Veqtor8

ROI RESEARCH BASED ON ENTERPRISE INTERVIEWS

The ROI of an AI Data Layer for Agentic AI

Turning agent deployments into business value faster, at lower cost and with consistent governance

Written by *Andrew Milroy, Chief Analyst*, Veqtor8



Contents

01 Introduction

02 Executive Summary

03 Deploying Trusted Agents Fast, at Scale and Under Governance

04 Modeled Benefits of Denodo's AI Data Layer

05 ROI Analysis

06 The Baseline Architecture and the Alternatives

07 Strategic Outcomes Beyond the Numbers

08 Conclusion

09 Methodology

Andrew Milroy
Chief Analyst, Veqtor8

01

Introduction

Enterprises are moving from experimenting with generative AI to deploying autonomous and semi-autonomous agents that work with live enterprise data. These agents retrieve records, call tools, update systems and complete multi-step workflows with limited human intervention. Two hyperscalers interviewed for this research reported that around twenty to twenty-five percent of enterprises were running agents in production, with most organizations still in pilot. The use of agents in production is increasing, although the number deployed and the complexity of the work they perform vary significantly between organizations.

Agents are only as effective as the data they can reach. To act reliably, an agent needs access to live, governed and semantically consistent enterprise data that is typically scattered across cloud platforms, on-premises systems, SaaS applications and operational databases. MCP servers can provide agents with access to these systems, but they do not by themselves create consistent enterprise context across them. Without a shared layer, much of that work must be repeated as new agents are deployed. In regulated industries a single production agent may touch thirty to fifty enterprise sources, and the value of a shared data layer rises with the number of systems a workflow has to orchestrate.

This paper compares the blended architecture most enterprises run today against a governed AI Data Layer sitting between agents and the underlying systems. In this research, we considered a governed AI Data Layer provided by Denodo. There is no single default architecture for enterprise agentic AI. Instead, the research identifies three architectural archetypes. These are lakehouse-centric, hyperscaler-centric and bespoke integration. Most organizations combine elements of all three, using the native AI and data capabilities of a primary cloud or lakehouse platform for some workloads, while custom integration, connectors and other platforms provide access to data and applications elsewhere. This blended architecture is therefore the baseline used in this paper. Section 6 compares the AI Data Layer with these archetypes and the combinations enterprises realistically use today.

The comparison is reflected in a composite enterprise ROI model, informed by structured interviews with enterprise practitioners, technology partners and hyperscalers, to quantify the impact of an AI Data Layer.

290%	68%	2.9x	3x
MODELED 3-YEAR ROI	COST AVOIDANCE	MORE BUSINESS VALUE GENERATED	FASTER AGENT DEPLOYMENT

Modeled over three years for a composite enterprise built from the seven organizations interviewed for this research.

02

Executive Summary

In this paper, an AI Data Layer refers to a logical layer through which agents can access and work with data across multiple enterprise systems.

As enterprises move from a handful of AI pilots to populations of agents, the cost of giving each agent reliable access to enterprise data becomes a material part of the AI business case. That cost grows with the number of agents, the number of data sources and the number of integrations required to keep them current.

Much of the economic benefit of an AI Data Layer comes from reuse. A source is connected, profiled and assessed once, while the integrations, semantics and governance applied through the layer become reusable across agents. A data product is built once and can then be used by every agent entitled to it. This reduces the amount of work that must be recreated for each new agent, driving faster deployment, greater integration agility and governance leverage while lowering the marginal cost of scaling agentic AI. A governed AI Data Layer removes most of the data-related work required to get a new agent into production, so agents reach the business faster and more of them can be deployed across the organization. Part of that value returns in the form of time released to the people working alongside agents. Another part of that time can be applied to work the business had no capacity to attempt before.

Every data point in this paper is derived from a composite enterprise, a single organization aggregated from the seven large enterprises interviewed for this research and set out in full in Section 9. This composite enterprise grows from managing 100 agents in Year 1 to 900 in Year 3 across 100 enterprise data sources and is measured against the baseline architecture described in Section 6. Two of the six categories in the model measure value created and four measure cost avoided, held visibly separate so that released capacity is not mistaken for cash saved.

FOUR FINDINGS DRIVE THAT RESULT:

01

$3.5
MILLION

Deployment velocity. With an AI Data Layer, agents reach production around three times faster, creating $3.5 million in modeled business value over three years.

02

$5.1
MILLION

Value created. Reduced data preparation effort, productive time returned to agent users and new work the business previously could not resource together add $5.1 million, the largest single benefit line.

03

$10.5
MILLION

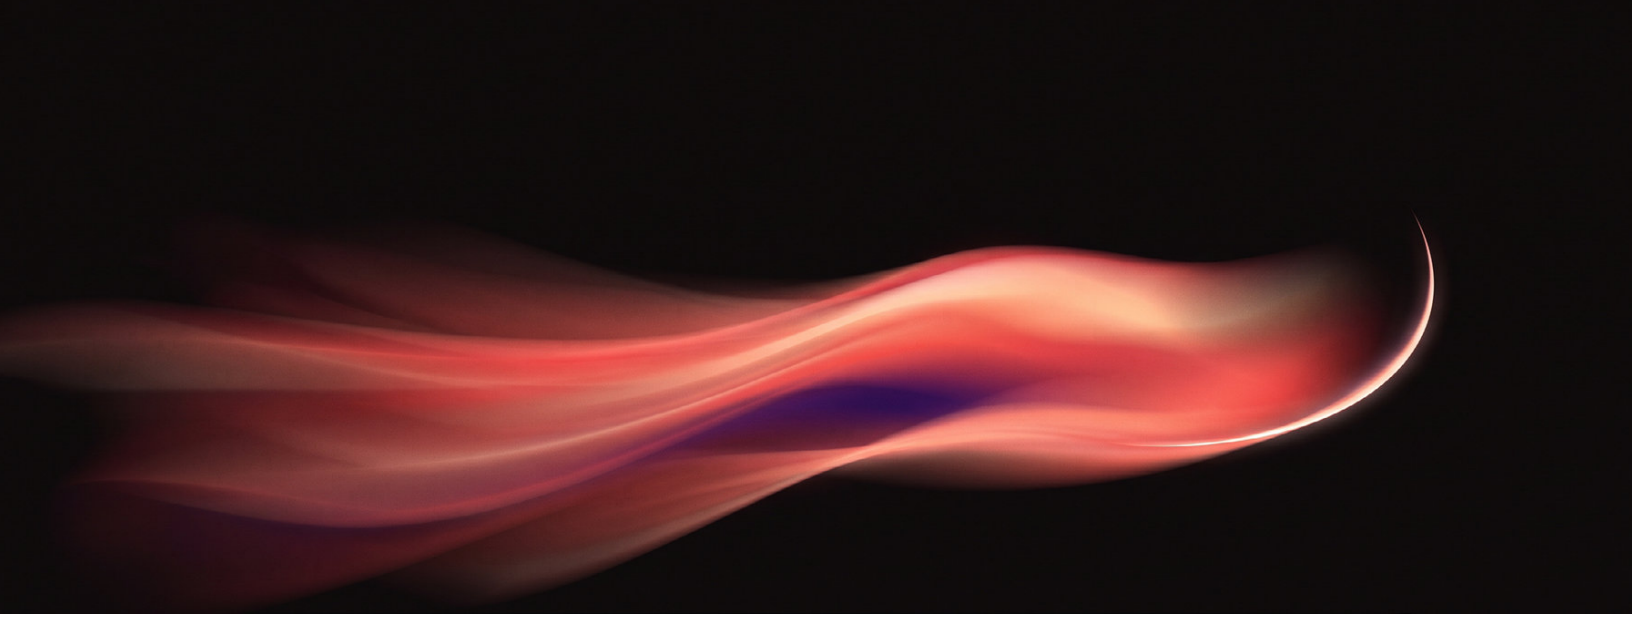
Cost avoided. Accuracy and trust, integration agility, governance and AI operating cost together avoid $10.5 million over three years, net of the AI Data Layer subscription.

04

↓68%
PERCENT

AI operating cost. The tokens needed for agents to ingest the necessary context, per task, falls from approximately 18,000 tokens to approximately 5,760, producing a 68 percent reduction in AI operating cost.

Individual outcomes vary with the maturity of the existing estate.



03

Deploying Trusted Agents Fast, at Scale and Under Governance

Most enterprises are now deploying agents. The challenge is increasingly the speed at which organizations can put reliable agents into production and scale them. Rapid deployment on its own is insufficient. An agent deployed quickly against data it cannot trust creates exposure rather than value, and an estate of agents assembled without common controls becomes increasingly difficult to govern as it grows.

Enterprises therefore need four things at once: Agents that reach production quickly, agents that act on data the business trusts, agents that run at a cost which scales with usage, and governance that holds as agent numbers grow. In the baseline architecture those four requirements pull against each other, because each is satisfied separately for every agent.

Speed depends on data readiness rather than model capability

Interviewees were consistent on this point. A senior AI leader at a global bank reported deployment times ranging from around ten days for simple agents to two months for complex deployments. He also noted that ready-made data products can reduce deployment effort from months to days or weeks. An AI Data Layer accelerates this process by speeding the creation of reusable data products and enabling their reuse across agents.

An APAC building materials group made the same observation from a different angle, noting that connecting any individual system is straightforward once integrated, and that the difficulty lies in orchestrating several systems into a coherent workflow.

> *"Data products with semantics and context make it easier and faster. It reduces deployment by months. It can reduce deployment to days or weeks."*
>
> SENIOR AI LEADER, GLOBAL BANK

Model capability alone does not determine reliability

An agent working with incomplete, inconsistent or outdated data can make an incorrect recommendation or take the wrong action. At a global custody and asset servicing bank interviewed for this research, agents repair missing routing and country codes on payment instructions, which requires resolving a reference value correctly against the authoritative source rather than retrieving a plausible one. The reliability of that task rests on the semantic consistency of the data beneath it, placing the burden on the data layer rather than the model.

> *"Deploying Denodo on the AWS infrastructure has allowed us to build an intelligent finance agent. By providing simple feedback, we're continuously training the model to quickly learn our unique business rules, traverse multiple tables and deliver accurate answers."*
>
> SENIOR FINANCE DIRECTOR, GLOBAL MANUFACTURING GROUP

Cost accumulates in a second layer beneath the agents

Automation at scale creates its own operating cost. An APAC telecommunications group runs approximately 100,000 localized models across its network in addition to its business workflow agents. Interviewees reported that production AI costs increasingly depend on model pricing, the size of the prompts and the contextual payloads supplied to agents. Controlling the volume of data work agents perform at runtime is therefore central to controlling that cost.

100,000 localized models

At this scale, controlling the volume of data work agents perform at runtime is central to controlling AI operating cost.

Governance load grows with the estate, not with the use case

Each new agent, workflow or data source adds integration and maintenance effort. Interviewees identified connector proliferation and MCP sprawl as emerging operational and security problems. This burden continues after deployment as data, schemas, pipelines and access policies change, requiring ongoing maintenance and testing across the agent estate.



Sources, data products and reuse

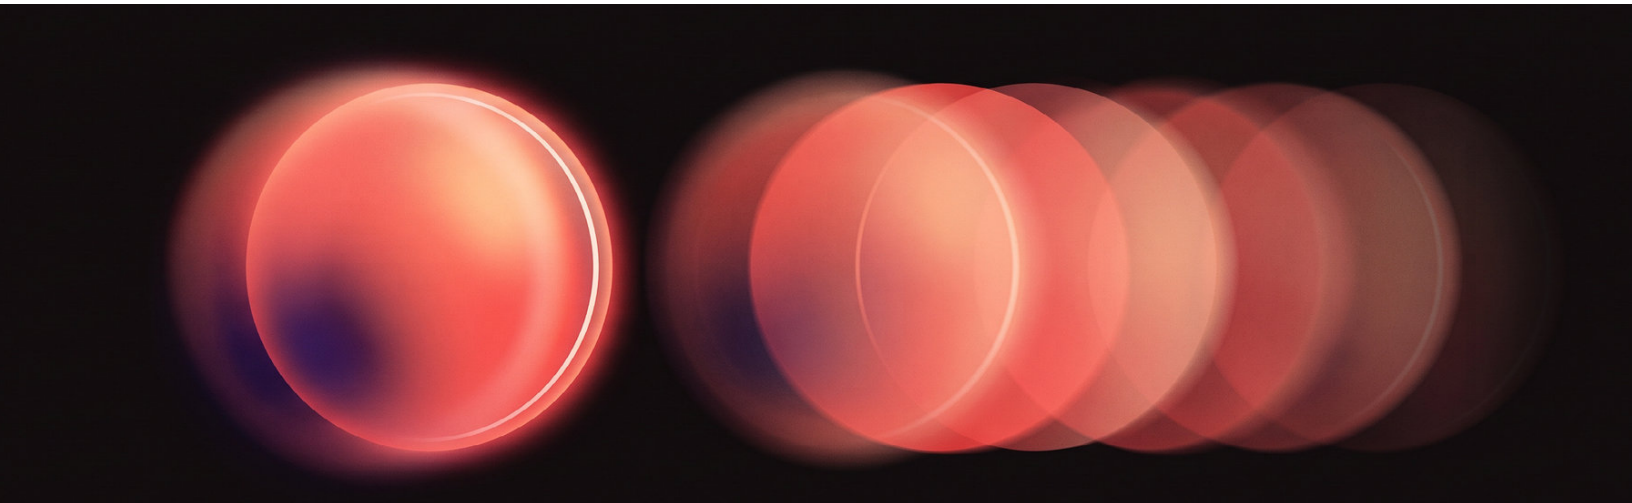
Agents access enterprise data through tools exposed through MCP servers, APIs or equivalent protocols. These can provide access to governed data products that combine reusable data and business meaning. A source is connected, profiled and assessed once, while data products built from it can be used by multiple agents.

This reuse is central to the economics. Instead of rebuilding data access for each new agent, a governed AI Data Layer allows agents to use existing data products across multiple underlying sources. This reduces repeated integration and maintenance as the number of agents grows.



"If you use Denodo as the universal data delivery layer across multiple systems like Salesforce and Snowflake, agents can come to one place to get the data they need."

GENERAL MANAGER, DATA AND ANALYTICS, APAC SYSTEMS INTEGRATOR



04

Modeled Benefits of Denodo's AI Data Layer

Speed and trust are the enablers, capacity is the outcome

Faster, more accurate access to trusted enterprise data enables an organization to deploy more agents and increase the work they can deliver. The model quantifies this as business value added across the composite enterprise over three years, comparing the baseline with an AI Data Layer.

Category	3-year baseline value	3-year value with an AI Data Layer	Net uplift	Multiple
Agent Deployment Velocity	$1,785,056	$5,314,400	$3,529,344	**3.0x**
Value Created	$2,802,800	$7,898,800	$5,096,000	**2.8x**
Total business value added	**$4,587,856**	**$13,213,200**	**$8,625,344**	**2.9x**

Table 1: Modeled three-year business value added by category, composite enterprise model

Value Created is the largest single benefit line at $5.1 million over three years. It combines productive time returned to people working alongside agents with new work the organization previously lacked the capacity to resource.

A global bank interviewed illustrates the productivity opportunity, with a team of five potentially delivering the output of forty. It also expects supervision-by-exception to reduce ongoing human oversight by around forty percent. A lending organization cited by one hyperscaler reduced human review by roughly seventy-five percent while increasing throughput approximately fourfold.

A global chemical and crop science group put its first production agent into service on a governed layer, giving supply chain staff a single route to information previously spread across separate systems and spreadsheets.

> *"We have built agents on top of Denodo, and the first one we put into production is a customer service assistant. Before that, our people had to go into many different systems and many different spreadsheets to answer one question. Now they ask the question in natural language and the answer comes back from one place."*
>
> DATA ARCHITECTURE LEAD, GLOBAL CHEMICAL AND CROP SCIENCE GROUP

Four further categories measure cost avoided, by comparing baseline costs with those under a governed AI Data Layer.

01 AGENT ACCURACY AND TRUST

The ongoing effort to review and correct agent outputs, and to investigate and recover from errors.

02 AGENT INTEGRATION AGILITY

The effort to set up and maintain MCP servers, pipelines and connectors, and to absorb change as source data and context evolve.

03 GOVERNANCE AND RISK REDUCTION

The effort to administer policy and access controls, and to prepare for and remediate audit findings, across a growing agent population.

04 AI COST OPTIMIZATION

The token and infrastructure cost of running agents, driven by how much context each agent task has to discover, retrieve and retry.



Category	3-year baseline cost	3-year cost with an AI Data Layer	Net benefit	Cost reduction
Agent Accuracy and Trust	$2,639,000	$837,200	$1,801,800	68%
Agent Integration Agility	$5,773,040	$1,228,209	$4,544,831	79%
Governance and Risk Reduction	$2,548,000	$618,800	$1,929,200	76%
AI Cost Optimization	$4,374,000	$1,399,680	$2,974,320	68%
AI Data Layer Subscription	$0	$750,000	($750,000)	n/a
Total cost avoidance	**$15,334,040**	**$4,833,889**	**$10,500,151**	**68%**

Table 2: Modeled three-year cost avoidance by category, using a composite enterprise model

The model shows benefits increasing as agentic AI adoption scales. As organizations deploy more agents across more workflows and data sources, the value of shared integration, governance and data access increases with them. AI cost benefits also grow with task volume, making the overall value of the shared data layer progressively larger as agentic AI use expands.

Production deployments demonstrate the effect at the workload level. A global manufacturing group using Denodo for an intelligent finance agent reports 95 percent accuracy in natural-language-to-SQL mapping and a reduction in reporting time from more than four hours to under five minutes.

Estimates gathered from interviews support both the direction and the approximate magnitude of the modeled results.

01

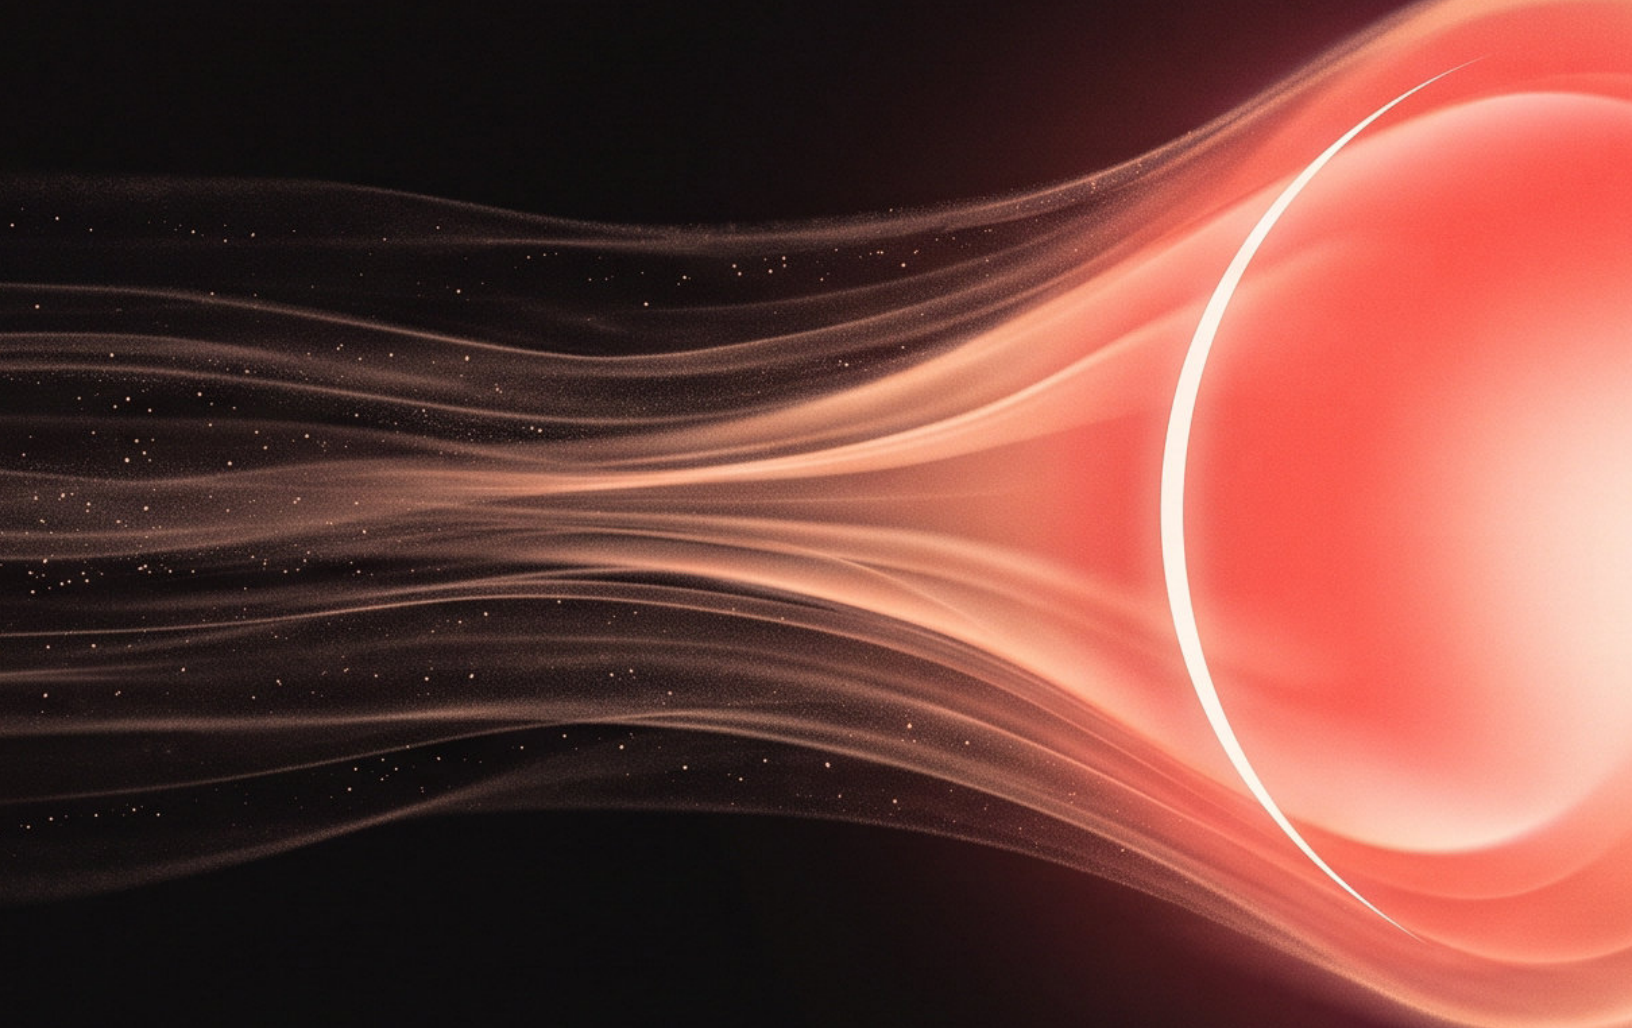
An executive at the APAC telecommunications group estimated approximately three times faster agent deployment, two to three times greater integration agility, around three times the governance benefit and two to three times lower token consumption if the layer were extended across the agent estate.

02

At the custody and asset servicing bank, uplift of approximately three to one was reported for agent deployment and four to one for integration agility.

03

An APAC systems integrator reported a customer onboarding a data source through the layer in approximately one week, against an estimated four to six weeks to reach the equivalent position through a lakehouse platform.



05

ROI Analysis

The model compares two scenarios for the composite enterprise described in Section 9 over three years. The baseline is the mixed architecture outlined in the Introduction and detailed in Section 6. The alternative scenario places a governed AI Data Layer between the agents and the same underlying estate.

The composite enterprise grows from 100 agents in Year 1 to 300 in Year 2 and 900 in Year 3, with each agent completing an average of 100,000 tasks a year across an estate of 100 enterprise data sources. Engineering effort, governance activity and AI operating costs all rise as deployment scales, and the analysis measures the impact of an AI Data Layer on those costs.

Token economics

01

Components

Token cost is modeled from the context an agent needs to complete a task, rather than from a flat cost per task.

02

Levers

Better metadata and semantic models can reduce the effort agents spend discovering relevant data.

03

Query delegation

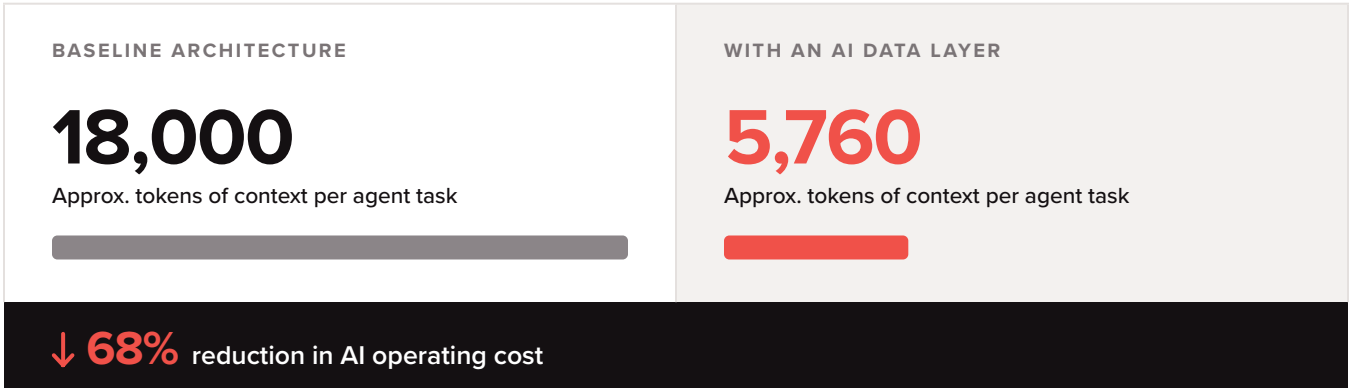
For structured data, one approach is to delegate the query to the data layer rather than retrieve large numbers of rows for the agent to process.

04

Modeled effect

the model produces a 68 percent reduction in AI operating cost.

How the AI Data Layer changes agent economics:



Components. Token cost is modeled from the context an agent needs to complete a task, rather than from a flat cost per task. Four components are counted: tokens spent discovering where relevant data sits, tokens consumed by the raw payload retrieved, tokens consumed by data retrieved that is irrelevant or outside policy, and tokens spent retrying failed attempts. Each is multiplied by a context compounding factor reflecting how many turns that context stays live, since context accumulates across the steps of an agentic workflow rather than being paid for once.

Levers. Several architectural choices can reduce this context. Better metadata and semantic models can reduce the effort agents spend discovering relevant data. Querying and filtering data before it reaches the model can reduce payload size and over-retrieval, while consistent definitions can reduce failed attempts and retries.

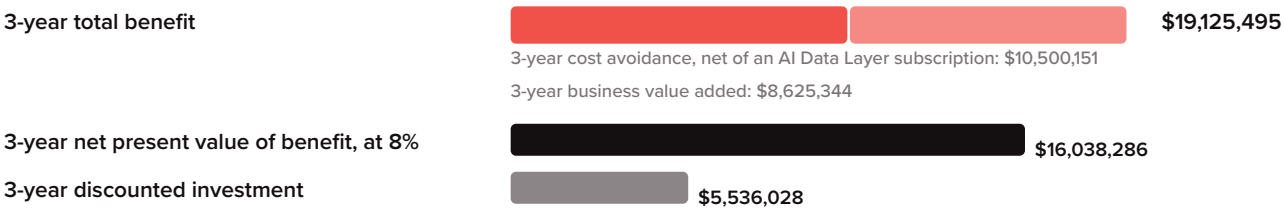
Query delegation. For structured data, one approach is to delegate the query to the data layer rather than retrieve large numbers of rows for the agent to process. Denodo describes its implementation of this pattern as Query RAG. Filtering, joining and aggregation take place before the result reaches the agent, reducing the amount of data entering the context window.

Modeled effect. In the baseline, the composite enterprise carries approximately 18,000 tokens of context per agent task once compounding is applied. With the governed AI Data Layer modeled here, the figure falls to approximately 5,760 tokens. Applied across the task volume of the composite enterprise, and including the cloud compute and storage overhead of running those workloads, the model produces a 68 percent reduction in AI operating cost.

Exclusions. Three deliberate exclusions make that figure conservative. The same input token price is applied to both scenarios, so every dollar of saving comes from consuming fewer tokens rather than from buying cheaper ones. Model routing and prompt caching are available to both scenarios and are not credited to the layer, even though one hyperscaler reported a customer reducing prompt payload by around eighty percent through caching alone. Infrastructure cost is also modeled separately from token cost, since the two are related but distinct, and several interviewees expect infrastructure to overtake token pricing as the dominant AI operating cost.

Three-year financial summary

Applying an 8 percent discount rate, consistent with rates commonly used in enterprise technology investment appraisal, produces the summary metrics below for the composite enterprise.



Metric	Value
3-year cost avoidance, net of an AI Data Layer subscription	$10,500,151
3-year business value added	$8,625,344
3-year total benefit	$19,125,495
3-year net present value of benefit, at 8%	$16,038,286
3-year discounted investment	$5,536,028
3-year ROI	**290%**

Table 3: Key financial metrics, composite enterprise model

"Denodo enables us to minimize the number of MCP servers we need, reducing server sprawl. This lowers complexity, improves security and trust, and controls token costs."

APPLIED AI AND PRACTICE LEAD, GLOBAL CUSTODY AND ASSET SERVICING BANK

06

The Baseline Architecture and Alternative Archetypes

The baseline reflects the mixed architecture typically found in enterprises today. Three architectural archetypes commonly contribute to that environment. These are custom integration, a lakehouse-native stack and cloud-native services. In practice, enterprises often combine these approaches rather than adopting one in isolation.

The alternative assessed in this paper is a governed AI Data Layer. For most enterprises, this does not replace existing lakehouse or cloud platforms but extends them by providing access to data across the wider enterprise. The three archetypes are outlined below.

01

Custom source-by-source integration. Bespoke MCP servers, direct API integrations, RAG pipelines, per-agent prompts and semantic mappings, with access control and audit handled separately in each place. Deployment can be relatively straightforward for the first few agents but becomes progressively more complex as agent numbers grow, because much of the integration, semantic and governance work has to be repeated. This archetype is common when data and AI ownership is distributed across business units, with individual teams optimizing for their own requirements. It can work well at that level but becomes harder to manage as agentic AI scales across the enterprise.

02

Lakehouse-native stack. Snowflake or Databricks as the primary platform, with ingestion through tools such as Fivetran, Airbyte or Informatica, transformation through dbt or equivalent tooling, and semantic modeling and governance handled within or alongside the platform. Custom APIs and MCP servers provide agent access. The effort required depends on data movement, semantic modeling, connector maintenance and policy enforcement as agent numbers grow, particularly if significant parts of the enterprise data estate remain outside the platform.

03

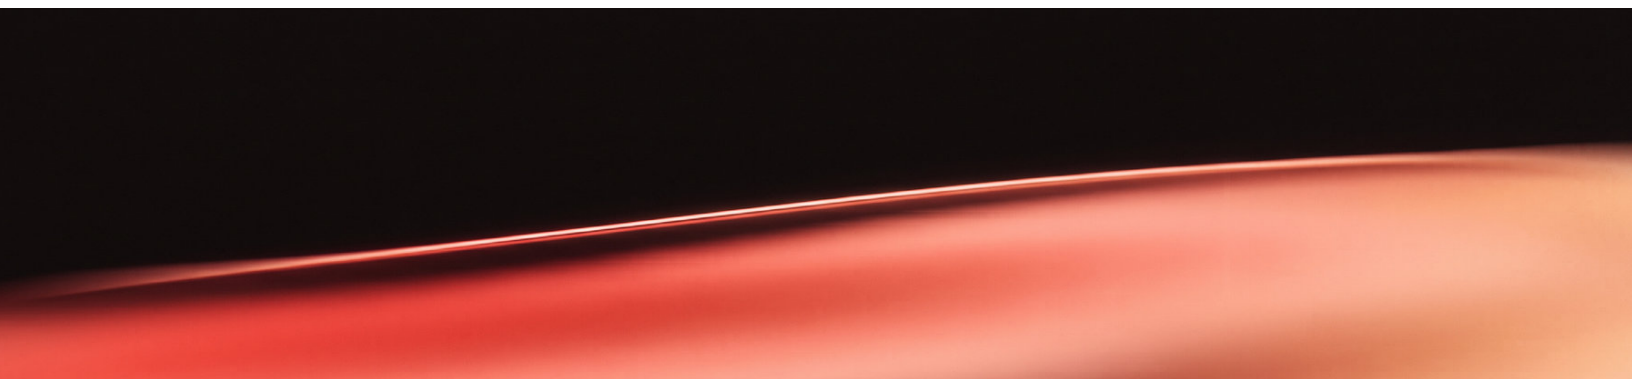
Cloud-native platform capabilities. Microsoft Fabric with Fabric IQ and Purview, AWS SageMaker and its underlying data services, or Google BigQuery with its Knowledge Catalog and cross-cloud capabilities. Enterprises may use the platform they already own to provide governed agent access without adding another layer.



Requirement	Custom integration	Lakehouse-native stack	Cloud-native platform	AI Data Layer
Live access to data held on the primary platform	Built per agent	Native access	Native access	**Access through the shared layer**
Live access to data held outside the primary platform	Built per agent, repeated each time	Requires ingestion or custom federation first	Connector or federation dependent outside the provider's ecosystem	**Federated access across supported external sources**
SaaS, operational and legacy system coverage	One connector per system, built and maintained separately	Connector-specific, usually with ingestion into the platform	Native within the provider's ecosystem, connector-dependent beyond it	**Broad prebuilt connectivity exposed through one access point**
Shared semantic layer	Usually per project	Possible through dbt and catalog	Provided through the platform catalog, within the platform	**Defined once and reused across every agent and platform**
Tool reuse across agents	Requires custom design	Requires custom design	Requires design, strongest inside one ecosystem	**Governed data products inherited by every entitled agent**
Centralized policy enforcement	Fragmented	Usually split across tools	Strong within the platform, variable beyond it	**Applied once at the access point, across all sources**
Minimal data movement	Depends on integration design	Variable	Variable	**Federation reduces the need for data movement**
Vendor neutrality	High, at the cost of effort	Partial	Ecosystem dependent	**Core architectural role**

Table 4: Qualitative comparison of architectural archetypes. Outcomes depend on the source estate, workload and governance requirements in each case.

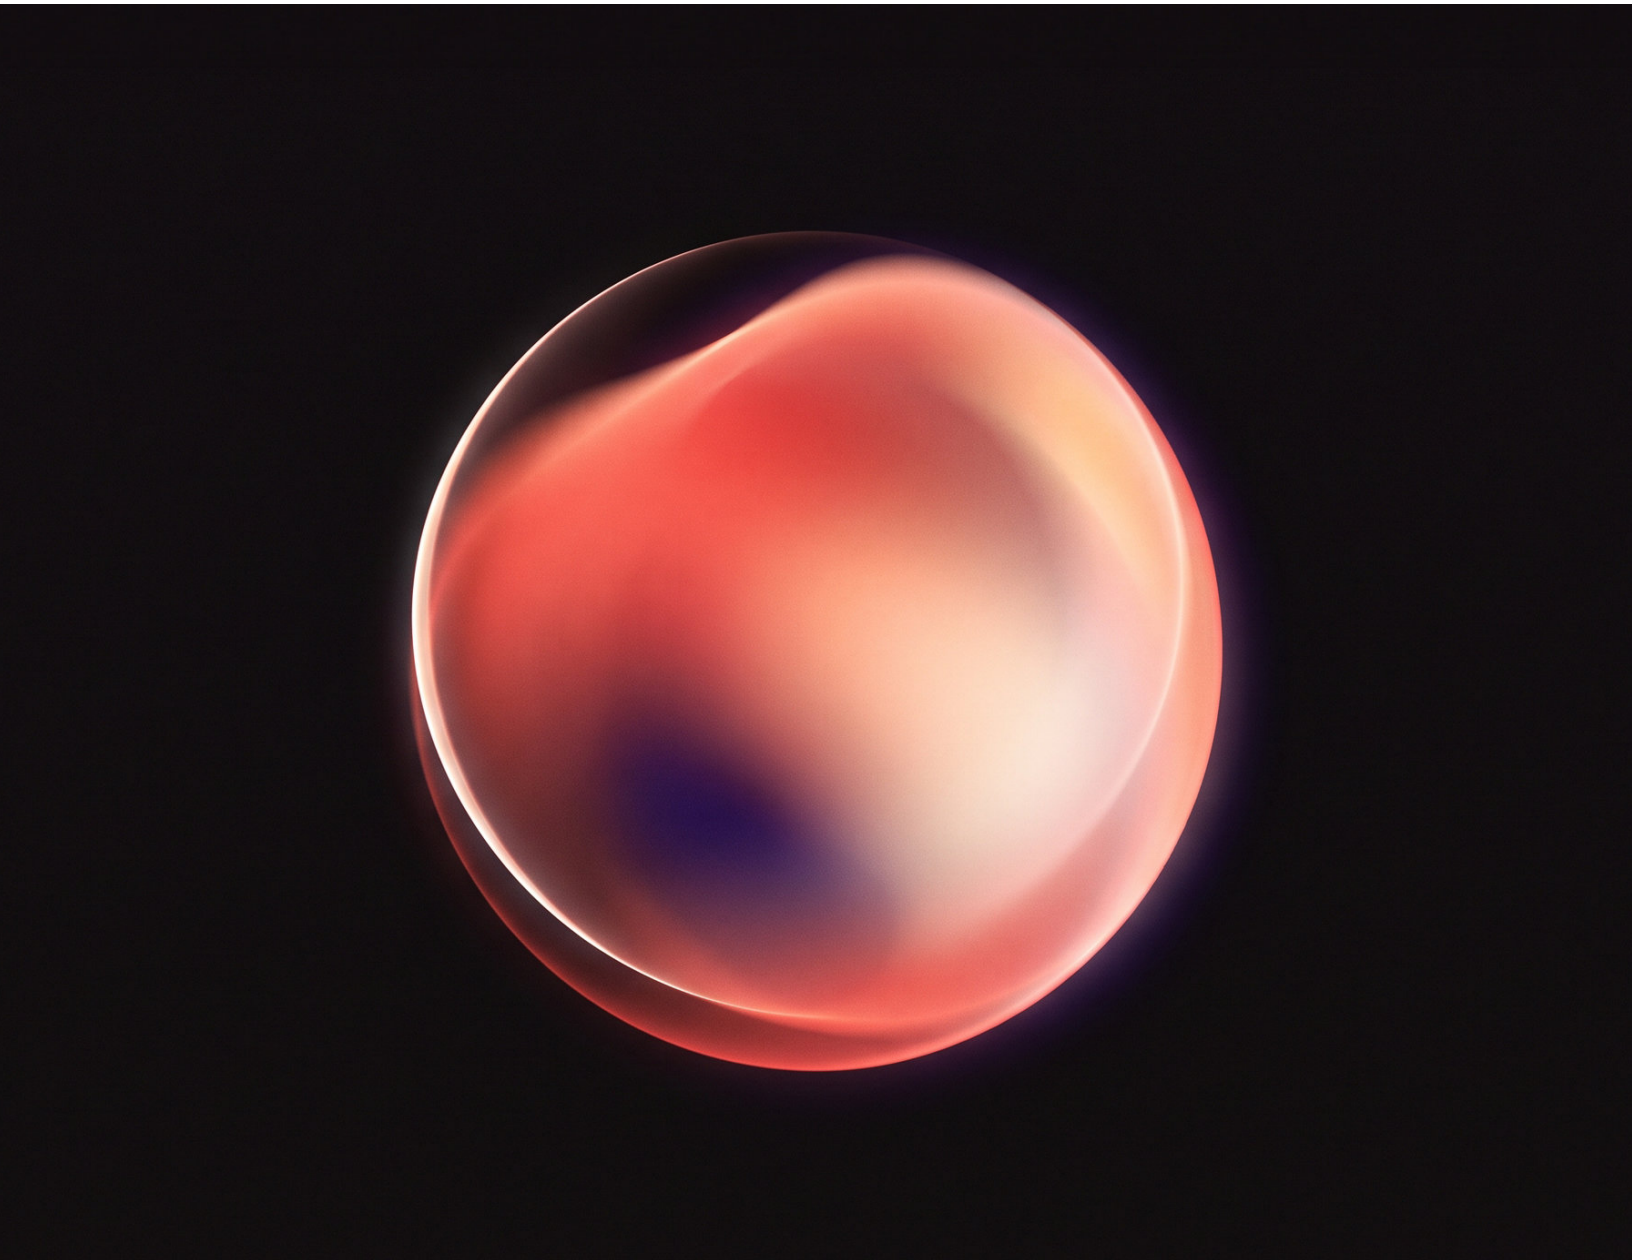
One of the main differences between the archetypes is how they handle data outside the primary platform. Lakehouse-native and cloud-native approaches are typically strongest where the data required by agents already sits within the platform or its ecosystem. Accessing data elsewhere may require additional connectors, federation or ingestion. An AI Data Layer is designed to work across these platform boundaries, extending consistent access, semantics and governance across distributed sources without requiring the data to be consolidated first. The relative value of this approach therefore increases as the data required by agents becomes more distributed across SaaS, operational, legacy and cloud environments.



Factors affecting returns

The modeled return rises with the proportion of agent-relevant data sitting outside the central platform, and with the number of domains a typical agent task has to cross. A single layer connecting a lakehouse, a CRM platform, an ERP system and a mainframe delivers considerably more value than the same layer operating across an already consolidated data environment. A hyperscaler interviewed for this research made a similar observation, noting that the value of a shared enterprise data layer increases with the number of systems involved in an AI workflow.

The converse also applies. An enterprise that has already consolidated most of the data its agents need onto one platform may see less incremental benefit from a shared data layer. However, previous Veqtor8 research found that achieving and maintaining equivalent integration, governance and data delivery capabilities through a lakehouse-centric architecture can require substantially greater engineering effort and cost. One organization in this research reported that roughly ninety percent of its enterprise data was held in a single cloud warehouse, but agents still required live access to operational data outside the central platform. An AI Data Layer can provide live access to that data where it sits, alongside data held in lakehouses and other platforms.



07

Strategic Outcomes Beyond the Numbers

Cost avoidance and modeled value creation are not the whole business case. Organizations further along in their agentic AI journey report four strategic benefits that cannot be quantified easily but increasingly influence the decision to invest.

01	02	03	04
Access to current data	**Faster agent deployment**	**Governance and auditability**	**Architectural flexibility**
Provides a common route to data across distributed sources, including live access where data remains in operational systems	Common connectivity and reusable data products reduce the engineering work needed for each new agent, accelerating time to deployment and improving the return on subsequent workflows	Policy enforcement is managed centrally rather than per agent or per source, and semantic consistency with lineage supports auditability of agent decisions and regulatory reporting	Decouples agents from underlying data platforms, reducing dependence on any single provider and supporting data portability and AI sovereignty

Table 5: The strategic role of an AI Data Layer in agentic AI enablement

Governance by design

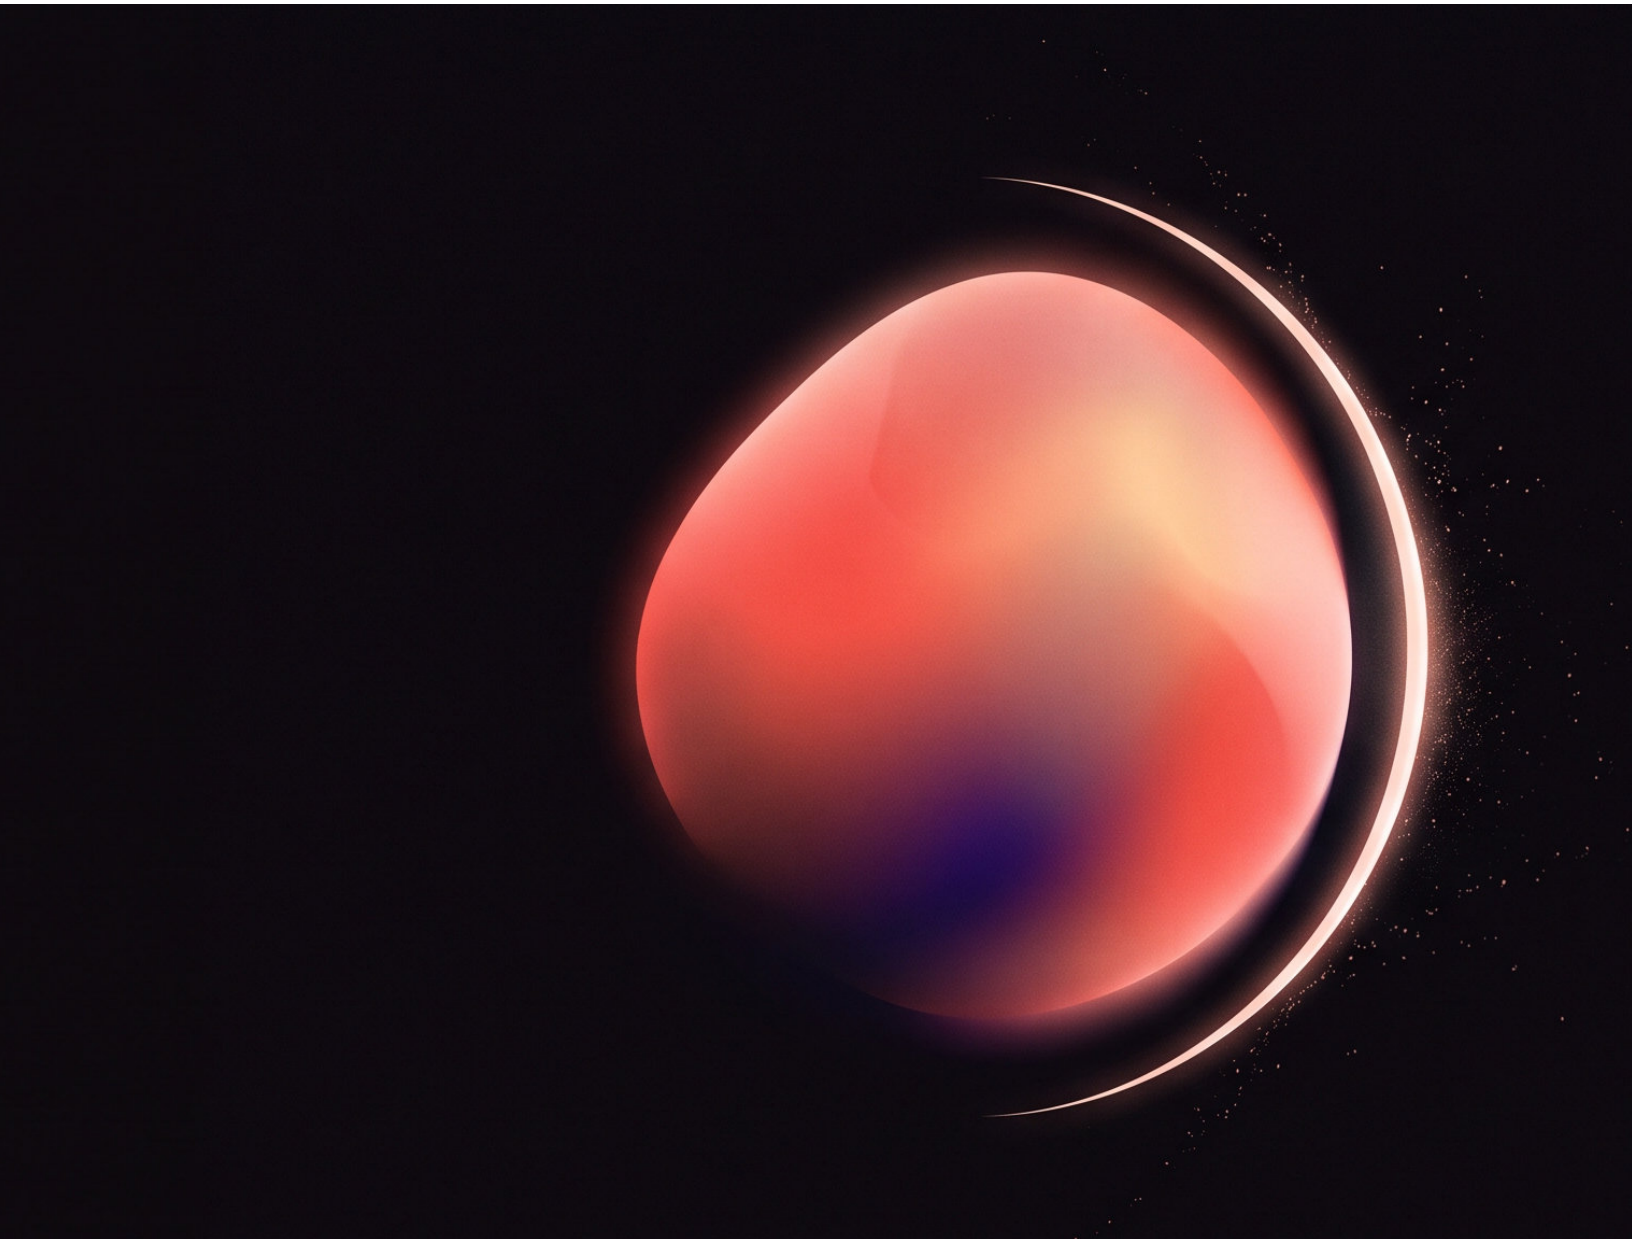
Several organizations describe governance as an architectural capability rather than a separate project. Behavioral monitoring identifies when an agent begins operating outside expected parameters, while governance at the data layer determines what the agent can see before it acts. Agent governance manages the behavior and lifecycle of the agent, while the AI Data Layer governs the data and context it can access. This separation of responsibilities becomes particularly important if agents hold enterprise credentials and report to human line managers, as they do at the custody and asset servicing bank profiled in this research, because a named individual remains accountable for the agent's actions.

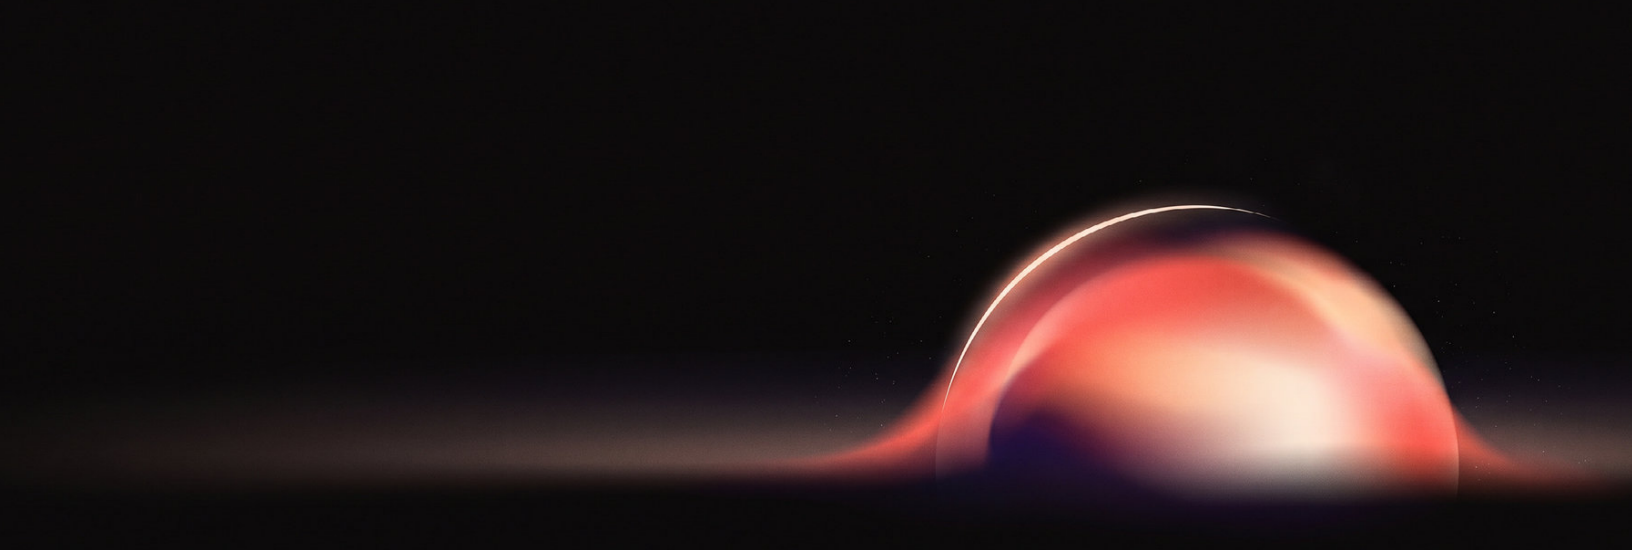
Classification and masking provide another level of control. Agents can be restricted to the data they need, with sensitive or irrelevant fields masked and access limited to their permitted scope. As more agents access enterprise data, applying these controls consistently becomes more important. A governed AI Data Layer allows access, classification, masking and data-sharing policies to be managed centrally rather than configured separately for each agent.

Semantic agility and architectural choice

A well-designed semantic model on a modern data platform can provide equally strong semantics for data already within that platform. The advantage of an AI Data Layer becomes clearer when agents need governed access to data distributed across multiple platforms and operational systems. Semantic views can be changed without first moving and transforming the underlying data.

Architectural choice is increasingly important as organizations use a mix of proprietary and open-weight models. Separating data access from the underlying model and platform allows organizations to change models without rebuilding data access. This supports AI sovereignty by giving organizations greater control over where models run, which models they use and how enterprise data is accessed.





08

Conclusion

As organizations expand their use of agentic AI, the route agents take to enterprise data becomes a key architectural choice. The choice is whether to keep rebuilding data access, semantics and governance as new agents are added, or establish these as reusable capabilities that can support agents across the enterprise. Without a shared data layer, much of this work is repeated as new sources and agents are added. Engineering effort increases, semantics are rebuilt wherever they are needed, governance is applied separately, and AI operating costs rise when agents must do more work to assemble their own context.

For enterprises whose AI initiatives depend on data distributed across cloud platforms, operational systems, SaaS applications and existing data platforms, a governed AI Data Layer provides reusable access to the live, governed and semantically consistent data those agents need. The model indicates a 290 percent three-year return for the composite enterprise, made up of $8.6 million in business value added and $10.5 million in cost avoidance, with benefit growing year on year as the agent population expands. The level of benefit will vary between organizations, though the trends are similar with large distributed enterprises.

09

Methodology

This ROI model was developed in two stages, consistent with the methodology used in Veqtor8's prior lakehouse ROI research. An initial prototype model was built from existing research, then developed and refined through a program of structured interviews with enterprise practitioners, hyperscalers and technology partners. The findings from the interviews were used to build a composite enterprise representing a typical large organization adopting agentic AI.

The seven organizations behind the composite

The composite enterprise used throughout this paper is built from seven large enterprises interviewed for this research. Each operates across multiple regions and each is either running agents in production or preparing to do so. Their data estates, agent populations and deployment experience are aggregated into a single organization, so the analysis presents one consistent set of figures rather than seven partial ones. The baseline architecture modeled in Section 6 is the aggregate of the estates these organizations run today.

GLOBAL **Banking** Agent deployment timelines, reuse of semantic data products, human oversight of agent decisions	GLOBAL **Custody and asset servicing** Agent reliability against authoritative reference data, MCP server sprawl, accountability for agent actions	GLOBAL **Chemicals and crop science** First production agents on a governed layer, business user access to distributed data	GLOBAL **Mining and resources** Fragmentation of the data estate and governed access to operational systems
NORTH AMERICA **Rail and transportation technology** Agent workloads drawing on distributed enterprise data	APAC **Building materials** Orchestration of multi-system workflows and the integration effort behind them	APAC **Telecommunications** AI operating cost at scale, deployment and integration uplift	

Table 6: The seven organizations aggregated into the composite enterprise

Findings were then tested with the major technology suppliers, each supplying a cross-customer view of agent adoption, architectural choice and AI operating cost. Participants are not named at their request.

Composite enterprise

Aggregating those seven organizations produces a composite enterprise with the following profile.

01

SIZE:

5,000 to 150,000 employees

02

REVENUE:

Greater than USD $1B annually

03

GEOGRAPHY:

Global or multi-region operations

04

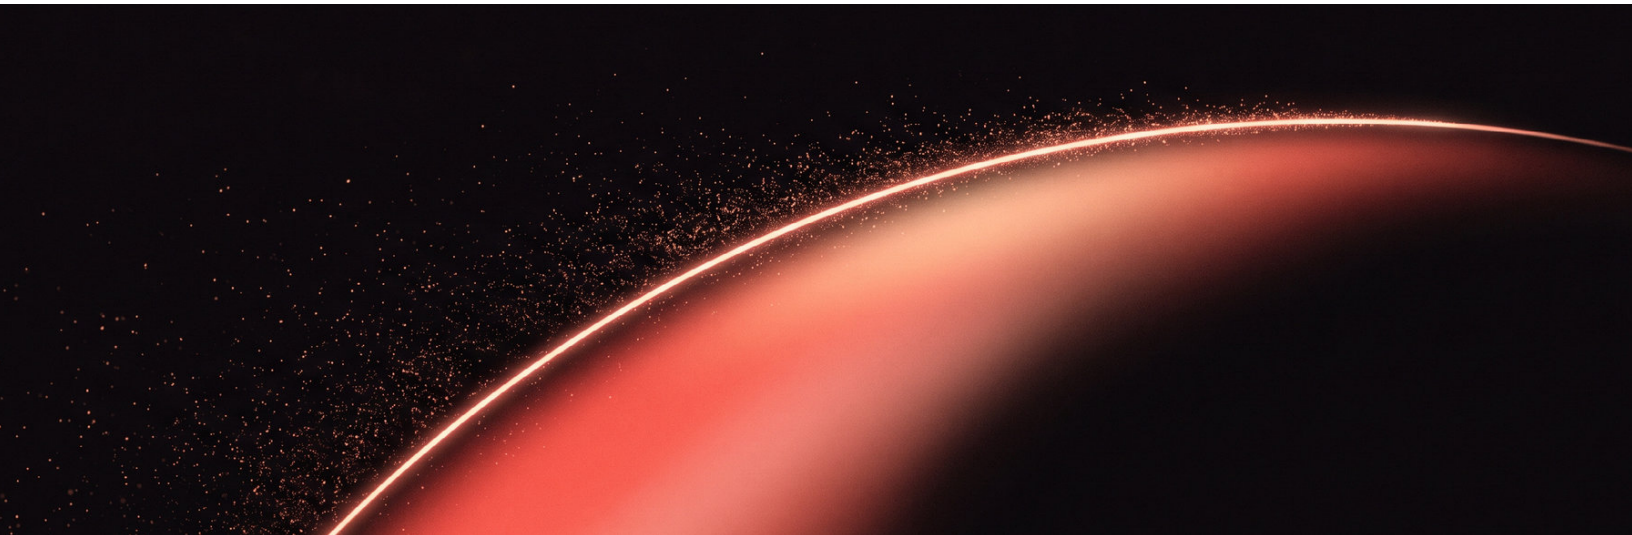
DATA ENVIRONMENT:

Diverse hybrid data ecosystem, with multiple on-premises, cloud and SaaS sources feeding agent workflows

The model quantifies the six sources of economic benefit set out in Section 4. Enterprise data sources are used as the unit of onboarding effort, while data products are the unit of consumption and reuse across agents. Interview evidence was used to test assumptions and calibrate modeled values, with estimates generally set toward the lower end of the ranges reported.

Assumptions and limitations

The analysis models the economic effect of making governed enterprise data reusable across a growing population of agents. It does not claim that every enterprise will achieve the headline return, that a logical layer removes the need for lakehouse or cloud data platforms, that federation is always preferable to data movement, that an AI Data Layer alone improves model quality, or that token cost is determined only by data access architecture. Results are indicative rather than predictive for any individual organization.

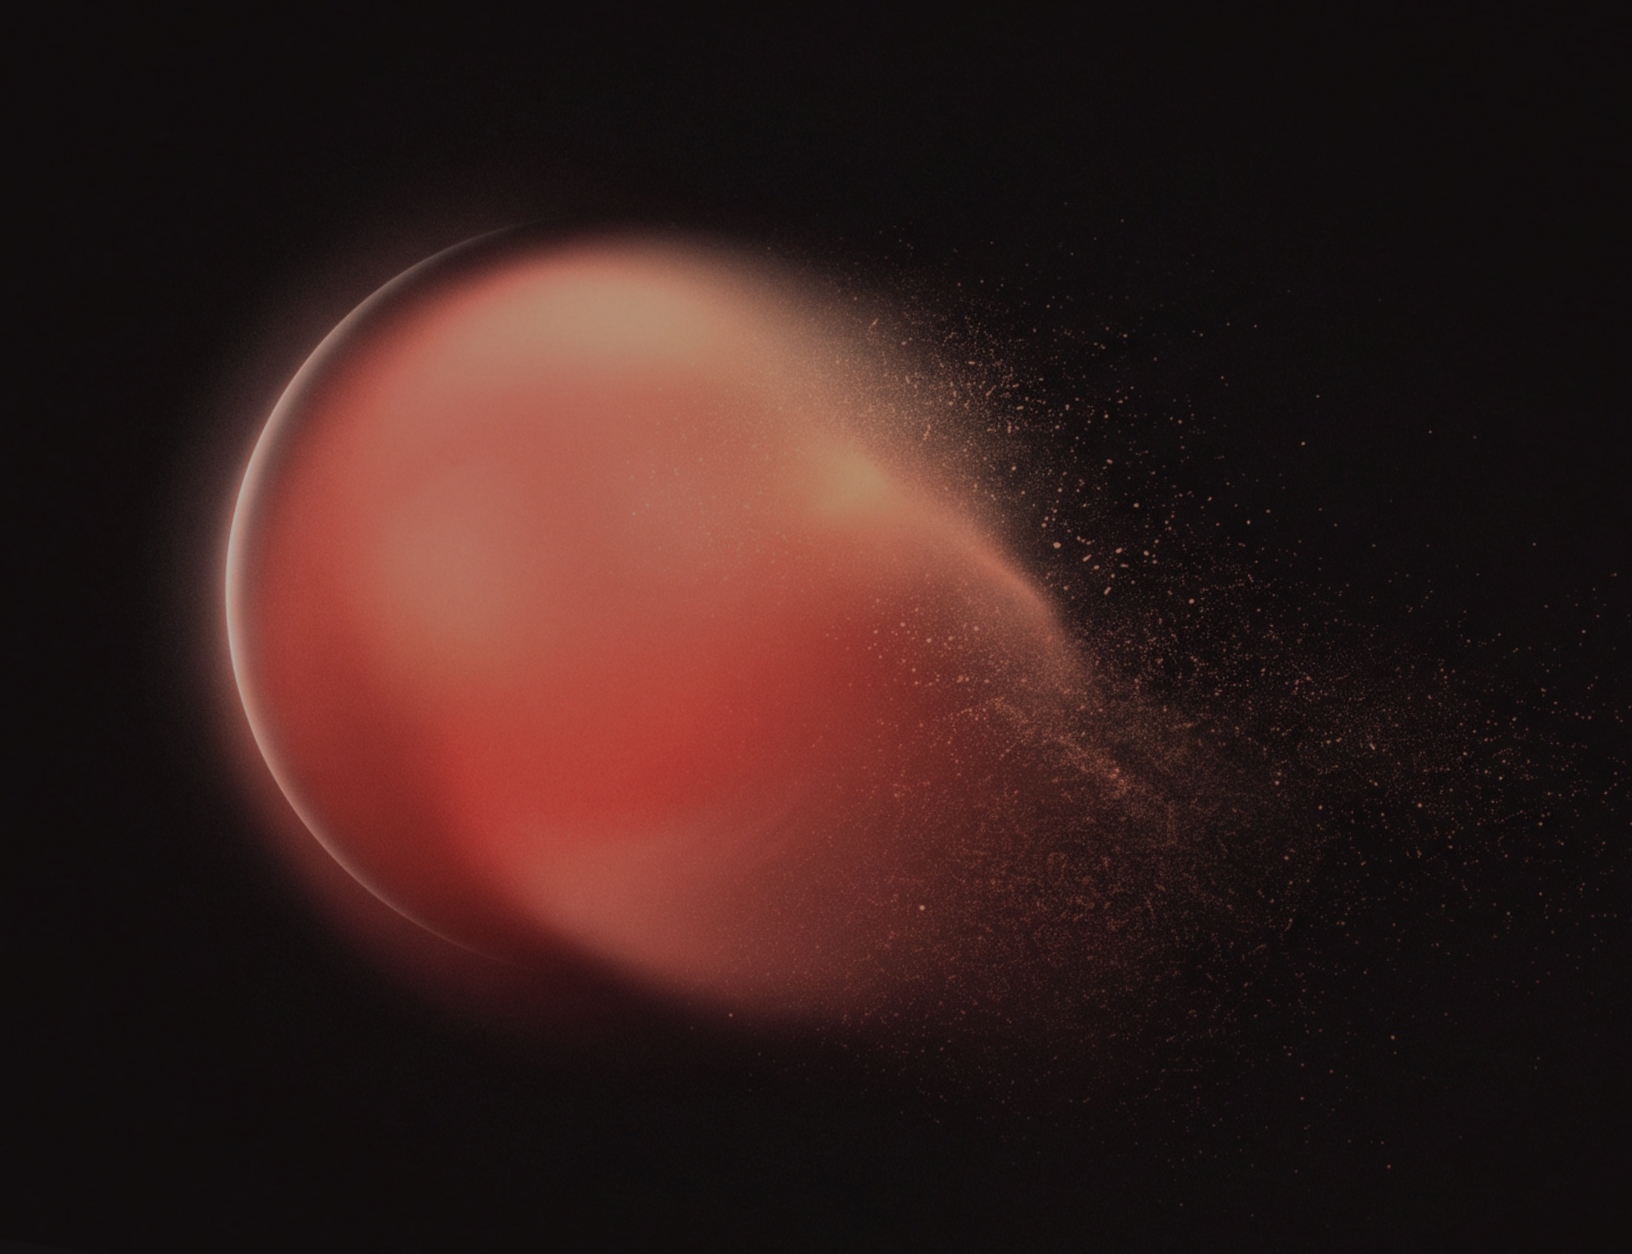


- **Sensitivity.** Outcomes are most sensitive to agent growth, task volume, integration reuse, context reduction and the treatment of released capacity as business value. The return rises with the proportion of agent-relevant data sitting outside the central platform and with the number of domains a typical agent task crosses, and falls where the estate is genuinely consolidated and agents work within it.
- **Agent adoption.** The model assumes continued growth in enterprise AI deployments, with accelerating growth in deployed agent numbers over time. Actual benefits will depend on the number of agents deployed, the complexity of their workloads and the pace of adoption.
- **AI operating costs.** Token pricing, cloud infrastructure costs and model efficiency are evolving rapidly. The model assumes token prices fall over the three-year period. Actual costs will vary with the organization's choice of infrastructure and AI models, and with the degree of portability between them.
- **Discount rate.** An 8 percent discount rate is applied, consistent with rates commonly used for enterprise technology investment appraisal.
- **Conservative benefit estimates.** Category-level assumptions sit at the lower end of the ranges reported in interviews and observed in comparable deployments, representing typical enterprise environments rather than best-case outcomes.
- **Value created.** Business value creation is modeled explicitly, covering semantic preparation time released, productive time returned to agent users, and net new work delivered by agents. These lines are kept separate so that work taken off people – and work the organization previously could not resource – are each counted once.
- **AI cost optimization.** Interviewees expressed differing views on the extent to which a governed data layer reduces AI operating costs. Token consumption depends on factors including model capability, workflow design, context size and retry behavior. The model therefore uses conservative assumptions and separates token and infrastructure costs.
- **Benefits still excluded.** The model does not attempt to quantify revenue growth, customer experience improvement, competitive positioning or cultural and organizational change. These are discussed qualitatively in Section 7.



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. Veqtor8 helps organisations become AI-ready, optimize their data strategies and strengthen digital resilience through practical, independent research, executive workshops, reports and presentations.

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

denodo

Denodo is the AI data layer company. The award-winning Denodo Platform powers trustworthy agents and applications, transforming enterprise data into reliable insights for analytics and self-service.

Visit: denodo.com | **Connect:** denodo.com/contact-us