DATA+AI SUMMIT BY databricks

RISE OF THE MEDALLION MESH

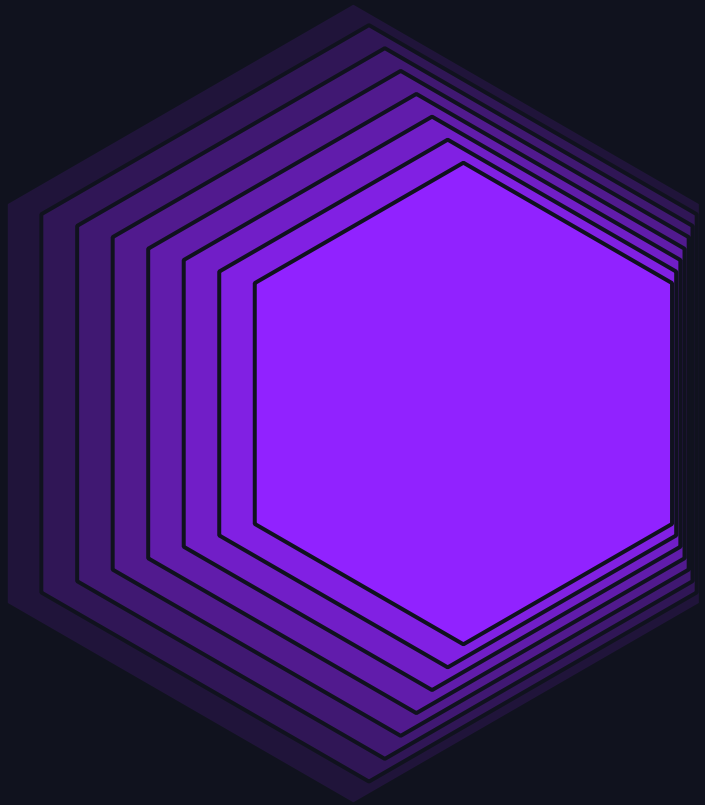

Franco Patano, Databricks, Strategic Data & AI Advisor

Top concerns in data and AI modernization*

For leaders and practitioners across industries

- Data security, governance, and quality
- Performance, scalability, and total cost of ownership
- Skills and expertise gap
- Data integration and migration
- Provide diverse tooling to support variety of users on a single version of the truth for all use cases

whoami

Franco Patano
Strategic Data and AI Advisor

- Background in ETL, Data Warehousing, BI, and Analytics
  - MSBI Stack (SQL Server, Reporting Services, Integrations Services, Analysis Services)
  - Informatica
  - Tableau
  - Cognos
  - Excel and VBA
- Grew up in Enterprises in Chicago Area
  - Career Education Corporation (Perdoceo Education Corporation)
  - Wintrust
  - JLL (Jones Lang Lasalle)
- Dreamed of joining a startup and changing the world
  - Joined Databricks in 2019 to change the world with Lakehouse!

- Likes
  - Time Travel fiction (and non-fiction)
    - Back to the future
    - Doctor Who
    - Bill and Ted
    - Predestination
  - Making Pizza from Scratch
  - EDM
    - Daft Punk, Alesso, Cazzette, DJ Caffeine
  - Rocky, The Princess Bride, and The Fifth Element are the most EPIC films of all time
- Dislikes
  - Data Warehouses
    - Ever since I first became an analyst, all I wanted to do was get rid of that pesky data warehouse

10,000+
global customers

$1.5B+
in revenue

$4B
in investment

Inventor of the **lakehouse**
&
Pioneer of **generative AI**

databricks
The data and AI company

Gartner-recognized Leader
Database Management Systems
Data Science and Machine Learning Platforms

Creator of

DELTA LAKE

mlflow™

APACHE Spark™

DBRX

Recognized as a leader in the industry

Gartner MQ Leader

Database Management Systems

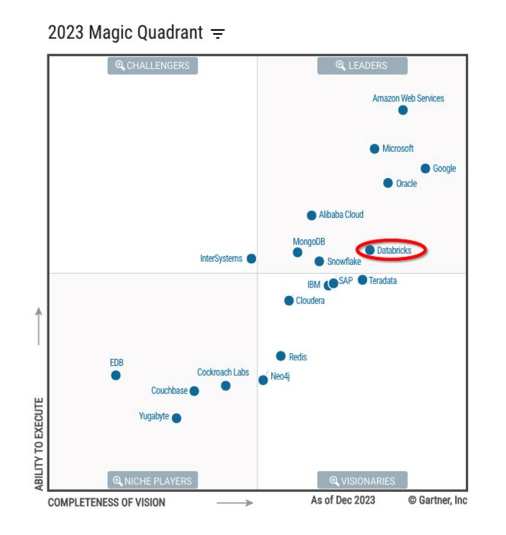

2023 Magic Quadrant

As of Dec 2023

© Gartner, Inc

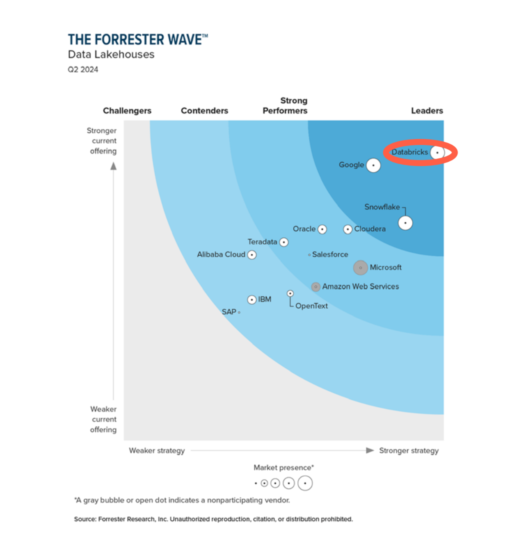

THE FORRESTER WAVE™

Data Lakehouses

Q2 2024

*A gray bubble or open dot indicates a nonparticipating vendor.

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

Forrester Wave Leader

Data Lakehouses

Databricks SQL

Intelligent data warehousing on the Data Intelligence Platform

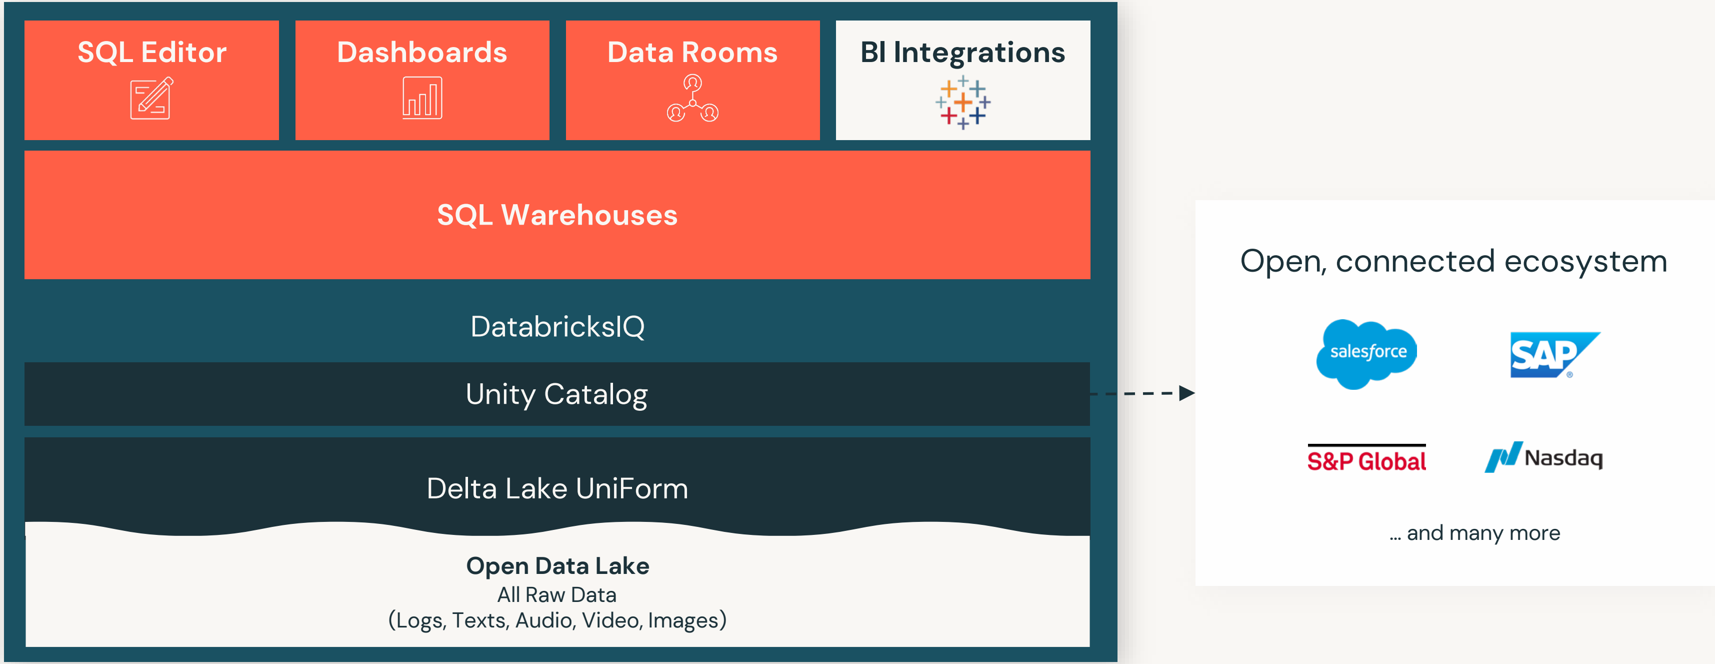

SQL Editor

Dashboards

Data Rooms

BI Integrations

SQL Warehouses

DatabricksIQ

Unity Catalog

Delta Lake UniForm

Open Data Lake
All Raw Data
(Logs, Texts, Audio, Video, Images)

Open, connected ecosystem

salesforce

SAP

S&P Global

Nasdaq

... and many more

Trusted by organizations of all sizes

4,600+ Databricks SQL customers across industries

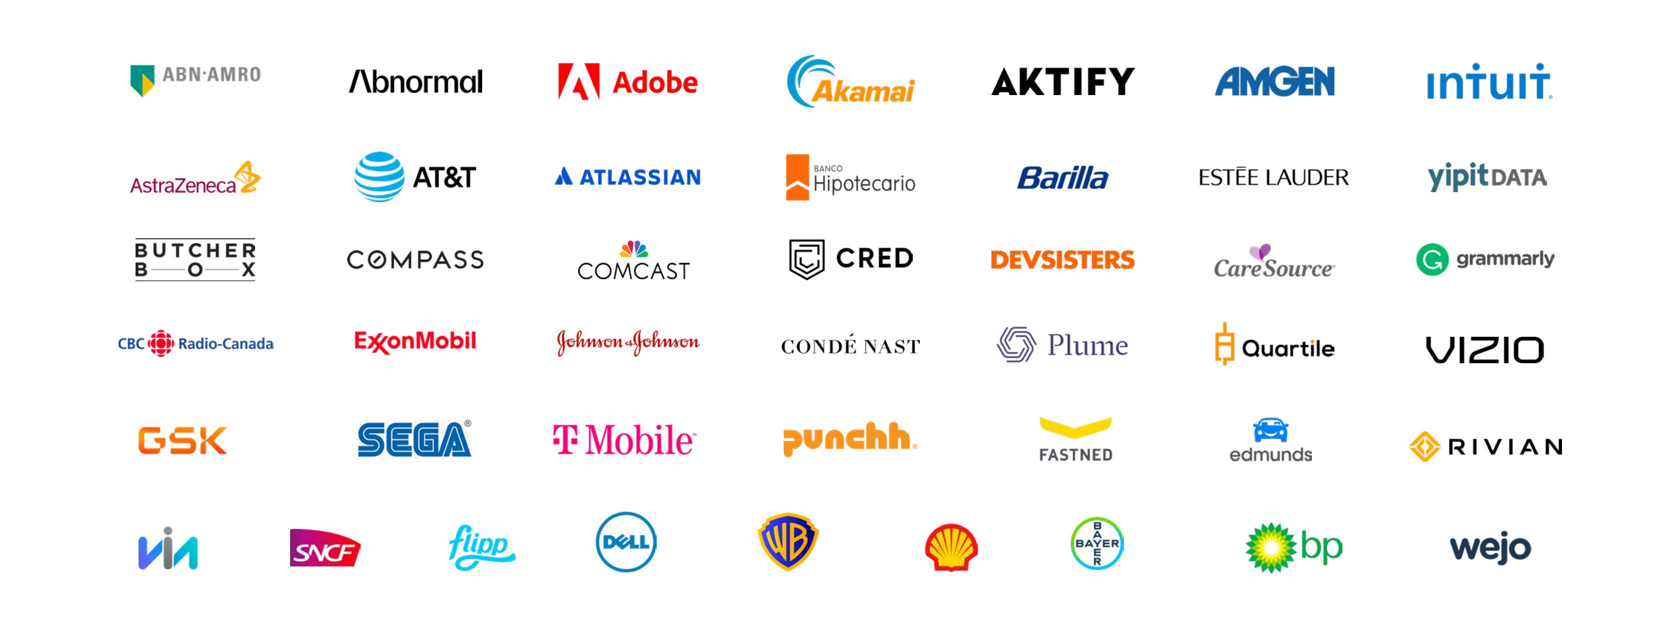

DATA SECURITY, GOVERNANCE, AND QUALITY

Governed and secured by Unity Catalog

Governance for all your data and AI assets

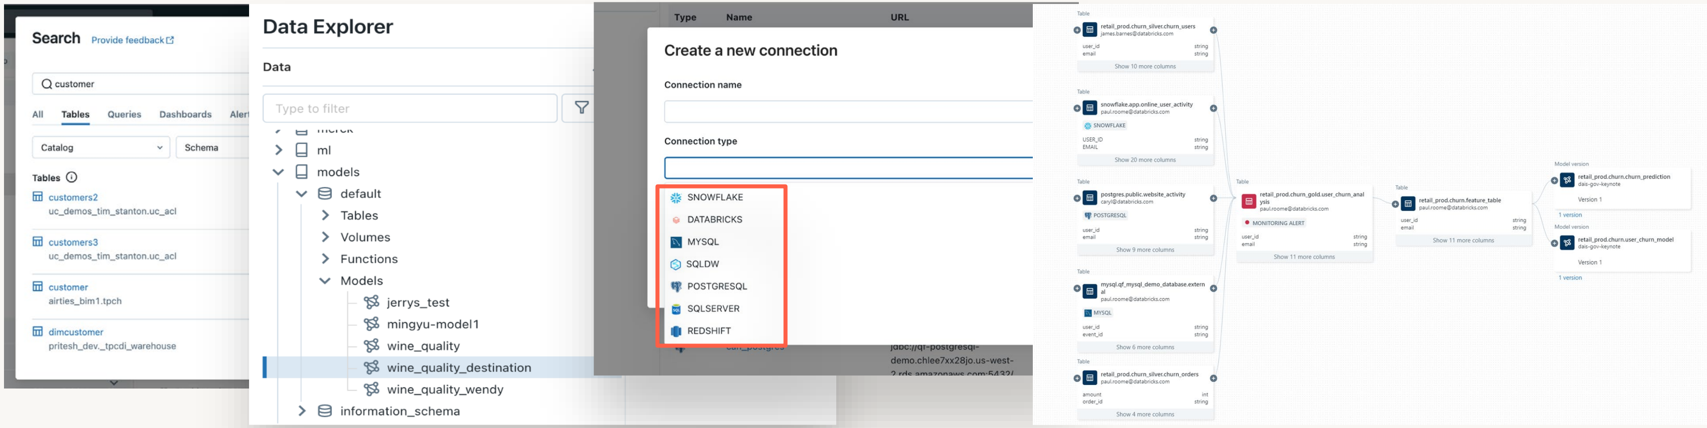

Simplified **data discovery**, **governance**, **federation**, **lineage**, and **compliance** with enhanced **security** and **auditing** with Unity Catalog and Databricks SQL

Simple security with private connectivity

Private Link or firewall support for stable VNET IPs

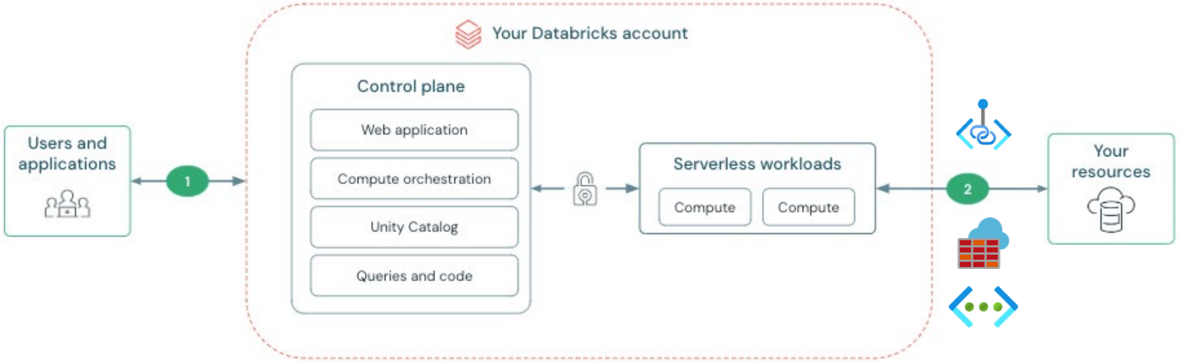

- Reduce data exfiltration risk
- Improved performance by minimizing network latency
- Meet compliance requirements for sensitive workloads
- PrivateLink included with Serverless SQL Warehouse cost on Azure

Simple AI-powered monitoring and observability

- Receive **proactive alerts** for quality issues with data and ML models
- Access **real-time data lineage** down to the column level for efficient root cause analysis and impact assessment
- Utilize **auto-generated** dashboards to easily share data and ML quality reports with stakeholders
- Achieve complete data and AI **observability** with operational intelligence for billing, auditing, lineage, and more

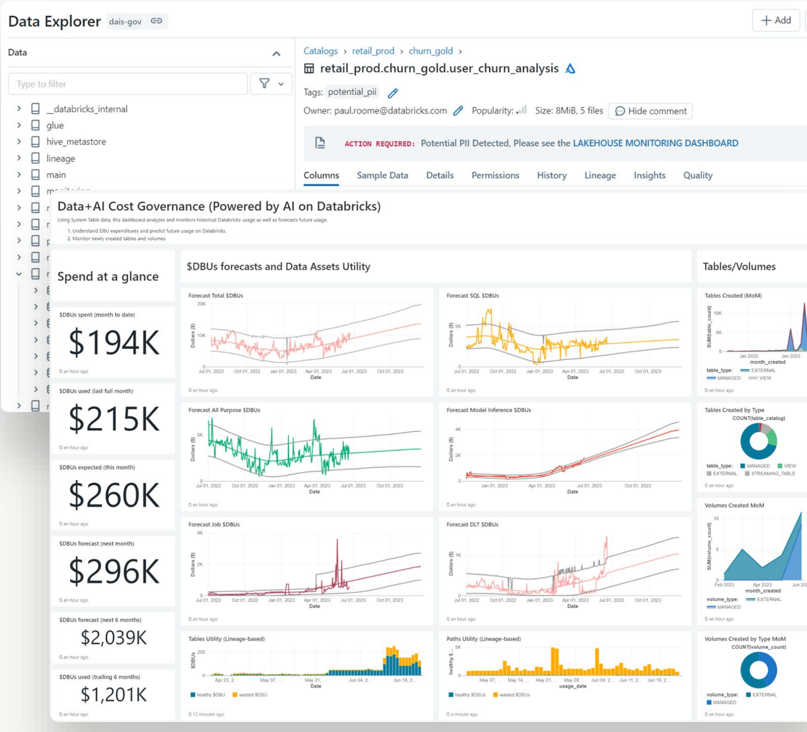

Data Explorer dais-gov

Catalogs › retail_prod › churn_gold ›

retail_prod.churn_gold.user_churn_analysis

Tags: potential_pii

Owner: paul.roome@databricks.com Popularity: Size: 8MiB, 5 files

ACTION REQUIRED: Potential PII Detected, Please see the LAKEHOUSE MONITORING DASHBOARD

Data+AI Cost Governance (Powered by AI on Databricks)

Using System Table data, this dashboard analyzes and monitors historical Databricks usage as well as forecasts future usage.

1. Understand DBU expenditures and predict future usage on Databricks.
2. Monitor newly created tables and volumes

Spend at a glance

$DBUs spent (month to date): $194K

$DBUs used (last full month): $215K

$DBUs expected (this month): $260K

$DBUs forecast (next month): $296K

$DBUs forecast (next 6 months): $2,039K

$DBUs used (trailing 6 months): $1,201K

Simple data monitoring

Lakehouse monitoring for reliable, insightful, and simple data-to-AI-BI pipelines

Simple: Log inference tables automatically, and generate metric tables and SQL dashboards.

Proactive: Automate alerts on table quality and custom metrics, and diagnose data or model issues.

Integrated: Track end-to-end lineage in the Unity Catalog for training data, feature tables, models, and inference logs, for simpler governance.

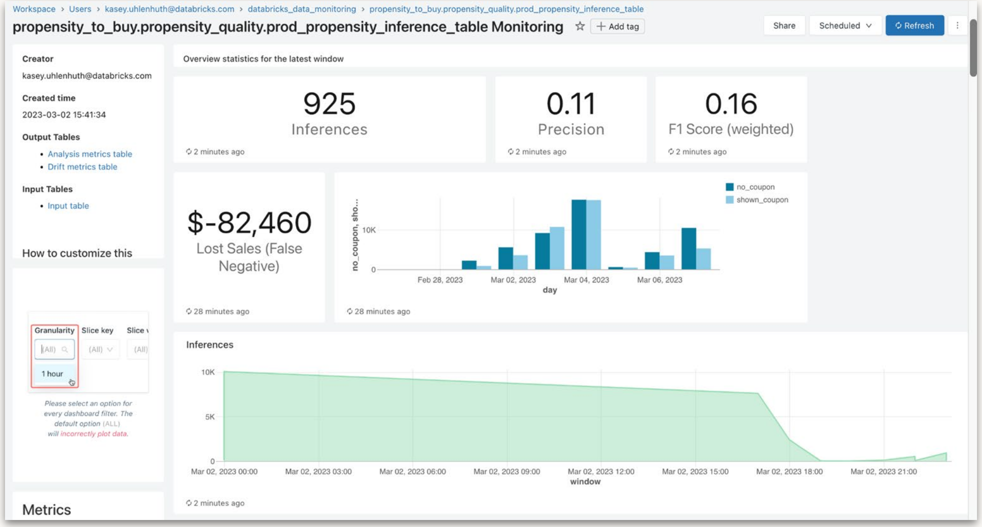

Workspace › Users › kasey.uhlenhuth@databricks.com › databricks_data_monitoring › propensity_to_buy.propensity_quality.prod_propensity_inference_table

propensity_to_buy.propensity_quality.prod_propensity_inference_table Monitoring ☆ + Add tag

Creator
kasey.uhlenhuth@databricks.com

Created time
2023-03-02 15:41:34

Output Tables
- Analysis metrics table
- Drift metrics table

Input Tables
- Input table

How to customize this

Granularity: 1 hour

Please select an option for every dashboard filter. The default option (ALL) will incorrectly plot data.

Metrics

Overview statistics for the latest window

925 Inferences — 2 minutes ago

0.11 Precision — 2 minutes ago

0.16 F1 Score (weighted) — 2 minutes ago

$-82,460 Lost Sales (False Negative) — 28 minutes ago

Inferences — 2 minutes ago

Simple and open data sharing

- **Avoid vendor lock-in** with open-source Delta Sharing for seamless data sharing across clouds, regions, and platforms, without replication
- Share **more than just data:** notebooks, ML models, dashboards, applications
- Explore and monetize data products through an **open marketplace**
- Collaborate securely on sensitive data with **scalable data clean rooms**

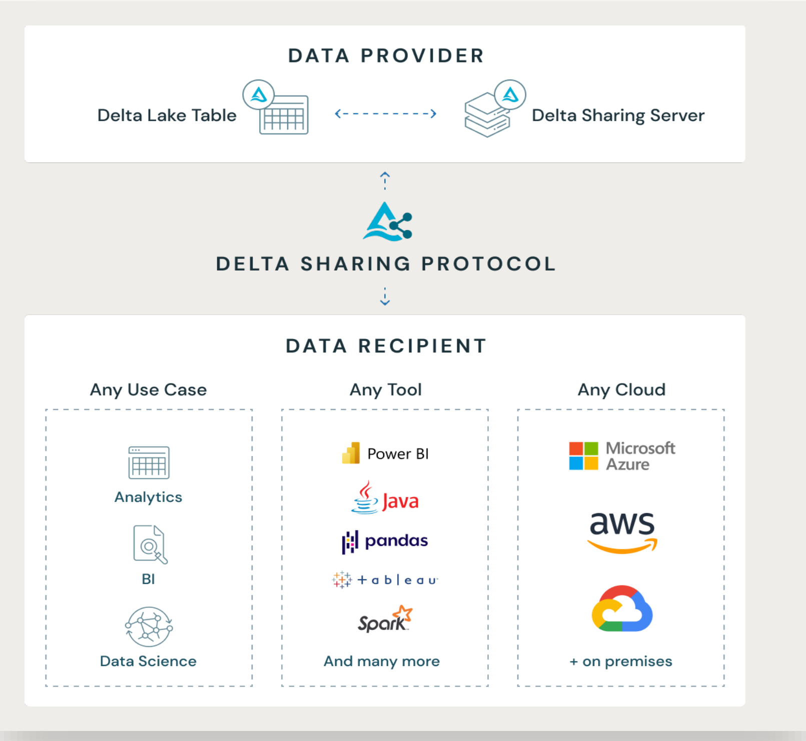

DATA PROVIDER

Delta Lake Table ↔ Delta Sharing Server

DELTA SHARING PROTOCOL

DATA RECIPIENT

Any Use Case: Analytics, BI, Data Science

Any Tool: Power BI, Java, pandas, tableau, Spark, And many more

Any Cloud: Microsoft Azure, aws, + on premises

PERFORMANCE, SCALABILITY, AND TOTAL COST OF OWNERSHIP

World-class performance & TCO

Meets or beats the price/performance of major CDWs across scales!

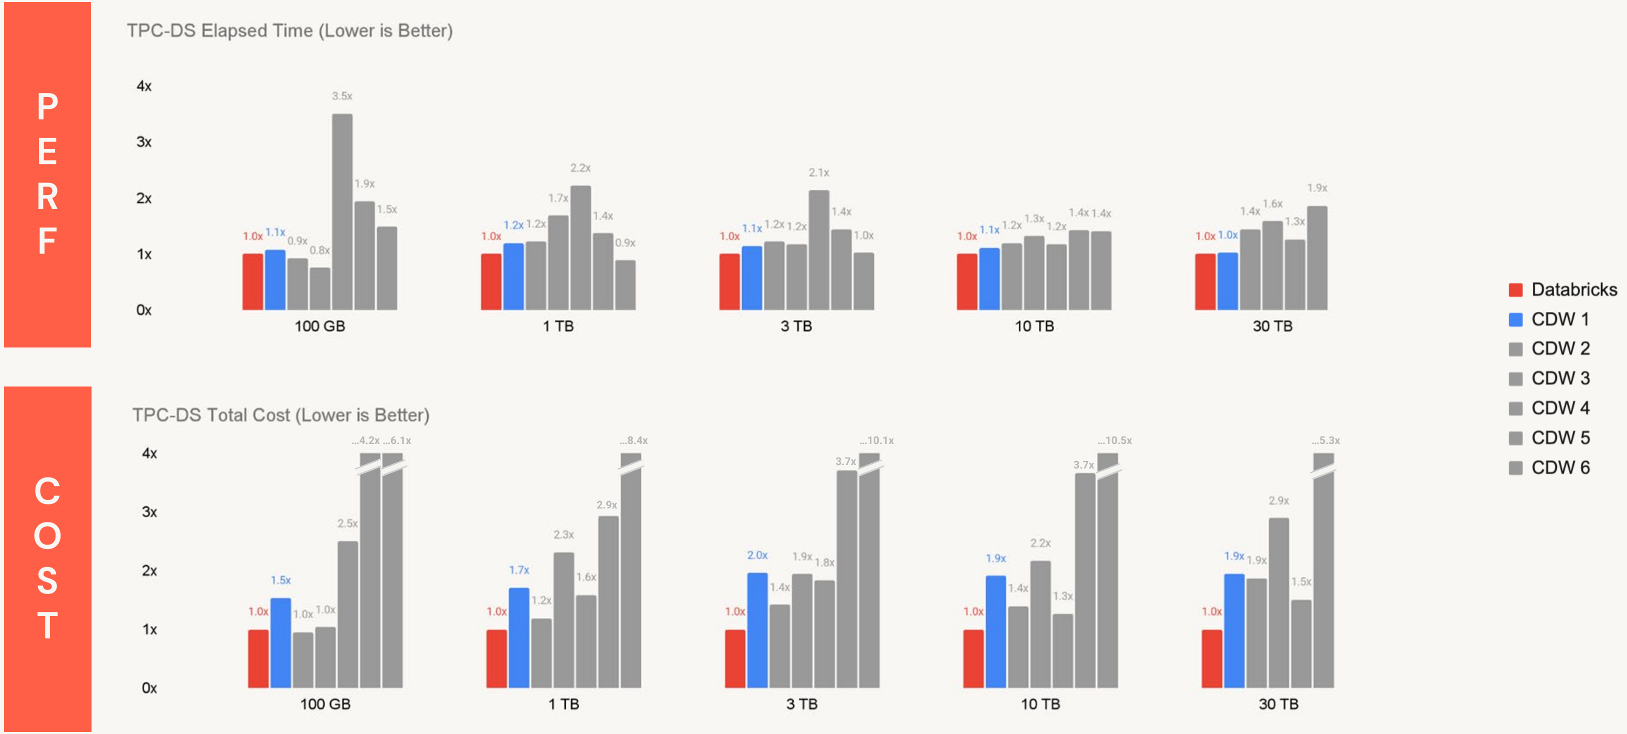

World-class performance for BI

High concurrency with low latency

Lower is better

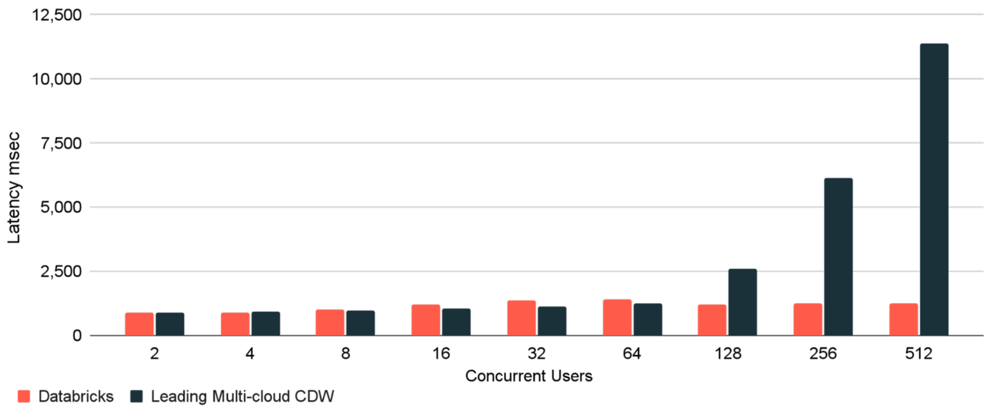

Adevinta

"Our analysts rely on Databricks SQL to derive **business intelligence . . . we have 30% better performance** and have reduced costs by 20% on average."

—Allard de Boer

Global Director of Analytics, Adevinta

Simple performance with Predictive Optimization

AI-optimized Delta table layouts for best price-performance

Runs OPTIMIZE, VACUUM, ANALYZE, Liquid clustering

AI model prioritizes tables to maximize ROI

Out-of-box observability with system tables

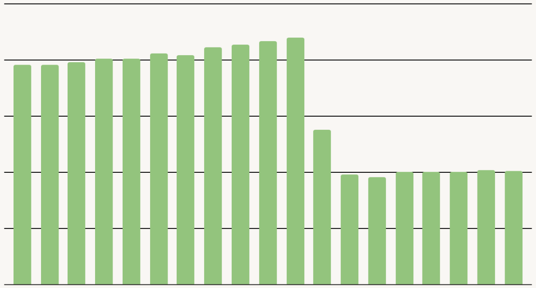

Customer storage costs
(actual data from Preview customer)

> “Databricks’ Predictive Optimizations intelligently optimized our Unity Catalog storage, which **saved us 50% in annual storage costs** while **speeding up our queries by >2x**. It learned to prioritize our largest and most-accessed tables. And, it did all of this **automatically,** saving our team valuable time.”
>
> **—Anker**

SKILLS AND EXPERTISE GAP

Intelligent experiences with natural language

Powered by DatabricksIQ

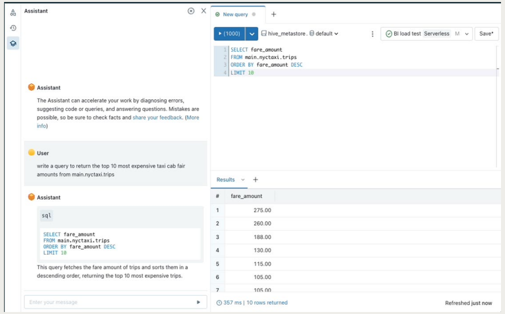

Assistant

The Assistant can accelerate your work by diagnosing errors, suggesting code or queries, and answering questions. Mistakes are possible, so be sure to check facts and share your feedback. (More info)

User

write a query to return the top 10 most expensive taxi cab fair amounts from main.nyctaxi.trips

Assistant

sql

```
SELECT fare_amount
FROM main.nyctaxi.trips
ORDER BY fare_amount DESC
LIMIT 10
```

This query fetches the fare amount of trips and sorts them in a descending order, returning the top 10 most expensive trips.

```
SELECT fare_amount
FROM main.nyctaxi.trips
ORDER BY fare_amount DESC
LIMIT 10
```

#	fare_amount
1	275.00
2	260.00
3	188.00
4	130.00
5	115.00
6	105.00
7	105.00

357 ms | 10 rows returned   Refreshed just now

SQL Editor

Data Science / Engineers

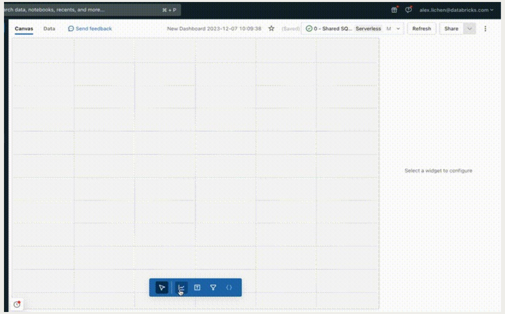
rch data, notebooks, recents, and more...
⌘ + P
alex.lichen@databricks.com
Canvas
Data
Send feedback
New Dashboard 2023-12-07 10:09:38
(Saved)
0 - Shared SQ... Serverless M
Refresh
Share
Select a widget to configure

Lakeview Dashboards

Analysts

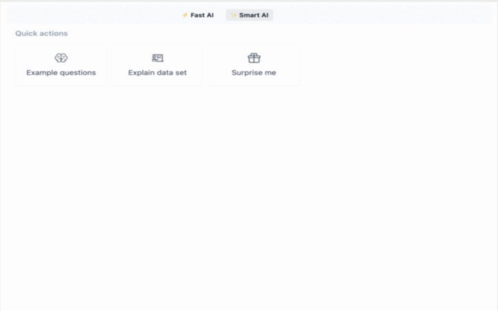

Project Genie

Coming Soon!

Business Users

Assistant in SQL Editor

Developers can create, explain and fix SQL code

Assistant in Query Editor

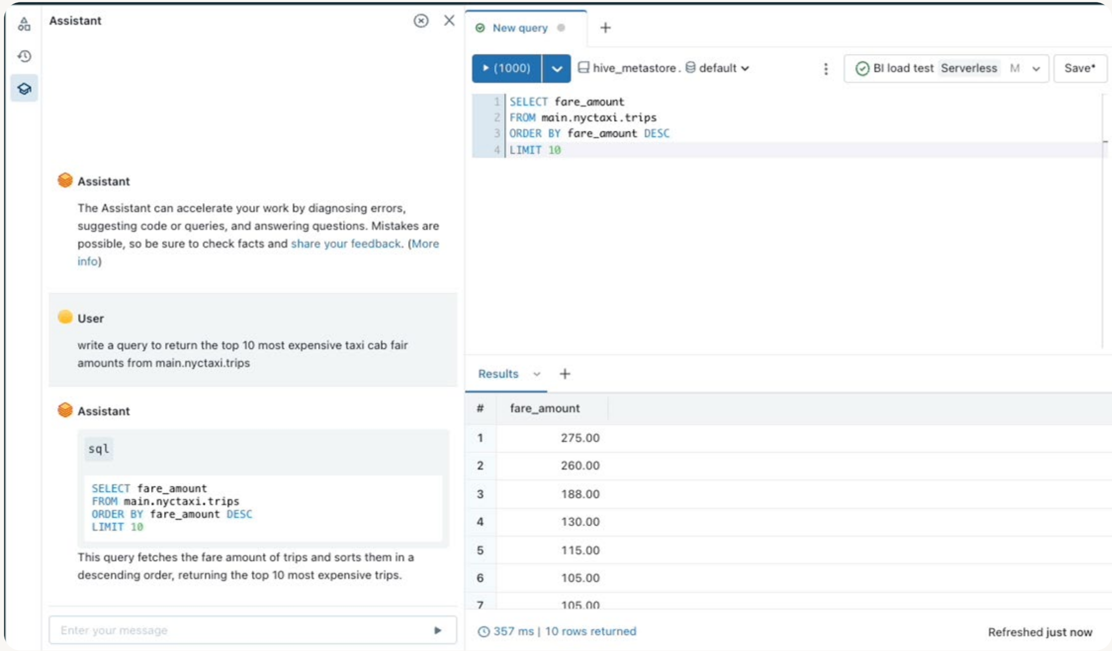

```
SELECT fare_amount
FROM main.nyctaxi.trips
ORDER BY fare_amount DESC
LIMIT 10
```

Assistant

The Assistant can accelerate your work by diagnosing errors, suggesting code or queries, and answering questions. Mistakes are possible, so be sure to check facts and share your feedback. (More info)

User

write a query to return the top 10 most expensive taxi cab fair amounts from main.nyctaxi.trips

Assistant

```sql
SELECT fare_amount
FROM main.nyctaxi.trips
ORDER BY fare_amount DESC
LIMIT 10
```

This query fetches the fare amount of trips and sorts them in a descending order, returning the top 10 most expensive trips.

#	fare_amount
1	275.00
2	260.00
3	188.00
4	130.00
5	115.00
6	105.00
7	105.00

357 ms | 10 rows returned

Refreshed just now

Assistant in Lakeview

Analysts can generate visualizations using natural language

Assistant
in Lakeview

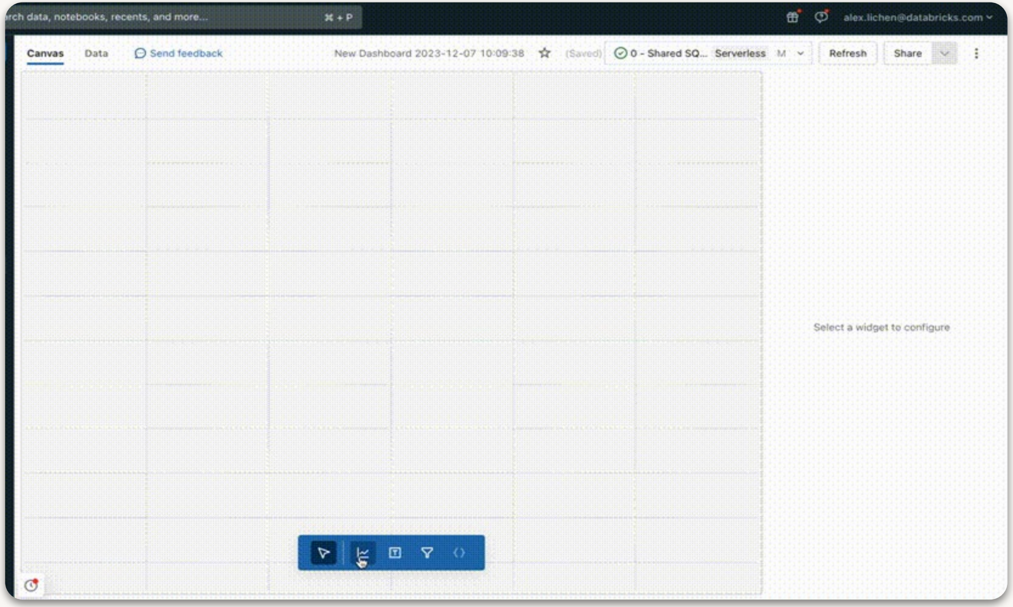
rch data, notebooks, recents, and more...
⌘ + P
alex.lichen@databricks.com
Canvas
Data
Send feedback
New Dashboard 2023-12-07 10:09:38
(Saved)
0 - Shared SQ... Serverless M
Refresh
Share

Select a widget to configure

Project Genie (Coming Soon)

Business users can query and visualize data using natural language

Project Genie

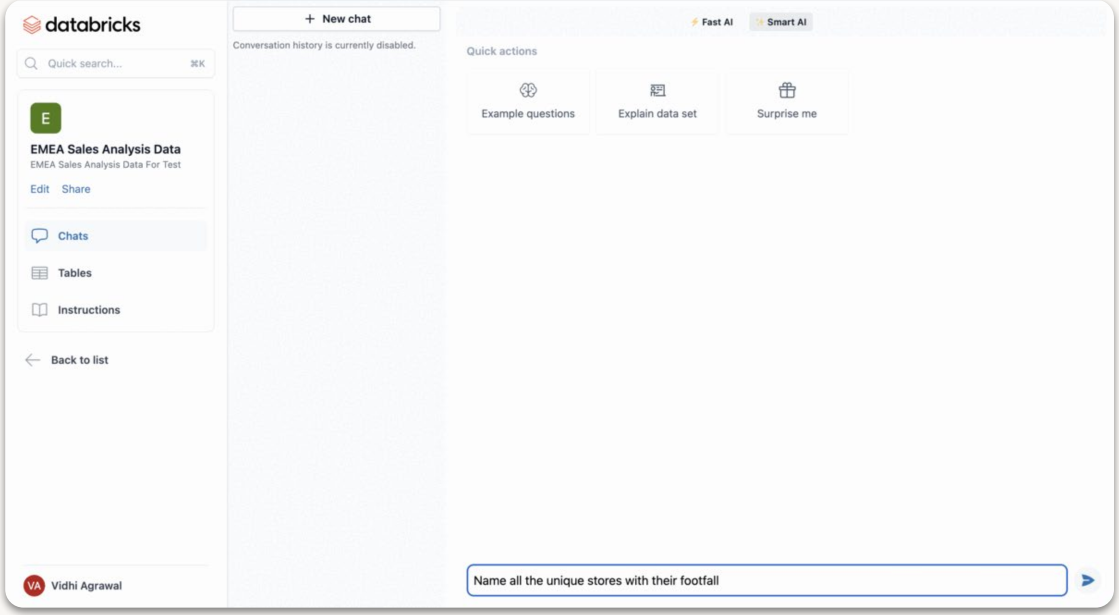

databricks

Quick search... ⌘K

E

EMEA Sales Analysis Data

EMEA Sales Analysis Data For Test

Edit Share

Chats

Tables

Instructions

Back to list

VA Vidhi Agrawal

+ New chat

Conversation history is currently disabled.

Fast AI Smart AI

Quick actions

Example questions

Explain data set

Surprise me

Name all the unique stores with their footfall

DATA INTEGRATION AND MIGRATION

Simple migrations at your pace

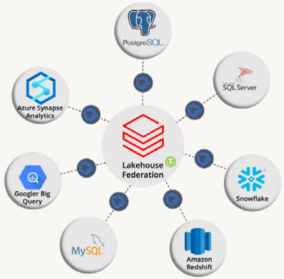

Federate

with Lakehouse Federation

Lower the barrier to entry and get started fast
Unity can secure and govern the sources, and track the consumption

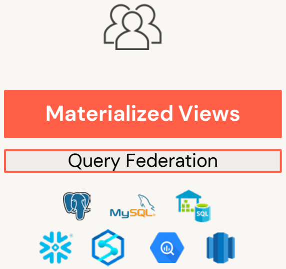

Materialize

with Materialized Views

When the users and business find value, lower the stress on federated data sources with Materialization in the cloud

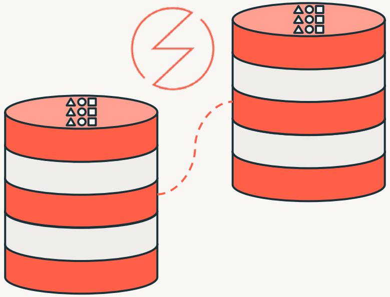

Ingest

with Change Data Capture

When the business and technical folks are aligned, lower the costs and migrate with low latency and high throughput incremental CDC

IN PRODUCTION TODAY

Announced at Summit

ALL OF THIS IS POWERED BY UNITY CATALOG

Simple data access with Federation

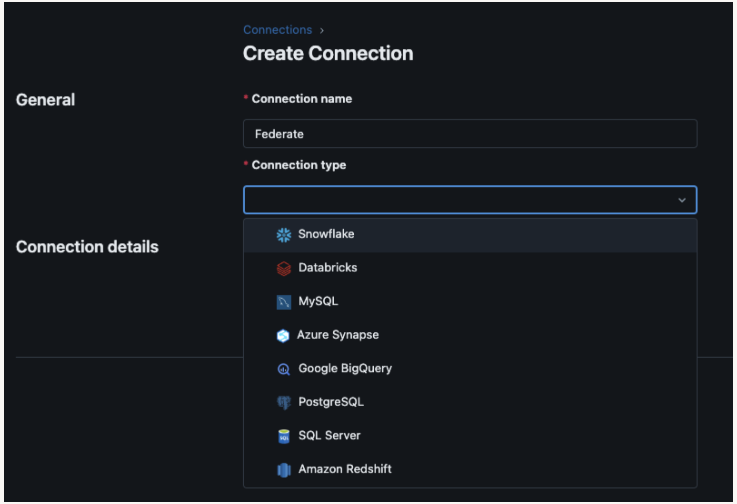

Connections ›

Create Connection

General

* Connection name

Federate

* Connection type

Connection details

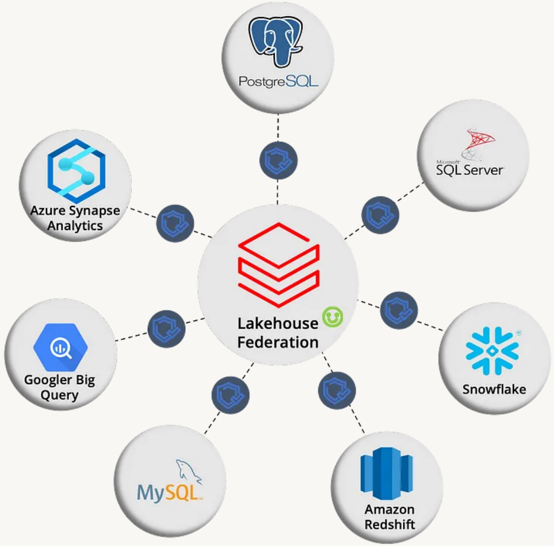

Discover, query, and govern ***all*** your data with **a unified view**, unified **engine** and **governance**—no matter where it lives

Simple and fast performance

Accelerating federated workloads

Federation 💛 Materialized views:

- Consistent latency & concurrency for data outside of the Lakehouse
- Accelerate cross-source joins and complicated transformation logic
- Offload access to underlying databases via materialized views to avoid high/concurrent loads on operational databases

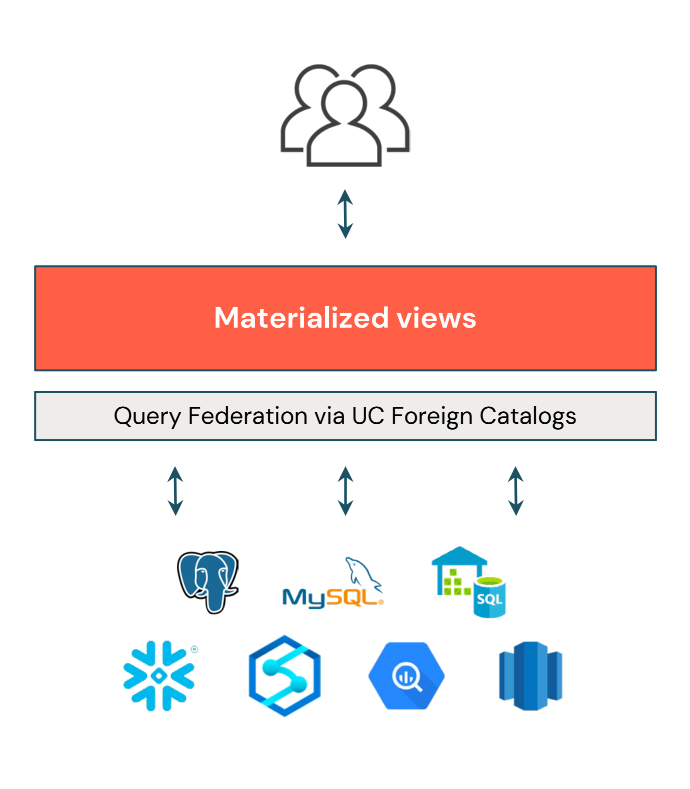

Simple data ingestion

Built-in connectors for common database and SaaS sources

Quickly and reliably replicate data across on-prem, cloud databases, and data platforms into the Lakehouse

Fully integrated into the Lakehouse

Low-maintenance API or no-code UI

Utilizing Arcion's advanced CDC technology

Private Preview in Q1

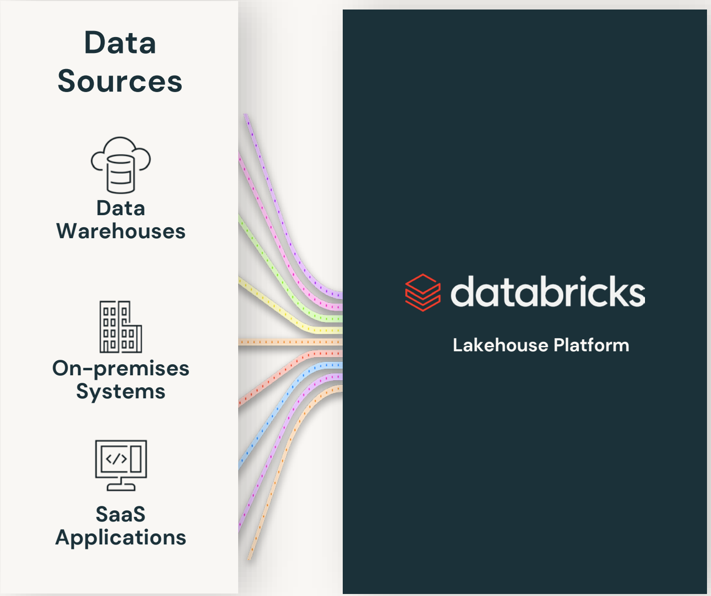

Warehousing technical debt

Stored procedures

- **Why**
  - Performance & cost
    - Runs on warehouse, compared to external ETL tools
- **What**
  - Powerful SQL script that encapsulates complex and reusable database operations
- **Pain**
  - Difficult to debug
  - Costly in the cloud
  - Migration pain with vendor lock-in
  - Anti-Patterns in the cloud

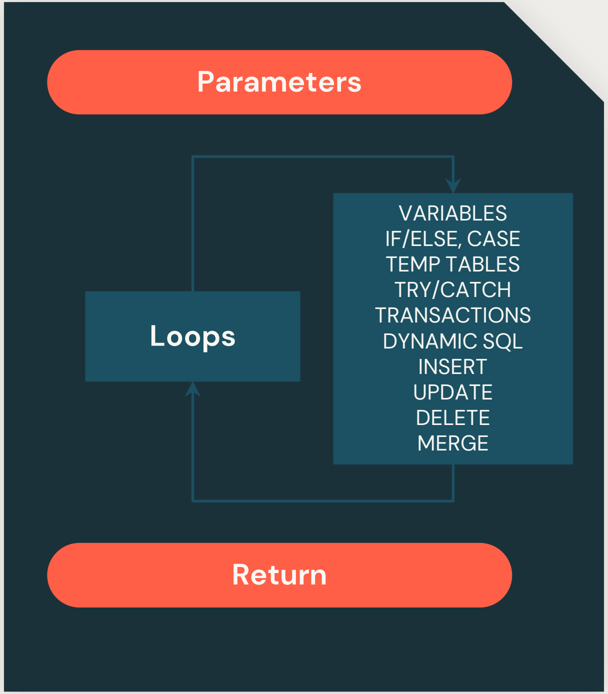

Refactor monolithic black box procedures to modularized Workflows

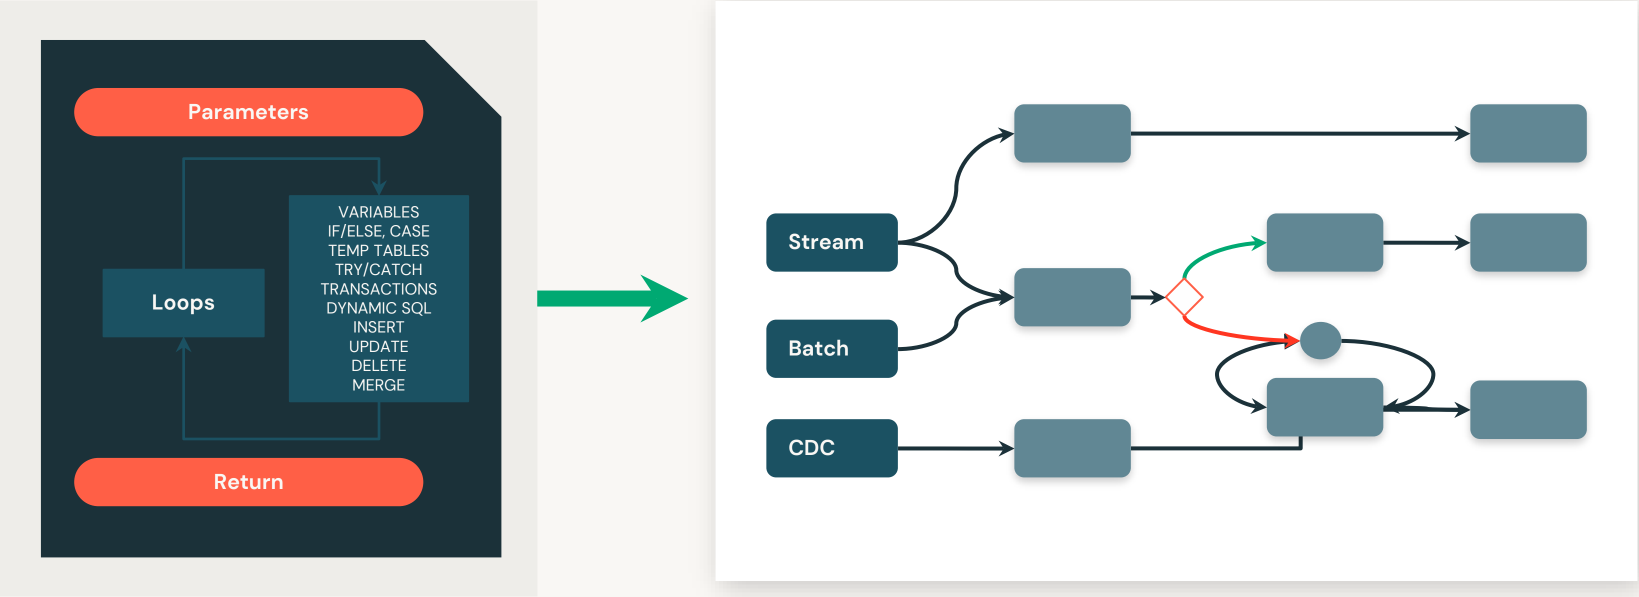

Simple authoring of Workflows

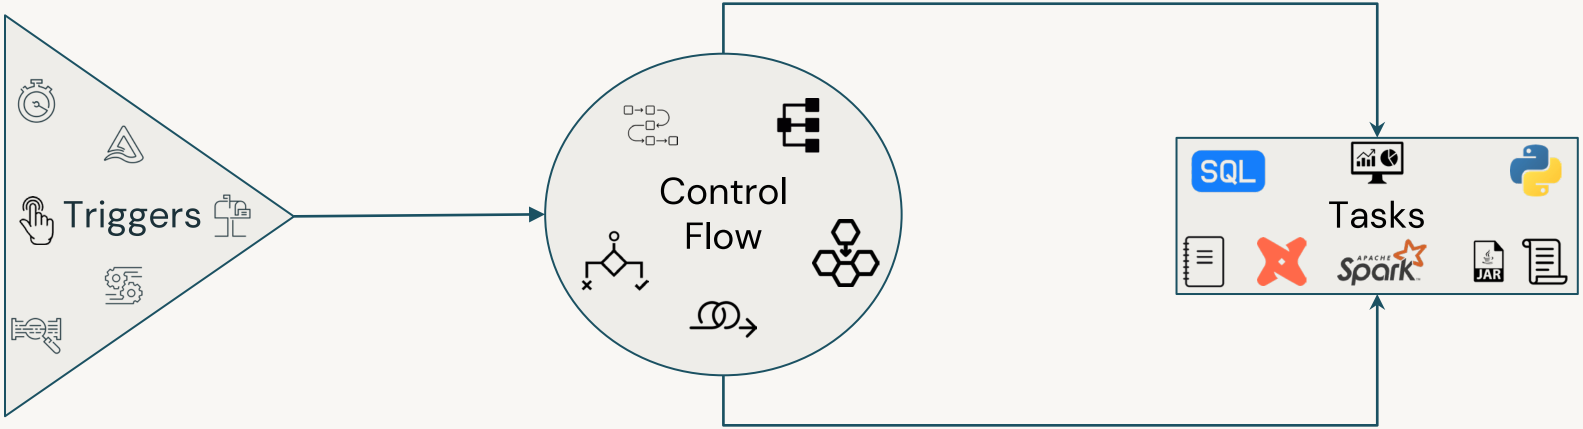

Manual
Scheduled
API
File Arrival
Delta table update
Continuous (Streaming)

Sequential
Parallel
Conditional (IF/ELSE)
Modular
Loops

SQL notebooks, files, queries
Python notebooks, files, scripts, wheels
Delta Live Tables
dbt
Java Jar
Spark Submit
Dashboards

Simple Migrations with AI

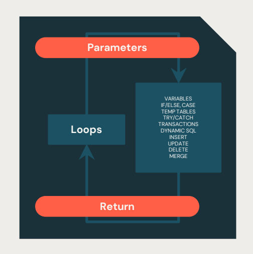

Use AI to extract Source to Target mappings and transformations and business logic from existing ETL + Warehouse

Source to Target

Transformations and Business Logic

Data Quality Rules

Map to a metadata driven framework

Use AI to generate workflow DAGs using Python support for Databricks Asset Bundles

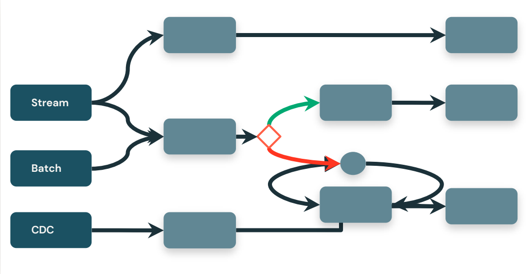

- Reduce risk of migration friction with LLM as judges to ensure transformation logic validity
- Accelerate time to value
- Backwards compatibility with Hive Federation, Delta Uniform or Delta Sharing to dependent systems
- Gain visibility with unified Governance and Orchestration
- Leverage existing Solutions Integrator frameworks

Simple real-time monitoring

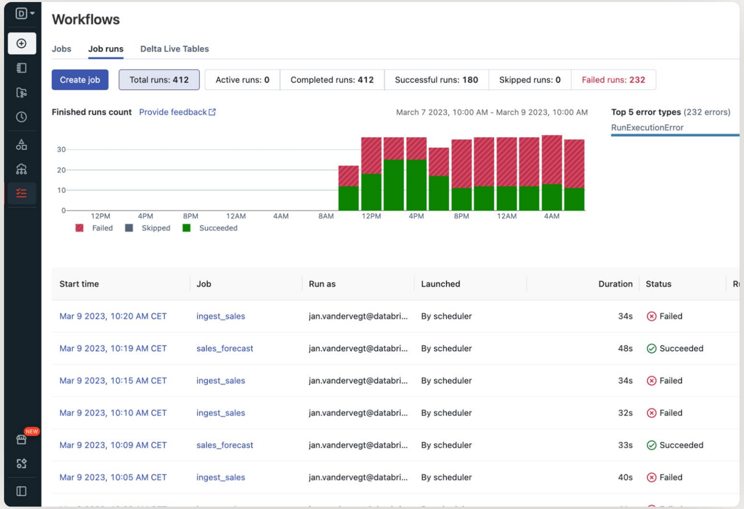

Start time	Job	Run as	Launched	Duration	Status
Mar 9 2023, 10:20 AM CET	ingest_sales	jan.vandervegt@databri...	By scheduler	34s	⊗ Failed
Mar 9 2023, 10:19 AM CET	sales_forecast	jan.vandervegt@databri...	By scheduler	48s	✓ Succeeded
Mar 9 2023, 10:15 AM CET	ingest_sales	jan.vandervegt@databri...	By scheduler	34s	⊗ Failed
Mar 9 2023, 10:10 AM CET	ingest_sales	jan.vandervegt@databri...	By scheduler	32s	⊗ Failed
Mar 9 2023, 10:09 AM CET	sales_forecast	jan.vandervegt@databri...	By scheduler	33s	✓ Succeeded
Mar 9 2023, 10:05 AM CET	ingest_sales	jan.vandervegt@databri...	By scheduler	40s	⊗ Failed

A simple and intuitive monitoring UI provides real-time metrics and detailed analytics for every workflow run

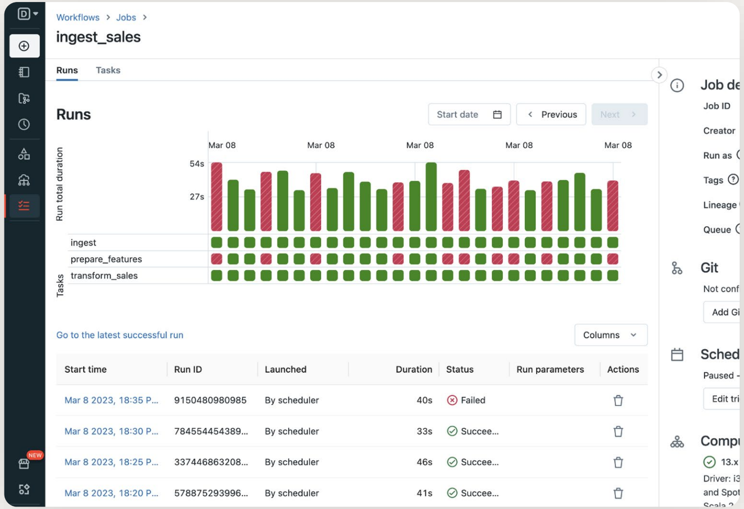

Start time	Run ID	Launched	Duration	Status	Run parameters	Actions
Mar 8 2023, 18:35 P...	9150480980985	By scheduler	40s	⊗ Failed		🗑
Mar 8 2023, 18:30 P...	784554454389...	By scheduler	33s	✓ Succee...		🗑
Mar 8 2023, 18:25 P...	337446863208...	By scheduler	46s	✓ Succee...		🗑
Mar 8 2023, 18:20 P...	578875293996...	By scheduler	41s	✓ Succee...		🗑

Drill down to understand which tasks are failing and why. Troubleshoot issues before your customers are impacted

Simple streaming with SQL

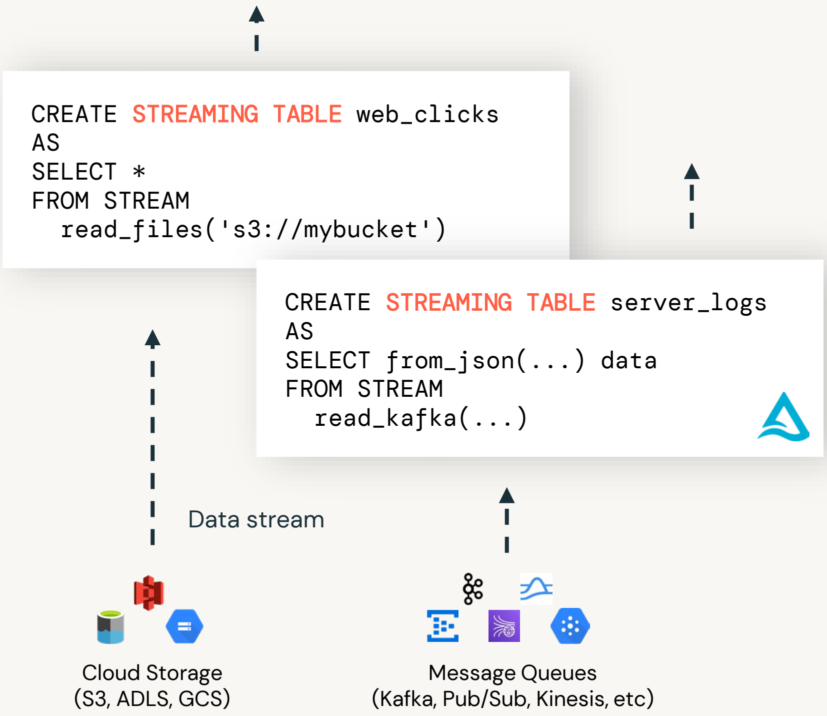

```
CREATE STREAMING TABLE web_clicks
AS
SELECT *
FROM STREAM
  read_files('s3://mybucket')
```

```
CREATE STREAMING TABLE server_logs
AS
SELECT from_json(...) data
FROM STREAM
  read_kafka(...)
```

Benefits:

1. **Unlock real-time use cases.** Ability to support real-time analytics/BI, machine learning and operational use cases with streaming data.
2. **Better scalability.** More efficiently handle high volumes of data via incremental processing vs. large batches.
3. **Enable more practitioners.** Simple SQL syntax makes data streaming accessible to all data engineers and analysts.

Simple and fast BI with Materialized Views

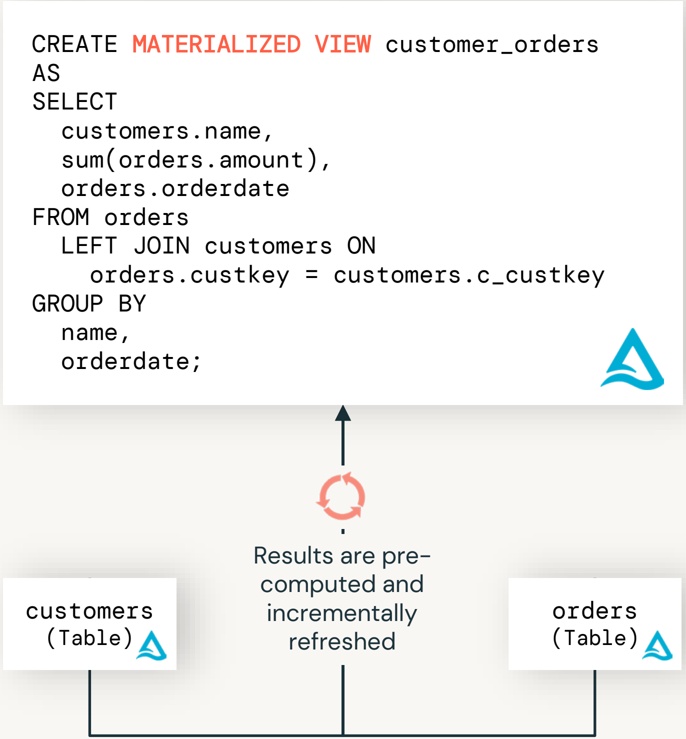

```
CREATE MATERIALIZED VIEW customer_orders
AS
SELECT
  customers.name,
  sum(orders.amount),
  orders.orderdate
FROM orders
  LEFT JOIN customers ON
    orders.custkey = customers.c_custkey
GROUP BY
  name,
  orderdate;
```

Benefits:

1. **Accelerate BI dashboards.** Much faster to query data that is pre-computed vs querying base tables.
2. **Reduce data processing costs.** MV results are refreshed incrementally avoiding the need to completely rebuild the view when new data arrives.
3. **Improve data access control.** More tightly govern what data can be seen by consumers by controlling access to base tables.

Ingest and Transform with Medallion Mesh

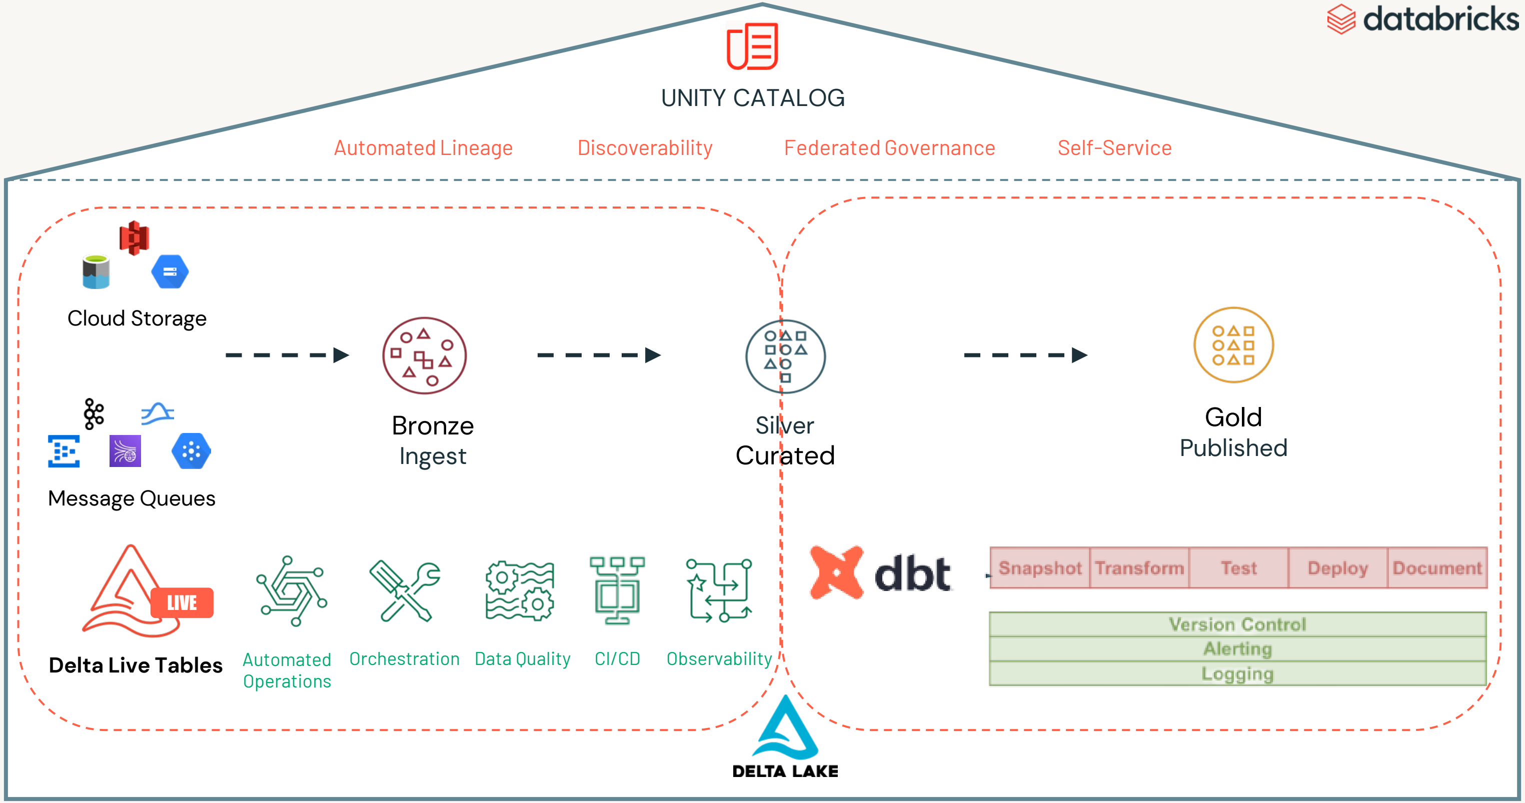

* https://future.a16z.com/emerging-architectures-modern-data-infrastructure/

Minimal data domain blueprint

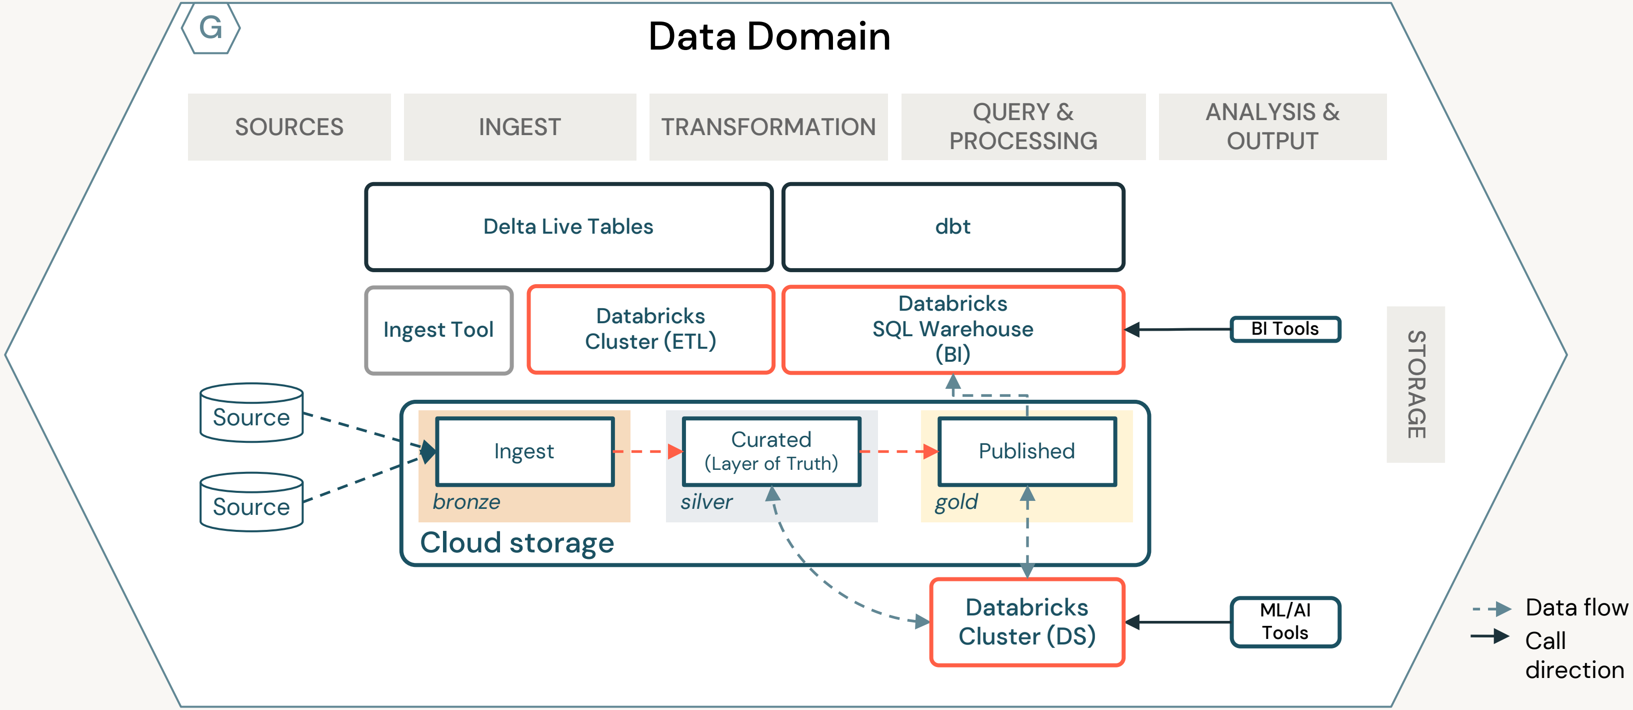

* https://future.a16z.com/emerging-architectures-modern-data-infrastructure/

Vertical view into the domains and layers

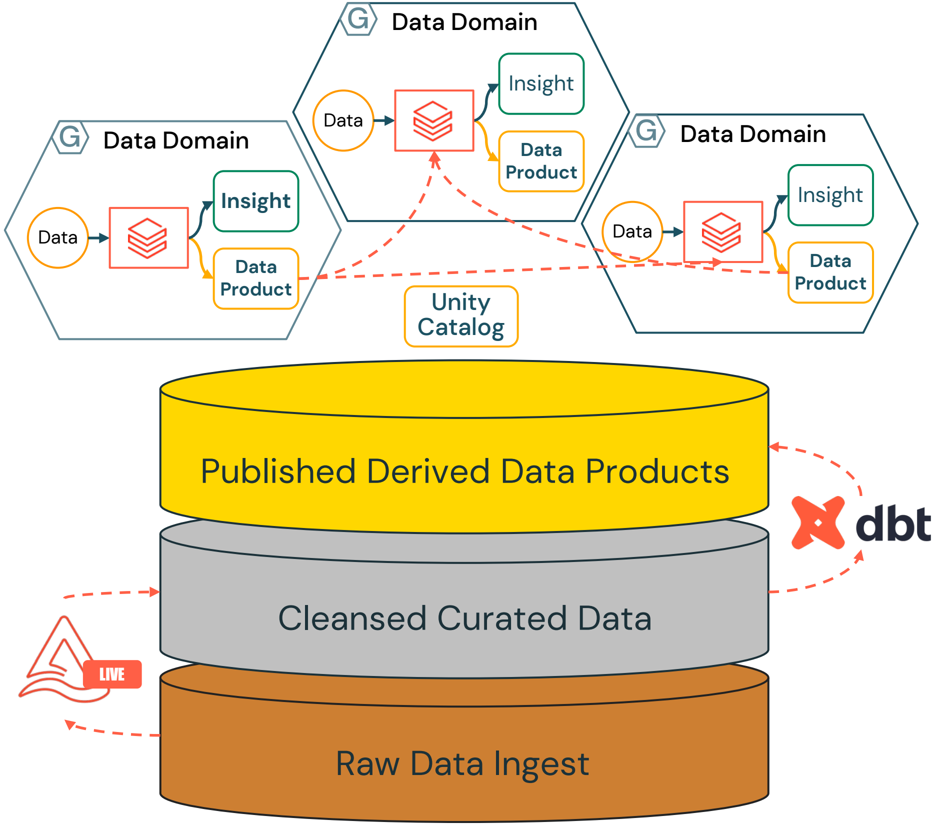

Unified data modeling in Medallion Mesh

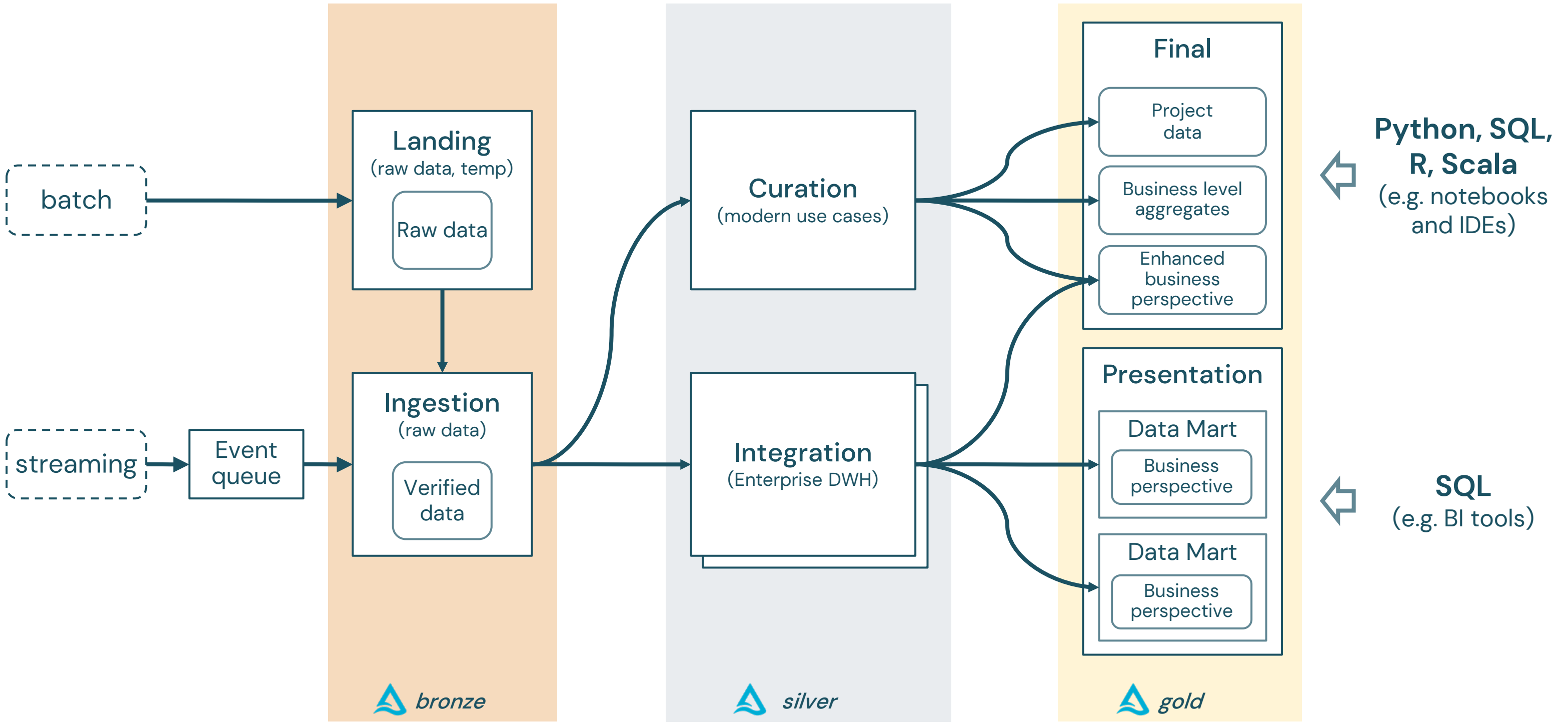

UNIFIED PLATFORM FOR ALL USERS AND USE CASES

Medallion Mesh on Data Intelligence Platform

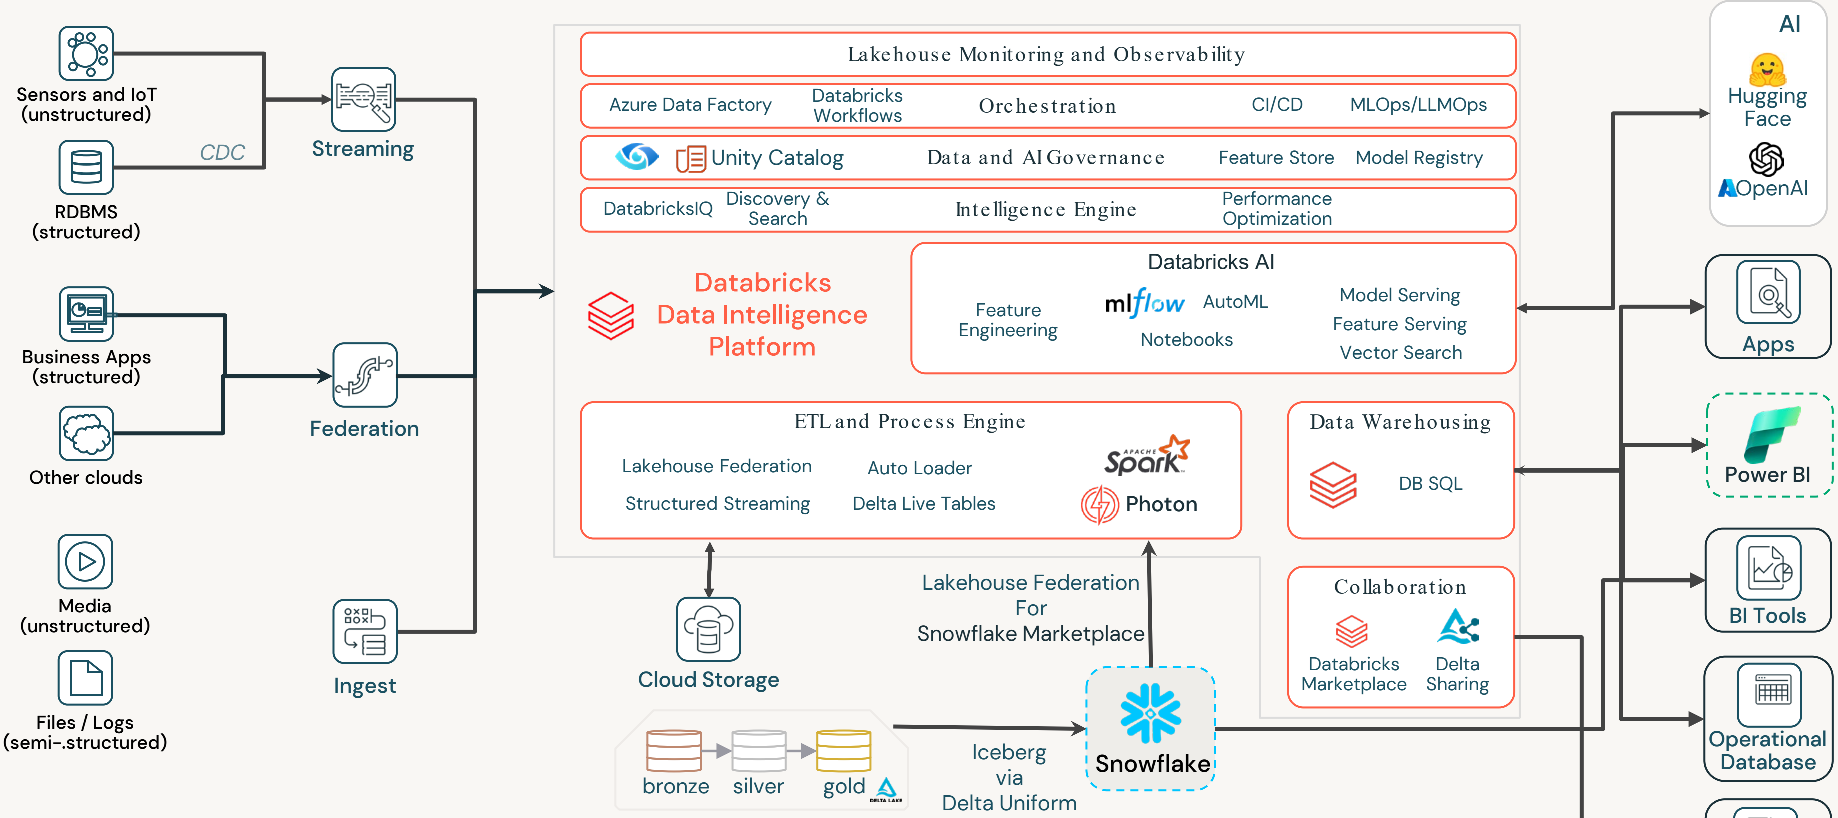

Simple data understanding

Real-time relational modelling visualization

Easily visualize **table-to-table** relationships leveraged by **BI tools** and **DatabricksIQ**

Modify and create PK/FK Constraints with UI

Visualize table to table relationships with a built-in Entity Relationship Diagram (ERD)

Built on Unity Catalog

View PK, FK columns via the schema browser

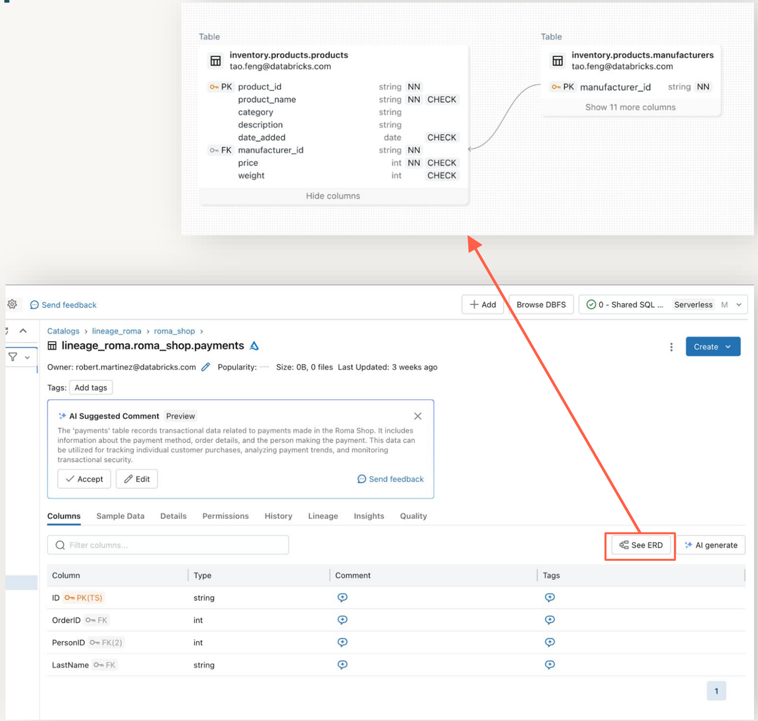

Table

inventory.products.products
tao.feng@databricks.com

PK	product_id	string	NN	
	product_name	string	NN	CHECK
	category	string		
	description	string		
	date_added	date		CHECK
FK	manufacturer_id	string	NN	
	price	int	NN	CHECK
	weight	int		CHECK

Table

inventory.products.manufacturers
tao.feng@databricks.com

PK	manufacturer_id	string	NN

Catalogs › lineage_roma › roma_shop ›

lineage_roma.roma_shop.payments

Owner: robert.martinez@databricks.com Popularity: Size: 0B, 0 files Last Updated: 3 weeks ago

Tags: Add tags

AI Suggested Comment Preview

The 'payments' table records transactional data related to payments made in the Roma Shop. It includes information about the payment method, order details, and the person making the payment. This data can be utilized for tracking individual customer purchases, analyzing payment trends, and monitoring transactional security.

Column	Type	Comment	Tags
ID PK(TS)	string		
OrderID FK	int		
PersonID FK(2)	int		
LastName FK	string		

Simple and unified visibility into data and AI

- **Discover and classify** structured and unstructured data, files, notebooks, ML models, and dashboards at one place
- Consolidate and query data from **other databases and data warehouses** using a **single point of access,** without moving or copying the data
- Build better **understanding of your data estate** with automated lineage, tags and auto-generated data insights
- Boost productivity by searching, understanding and gaining insights from your data and AI assets, using **natural language**

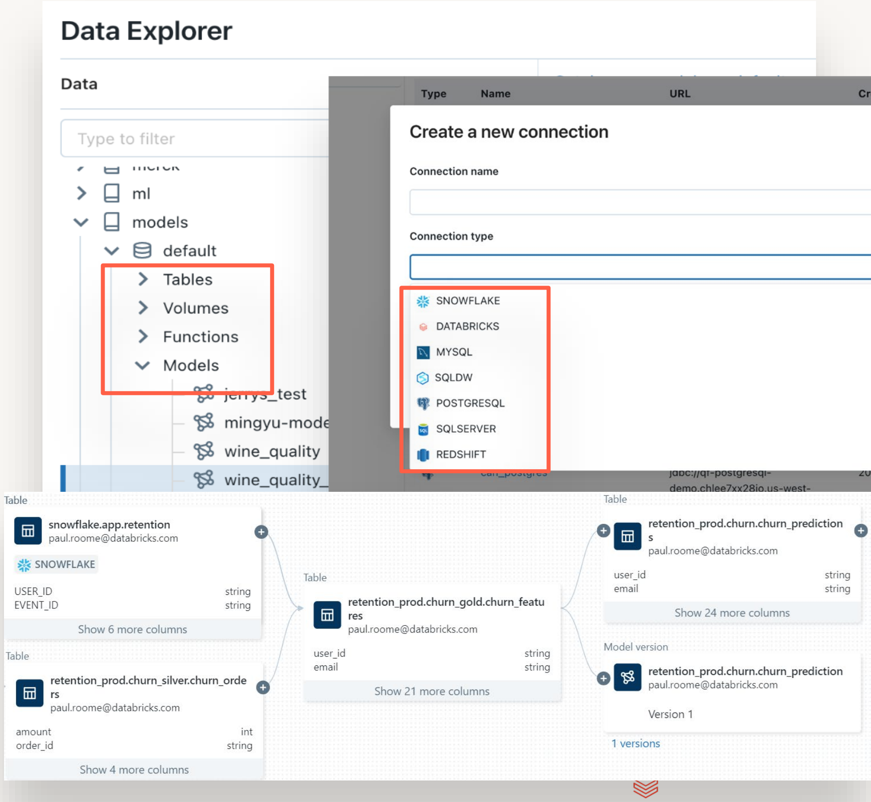

Simple help from AI in your SQL

Public Preview

Write SQL to get insight from unstructured text data via LLMs

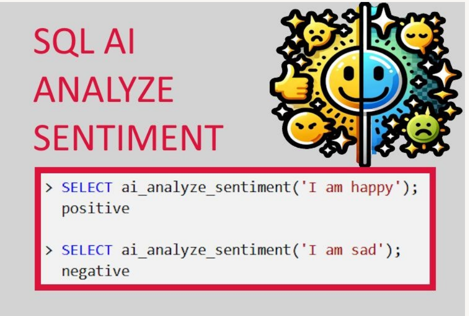

SQL AI ANALYZE SENTIMENT

```
> SELECT ai_analyze_sentiment('I am happy');
  positive

> SELECT ai_analyze_sentiment('I am sad');
  negative
```

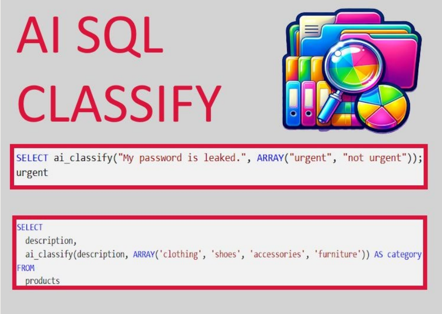

AI SQL CLASSIFY

```
SELECT ai_classify("My password is leaked.", ARRAY("urgent", "not urgent"));
urgent
```

```
SELECT
  description,
  ai_classify(description, ARRAY('clothing', 'shoes', 'accessories', 'furniture')) AS category
FROM
  products
```

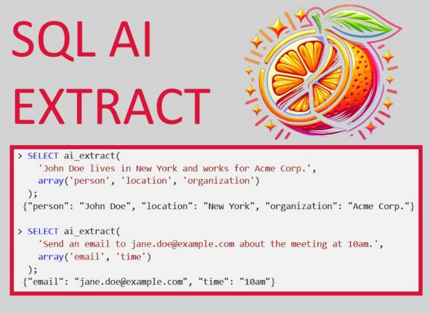

SQL AI EXTRACT

```
> SELECT ai_extract(
    'John Doe lives in New York and works for Acme Corp.',
    array('person', 'location', 'organization')
  );
 {"person": "John Doe", "location": "New York", "organization": "Acme Corp."}

> SELECT ai_extract(
    'Send an email to jane.doe@example.com about the meeting at 10am.',
    array('email', 'time')
  );
 {"email": "jane.doe@example.com", "time": "10am"}
```

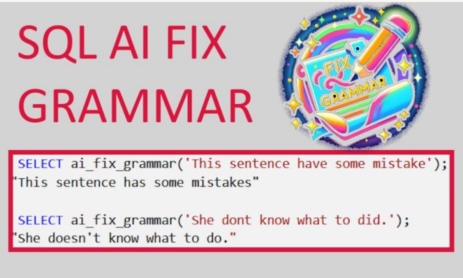

SQL AI FIX GRAMMAR

```
SELECT ai_fix_grammar('This sentence have some mistake');
"This sentence has some mistakes"

SELECT ai_fix_grammar('She dont know what to did.');
"She doesn't know what to do."
```

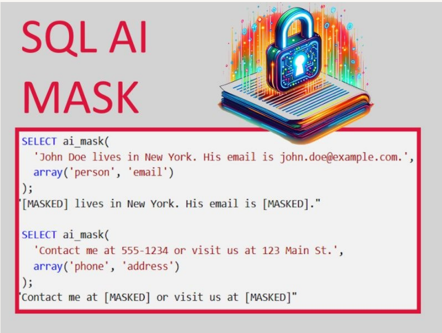

SQL AI MASK

```
SELECT ai_mask(
  'John Doe lives in New York. His email is john.doe@example.com.',
  array('person', 'email')
);
"[MASKED] lives in New York. His email is [MASKED]."

SELECT ai_mask(
  'Contact me at 555-1234 or visit us at 123 Main St.',
  array('phone', 'address')
);
"Contact me at [MASKED] or visit us at [MASKED]"
```

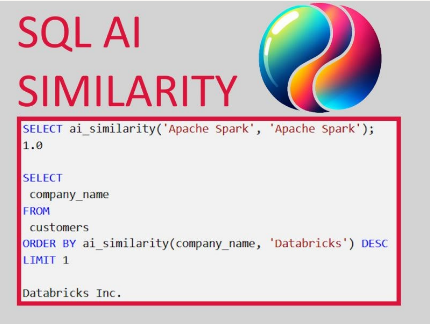

SQL AI SIMILARITY

```
SELECT ai_similarity('Apache Spark', 'Apache Spark');
1.0

SELECT
  company_name
FROM
  customers
ORDER BY ai_similarity(company_name, 'Databricks') DESC
LIMIT 1

Databricks Inc.
```

Simplify your SQL

Iterate faster with AI

Assistant does all the heavy lifting

Cell in focus

Focus on a cell to turn it into a SQL Editor

Focus back to notebook mode

Run Cells in parallel

Use Run Now to execute SQL cells in parallel, no more waiting for execution

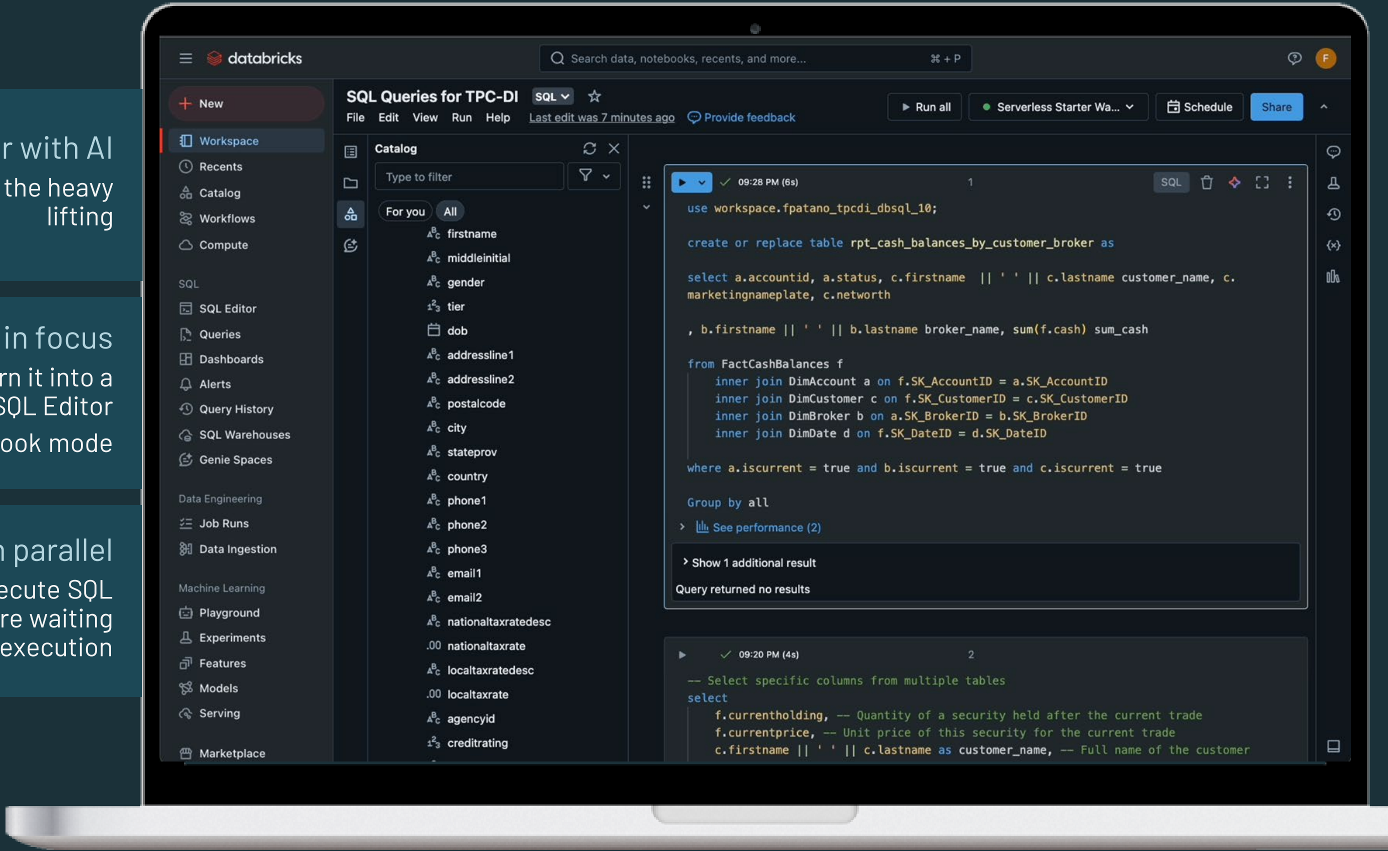

```
use workspace.fpatano_tpcdi_dbsql_10;

create or replace table rpt_cash_balances_by_customer_broker as

select a.accountid, a.status, c.firstname  || ' ' || c.lastname customer_name, c.
marketingnameplate, c.networth

, b.firstname || ' ' || b.lastname broker_name, sum(f.cash) sum_cash

from FactCashBalances f
    inner join DimAccount a on f.SK_AccountID = a.SK_AccountID
    inner join DimCustomer c on f.SK_CustomerID = c.SK_CustomerID
    inner join DimBroker b on a.SK_BrokerID = b.SK_BrokerID
    inner join DimDate d on f.SK_DateID = d.SK_DateID

where a.iscurrent = true and b.iscurrent = true and c.iscurrent = true

Group by all
```

Query returned no results

```
-- Select specific columns from multiple tables
select
    f.currentholding, -- Quantity of a security held after the current trade
    f.currentprice, -- Unit price of this security for the current trade
    c.firstname || ' ' || c.lastname as customer_name, -- Full name of the customer
```

databricks

Simplify your Data Science and ML

Provide a collaborative environment for Unified Machine Learning & Data Analytics

Multi-Language

Scala, SQL, Python, R:
All in one notebook

Visualizations

Built-in visualizations and support for the most popular visualization libraries (e.g. matplotlib, ggplot)

Experiment Tracking

Built-in tracking of Data Science and ML experiments, with metrics, parameters, artifacts, and more

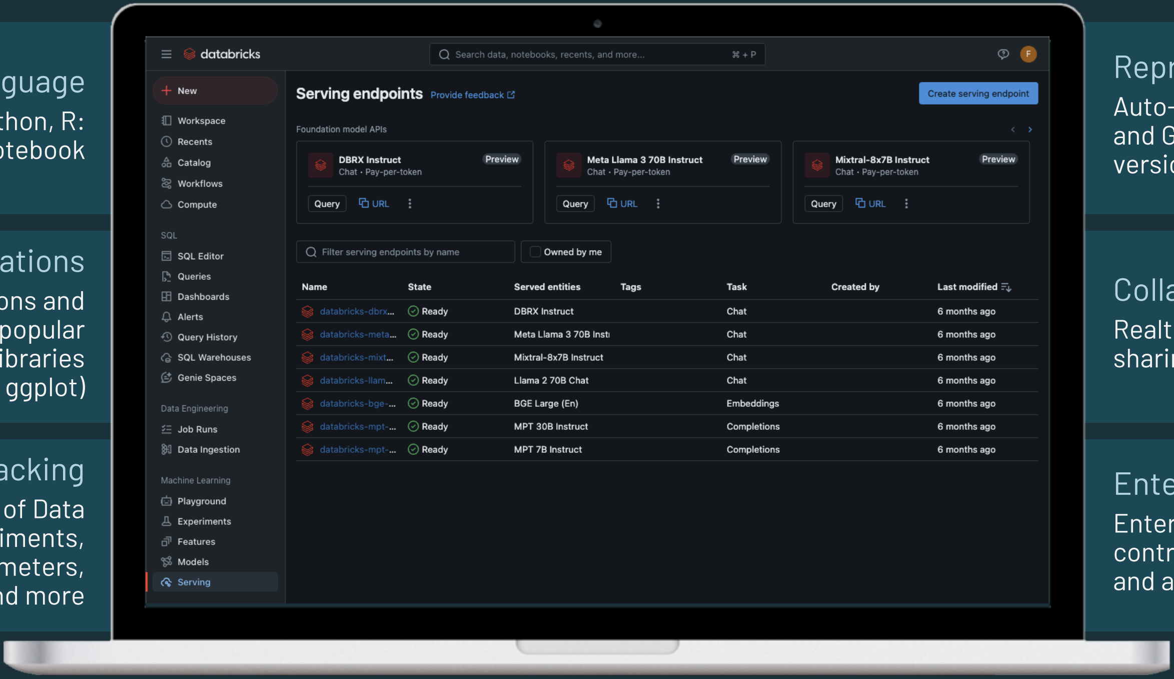

Serving endpoints

Foundation model APIs

DBRX Instruct Chat · Pay-per-token Preview	Meta Llama 3 70B Instruct Chat · Pay-per-token Preview	Mixtral-8x7B Instruct Chat · Pay-per-token Preview

Name	State	Served entities	Tags	Task	Created by	Last modified
databricks-dbrx...	Ready	DBRX Instruct		Chat		6 months ago
databricks-meta...	Ready	Meta Llama 3 70B Insti		Chat		6 months ago
databricks-mixt...	Ready	Mixtral-8x7B Instruct		Chat		6 months ago
databricks-llam...	Ready	Llama 2 70B Chat		Chat		6 months ago
databricks-bge-...	Ready	BGE Large (En)		Embeddings		6 months ago
databricks-mpt-...	Ready	MPT 30B Instruct		Completions		6 months ago
databricks-mpt-...	Ready	MPT 7B Instruct		Completions		6 months ago

Reproducible

Auto-logged revision history and Git integration for version control

Collaborative

Realtime co-editing, with sharing and permissions

Enterprise Ready

Enterprise grade access controls, identity pass-through, and auditability

Simplify DEVOPS with Repos

Repos API to automate CI/CD

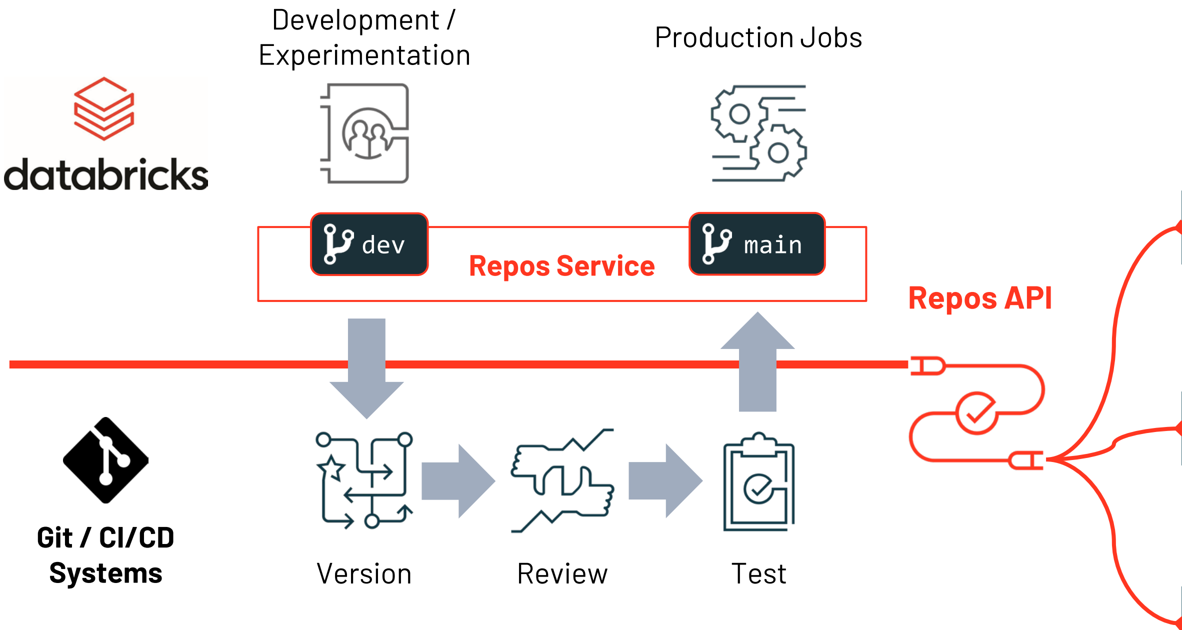

Programmatically manage Repos

E.g., set up automation to automatically create Databricks Repos for all of your repositories

Perform Git operations on Repos

E.g., set up automation to always keep Repos in a production folder at a specified version

Specify Git branch, tag, hash for Jobs

E.g., update the Repo to a specific branch and commit and run Jobs against the notebooks in the repo

Simplify MLOPS with mlflow

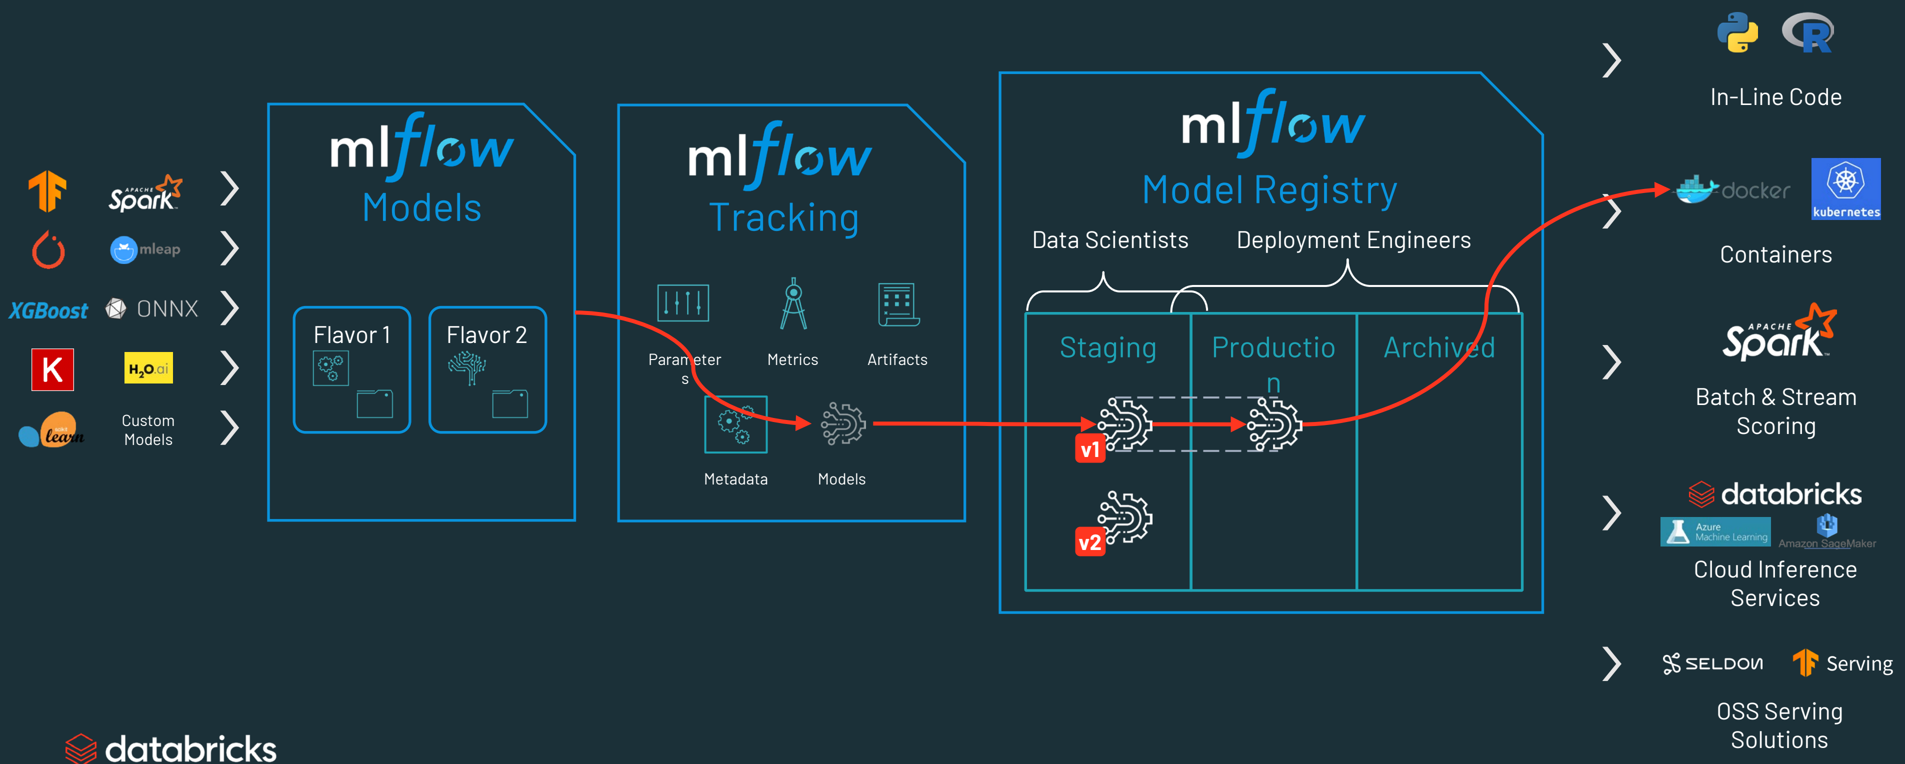

databricks

Simple Prompt Engineering UI with mlflow

Adjust text prompts for Gen AI models to elicit better responses

MLflow Prompt Engineering lets you compare and analyze many models and prompts, across many inputs.

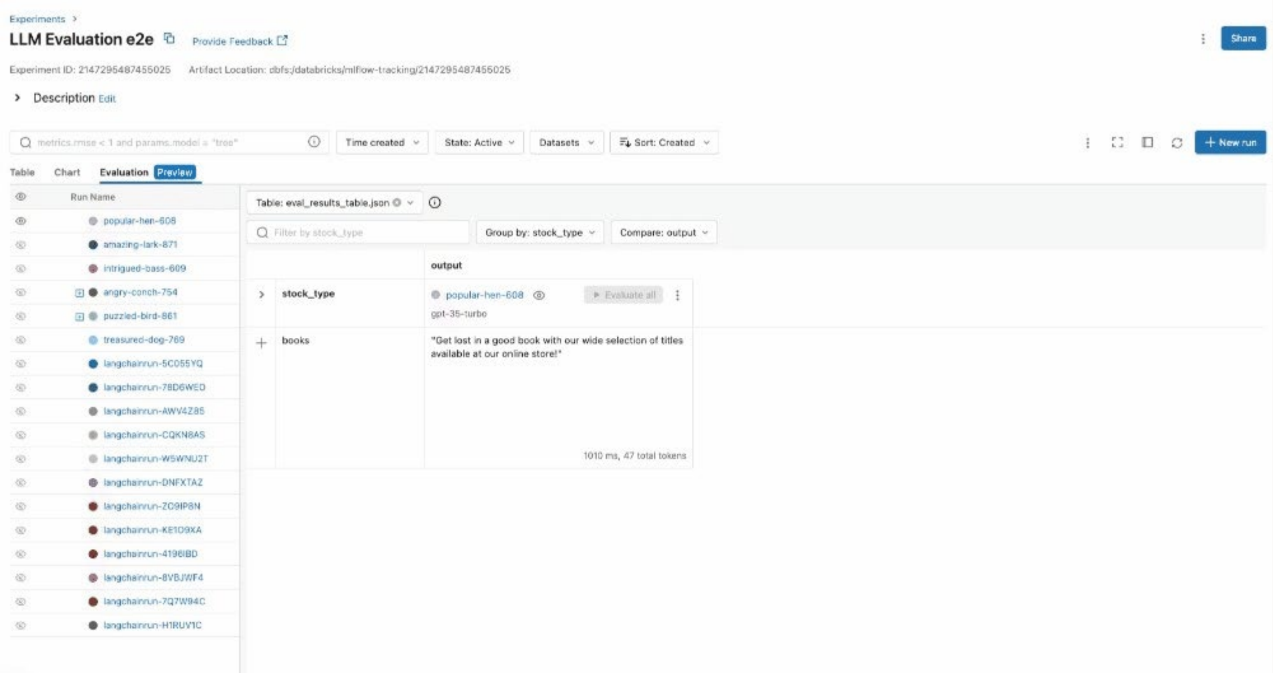

Experiments >

LLM Evaluation e2e

Experiment ID: 2147295487455025   Artifact Location: dbfs:/databricks/mlflow-tracking/2147295487455025

> Description Edit

Run Name
popular-hen-608
amazing-lark-871
intrigued-bass-609
angry-conch-754
puzzled-bird-861
treasured-dog-769
langchainrun-5C055YQ
langchainrun-78D6WED
langchainrun-AWV4Z85
langchainrun-CQKN8AS
langchainrun-W5WNU2T
langchainrun-DNFXTAZ
langchainrun-ZC9IP8N
langchainrun-KE1D9XA
langchainrun-4196IBD
langchainrun-8VBJWF4
langchainrun-7Q7W94C
langchainrun-H1RUV1C

Table: eval_results_table.json

	output
stock_type	popular-hen-608 gpt-35-turbo
books	"Get lost in a good book with our wide selection of titles available at our online store!" 1010 ms, 47 total tokens

RISE UP AND UNIFY WITH THE MEDALLION MESH

fin