




How to use Databricks SQL for Analytics on Your Lakehouse

Pearl Ubaru

Sr. Solutions Engineer



Meet the Presenter



Pearl Ubaru
Sr. Solutions Engineer
pearl.ubaru@databricks.com



Agenda

- Introduction & Key Capabilities of Databricks SQL (15 minutes)
- Hands-on Demo (50 minutes) + Break (10 minutes)
- Q&As (5 minutes)



Databricks

The Data + AI Company



Inventor and pioneer of the **data lakehouse**

Gartner recognized leader in both

- Database Management Systems
- Data Science and Machine Learning Platforms

Creator of highly successful OSS data projects: Delta Lake, Apache Spark, and MLflow

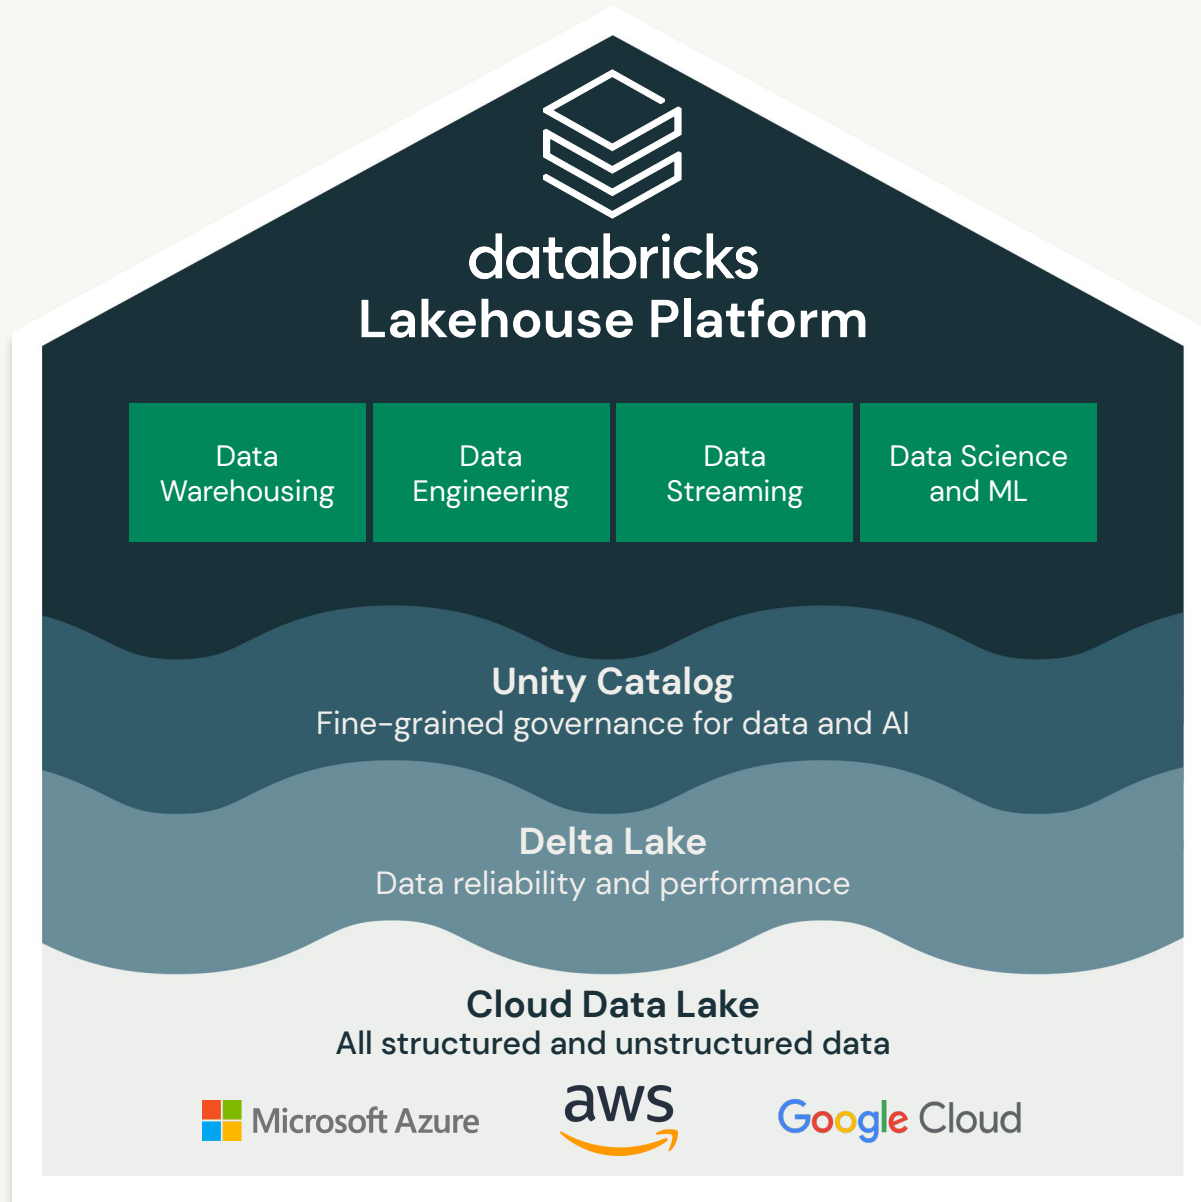
Raised over \$3B in investment

3000+ employees across the globe

Global adoption

Over 6000 customers, from F500 to unicorns





Databricks Lakehouse Platform

Simple

Unify your data warehousing and AI use cases on a single platform



Open

Built on open source and open standards

Multi-cloud

One consistent data platform across clouds

Introducing Databricks SQL

Data Warehousing on the Lakehouse




Serverless Compute

Lower costs, get the best price/performance, and eliminate the need to manage, configure or scale cloud infrastructure with serverless.



Bring your Own Tools

DB SQL works with your favorite tools. Easily connect Fivetran, dbt, PowerBI, or Tableau to ingest, transform and visualize all your latest data in-place.



Native SQL Experience

Enable every analyst in your organization to quickly discover and collaboratively query the latest data, mock-up interactive dashboards, and find new insights.



Built-in Governance

Confidently discover, access, and govern all your data using fine-grained governance and standard SQL, with built-in security and compliance.



Key Capabilities




Eliminate resource management with serverless

Instant, elastic SQL compute decoupled from storage

Quickly setup instant, elastic SQL **serverless** compute – decoupled from storage – to free up time and resources. Databricks automatically determines instance types and configuration for the **best price/performance** through Photon.

100TB TPC-DS Price/Performance
Lower is better



The screenshot displays the Databricks SQL Endpoints management interface. A modal window for creating a 'New Serverless SQL Endpoint' is open, showing the following configuration:

- Name: Shared Endpoint
- Cluster Size: Medium (24 DBU / cluster)
- Auto Stop: Enabled (After 10 minutes of inactivity)
- Multi-cluster Load Balancing: Enabled (Cluster Count: Min 1, Max 3)
- Photon: Enabled

The background shows a list of existing endpoints and a monitoring dashboard for a 'Shared Endpoint'. The monitoring dashboard includes a 'Queries' bar chart and a 'Cluster Scale' step chart for August 24th.

Queries Monitoring (Aug 24):

- Peak Running: 60
- Peak Scheduled: 30
- Starting up: 10:00 AM - 10:05 AM

Cluster Scale Monitoring (Aug 24):

- Cluster Scale 6 - 8 (10:55 AM - 11:00 AM)

Built from the ground up for best performance

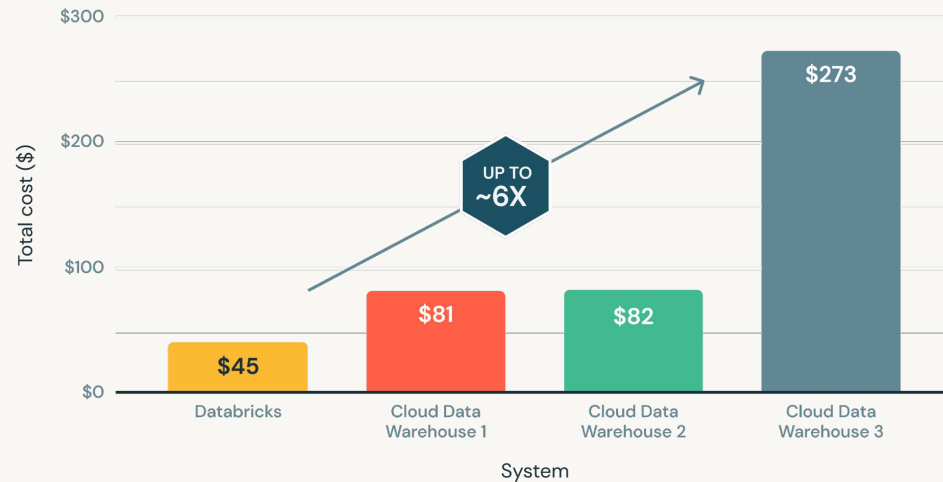
Lightning fast analytics

Query and analyze your most complete and freshest data with **up to 12X better price/performance** than cloud data warehouses.



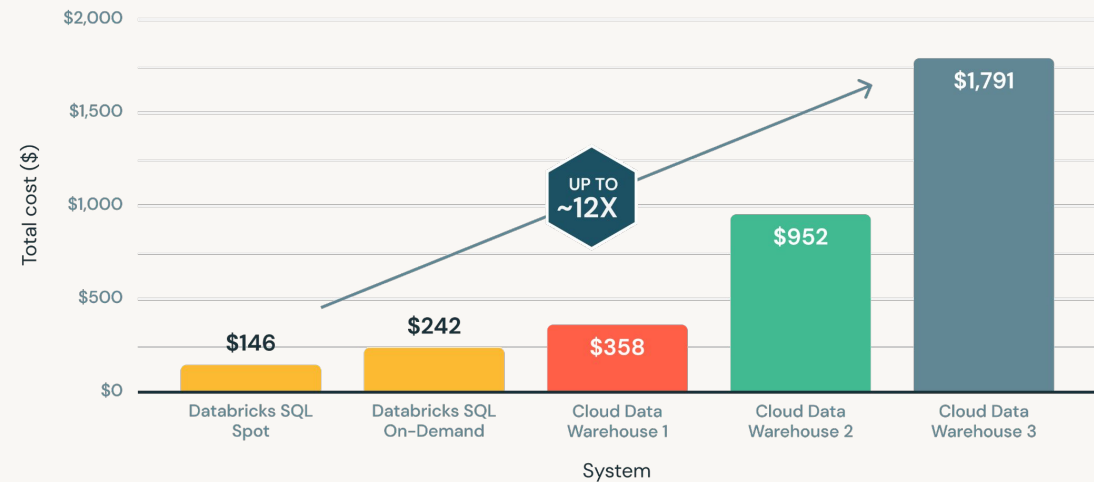
30TB TPC-DS Price/Performance

Lower is better



100TB TPC-DS Price/Performance

Lower is better



Source: Performance Benchmark with Barcelona Supercomputing Center

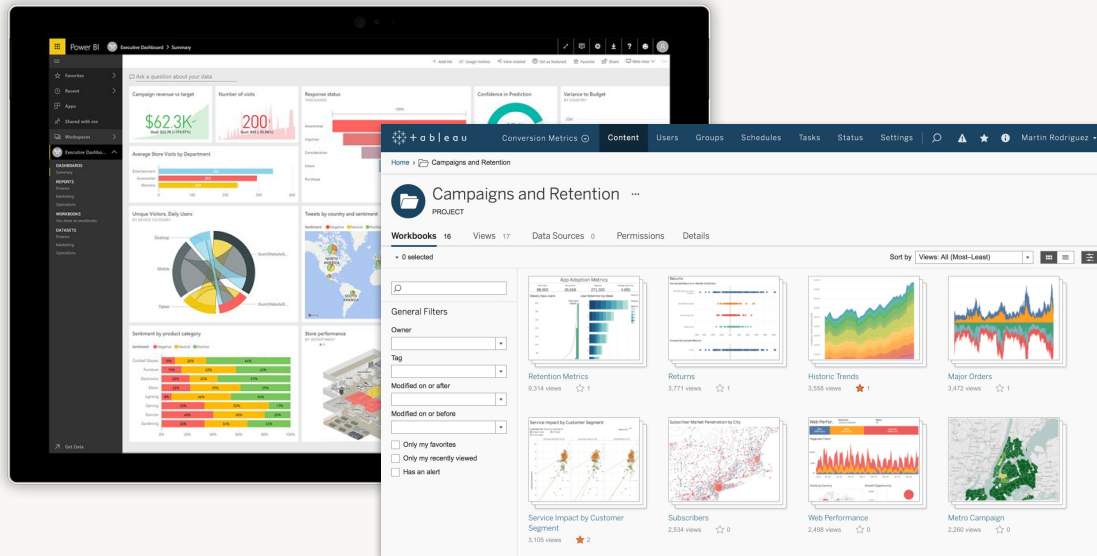


Modern analytics with your tools of choice

Seamless integration with the modern data stack



Easily **ingest** and **transform** data from cloud storage or business applications, and leverage **your existing BI tools** to query, analyze, and visualize new insights.



Ingest & ETL



BI & Analytics




ThoughtSpot



First-class SQL development experience

Collaboratively query, explore, and transform data in-place

Query all your data using familiar **ANSI SQL**, and collaboratively find and share new insights faster with the built-in SQL query editor, alerts, visualizations, and interactive dashboards.



Demo: Databricks SQL



Databricks SQL Hands-on Demo



- Section 1: The Databricks SQL workspace
- Section 2: Import sample dashboard
- Section 3: SQL Warehouses & Understand Computation Resources
- Section 4: Explore database and tables, and data access controls
- Section 5: Create & execute queries and visualizations
- Section 6: Monitor a SQL Endpoint, query history, query performance
- Section 7: Configure alerts
- Section 8: Creating Dashboards



Q&A



Next steps

- Get started on Databricks SQL:
<https://databricks.com/product/databricks-sql>
- Keep learning with self-paced training (free for customers)
<https://academy.databricks.com/>
- Connect with me on LinkedIn: Pearl Ubaru



