MCKNIGHT
CONSULTING GROUP

Comparative Performance and Cost Analysis of Modern Analytical Data Platforms

High concurrency vs high demand

EDB Postgres AI for WarehousePG
Snowflake
Databricks
Redshift
Hive on Iceberg

McKnight Consulting Group

February 2026

Table of Contents

Executive Summary

The primary objective of this testing was to evaluate what a modern analytics strategy looks like through extensive testing and usage scenarios on EDB Postgres AI for WarehousePG, Snowflake, Databricks, Redshift, and Hive on Apache Iceberg (AWS EMR).

Currently, many organizations are trapped in a cycle of operational friction, facing system instability during peak reporting periods or scaling back their data science ambitions to stay within budget. This is the "Flexibility Tax" in action—where the cost of maintaining a fragmented, best-of-breed stack outweighs the insight it provides. As enterprises simplify their stacks and bring warehouses into their data and AI platforms, the pressure on high-concurrency workloads has never been greater. We sought out to find a solution.

We extended the TPC-DS benchmark by adding approximately 20% additional queries that are simpler than standard TPC-DS queries, reflecting the reality of modern enterprise data usage: a mix of heavy, complex queries and high-frequency, routine reports and dashboards. TPC-DS is the industry-standard benchmark for evaluating how well a data platform performs complex, real-world analytical workloads, measuring not just speed but also cost-effectiveness under business-critical conditions.

Under concurrency and mixed workload testing, cloud-only analytics platforms show clear throughput advantages, but they incur coordination and planning overhead. In contrast, WarehousePG handles concurrent users with a lower performance differential, providing a more consistent experience when predictability is a priority.

The results demonstrate that cloud warehouses suit high-performance analytics for the most demanding queries, justifying their pricing. Conversely, WarehousePG works efficiently for the high-concurrency analytics that power daily operations, providing consistent performance with better cost efficiency. This reveals the merits of a hybrid approach: cloud-native for high-scale analytics, and self-hosted systems for always-on business intelligence (BI).

Introduction

This study assesses the scaling and cost characteristics of leading analytical data platforms under realistic usage scenarios, moving beyond query speed to examine workloads representative of enterprise environments. The focus is on understanding how different architectural approaches handle analytical queries, concurrency, and mixed workloads combining reporting and dashboarding.

Using an extended version of the industry-standard TPC-DS benchmark (10 TB dataset, official kit, full specification compliance), we evaluated WarehousePG, Snowflake, Databricks, Amazon Redshift, and Hive with Trino on Apache Iceberg (AWS EMR), representing a range of architectures.

The testing methodology includes concurrent and mixed workloads, designed to reveal predictability under load and cost efficiency. Results provide insights into platform strengths, trade-offs, and alignment with workload priorities and budget constraints.

As data becomes embedded in everyday enterprise workflows — and agentic workloads grow — Data Scientists and infrastructure leaders will face mounting pressure to balance peak performance with the practical demands of high concurrency.

Testing Methodology and Workload Design

The primary objective of this testing was to evaluate what a modern analytics strategy looks like through extensive testing and usage scenarios on WarehousePG, Snowflake[1], Databricks, Redshift, and Hive on Apache Iceberg (AWS EMR).

TPC-DS

We tested workloads based on the TPC-DS[2].

The TPC-DS benchmark is the industry standard for measuring how well a data warehouse handles complex analytical tasks. It simulates the types of queries and workloads an organization might run—spanning sales, inventory, customer behavior, promotions, and forecasting—using 99 queries that test everything from simple aggregations to highly complex joins and subqueries. Unlike synthetic tests that only measure raw speed, TPC-DS reflects the challenges customers face in running business-critical workloads at scale. Strong TPC-DS results demonstrate a platform's ability to handle the most demanding, large-scale analytics needs.

While TPC-DS is excellent for evaluating these complex scenarios, the most common queries in an enterprise are often simpler and run with much higher concurrency, such as frequent dashboard refreshes. To more closely match these real-world enterprise needs, we crafted a mixed workload that incorporates the complexity of TPC-DS while increasing the volume of simpler queries.

For all of our workloads, the following testing conditions were used:

- Scale factor (data size) of 10TB
- Data generated by official TPC-DS kit
- Query order per TPC-DS specifications
- Query syntax compliant with TPC-DS specifications
- Tables used cluster keys (Snowflake and Databricks), distribution/sort keys (Redshift), or partitions (Hive and WarehousePG)

---

[1] We tested Snowflake with and without multi-cluster warehouses. For multi-cluster testing, according to Snowflake:

> By default, a virtual warehouse consists of a single cluster of compute resources available to the warehouse for executing queries. As queries are submitted to a warehouse, the warehouse allocates resources to each query and begins executing the queries. If sufficient resources are not available to execute all the queries submitted to the warehouse, Snowflake queues the additional queries until the necessary resources become available.

[2] This was not an official TPC-DS. For more information, see the official site at: https://www.tpc.org/tpcds/

Workloads

We tested two different workloads:

- Concurrency Workload (5 concurrent users, TPC-DS queries)
- Mixed Workload (multiple users, mixture of TPC-DS and BI queries)

Mixed Workload (multiple users, TPC-DS and BI queries)

The analytics workload objective was to simulate a realistic mixed workload that combines occasional heavy analytics (complex, long-running reports); steady, moderate-complexity BI and reporting queries; and high-frequency dashboard pings consisting of simple queries that complete in around one second. This workload used the standard TPC-DS schema and data and executed the 99 official TPC-DS queries, divided into groups: 14 “complex” queries and a “medium” group of 85 queries. These 2 groups constituted the 99 TPC-DS standard queries.

We also wanted roughly 20% of our workload queries to be simpler than any in TPC-DS, reflecting environments where users rely heavily on routine or mundane queries. So we included 25 new, custom “simple” queries, each designed with zero or one join, two to three WHERE filters, and small result sets appropriate for sub-second response times[3].

The following list details the testing conditions and the query groups:

Complex TPC-DS queries (14 queries):
- Q4, Q11, Q23, Q24, Q64, Q65, Q67, Q74, Q78, Q87, Q88, Q93, Q95, Q97
- Only 2 executed per hour
- Represent “end-of-day” reports, complex cube-like analyses, multi-join queries with large scans

Medium TPC-DS queries (85 queries):
- All other TPC-DS queries (Q1–Q99 excluding the 14 above)
- Represent typical dashboard/report queries, ad-hoc analysis, and more complex BI tiles

Simple BI queries (25 queries crafted by MCG):
- Queries were parametrized with 100 randomly generated parameters
- Each uses one base fact or dimension table (optionally one join)
- Max one join, 2-3 filters, small group by result sets
- Represent users hitting dashboards or KPIs every few seconds

Testing conditions:
- Query queues were shuffled in random order once, so all platforms received the same queries in the same order
- The throughput in queries per hour (QpH) was the primary performance metric

[3] See https://github.com/mcknightconsulting/MixedWorkload.git for the simple queries

Concurrency Workload (5 concurrent users, TPC-DS queries)

The concurrency workload was a concurrent run of all 99 TPC-DS queries with 5 concurrent users. The following list details the testing conditions:

- Five users each running a stream of the 99 TPC-DS
- Query order by stream determined by the [TPC-DS specifications]

The total elapsed time for all 5 users to finish their 99 queries in seconds was the primary performance metric

Platforms Under Test

The following table details the systems we tested and their sizing, configuration, and cost. All testing was performed in the AWS cloud and US East 1 region.

Configuration	**Snowflake Enterprise**	**Redshift Provisioned**	**Databricks SQL Serverless**	**Hive+Trino on Iceberg (Amazon EMR)**	**WarehousePG Single Node**
Cloud/Region	AWS US East 1	AWS US East 1	AWS US East 1	AWS US East 1	AWS US East 1
Version	9.37.1	1.0.162991	2025.35	Hive 3.1.3 Trino 476	7.2.1
Compute Configuration	Generation 2 Large (multi-clusters on and off)	8 nodes of ra3.4xlarge	Large	EC2 i4i.4xlarge 1 Primary 1 Core 10 Workers	VM: 96 cores, 1536GB RAM, 45 TB of NVMe
Storage	Native	Managed	Delta	Iceberg+NVMe	NVMe
Total Hourly Cost	**$32.40**[4]	**$26.08**[5]	**$28.00**[6]	**$22.65**[7]	**$25.44**[8]

Table 1. Platforms under test and their configurations

[4] See https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf for Snowflake pricing
[5] See https://aws.amazon.com/redshift/pricing/ for Redshift pricing
[6] See https://www.databricks.com/product/pricing/product-pricing/instance-types for Databricks pricing
[7] See https://aws.amazon.com/emr/pricing/ for EMR pricing
[8] Contact EnterpriseDB for WarehousePG pricing

Test Results

Concurrency Workload (5 users, TPC-DS queries)

The following table shows the query slowdown from 1 user to 5 users for a concurrent run of 99 TPC-DS queries, attained by dividing the 5-user time by the single-user time.

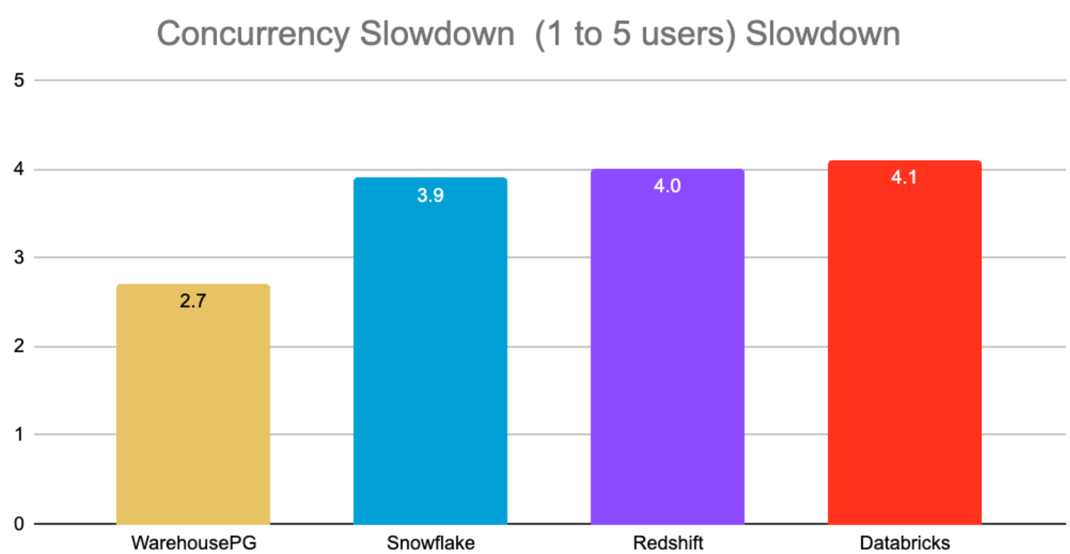

Figure 1: Concurrency slowdown (1 user to 5 users)

Platform	Compute Configuration	Slow Down
WarehousePG	Single Node with 96 cores, 1.5TB RAM, 45 TB of NVMe SSDs	**2.7x**
Snowflake	Generation 2 Large Warehouse	**3.9x**
Redshift	Provisioned cluster with 8 nodes of ra3.4xlarge	**4.0x**
Databricks	Serverless Large SQL Warehouse	**4.1x**

Table 2. Concurrency slowdown (1 user to 5 users)

The slow-down comparison reveals interesting trade-offs. WarehousePG, running with segments split across 96 cores, 1.5 TB of RAM, and 45 TB of NVMe storage, has the lowest slow-down factor at 2.7×, indicating predictable performance for this workload due to its raw compute and memory.

In contrast, cloud-only platforms—Snowflake Gen 2 (3.9×), Redshift with an 8-node ra3.4xlarge cluster (4.0×), and Databricks Serverless (4.1×)—experience higher slow-downs. While Snowflake slightly outperforms Redshift and Databricks, all three show similar scaling behavior under heavier workloads.

Their somewhat higher slow-downs reflect the coordination and planning overhead inherent in automated, cloud-only systems, which can make simpler architectures a better fit for certain workloads. At the same time, cloud-only platforms provide the elasticity and fault tolerance required for complex analytical workloads such as predictive modeling, large-scale data science, cross-data science, and cross-dataset joins.

Depending on the workload and operational priorities, either approach can be appropriate. Systems like WarehousePG excel at providing consistent performance for routine, high-concurrency analytics—such as operational reporting, dashboard queries, and real-time BI—where predictability is more important than sub-second latency, while cloud-only systems offer broader scalability for the most demanding workloads, such as ad-hoc exploratory analysis, machine learning feature engineering, and multi-tenant analytics environments.

Annual Cost

The following are the total annual cost of the platforms under test.

For Snowflake multi-cluster, we assumed 20 hours per week of heavy use where it would be scaled out to 3 multi-cluster warehouses (1,040 hours per year). The rest of the year, it is scaled down to the default single warehouse. We assumed 10TB of raw data compressed to 2.67TB.

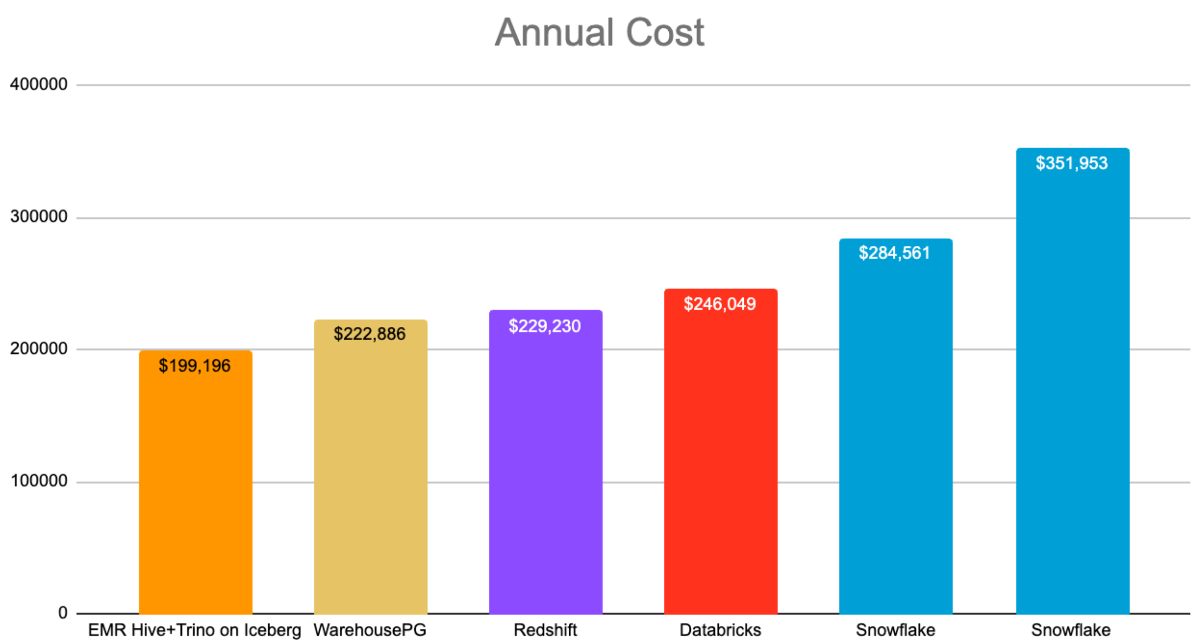

Figure 3: Annual cost of platforms tested

Platform	Compute Configuration	Annual Cost
EMR Hive + Trino on Iceberg	EC2 i4i.4xlarge with 1 Primary, 1 Core, 10 Task Workers	$199,196
WarehousePG	Single Node with 96 cores, 1.5TB RAM, 45 TB of NVMe SSDs	$222,886
Redshift	Provisioned cluster with 8 nodes of ra3.4xlarge	$229,230
Databricks	Serverless Large SQL Warehouse	$246,049
Snowflake	Generation 2 Large Warehouse (No Multi-clusters)	$284,561
Snowflake	Generation 2 Large Warehouse (Multi-clusters max=3)	$351,953

Table 4. Annual cost of platforms tested

The cost comparison highlights the trade-offs between different data warehouse platforms.

WarehousePG ($222,886 annually) and EMR Hive + Trino ($199,196) are the most cost-efficient, while Snowflake multi-clusters is the most expensive ($351,953), reflecting the price of scaling for maximum performance and throughput.

Though inexpensive, Hive is only built for batch processing while Trino excels at interactive queries. This option can also cause performance and consistency issues. We were unable to complete the concurrency workload on EMR Hive + Trino on Iceberg due to insufficient memory errors.

For these standard business queries, a scaled Snowflake deployment costs approximately 58% more than WarehousePG. This represents an opportunity to optimize costs, potentially saving over $129,000 per year per instance by aligning always-on analytics workloads with WarehousePG and its core-based pricing model.

Overall, cloud platforms deliver high performance, elasticity, and scalability for large, performance-sensitive workloads, justifying their premium for mission-critical use cases, whereas WarehousePG provides a consistent and predictable performance profile at a fraction of the cost, making it attractive for routine, high-concurrency analytical operations.

Conclusion

In today’s market, many enterprises no longer want their BI strategy constrained by an unpredictable consumption tax or perpetual rental model. In this “unbundled” DIY era, we’ve observed a “broken meter” effect: as your user base expands and AI agents proliferate, the outcome isn’t merely a larger data warehouse invoice. It becomes a performance ceiling that compels teams to ration analysis to avoid financial volatility.

WarehousePG offers superior predictability with the lowest slow-down factor (2.7×) for routine, high-concurrency analytics, while cloud-only platforms provide the raw performance and elasticity necessary for large-scale, complex data science workloads. Ultimately, the choice depends on operational priorities: sovereign architectures excel at consistent reporting and cost control, whereas cloud-native systems like Snowflake and Databricks dominate in ad-hoc exploratory analysis and massive scalability.

The trade-off is raw performance versus efficiency: cloud platforms maximize throughput for the most complex workloads, but WarehousePG can meet daily analytics needs while keeping infrastructure spend low.

An analytics team could effectively split workloads with mission-critical, large-scale, or infrequently run queries on a cloud-only platform to take advantage of performance and elasticity, while using WarehousePG for routine, high-concurrency analytics to ensure predictability and control spend. This “hybrid” approach is increasingly the reality in enterprises.

Transitioning to sovereign data infrastructure like WarehousePG can empower your developers and AI agents to execute intensive, real-time queries without hitting reliability and cost bottlenecks. This opens up much-needed space for enterprises to adopt agentic analytics alongside their daily business-critical operations, without spiraling costs or risking downtime for critical systems.

About EDB

EDB Postgres® AI (EDB PG AI) is the first open, enterprise-grade sovereign data and AI platform—secure, compliant, and scalable, on-premises and across clouds. Built on Postgres, the world's leading database, EDB PG AI unifies transactional, analytical, and AI workloads, enabling organizations to operationalize their data and LLMs while maintaining control over sovereign environments. EDB PG AI is supported by a global partner network and delivers up to 99.999% availability as well as hybrid management, Postgres-based data warehousing, and a built-in AI factory. As one of the most active contributors to the PostgreSQL project, EDB is deeply invested in the vitality of the global community. To learn more, visit www.enterprisedb.com.

About McKnight Consulting Group

Information Management is all about enabling an organization to have data in the best place to succeed to meet company goals. Mature data practices can integrate an entire organization across all core functions. Proper integration of that data facilitates the flow of information throughout the organization which allows for better decisions – made faster and with fewer errors. In short, well-done data can yield a better run company flush with real-time information... and with less costs.

However, before those benefits can be realized, a company must go through the business transformation of an implementation and systems integration. For many that have been involved in those types of projects in the past – data warehousing, master data, big data, analytics - the path toward a successful implementation and integration can seem never-ending at times and almost unachievable. Not so with McKnight Consulting Group (MCG) as your integration partner, because MCG has successfully implemented data solutions for our clients for over a decade. We understand the critical importance of setting clear, realistic expectations up front and ensuring that time-to-value is achieved quickly.

MCG has helped hundreds of enterprises with analytics, big data, master data management and “all data” strategies and implementations across a variety of industries and worldwide locations. MCG offers flexible implementation methodologies that will fit the deployment model of your choice. The best methodologies, the best talent in the industry and a leadership team committed to client success makes MCG the right choice to help lead your project.

MCG, led by industry leader William McKnight, has deep data experience in a variety of industries that will enable your business to incorporate best practices while implementing leading technology. See www.mcknightcg.com.

Disclaimer

McKnight Consulting Group (MCG) runs all its tests to strict ethical standards. The results of the report are the objective and unbiased results of the application of queries to the simulations described in the report. The report clearly defines the selected criteria and the process used to establish the field test. The report also clearly states the data set sizes, the platforms, the methods, etc. that were used. The reader is left to determine for themselves how to qualify the information for their individual needs. The report does not make any claims regarding third-party certification and presents the objective results received from the application of the process to the criteria as described in the report. The report strictly measures performance and cost and does not purport to evaluate other factors that potential customers may find relevant when making a purchase decision. This is a sponsored report. The client chose its configuration, while MCG chose the test, configured the databases and testing application, and ran the tests. MCG also chose the most compatible configurations for the other tested platforms. Choosing compatible configurations is subject to judgment. The information necessary to replicate this test is included. Readers are encouraged to compile their own representative configurations and test for themselves.