



# Starburst Galaxy

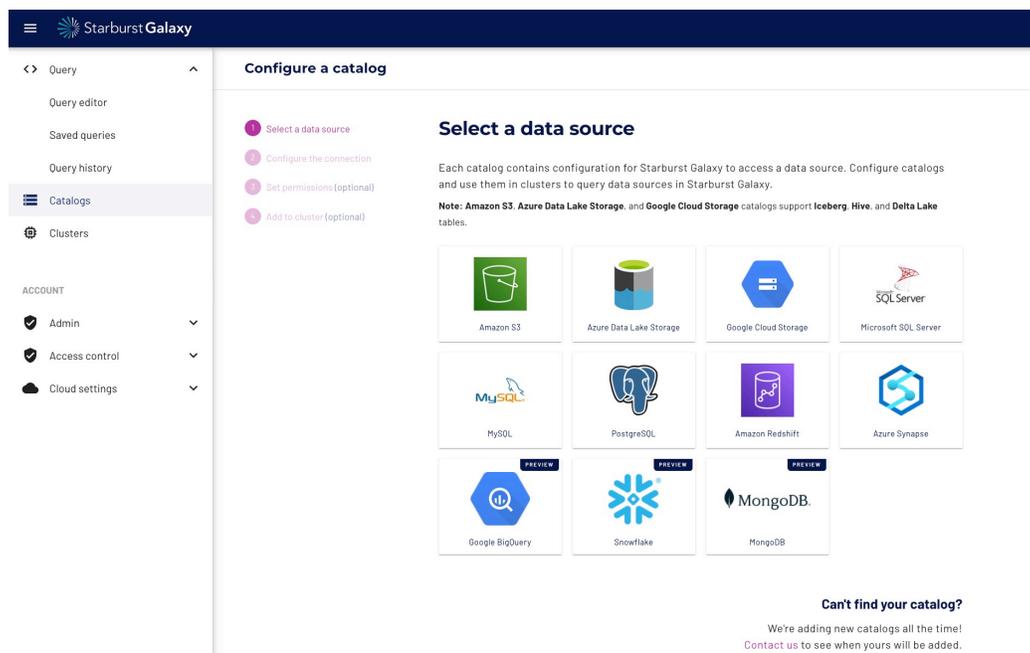
## Future proof your analytics

### The freedom to be curious

Traditionally organizations with new analytic needs move their data from its underlying sources in databases and data lakes into a data warehouse. The time commitment and cost associated with migrating data can be immense.

Starburst Galaxy is a fully managed service designed for running fast, interactive analytic queries against distributed data sources

like data lakes and data warehouses ranging in size from terabytes to petabytes with 24x7 support from the experts in large scale SQL engines and analytics. It enables immediate analysis of siloed data without expensive data warehouse appliances.



**Configure a catalog**

- Select a data source
- Configure the connection
- Set permissions (optional)
- Add to cluster (optional)

### Select a data source

Each catalog contains configuration for Starburst Galaxy to access a data source. Configure catalogs and use them in clusters to query data sources in Starburst Galaxy.

**Note:** Amazon S3, Azure Data Lake Storage, and Google Cloud Storage catalogs support Iceberg, Hive, and Delta Lake tables.

- Amazon S3
- Azure Data Lake Storage
- Google Cloud Storage
- Microsoft SQL Server
- MySQL
- PostgreSQL
- Amazon Redshift
- Azure Synapse
- Google BigQuery
- Snowflake
- MongoDB

**Can't find your catalog?**  
We're adding new catalogs all the time!  
Contact us to see when yours will be added.

### The fully-managed experience

Spend less time managing your data and more time analyzing it: Starburst Galaxy is a fully-managed, enterprise-ready SaaS platform providing you with the easiest way to access the power of Starburst's best-in-class MPP SQL engine. Connect your data sources, start a cluster and start querying using the analytics tools you already know in minutes.

Instantly scale as your analytics needs grow without compromising performance. Add a cluster for whoever in your organization needs it. Choose a t-shirt sized cluster to match your desired performance and cost. No more queuing behind slow queries or noisy neighbors. And with autoscaling you can easily scale up and down with demand. When you're done, Starburst Galaxy's idle shutdown feature automatically shuts down clusters after a set period.

## Accelerate your time to insight

In addition to the handling all the complex processes, Starburst Galaxy gets your organization up and running in minutes, not months:



**Start fast:** From login to production-ready in minutes with a guided user experience



**Flexible:** Pay-as-you-go pricing and the ability to manage for cost and performance



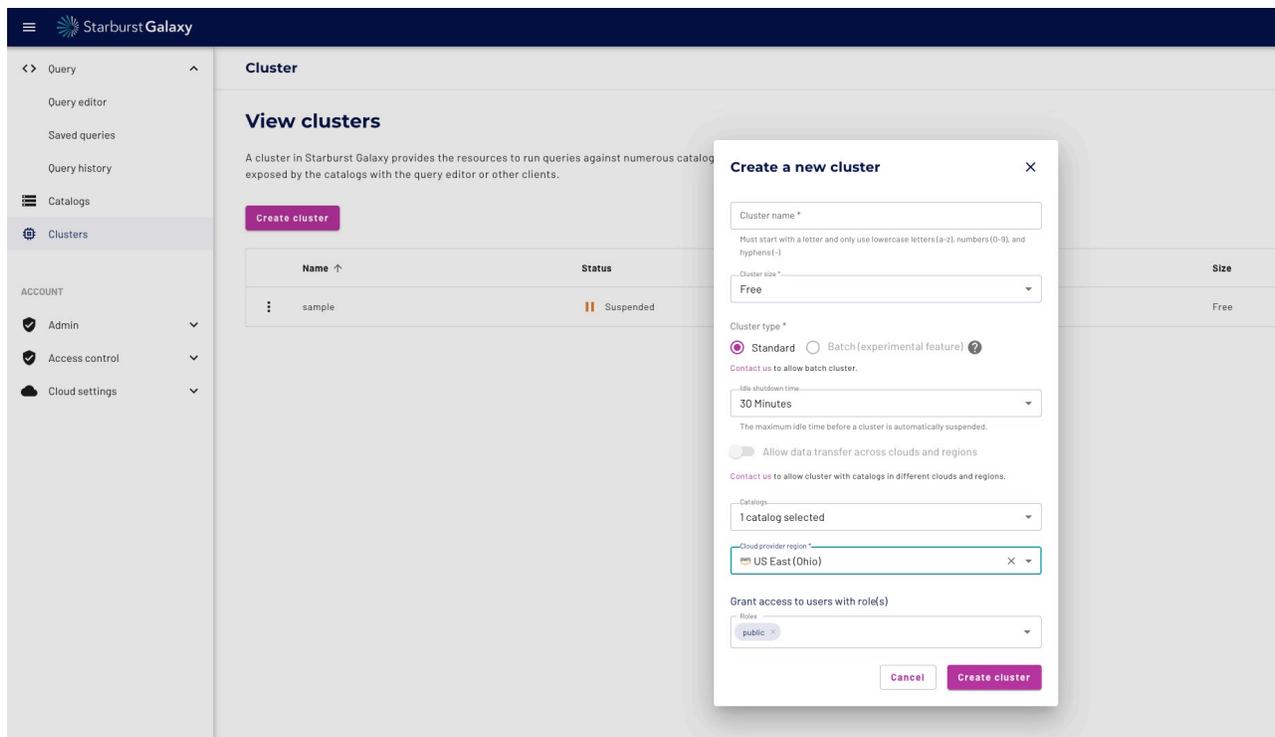
**Intuitive:** Designed for anyone in your organization to start querying data instantly

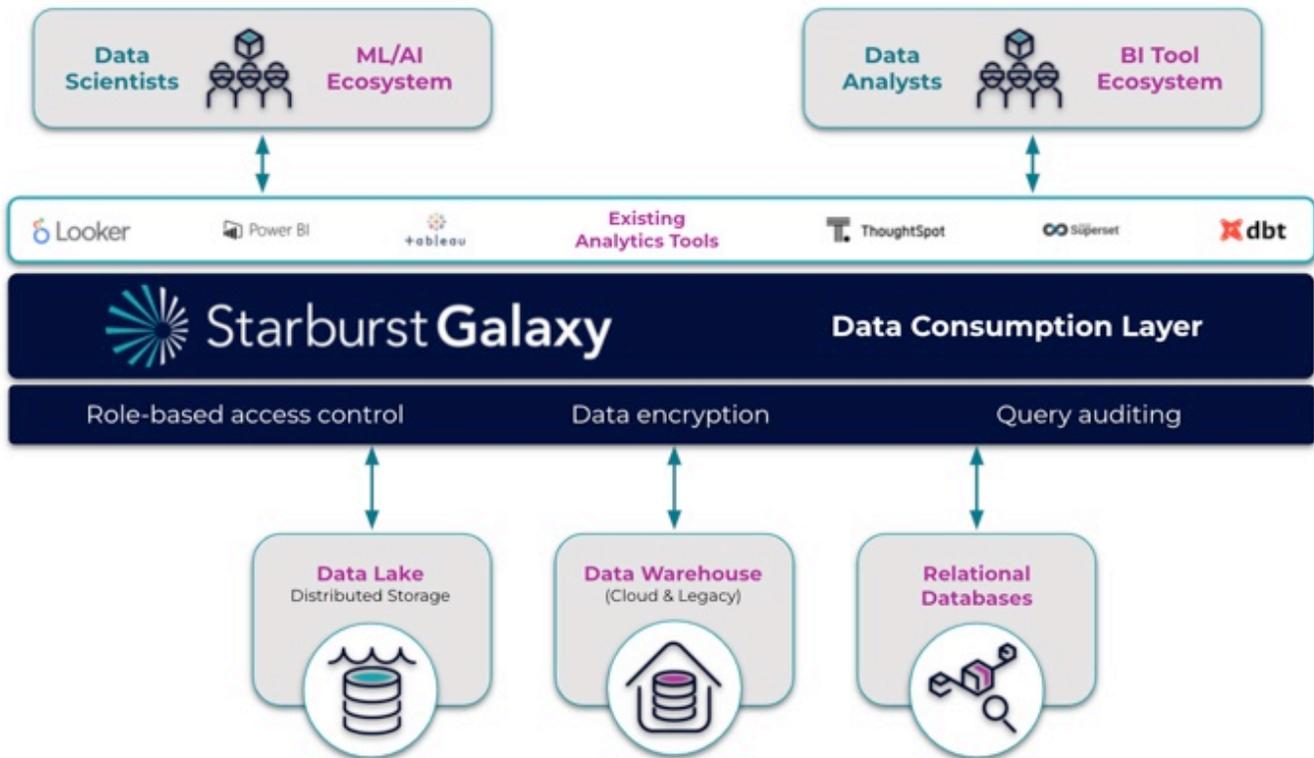


**Support:** 24x7 support from the Trino experts; in a fully-managed environment



**Self-Service:** Seamless integration with your existing analytics tools and workflows





Starburst Galaxy dramatically simplifies deployment and application development handling complex processes like configuration, upgrades and scaling seamlessly. Our expertise and 24x7 support ensures Starburst Galaxy is operating optimally allowing you to focus on delivering actionable insights to drive better business outcomes.

### The analytics engine for the data lake

Starburst Galaxy offers the flexibility to run interactive and ELT workloads in one engine on the data lake.

- Customers can easily transform, join and enrich data using standard SQL, without being strained like traditional engines.
- Future proof your architecture by eliminating vendor lock-in with a query engine designed for the most popular data lake open table formats like Iceberg, Delta Lake and Hive.
- Available on all major cloud platforms enabling you to instantly query data in Amazon S3, Azure Data Lake Storage and Google Cloud Storage

The screenshot shows the Starburst Galaxy interface. On the left is a navigation menu with options like 'Query', 'Query editor', 'Saved queries', 'Query history', 'Catalogs', and 'Clusters'. The main area displays a SQL query being executed on a cluster named 'Cluster explo...'. The query is as follows:

```

1
2
3 select
4   province_state,
5   sum(confirmed) as confirmed_cases,
6   sum(active) as active_cases
7 from
8   ejhu_stg_australia
9 where
10  last_update = '2020-05-30T02:32:48'
11 group by
   province_state;
    
```

The query has finished successfully. The results table shows the following data:

province_state	confirmed_cases	active_cases
Australian Capital Territory	107	0
Northern Territory	29	0
Queensland	1058	6
South Australia	440	1

## The Biggest Brands in the World Rely on Starburst and Trino

Comcast runs over **350,000**  
Trino queries daily

**Netflix, Lyft, Salesforce,  
LinkedIn, and Amazon**  
all depend on Trino

More and more of the largest,  
best-known enterprises are already  
using Starburst, including **Slack,  
Comcast, Zalando, and FINRA**

Common use cases include interactive  
data investigation, BI dashboard and  
reports, data science, ETL, and as a high-  
performance data lake query engine

## Starburst Galaxy

- ✓ **Kickstart your data lake strategy**
- ✓ **Join across data sources**
- ✓ **Directly query distributed object storage**
- ✓ **Native enterprise-ready security features**
- ✓ **Fully managed environment with automated maintenance**
- ✓ **24x7 support from the experts**
- ✓ **Easy access from the most popular BI tools**
- ✓ **Reduced big data costs by optimizing cloud spend**