

# Starburst Insights

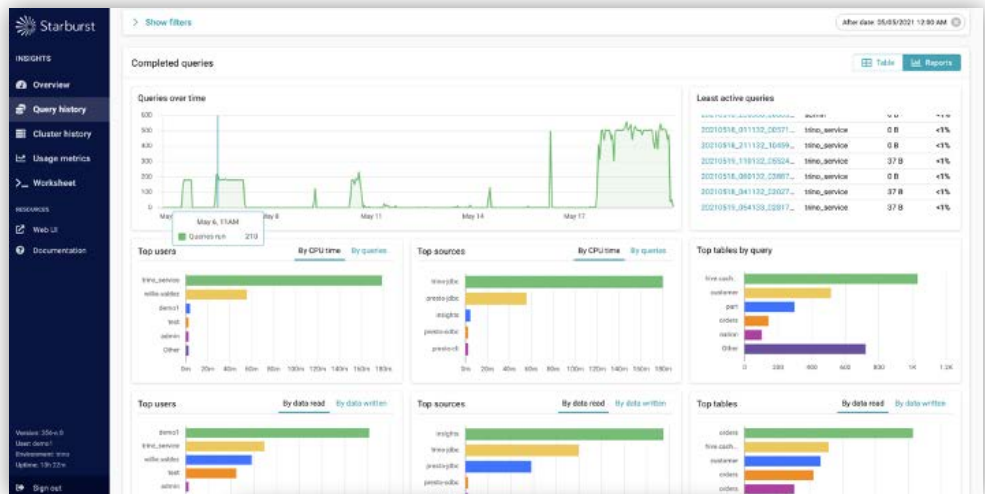
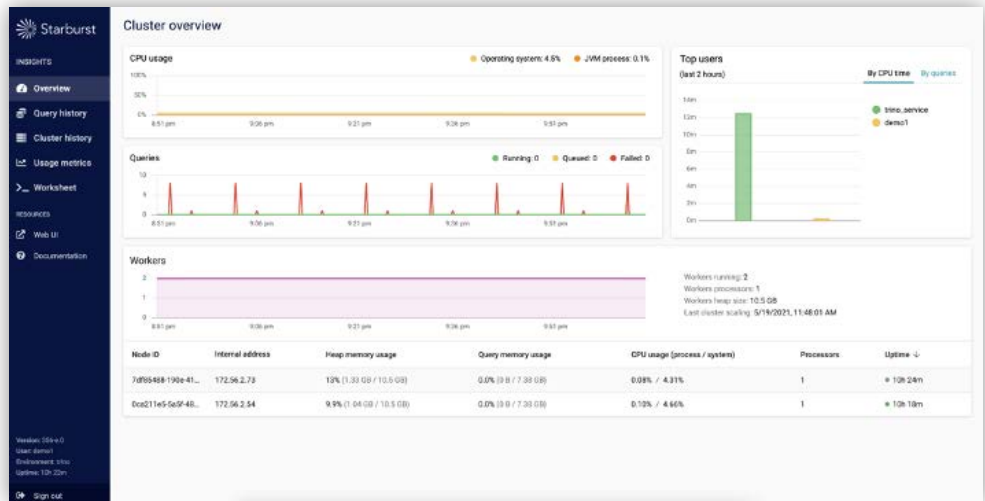
A holistic view of your cluster and query environment

Seeing the entire picture is essential for organizations to make accurate and timely decisions to drive business outcomes. With numerous disparate data sources, and many concurrent users running different workloads, gaining the visibility into overall performance and usage is an enormous benefit to our customers. Starburst Insights is a web application for cluster and query reporting on Starburst's environment.

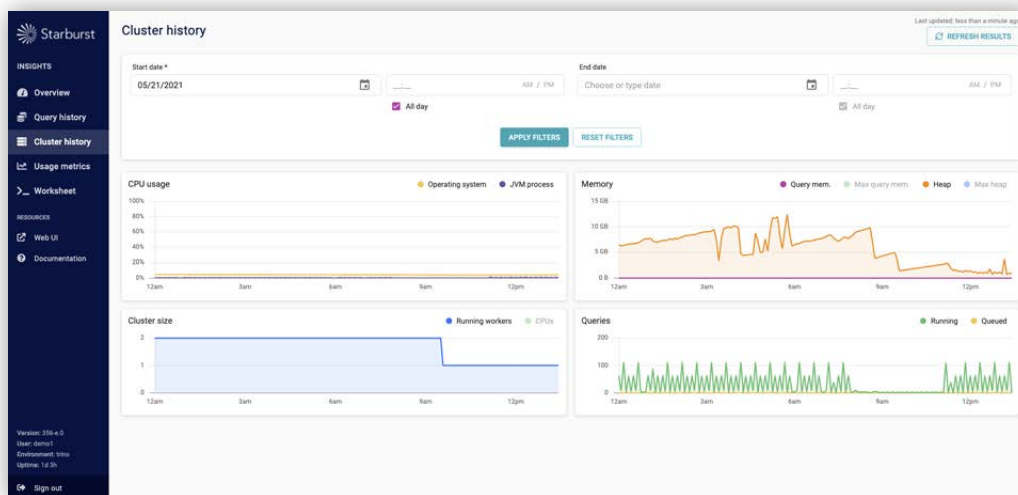
Starburst Insights provides a visual overview of important metrics about your Starburst Enterprise cluster for all types of users, from platform administrators to data consumers. From the Insights interface, you can access detailed query history, including single-query statistics and query plans, and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

**Additionally, users can drill down into query history, including:**

- Most accessed tables, and most active users by CPU time and queries
- Most and least active queries
- Completed query history
- Query details, including query statistics and execution plan



Data Teams need insights into their clusters operation and query execution so that they can identify any opportunity for cluster and query optimization. This way an analyst can optimize their queries to reach a meaningful business decision faster, and administrators can drill down into the usage pattern to perform capacity planning by analyzing historical resource usage, performance, and costs.



Starburst helps you increase your data efficiencies by providing a single source of access to all of your data sources no matter where they live. Starburst Insights provides a holistic view for managing your environment, including filterable analysis of cluster and query reporting. This includes the ability to search through the history of queries to understand the volume of queries by a user, or group of users. The analysis includes CPU, Memory, and query costs over multiple time dimensions. Customers can view usage metrics, monitor cluster health, explain queries, and proactively take actions to improve performance and reduce costs.

### Worksheet: A visual workbench to run ad hoc queries and explore configured data sources

With Worksheet querying a Trino cluster has never been easier. Instead of identifying a tool such as DBeaver or Superset, and a corresponding ODBC or JDBC driver, then downloading, installing, and configuring a connection, customers can simply utilize the Starburst Insights Worksheet to run queries and explore data. Tools like DBeaver, while flexible, are complex to set up, and require local desktop installation and sometimes organizational approval.

The Worksheet interface allows for running ad-hoc queries. It features a 'RUN' button, a 'Select catalog' dropdown, and a SQL editor. The query executed is:
 

```

    1 SELECT
    2   a.name,
    3   sum(l.extendedprice * (1 - l.discount)) AS revenue
    4 FROM
    5   "delta"."sf1"."customer" AS c,
    6   "delta"."sf1"."orders" AS o,
    7   "delta"."sf1"."lineitem" AS l,
    8   "delta"."sf1"."supplier" AS s,
    9   "delta"."sf1"."nation" AS n,
    10  "delta"."sf1"."region" AS r
    11 WHERE
    12  c.custkey = o.custkey
    13  AND l.orderkey = o.orderkey
    14  AND l.supplykey = s.supplykey
    15  AND c.nationkey = s.nationkey
    16  AND s.nationkey = n.nationkey
    17  AND n.regionkey = r.regionkey
    18  AND r.name = 'ASIA'
    19  AND o.orderdate <= DATE '1994-01-01'
    20  AND o.orderdate >= DATE '1994-01-01' + INTERVAL '1' YEAR
    21 GROUP BY
    
```

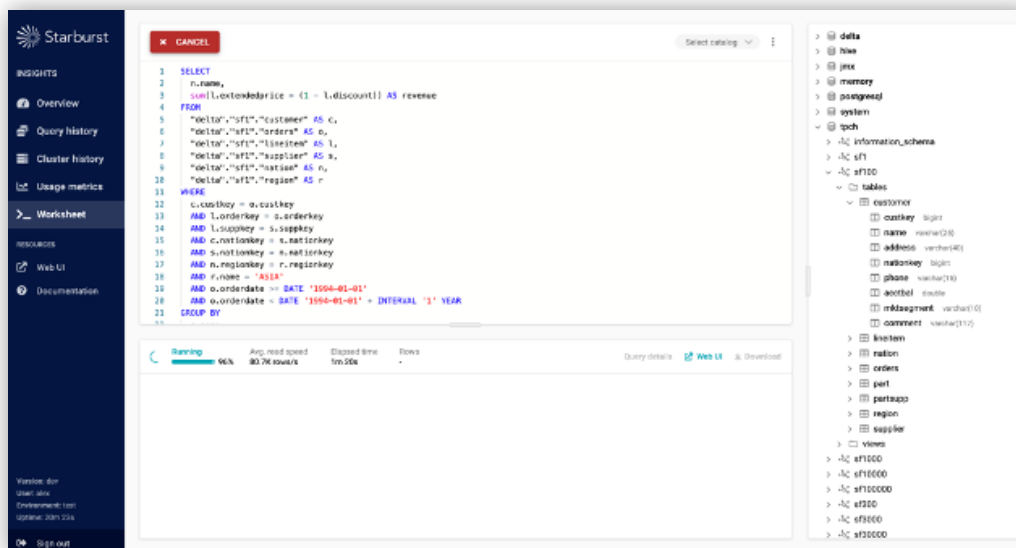
 The execution status is 'Finished' with an average read speed of 80.4K rows/s, an elapsed time of 1m 22s, and 5 rows returned. The results table is as follows:
 

name	revenue
INDONESIA	53358544.1697
Vietnam	53162101.99670004
CHINA	53581911.2566
INDIA	51889019.0002
JAPAN	45287756.69540001

 The right sidebar shows a tree view of the data catalog, including tables like customer, nation, orders, and supplier, and views like sf1000 through sf30000. A sidebar on the left provides navigation options similar to the Cluster history dashboard.



Worksheet is a SQL Integrated Development Environment (IDE) with an editor pane, catalog browser pane, and results pane. The editor pane allows users to run queries with menu options that include a copy to clipboard option, and a 'prettify' option to format queries as needed. The catalog browser pane displays an expandable, hierarchical view of catalogs, schemas, tables and views configured in Starburst Enterprise. When the query is finished executing, the result set appears in the results pane.



Worksheet does not require any installation and has no network latency. Data administrators and consumers are enabled with details in the IDE about queries, schemas, and query history. Furthermore, users can triage and test any queries with the built-in editor and obtain more detailed information such as the query execution plan and its explain plan. While Worksheet cannot replace the benefits and robust capabilities of many business intelligence tools, it empowers data teams and dramatically helps organizations increase the adoption of Starburst Enterprise.

## The Starburst and Trino difference

Starburst Enterprise leverages Trino, formerly PrestoSQL, an open-source, distributed SQL query engine, to make better decisions with lightning-fast access to all data, no matter where it lives. Trino eliminates the need to maintain a traditional data warehouse and separates storage from compute, allowing organizations to leverage low-cost storage without sacrificing insights. In addition to tremendous cost savings, Trino improves productivity and time to insight on actionable data, enabling decision makers to impact the business with faster, smarter, data-driven outcomes.

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



**Performance:** Includes the latest optimizations; Starburst Cached Views available for frequently accessed data; stable code that minimizes failed queries



**Connectivity:** 40+ supported enterprise connectors; high-performance connectors for Synapse, Oracle, Teradata, Snowflake, IBM DB2, Google BigQuery, 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, Starburst Insights monitoring dashboards; high availability; and easy deployment on anywhere



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

Unlock the value of your data with Starburst today.  
For more information, contact us at [starburst.io](https://starburst.io)

