kyvos

Supercharging Databricks with Kyvos Semantic Layer

10x
faster analytics
on any scale, any concurrency

90%
lower cost
with no performance slowdowns

~90%
fewer analyst hours
improved productivity

With Kyvos semantic layer on top of Databricks, enterprises get instant, scalable analytics while drastically reducing compute consumption and query costs, when compared to native Databricks performance.

What You Gain with Kyvos on Databricks

Single Source of Semantic Truth

Universal, consistent, trusted semantic intelligence for agents, LLMs and BI.

Speed & Scale

Fastest analytical engine for sub-second query results on petabyte-scale data.

Lower Costs

AI-powered smart aggregation and dynamic resource adjustment for reduced costs.

High Concurrency

Fully distributed, elastic architecture to support thousands of users simultaneously.

© 2026 Kyvos Insights. All rights reserved.

Semantic Layer

kyvos

Benchmarking Environment

500 Analysts | **250** Days | **50** Cold queries / Analyst / Day | **6.25M** Queries executed

Speed and Savings at Scale

	Databricks	Supercharging Databricks with Kyvos
Average Query Time (1.2B rows)	11.61 sec	**1.19 sec** ↓ 89.8%
Average Query Cost	$0.18	**$0.02** ↓ 88.9%
Annual Cost	$ 1,125,000	**$ 125,000** ↓ 88.9%
Query Push-Down	Sends requests to the underlying table every time	Reads data only once from the source
Query Type	Served from raw tables	Served from smart aggregates
Performance Consistency	Variable, depends on compute and resources	Consistently fast
Partitions	Physical, manual tuning, over-partitioning can lead to slow performance	Semantic, logical partitions, minimal rebuilds
Compute Consumption	High, rises sharply with data growth	Low, regardless of query volume or data size

© 2026 Kyvos Insights. All rights reserved.