Unlock the potential of data lakehouses for AI and analytics

Data lakehouses are powerful platforms, but an enhanced approach to migration is required to unlock their full potential.

Highlights of the Quest solution

Quest provides solutions that facilitate a model then move approach that focuses on schema transformation, workload optimization and governance-first migration strategies. By aligning the logical source model with the target physical model, your organization can maximize the performance potential of your target data lakehouse instead of carrying forward old database constraints.

- Accelerate adoption with an automated, governed migration framework
- Ensure real-time, zero-downtime data movement from Oracle and PostgreSQL to your data lakehouse
- Optimize workloads for performance, cost efficiency and governance
- Enable end-to-end data governance with metadata tracking, lineage and compliance
- Monitor and optimize query performance to maximize cost savings

Understanding the barriers to data lakehouse success

A successful migration to a data lakehouse like Snowflake or Databricks isn't just about transferring data — it's about transforming how organizations manage, govern and optimize their data for future growth. While the platform providers sport a modern, cloud-native architecture, the differences between

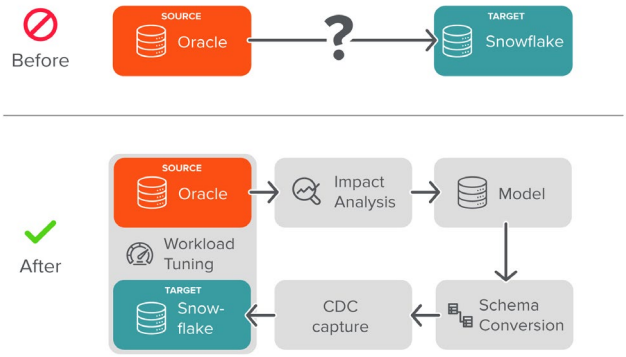

Assuring data visibility, integrity, governance and performance when stocking your cloud data lake.

A smarter way to ingest Oracle/Postgres data into Snowflake.

legacy databases like Oracle and PostgreSQL can introduce inefficiencies if not carefully addressed. Poorly structured migrations can result in you suffering from cost overruns, governance gaps and performance bottlenecks that limit the full potential of your data lakehouse for AI and analytics.

A common pitfall practitioners face is assuming that the target platform will automatically optimize performance for them. But migrating 'as is'— without reshaping schemas and tuning workloads for the architecture of the data lakehouse — means they're bringing old inefficiencies into a new system. The platform providers don't reconfigure legacy structures to fit their columnar, cloud-native design; that responsibility falls on your organization. Fortunately, Quest® can help.

Why now?

Surging data lakehouse adoption

A recent survey indicates that 67% of organizations plan to run the majority of their analytics on data lakehouses within the next three years, up from 55% today.

AI integration driving growth

The integration of AI technologies is accelerating the adoption of data lakehouses, with 85% of firms enhancing their AI readiness through this architecture.

Expanding cloud data warehouse market

Organizations modernizing their data ecosystems are driving cloud data warehouse adoption, pushing the market to $10.42 billion by 2026 (15% CAGR). The shift to cloud-native analytics and AI-ready platforms fuels this growth.

Surging demand for real-time data processing

The need for instant insights and AI-driven analytics is accelerating active data warehousing, expected to grow from $12.6 billion in 2023 to $24.2 billion by 2030. Businesses are prioritizing real-time, high-performance platforms like Snowflake and Databricks to stay competitive.

AI as a business imperative

Enterprises employing AI and big data analytics are 23 times more likely to acquire customers, 19 times more likely to be profitable and experience 7 times faster growth compared to their peers.

Quest solutions for data lakehouses

Overview

Many migrations fail because they don't account for the differences between legacy databases and the cloud-native architecture of a data lakehouse platform. This is where the Quest model then move approach and solutions come in — delivering a structured, five-step approach to ensure a smooth, risk-free migration.

Key capabilities

Many organizations take a "lift and shift" approach to a data lakehouse — only to find themselves battling performance bottlenecks, governance gaps and unexpected costs. Quest takes a different approach:

- **Assess with impact analysis —** Identify dependencies, risks and optimization opportunities before migration begins
- **Model first —** Reverse-engineer, optimize, and govern data schemas before migration
- **Move with confidence —** Ensure your data lakehouse runs at peak efficiency with proactive monitoring and cost controls
- **Monitor and optimize —** Ensure your data lakehouse runs at peak efficiency with proactive monitoring and cost controls
- **Govern for long-term success —** Maintain compliance, security and lineage tracking across all data assets

By aligning migration with business strategy and technical architecture of your data lakehouse, Quest delivers faster adoption, lower risk and better outcomes than traditional approaches.

Result: A continuously optimized data lakehouse environment that delivers maximum ROI — without the baggage of legacy inefficiencies.

The Quest five-step approach for data lakehouse migration

1. Impact analysis: Plan with precision

Before any migration begins, organizations need full visibility into their current data landscape. Without an upfront impact analysis, hidden dependencies, schema mismatches and inefficient workloads can introduce costly delays and errors. The Quest approach ensures that every migration is informed by a detailed assessment of workloads, data relationships and performance requirements before any data is moved.

By analyzing existing database structures, Quest helps organizations map their data lakehouse architecture, ensuring that data flows, schema

relationships and governance policies align with business goals. This minimizes surprises and prevents costly rework down the line.

2. Model: Optimize before you move

One of the biggest mistakes organizations make is assuming that legacy schemas can simply be lifted and shifted into a data lakehouse without modification. But the cloud-native, columnar architecture of a data lakehouse platform is fundamentally different from traditional relational databases like Oracle and PostgreSQL.

What works in a row-based database may cause serious inefficiencies in a data lakehouse — resulting in slow queries, excessive compute costs and storage waste. The platform providers don't restructure legacy schemas for you — you have to do it before the move.

erwin® Data Modeler by Quest ensures that organizations don't fall into this trap. By reverse-engineering legacy schemas and restructuring them for the architecture of your data lakehouse, Quest enables your organization to:

- Eliminate unnecessary joins and redundant structures that slow down performance
- Optimize indexing and partitioning to align with the distributed query processing of your data lakehouse
- Ensure efficient data compression and storage strategies to reduce compute costs
- Automate metadata and governance tracking so organizations maintain full visibility

By transforming the source logical model into a data lakehouse-ready physical model, your organization can unlock the true performance potential of the data lakehouse platform — rather than forcing a square peg into a round hole.

3. Move: Seamless, zero-downtime data replication *

Data migration is one of the most critical — and riskiest — phases of any data lakehouse adoption. Traditional "bulk move" migrations often result in downtime, data inconsistencies and business disruption.

SharePlex® by Quest* provides real-time, zero-downtime replication for Oracle and PostgreSQL databases, allowing organizations to keep legacy systems running while gradually transitioning to a data lakehouse platform. This approach enables a phased migration strategy, where organizations can test workloads in your data lakehouse before fully decommissioning their legacy environments.

The replication technology in SharePlex ensures that all transactions remain synchronized, preventing data drift and maintaining operational continuity during the migration process.

4. Monitor: Optimize performance and costs *

Once data is in a data lakehouse, organizations often struggle with query performance issues, runaway compute costs and inefficient resource allocation. Without proper monitoring, businesses risk spending more on the data lakehouse than anticipated while failing to maximize its analytical power.

Foglight® by Quest* provides real-time query monitoring, credit usage tracking and security insights to ensure organizations maintain both performance and cost control. By proactively identifying slow-running queries, excessive compute consumption and unused workloads, Foglight enables businesses to continuously optimize their data lakehouse environment and prevent unnecessary expenses.

5. Govern: Ensure compliance and security

A data migration is only successful if organizations can trust their data in the new environment. Without proper lineage tracking, metadata management and security enforcement, businesses risk compliance failures and governance breakdowns.

The governance-first approach of Quest ensures that all migrated data is classified, auditable and secure. Through automated metadata tracking, role-based access controls and regulatory compliance alignment (GDPR, HIPAA, and industry-specific mandates), Quest ensures that data lakehouse environments remain governed, compliant and secure from the moment they go live.

Key business outcomes

The right data lakehouse migration strategy can drive significant business value — but only when executed correctly. With the Quest model then move approach and solutions, organizations benefit from:

- **50% faster migration timelines —** By modeling before moving, organizations avoid time-consuming rework and accelerate data lakehouse adoption.
- **30% cost reduction on compute and storage —** Optimized schema designs and workload monitoring prevent unnecessary data lakehouse platform expenses.
- **2x faster query performance —** Performance-tuned schemas ensure high-speed data access for AI, analytics and business intelligence.
- **Enhanced security and compliance —** Automated governance tracking ensures data integrity, lineage visibility and regulatory compliance across platforms.

Final thought: A data lakehouse platform is powerful — but when paired with the right solutions, it becomes truly transformative!

Organizations that assume Snowflake or Databricks* will automatically handle performance optimization for them are setting themselves up for frustration. Moving into their platforms without restructuring schemas and tuning workloads is like trying to run a race in the wrong kind of shoes — it won't get you where you need to go efficiently.

Quest ensures that organizations don't just move to a data lakehouse platform — they move smartly. By transforming data structures before migration, ensuring real-time synchronization and continuously optimizing performance and governance, organizations can unlock the full potential of these platforms and avoid the pitfalls of a poorly planned transition.

Why choose Quest?

Successful adoption of a data lakehouse platform requires more than just migration tools — it demands a strategic, structured approach that ensures data integrity, performance and governance from day one.

Quest delivers:

- **End-to-end data lakehouse platform expertise —** A full-stack approach covering modeling, replication, performance tuning and governance*
- **Industry-leading technology —** erwin Data Modeler + SharePlex* + Foglight* provide a comprehensive migration and optimization toolkit
- **Optimized for AI and analytics —** Ensure your data lakehouse is AI-ready and primed for advanced machine learning workloads by managing and monitoring data products and data quality
- **Cost-effective cloud adoption** – Prevent runaway costs and right-size data lakehouse workloads for maximum efficiency

About Quest Software

Quest Software creates technology and solutions that build the foundation for enterprise AI. Focused on data management and governance, cybersecurity and platform modernization, Quest helps organizations address their most pressing challenges and make the promise of AI a reality. Around the globe, more than 45,000 companies including over 90% of the Fortune 500 count on Quest Software. For more information, visit www.quest.com or follow Quest Software on X (formerly Twitter) and LinkedIn.

* SharePlex and Foglight currently support Snowflake and are planned to support Databricks in a soon to be forthcoming release.