Veqtor8

Whitepaper

The ROI of Using the Denodo Platform alongside the Modern Data Lakehouse

Save $3.6 million in costs while delivering insights 4x faster with logical data management

Written By **Andrew Milroy, Chief Analyst, Veqtor8**

Table of Contents

Introduction

Enterprises have invested heavily in platforms like **Snowflake and Databricks** to unify both storage and compute. The challenge they face is delivering the right data to the right users in real time.

This white paper examines how adding a logical data management platform to a lakehouse-centric data architecture addresses this challenge. Drawing on customer interviews and an ROI model, it shows how the Denodo Platform enables organizations to:

- ✓ **Integrate and govern distributed data without requiring replication**
- ✓ **Reduce engineering complexity and infrastructure cost**
- ✓ **Deliver AI-ready, business-ready data faster and more consistently**
- ✓ **Enable self-service and real-time decision-making at scale**

By introducing logical data management with a common semantic layer to one's data architecture, Denodo can extend a data lakehouse into a complete, intelligent data platform that delivers trusted insights for AI, analytics, and operations, faster and at a lower cost.

Executive Summary

In today's data-driven environment, success depends not only on integrating and managing data but also on delivering it to end users with speed, trust, and relevance. **When paired with a modern data lakehouse, Denodo can provide the integration, governance, and delivery capabilities that data lakehouses lack.**

Working together, data lakehouses and Denodo:

- ✓ Accelerate AI adoption through governed access to diverse, high-quality data
- ✓ Enable business self-service without overburdening IT
- ✓ Optimize infrastructure by reducing redundant pipelines and compute overhead
- ✓ Support real-time decisions by exposing live operational data to BI tools, AI agents, and applications

Denodo's logical data management platform can bridge the gap between raw data storage and business-ready intelligence, enabling enterprises to extract more value from existing data lakehouse investments.

When using Denodo alongside a data lakehouse, customers realized...

75%

Reduction in data integration time

3-4x

Faster time to insights

345%

ROI over 3 years

$3.6M

Savings, payback in 6 ½ months

Delivering Trusted Data Products for the Enterprise

Enterprises need data that is timely, trusted, and ready for action, whether for AI models, business analytics, or operational decisions. ***Data lakehouses provide scalable storage and compute, but as mentioned, often struggle with the last mile of data delivery: integrating distributed sources, applying consistent governance, and enabling real-time consumption.***

Denodo extends data lakehouse architectures to close this gap. The use cases on the following page illustrate how.

DENODO PLATFORM ARCHITECTURE

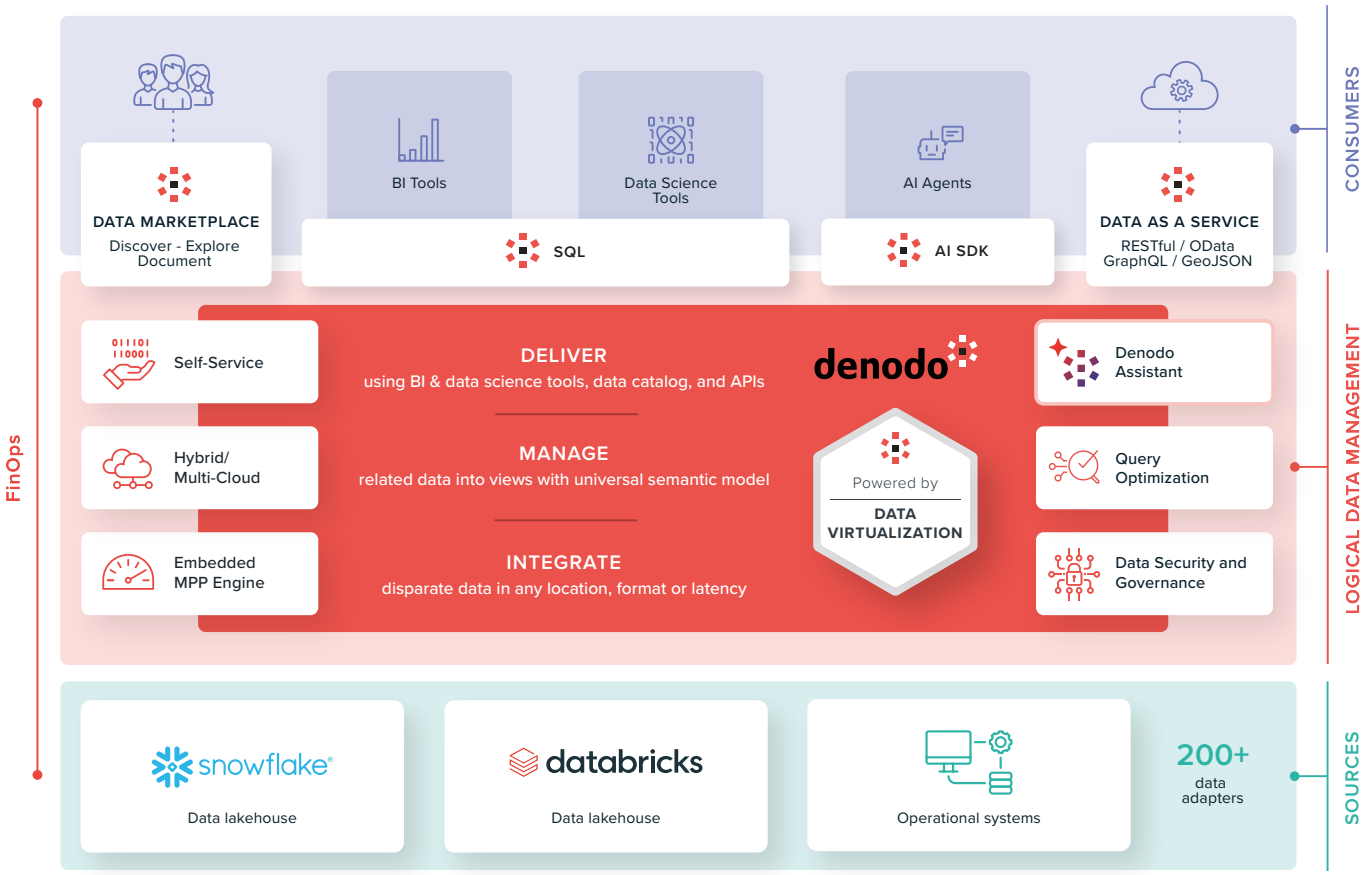

SCALING AI INITIATIVES WITH TRUSTED, REAL-TIME DATA

AI models require accurate, explainable, and governed data—but often this data is spread across on-premises, cloud, and SaaS systems. Denodo delivers this through a logical layer that unifies metadata and provides live, federated access without requiring replication. This accelerates GenAI and agentic AI projects while improving model accuracy, explainability, and compliance.

ENABLING BUSINESS SELF-SERVICE AT SCALE

Centralizing your data alone does not deliver democratization. Business users need intuitive, governed access without depending on IT. Denodo abstracts complexity into business-friendly views enriched with common semantics, enabling users to find, combine, and consume trusted data for analytics, applications, and composable experiences, on demand.

OPTIMIZING INFRASTRUCTURE AND REDUCING COST

Data movement, duplication, and over-provisioned compute inflate costs. Denodo minimizes replication and intelligently federates queries so data is processed where it lives, reducing storage and compute overhead while scaling infrastructure more efficiently.

POWERING REAL-TIME OPERATIONAL INTELLIGENCE

Dashboards, fraud detection, logistics, and customer engagement depend on live data, not ageing extracts. Denodo integrates real-time and historical data into unified views, delivered via APIs, BI tools, or event-driven systems. This enables faster, smarter operational decisions.

Use Case	Business Objective	Common Challenge	Role of Logical Data Management
Scalable AI Initiatives	Accelerate AI adoption and improve model performance and trust	Lack of real-time, explainable, governed data	Powers trusted, governed data delivery with semantic consistency
Business Self-Service	Empower business users	Inconsistent access and reliance on IT for data delivery	Delivers business-ready data through curated, governed views
Infrastructure Efficiency	Reduce compute/storage requirements and avoid duplication	Costly replication and compute	Minimizes infrastructure load by federation, query acceleration, and reducing the need for data movement
Real-Time Operations	Enable faster, data-driven decisions	Batch delays, inconsistent integration	Live access to operational and historical data from distributed systems

Table 1: Logical Data Management Use Cases

The Realized Benefits of Denodo Alongside a Data Lakehouse Architecture

Based on a series of structured interviews with data and technology leaders, **Veqtor8 identified six key categories in which organizations consistently realized value from combining Denodo with a modern data lakehouse architecture.** These categories were selected because they represent both strategic business outcomes and operational efficiencies that can be quantified across a wide range of industries.

Each category was broken down into variables that show measurable improvements when Denodo is implemented and operated as part of a data lakehouse architecture. These benefits stem directly from Denodo's capabilities, particularly its ability to abstract, integrate, and govern data, across multiple sources, in real time without requiring physical replication. The improvements listed below are drawn from real-world implementations and provide a practical framework for assessing Denodo's impact. These results reflect Denodo's ability to reduce effort and complexity while improving access, consistency, and responsiveness across enterprise data environments.

Benefit Area	Improvement Realized	Enabled By
Time-to-Insight	3-4x faster	Virtualized access and semantic abstraction layer
Implementation	78% reduction in cost	Streamlined setup with reusable connectors, reduced custom coding
Data Engineering Productivity	75% less effort	No-code integration and reusable views
Governance and Consistency	77% reduction in cost	Centralized metadata and policy enforcement
Business Agility	5-10x faster rollout	Unified access to real-time and historical data. Much faster time-to-value.

Table 2: The Realized Benefits of Denodo alongside a Data Lakehouse Architecture

These benefits highlight the synergies between Denodo and data lakehouses. **Denodo's logical data abstraction and federated access can unlock the full potential of a scalable lakehouse infrastructure.** This enables organizations to rapidly deploy AI and analytics solutions across silos without compromising governance or data quality.

"With Denodo, we can deliver governed, business-ready data across all our platforms without duplicating it. This means AI and analytics initiatives go live much faster, and we don't compromise on governance or data quality."

- Head of Enterprise Data Architecture at a Large Bank.

ROI Analysis

The ROI model underpinning this white paper is based on structured inputs derived from stakeholder interviews, industry benchmarks, and Denodo data. **It evaluates the financial impact of deploying Denodo as a logical data management layer on top of a modern data lakehouse architecture over a three-year period.**

The analysis reflects a composite organization constructed from insights provided by eight enterprises as well as data lakehouse providers. **The enterprises included five Denodo customers, and three large enterprises that are not using Denodo.** These firms shared detailed input on their integration timelines, cost structures, and AI project goals to help shape a representative financial model. Specifically, the ROI model compares a Denodo augmented data lakehouse architecture with an alternative approach leveraging native Snowflake and/or Databricks capabilities, often supplemented by widely used third party tools such as Fivetran for data integration, dbt for semantic views, and Collibra for governance. Such an approach can also involve the creation of custom pipelines, point-to-point integrations, and manual data preparation.

The ROI model reflects multi-year value realization, with implementation and enablement costs embedded in Year 1, followed by recurring benefits across Years 2 and 3. Financial analysis applies an 8% discount rate to calculate Net Present Value (NPV), compares Total Cost of Ownership (TCO) across scenarios, and computes ROI and payback period based on cumulative cost avoidance.

The ROI analysis begins by comparing the costs of achieving equivalent integration, governance, and delivery outcomes with and without Denodo in a modern data lakehouse environment. To ground these figures in reality, we constructed a composite enterprise based on inputs from five Denodo enterprise customers, three large non-customers, and leading data lakehouse vendors.

This composite represents a large, global organization (5,000–150,000 employees) operating across multiple regions, with an established Snowflake, Databricks, or equivalent data lakehouse investment. The environment already supports multiple use cases, AI/ML, analytics, self-service, and operational intelligence, ensuring that cost efficiencies and governance gains can be shared across the organization from day one.

The model is built around six outcome categories:

1. **Implementation** – Costs and time associated with initial deployment, including integration, configuration, enablement, and training
2. **Infrastructure and Tooling** – Reduction in infrastructure/tooling, legacy integration software, and manual data preparation
3. **Data Engineering Productivity** – Improvements in developer efficiency, reuse of data views, and reduction in redundant data wrangling
4. **Governance and Risk Reduction** – Better policy enforcement, fewer data quality issues, and simplified compliance
5. **Acceleration of Time-to-Insight** – Speeds up access to actionable insights, enabling timely decisions that drive revenue, efficiency, and competitive advantage
6. **Value Created from New Business Initiatives** – Unlocks and creates new value by providing rapid, governed access to data across silos, fueling scalable AI/ML and analytics

Each category incorporates real variables captured from client discussions, including:

- Time to integrate new data sources
- FTE effort reduction across data engineering and analytics teams
- Time-to-market acceleration for AI/ML projects and other data-driven initiatives
- Governance workload reduction due to unified metadata and policy controls

In the "without Denodo" scenario, enterprises rely on a patchwork of custom pipelines, point integrations, bespoke semantic layers, and separate governance tools. **This drives higher engineering effort, longer timelines, and greater infrastructure costs**. **In the Denodo-enabled scenario, these capabilities are consolidated into a single logical data layer, reducing duplication, avoiding unnecessary replication, and streamlining operations**.

The table below (Table 3) quantifies these differences across key outcome categories, highlighting where Denodo delivers the most substantial cost advantages for our composite enterprise.

Category	Cost Without Denodo (USD)	Cost With Denodo (USD)	Net Benefit (USD)	Cost Saving	Value Driver
Implementation (Year 1)	2,126,400	466,800	1,659,600	78%	Reduced manual integration and setup effort
Infrastructure and Tooling*	426,400	78,400	348,000	82%	Infrastructure/tooling reduction
Data Engineering Productivity*	601,200	150,600	450,600	75%	Time/effort saved
Governance, Consistency and Risk Reduction*	57,600	13,200	44,400	77%	QA and consistency uplift
Denodo Licensing*	0	200,000	-200,000	—	Denodo subscription cost
Total Investment (Year 1)	3,211,600	909,000	2,302,000	72%	—
Total Operational Cost* (Ongoing, excluding implementation)	1,085,200	442,200	643,000	60%	—

** Annual*

Table 3: Cost Benefits of Using Denodo – Composite Enterprise Model

While cost savings alone provide a strong justification, Denodo also generates measurable business value by enabling faster insights and unlocking entirely new initiatives that would otherwise be delayed or impossible. These benefits are particularly impactful in AI/ML, analytics innovation, and cross-domain decision-making, where timely, trusted data drives competitive advantage. The table below (Table 4) captures these incremental gains including revenue opportunities, operational improvements, and innovation capacity, resulting from Denodo's ability to deliver governed, real-time, reusable data products at scale. Again, these figures are drawn from our composite enterprise model.

> **"With Denodo, governance isn't a separate, manual exercise, it's built into how we deliver data. That's a big reason for choosing Denodo. Without it, we'd be stitching together controls across multiple tools, increasing both delivery time and compliance risk."**
>
> - Lead Business Analyst at a Global Bank

> **"A base view in Denodo takes minutes; even our most complex derived views rarely take more than a day. That's a fraction of the time we'd need with manual pipelines."**
>
> - Head of Data Architecture at a Global Manufacturing Company

Category	Value Without Denodo (USD)	Value With Denodo (USD)	Net Uplift (USD)	Multiplier	Value Driver
Accelerated Time-to-Insight*	95,000	400,000	305,000	3-4x	Faster delivery of insights
Value Created from New Business Initiatives*	367,200	928,800	561,000	2-3x	Enablement of new AI/ML and data-driven initiatives
Total Business Uplift (ongoing)*	462,200	1,328,800	866,600	2-3x	New value beyond cost savings: enabling new AI/ML& data initiatives and faster time-to-market

** Annual*

Table 4: Additional Business Benefits of Using Denodo – Composite Enterprise Model

Building on the comparisons in the above table, the financial summary below presents a consolidated view of the overall impact over a three-year period. **It shows how the efficiencies and business value enabled by the Denodo Platform, implemented alongside a data lakehouse, translate into significant financial returns when compared to achieving the same outcomes without Denodo.** This includes cumulative cost savings and business uplift, the net present value (NPV) of those benefits, and key performance indicators like payback period and ROI. For data lakehouse users, these figures provide clear, data-backed justification for Denodo investment.

Taken together, the key financial metrics below illustrate just how impactful Denodo can be from a business value perspective.

The combination of strong cost avoidance, short payback period, and a high ROI presents a strong case for moving forward with Denodo as a core part of any data lakehouse architecture.

These figures are representative of our composite enterprise.

> “With Denodo, our time-to-insight has gone from months to days for key high-value use cases.”
>
> - Head of Technology Procurement for a Global Mining Company

Metric	Value (USD)
3-Year Total Cost of Ownership – Without Denodo	$5,382,000
3-Year Total Cost of Ownership – With Denodo	$1,793,400
Cost Avoidance	$3,588,600
3-Year Total Business Benefits	$2,599,800
3-Year Net Present Value (NPV)	$5,427,052
Breakeven Point (Cost)	$1,571,815
3-Year ROI (Value / Cost)	345%
Payback Period	6.42 months

Table 5: Key Financial Metrics – Composite Enterprise Model

Strategic Benefits Beyond the Numbers

While the quantified ROI provides a compelling business case, Denodo's impact on AI initiatives is particularly important. As AI/ML models become key differentiators, Denodo keeps data accessible, consistent, trusted, and delivered in real time. This makes it an essential enabler of enterprise AI.

Some of the key strategic advantages reported by Denodo customers include:

Faster Experimentation and Innovation: Real-time data access reduces the lag between hypothesis and insight, enabling AI and analytics teams to iterate rapidly and test new ideas with minimal friction.

Improved Data Democratization: Denodo empowers business users with self-service access to governed, trusted data, reducing reliance on central IT and accelerating decision-making.

AI/ML Enablement: Denodo plays a foundational role in modern AI pipelines, feeding high-quality, consistent data into machine learning and large language models (LLMs), and promoting trust in automated decisions.

Architectural Flexibility: By abstracting underlying sources, Denodo reduces vendor lock-in and enables organizations to evolve their data platforms without breaking data delivery chains.

> **"Denodo has transformed the way our data science team experiments. Instead of spending weeks preparing datasets, they now focus on training models and generating business value."**
>
> - Lead Business Analyst at a Global Bank

These benefits illustrate that Denodo can be a strategic enabler of agility, innovation, and long-term competitiveness.

AI Impact Area	Denodo's Contribution
Data Availability for ML	Unified access to diverse structured and unstructured sources
Feature Engineering Efficiency	Live query access reduces prep time and increases reuse
Model Training and Retraining Speed	Real-time pipelines eliminate the need for staging layers
Explainability and Model Trust	Semantic consistency and lineage tracking for auditability
Model Deployment Agility	Integration via APIs, views, and governed endpoints

Table 6: Denodo's Strategic Role in AI Enablement

> **"We use Denodo to feed our fraud models with real-time transaction data across systems. This shortens deployment by months."**
>
> - Head of Technology from a Global Bank

Conclusion

When paired with a modern data lakehouse, Denodo can deliver significant improvement in both business agility and return on data investments. It bridges the gap between raw data storage and business-ready intelligence, enabling faster AI deployments, stronger governance, and real-time responsiveness across the enterprise.

Our composite enterprise model shows:

- **$3.6M in cost avoidance**
- **345% ROI over three years**
- **Payback in 6.42 months**

Denodo can achieve this without replacing existing platforms and it works seamlessly with Snowflake, Databricks, and others, to unify data delivery and governance while minimizing infrastructure drag.

In a market where trusted, governed access to data is as critical as storage and compute, Denodo's logical data management platform addresses a key architectural gap in modern data lakehouse strategies. By enabling faster AI deployment, broader self-service, and more efficient use of infrastructure, it can provide a pragmatic path to extracting greater value from existing data investments.

Methodology

The findings in this report are based on structured research conducted by Veqtor8, including in-depth interviews with five Denodo enterprise customers, three large enterprises not currently using Denodo, and leading data lakehouse platform vendors. The objective was to quantify the technical, operational, and financial impact of deploying Denodo alongside a modern data lakehouse architecture, compared with achieving equivalent outcomes using alternative tools and integration methods.

As part of this process, the analysis included direct cost and benefit benchmarking between Denodo-enabled architectures and alternative approaches using native Snowflake and Databricks features, as well as third-party tools for data integration, virtualization, and governance. This ensured the financial and operational comparisons were grounded not only in vendor data but also in real-world alternative design patterns.

The study compares two primary scenarios:

1. **Denodo-enabled architecture** – A modern data lakehouse enabled with Denodo's logical data management platform.

2. **Alternative approach** – A modern data lakehouse using a combination of alternative integration, governance, and delivery tools, without Denodo.

A composite enterprise was defined to represent the organizations interviewed and to provide a realistic basis for financial modelling:

- Size: 5,000–150,000 employees
- Revenue: Greater than USD $1B annually
- Geography: Global or multi-region operations
- Industry: Representative of sectors with high data volume and complexity (e.g., banking, manufacturing, logistics, resources)
- Data environment: Hybrid, with multiple on-premises, cloud, and SaaS sources integrated into a central lakehouse

ASSUMPTIONS

The model provides a view of how a typical global enterprise integrates, manages, and delivers data. However, it is important to highlight key limitations and assumptions, as follows:

- **Multiple Use Cases in Scope** – The cost efficiency gains modeled in this study assume that Denodo is implemented to support multiple use cases in production from the outset. This is reinforced by the fact that all enterprises in scope already have substantial data lakehouse investments (e.g., Snowflake, Databricks), meaning the underlying infrastructure, data assets, and governance frameworks are already in place to support broad deployment across AI, analytics, self-service, and operational needs. This enables infrastructure, governance, and productivity benefits to be shared and scaled from day one.

- **Definition of New Business Uplift** – "New Business Uplift" refers to new value creation above and beyond the baseline cost efficiencies—such as enabling new AI/ML initiatives, accelerating time-to-market for products and services, and creating revenue opportunities that would not be achievable without Denodo.

- **Top-Loaded Implementation Costs** – Implementation and enablement costs are treated as largely occurring in Year 1. This approach reflects the real-world pattern where the majority of integration, configuration, training, and onboarding efforts are concentrated at deployment. While some ongoing tuning and incremental onboarding occurs, these costs are relatively small compared to initial setup, making top-loading a more accurate reflection of actual investment patterns.

- **Conservative Benefit Estimates** – Gains in productivity, time-to-insight, and AI enablement are modeled at the lower end of observed ranges from interviews to avoid overstating impact. Interviews revealed significant variation in realized benefits across enterprises, driven by differences in data maturity, industry, use cases, governance practices, number of integrated sources, and number of views. For the composite organization, we selected benefit values that sit within the lower-middle range of reported outcomes, so they remain realistic for most large enterprises while avoiding best-case bias. The values were also validated by sharing them with four of our respondents.

LIMITATIONS

The ROI model is based on a representative composite organization and may differ from individual enterprise results depending on scope, maturity, use case, and operational context. This analysis includes direct cost and benefit benchmarking between Denodo-enabled architectures and alternative approaches using native Snowflake or Databricks features and third-party tools for data integration, virtualization, and governance. However, results for each approach may still vary according to enterprise architecture, deployment scale, and the maturity of the surrounding data environment. The comparative findings are based on representative scenarios and may not capture all nuances of alternative solutions or emerging features introduced after the study period. The model does not capture intangible or difficult-to-quantify benefits such as employee satisfaction, cultural impact, or competitive signalling.

About **Andrew Milroy**

Andrew Milroy is the chief analyst at Veqtor8, a technology advisory firm. He is one of the world's most experienced industry analysts, with an accomplished track record spanning the U.S. Europe and Asia Pacific. His work focuses on helping organizations unlock value from emerging technologies, with a particular emphasis on AI, data management, cloud platforms, and cybersecurity.

Andrew has held leadership roles with Frost & Sullivan, Omdia, IDC, and NelsonHall and is often invited to keynote, chair, and moderate at major technology events.

Veqtor8

Veqtor8 is an independent advisory and insights platform focused on AI, data management, cybersecurity, and cloud infrastructure. We help organizations become AI-ready, optimize their data strategies, and strengthen digital resilience. Our approach is practical and independent, drawing on real-world projects and global best practices to deliver clear, actionable guidance. Through research, executive workshops, reports, and presentations, we equip leaders to manage digital risk, innovate with data, and harness AI responsibly.

Visit: veqtor8.com | **Email:** info.enquiries@veqtor8.com

denodo

Denodo is a leader in data management. The award-winning Denodo Platform is the leading logical data management platform for transforming data into trustworthy insights and outcomes for all data-related initiatives across the enterprise, including AI and self-service. Denodo's customers in all industries all over the world have delivered trusted AI-ready and business-ready data in a third of the time and with 10x better performance than with lakehouses and other mainstream data platforms alone.

Visit: denodo.com | **Email:** info@denodo.com