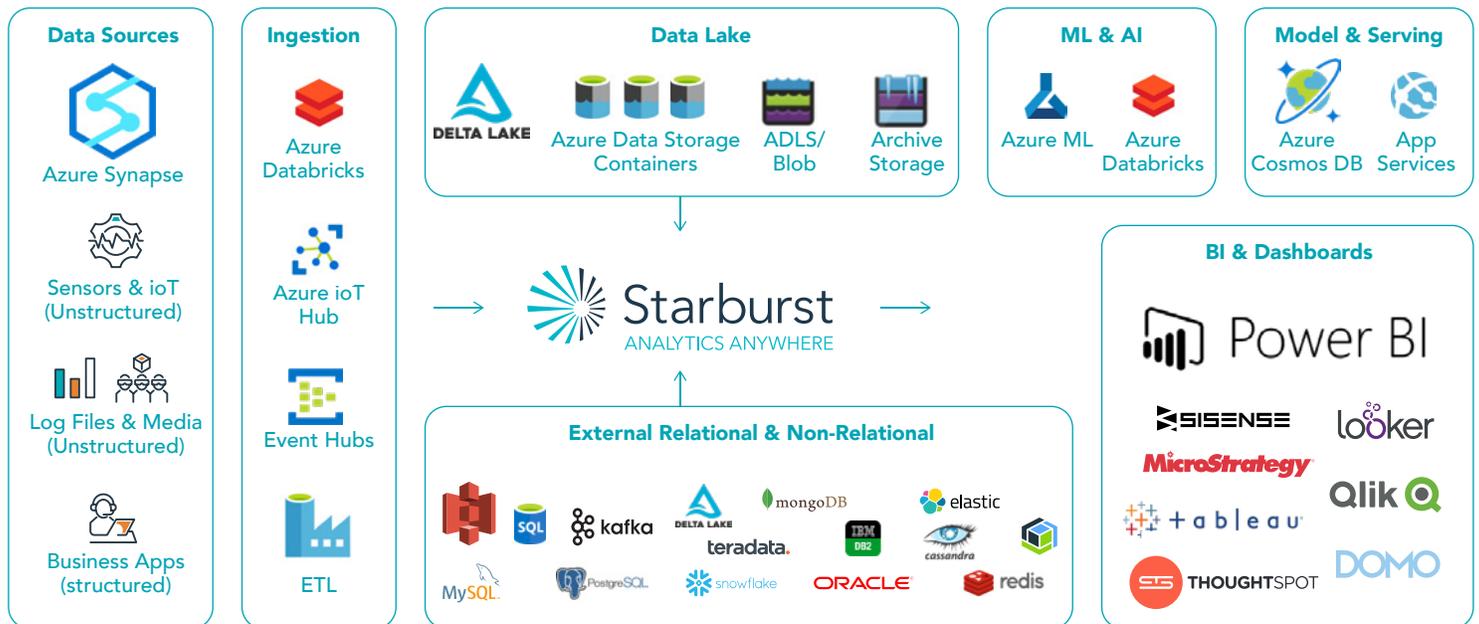
Management:

Enhanced tools for configuration, autoscaling, and monitoring; high availability; easy deployment on Azure platforms

Support:

24/7/365 enterprise-grade support from the largest team of Presto experts in the world; fully tested, stable releases; hot fixes & security patches

What these and other features add up to is a more powerful, stable, cost-effective, and customizable query engine that drives adoption within organizations, enabling companies to generate more value from their Azure resources.



Common Enterprise Use Cases

Organizations are using Starburst Presto on Azure in a variety of ways, but the following five use cases are the most popular.

✓ **Interactive Data Investigation**

The Starburst platform allows for rapid ad hoc analysis and data exploration. Data scientists can use programming libraries such as R, Python and Java to examine data where it resides before running ETL jobs. Business analysts can quickly glean insights from data to find value in unexpected market disruptions or other time-sensitive events. Generally, end users can explore data in ways they were never able to previously.

✓ **BI Dashboard and Reports**

Getting data to a dashboard typically demands a tough decision – you either suffer the latency or you sample a smaller, less accurate subset of your data to yield results faster. Starburst Enterprise Presto gives customers access to all their data in a latency-free environment – ensuring faster, more accurate reporting and better results.

✓ **Data Science**

Jobs that used to take hours or days to run can be shrunk down to minutes, which allows you to extract value out of your data faster and enables your data scientists to work at speed. Starburst Enterprise Presto has libraries/package support for Java, Python, R, Spark, etc. – familiar data science tools such as Jupyter Notebook can talk directly to Starburst, which then queries data where it lies.

✓ **ETL**

The Starburst platform can remove the need for ETL and the setup of data marts by querying data in place. But ETL is still essential in some cases. Many Starburst enterprise customers enhance their large ETL jobs by eliminating the need to write data to a temporary area before moving it to the final destination, such as an Azure data lake. Instead, the ETL job writes a query to Starburst, which returns it as a consolidated set, then writes it directly to the target.

✓ **Data Lake Query Engine**

Presto has become the de facto standard data lake query engine. A Starburst deployment enhances its functionality for the enterprise with role-based access control, autoscaling, high concurrency, ANSI SQL 2018 compatibility, and other benefits.



BUSINESS OUTCOMES

Overall, enterprises and large organizations that partner with Starburst to optimize their Azure investment benefit from:

→ Shorter Time to Value

→ Increased User Adoption

→ Improved Price vs. Performance

→ Enhanced Security



Companies using the Starburst Enterprise Platform on Azure discover and extract value from data faster — and turn these insights into actionable business initiatives that drive significant revenue gains.

Unlock the value of your data with Starburst and Azure today.

For more information, contact us at starburstdata.com