

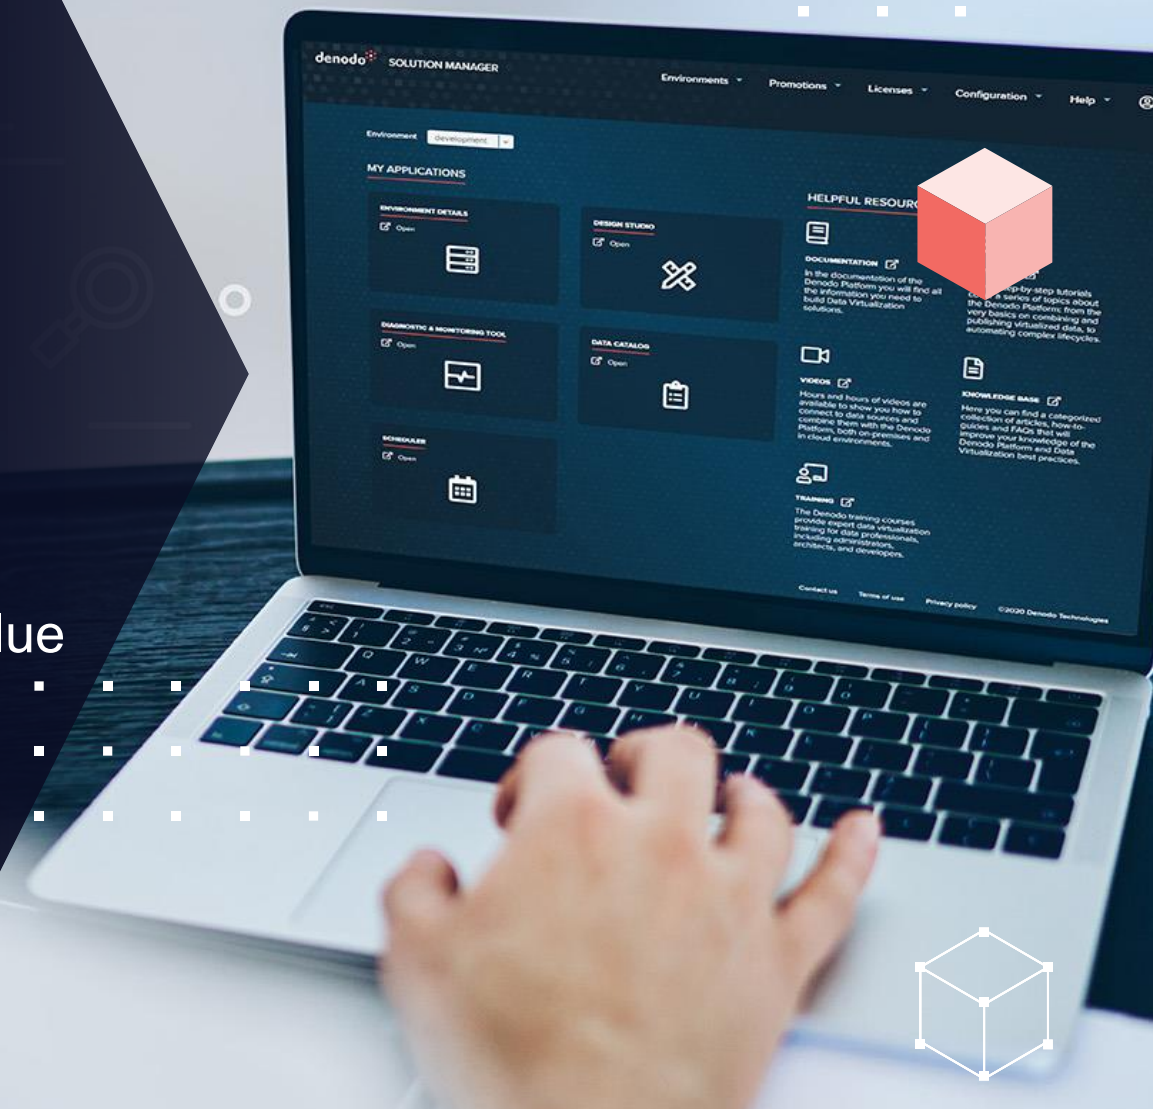



The AI Governance Imperative:

Balancing Innovation, Compliance, and Business Value







Terry Dorsey
Evanvelist NA | Denodo





Building a Modern Data Foundation



Modernization Has Entered a New Phase

AI now changes the architectural question

Modernization Efforts Included:

- Cloud migration has improved scalability and resilience.
- Lakehouse architectures have improved storage flexibility.
- Data products have improved organizational ownership.
- AI pilots are accelerating experimentation.

AI now changes the architectural question



The New Success Criteria

AI initiatives can no longer be just about technology implementation.

Initiatives are now measured by:

- Speed to deliver new use cases, reducing time-to-market.
- Ability to scale across domains, moving beyond isolated pilots
- Cost-to-value over time, ensuring sustainable economics as usage grows
- Risk and compliance readiness
- Business adoption

Success is defined by business impact and scalability, not just technical capability.

AI Changes the Data Foundation: From Unified Storage to Unified Access

In an AI-driven enterprise, unification must come from shared semantics and governed access — not from physically moving all data into one platform.



Live Data, Zero Copy

AI needs **live data** from original source systems.



Unified Business Context

AI needs more than good data. It needs **the right data**.



In-Place Data Governance and Access

AI needs consistent enforcement of **guardrails** and business SLAs.



Cost and Performance Optimized

AI is more **compute-intensive and data-hungry** than anything that preceded it.

Across All Sources

Multi-cloud, Hybrid, On-premise, SaaS, 3rd party.
Structured & Unstructured.

The Core Architectural Question

How do we architect **a trusted data foundation** that ensures consistent meaning, controlled access, and predictable performance...

...so new use cases can scale without rework or risk?

The Three-Part Answer




Universal Semantic Layer

To standardize meaning, access, and control across the entire enterprise.




Data Products

Turn trusted data into reusable, outcome-aligned assets that accelerate delivery across the enterprise.




AI Delivery Model




To remove data complexity from builders, enabling scalable AI adoption.



The Universal Semantic Layer



To standardize meaning, access, and control across the entire enterprise.



Think of the Universal Semantic Layer as an Architectural Boundary

It is the control point in the architecture where meaning, governance, and consumption are standardized once for the entire enterprise.

- Separates business meaning from physical data structures
- Decouples governance from individual platforms and tools
- Defines the contract between data producers and consumers
- Prevents logic from being embedded in pipelines, BI, and AI



What the Universal Semantic Layer Changes

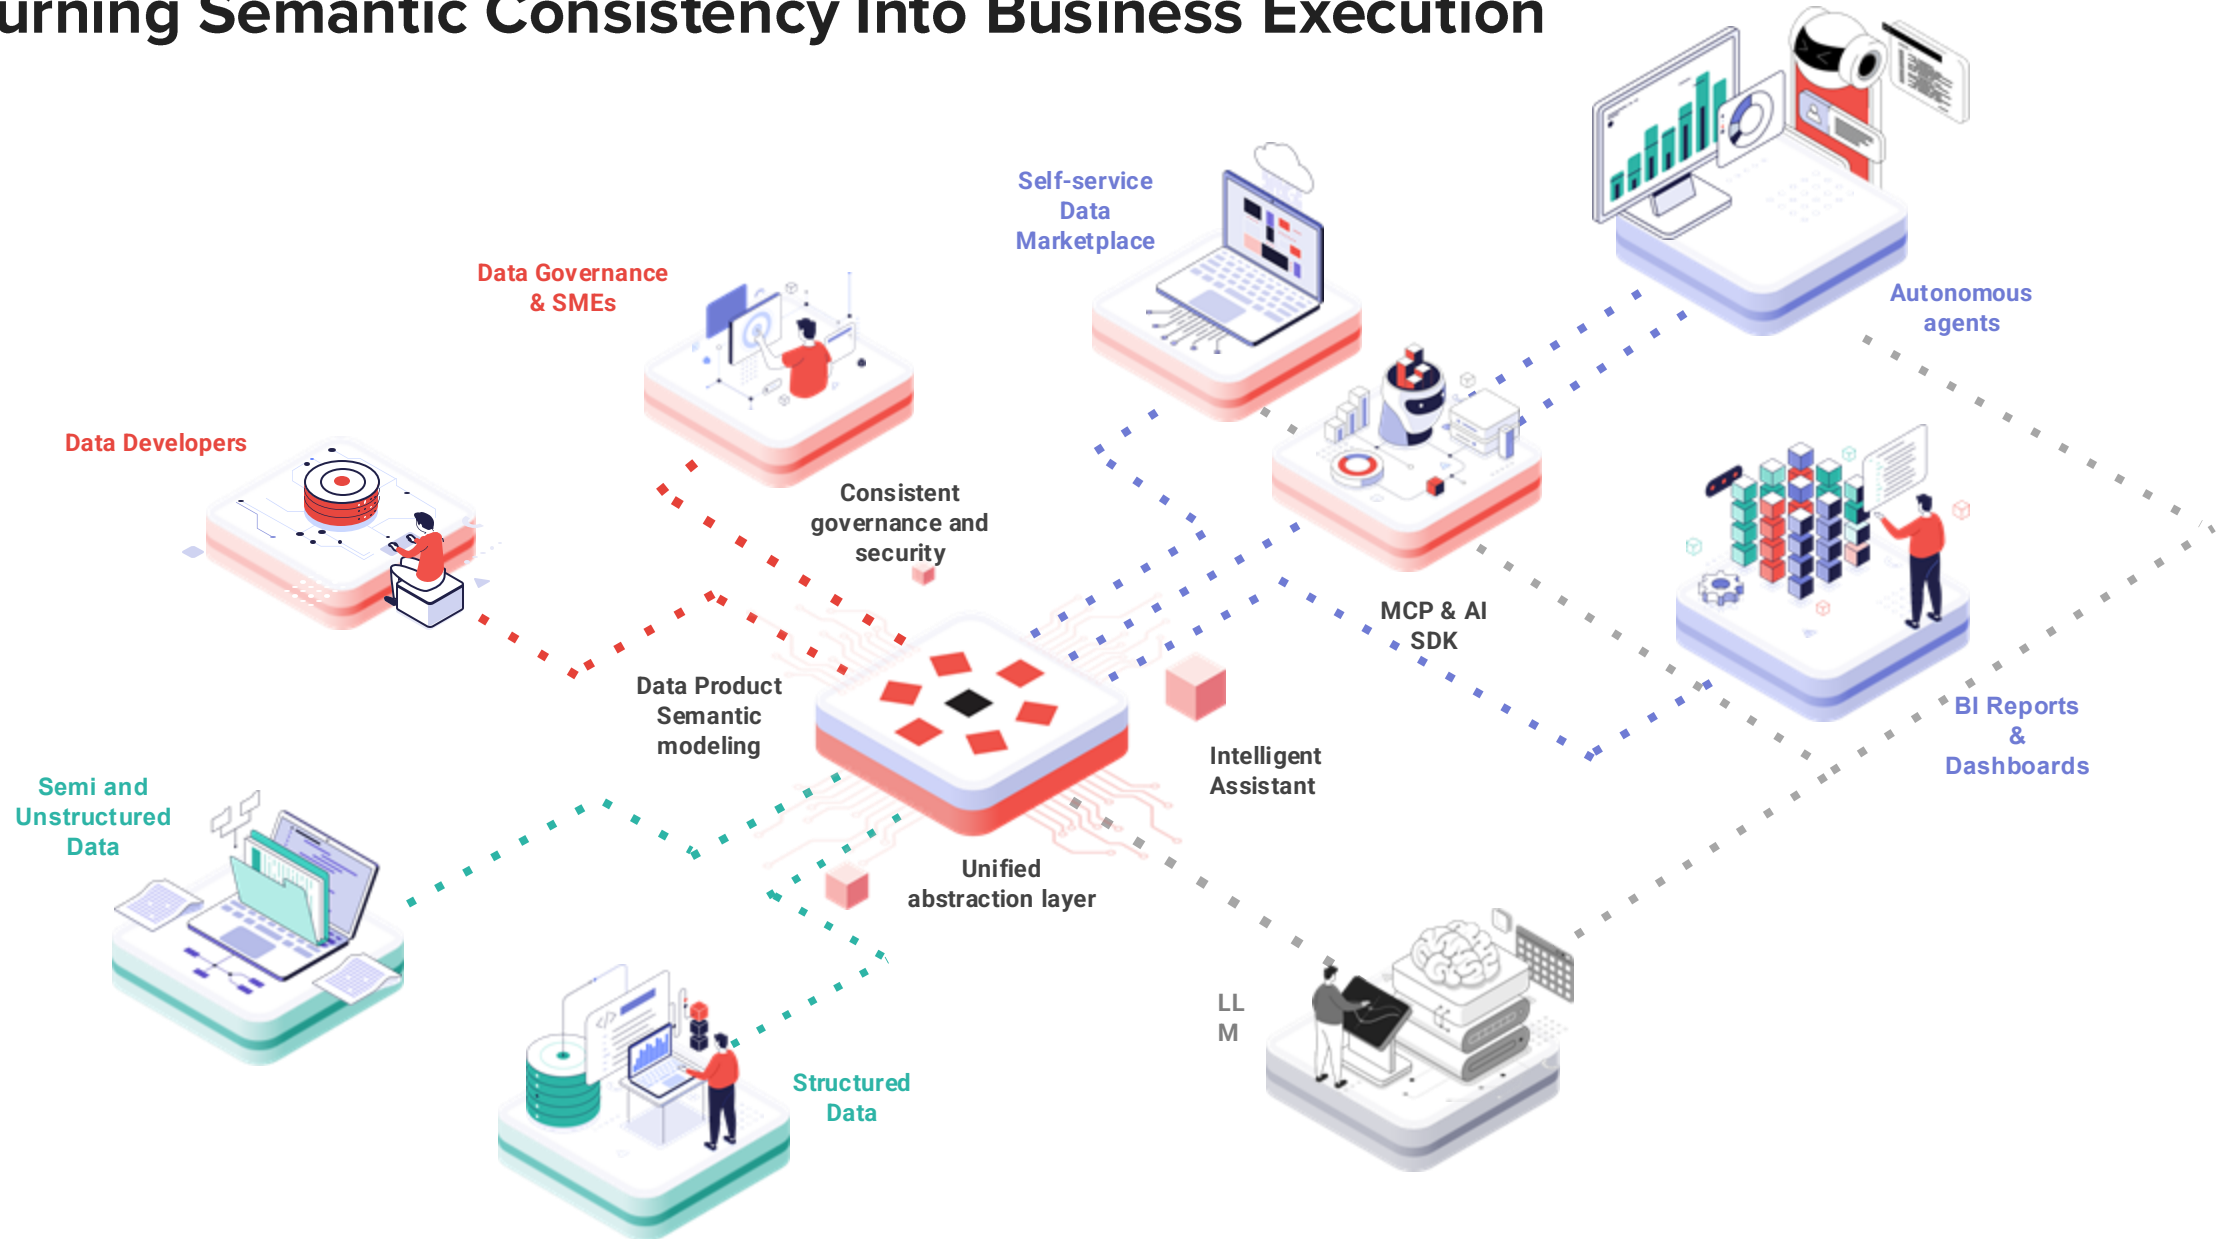
By abstracting data complexity, we shift from fragmented logic to a unified architectural standard.

- **One business definition reused across all use cases**
 - Ensure every report, model, and application speaks the same language
- **One governed access model for all consumers**
 - Apply security policies centrally and enforce them globally across the enterprise
- **One logical access point to all data**
 - Connect to data where it lives without forcing physical centralization

Core Shift

Define once, secure once, deliver everywhere.

Turning Semantic Consistency Into Business Execution



The Universal Semantic Layer as the Foundation for Scalable Data and AI

It creates the trusted, governed, and reusable data foundation required to deliver new business initiatives with speed and consistency.




One Shared Business Meaning

To define data once and deliver trusted, consistent metrics across all use cases.




One Governed Access Model

To enforce security and policy centrally for every user, tool, and AI workload.



One Logical Access Layer




To connect to all data in real time without physical consolidation.



Enterprise Data Products



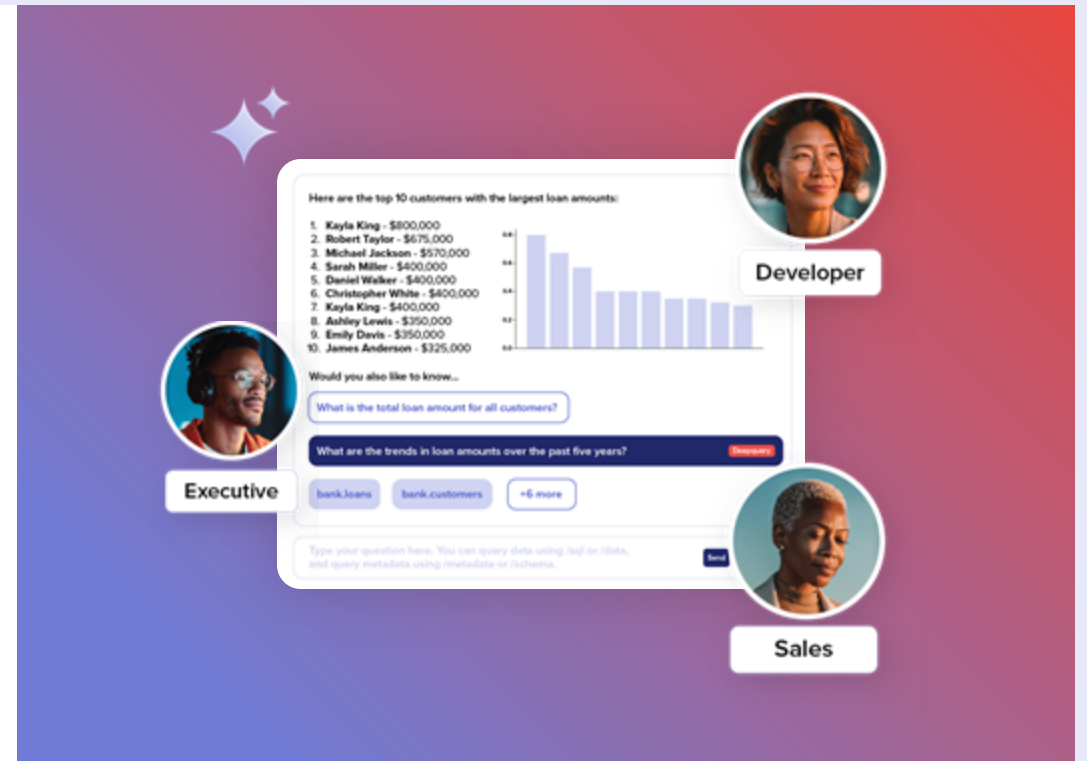
Turn trusted data into reusable, outcome-aligned assets that accelerate delivery across the enterprise.



Data Products as the Engine for Enterprise Scale

Data products shift delivery from isolated projects to a scalable model built on reusable assets.

- Data products package data, definitions, quality expectations, and policies into reusable assets
- Teams do not need to be data experts to create and maintain trusted data products
- Ownership is distributed to the teams closest to the business context, not centralized in a committee
- Governance becomes embedded in day-to-day work rather than enforced after the fact



Data Products Scale the Rate of Business Delivery

Data products change the delivery model from project-by-project execution to a scalable, repeatable business capability.

Shift to Reusable Assets



Shift from one-off projects to reusable, outcome-aligned assets

Accelerate Use Cases

Deliver new use cases without rebuilding logic, policy, and access

Scale Success

Increase production success without expanding central bottlenecks



Outcomes Define the Data Product

Value starts by anchoring every initiative to measurable business outcomes—data products package the trusted data required to deliver them.



Business-Driven Scope

Scope is driven by what the business must improve – not by available data



Aligned Success Criteria

Business and technical teams align on shared success criteria



Value-First Investment

Investment prioritizes what delivers impact first



The Reusable Unit Is the Data Product

A data product packages everything required to reuse trusted data across the enterprise.



Unified Asset Delivery

Data, definitions, quality rules, and policies delivered as one asset.



Immediate Consumption

Business-ready and immediately consumable for analytics and AI.



Build Once, Reuse Everywhere

Build once — reuse across domains and use cases.



A Delivery Model — Not an Organizational Mandate

Data products define how data is delivered — not how organizations must structure their teams.



Flexible Operating Models

Works with centralized, federated, and hybrid operating models seamlessly.



Business-Aligned Ownership

Ownership aligns to business context without forcing a reorganization.



Consistent Governance

Governance remains consistent regardless of structure.



What Changes in the Delivery Lifecycle




The biggest impact of data products is how they change the delivery lifecycle for new initiatives.

Traditional Delivery

-  **Rediscover & Prepare** Rediscover and prepare data from scratch for every new initiative.
-  **Rebuild Logic** Rebuild transformations and definitions per specific use case.
-  **Reapply Governance** Reapply governance rules and access controls each time.



Data Product Delivery

-  **Consume Trusted Assets** Consume trusted, ready-to-use assets immediately.
-  **Reuse Context** Reuse logic, policy, and semantic context across projects.
-  **Built-in Governance** Deliver faster with governance built-in by default.

Self-Service on Governed Data – Not Raw Access

Self-service only works when users can safely consume trusted data products instead of raw data sources.



Business Semantic Layer

Consumers use trusted data in business language.



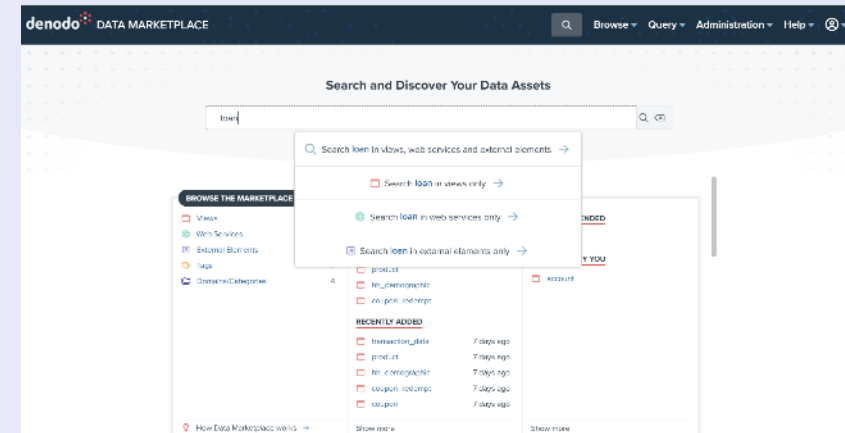
Eliminate Manual Prep

No schema discovery or manual data preparation.



Built-in Governance

Access, security, and SLAs enforced by default.



Works With Your Operating Model

Data products scale across organizations regardless of how data teams are structured.



Centralized

Shared platform team delivers domain data products



Federated

Domains build and own their products independently



Hybrid

Standards are centralized – ownership is distributed

✔ The delivery model remains the same across all three.

From Data Projects to a Scalable Business Capability

When data products become the standard delivery model, organizations turn isolated successes into repeatable capabilities.



Accelerate Delivery

Faster onboarding of new initiatives with reusable assets



Clear Ownership

Accountability defined by business domain context



Controlled Self-Service

Empower users at enterprise scale without chaos

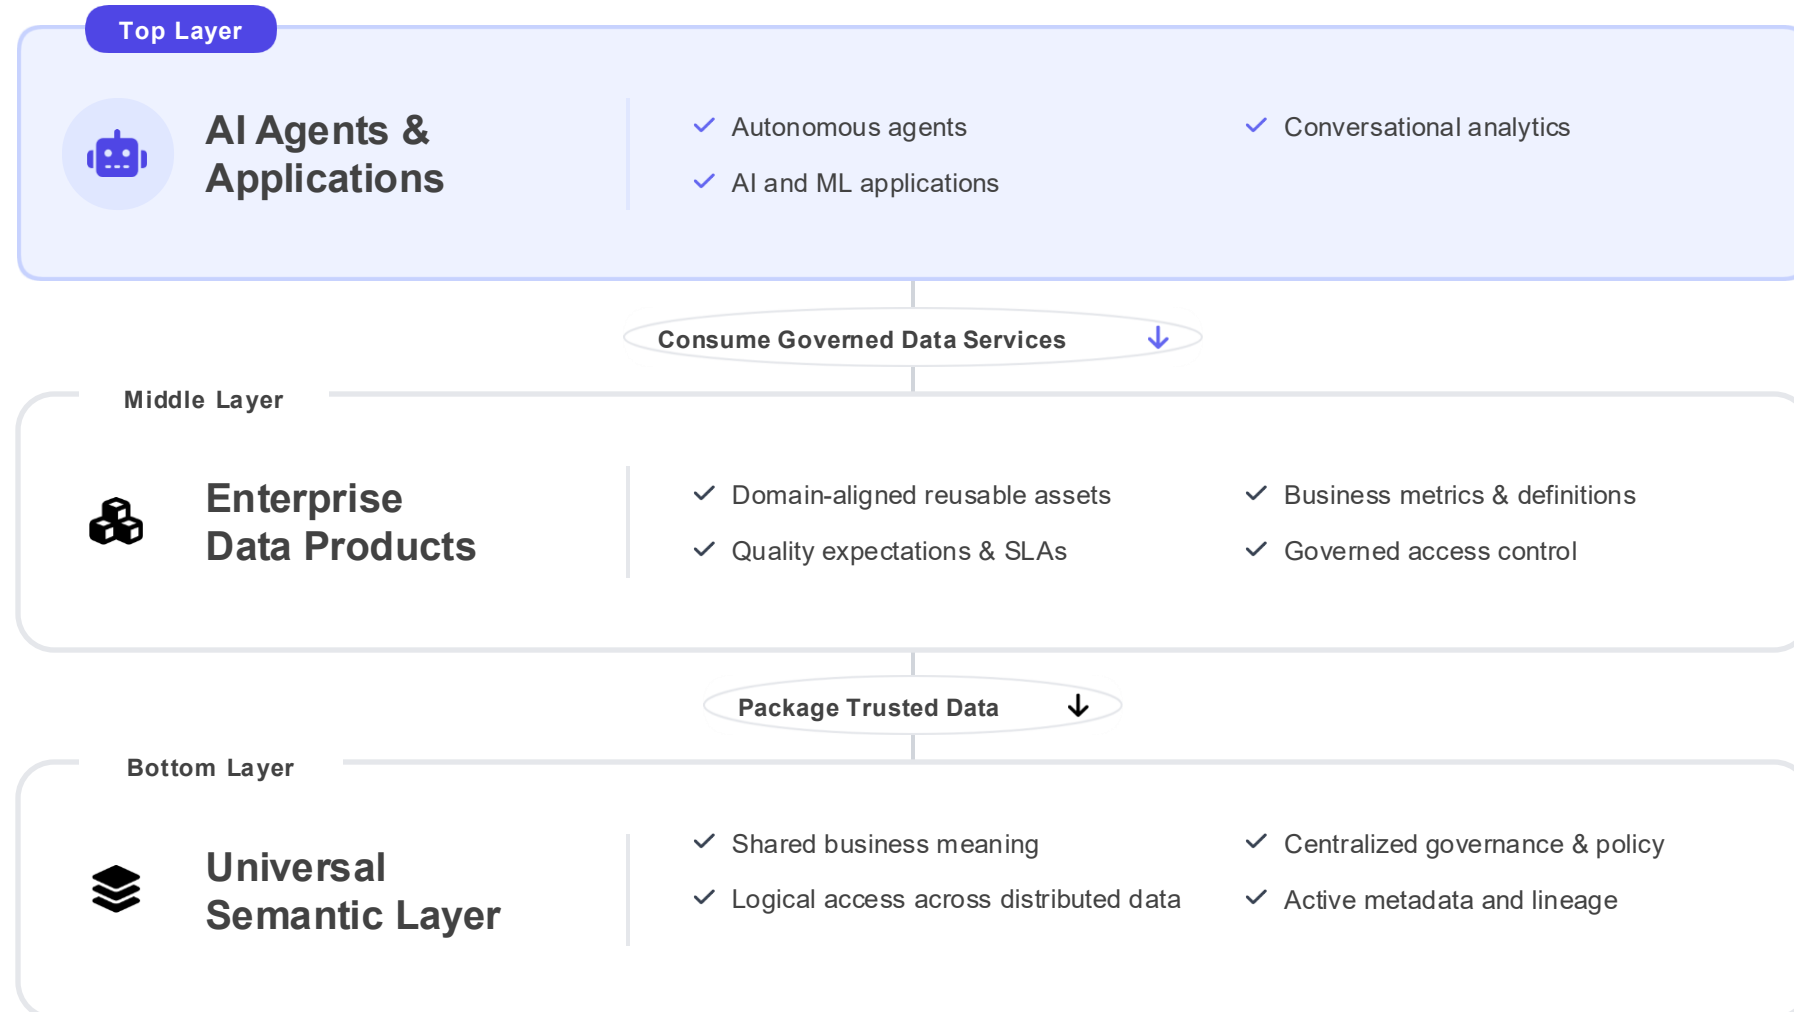


Trusted AI & Analytics

Built on trusted, reusable data foundations

From Semantic Foundation to AI Execution

A scalable AI architecture separates data foundation, reusable assets, and intelligent applications.





Accelerating AI Delivery

AI scales when teams can build on trusted data without becoming data engineers.



AI Delivery: The Business Challenge

While Generative AI and ML prototypes are increasingly easy to build, the path to production is blocked by data friction.

Why production takes too long:

- **AI prototypes succeed — production stalls**

The "last mile" of integrating with enterprise data becomes a permanent roadblock.

- **Heavy dependency on specialized teams**


Data scientists and AI builders wait weeks for data teams to prepare and pipeline data.

- **Governance and access slow deployment**

Security reviews, PII masking, and access policies must be manually re-implemented for every model.

💡 Core Challenge

AI scales when teams can build on trusted data without becoming data engineers.



The Architectural Shift

To scale AI, we must create a clear separation between reasoning and governed data access.

Fundamental changes to the delivery model:

- **Separate AI orchestration and reasoning**

Focus AI teams on model logic and application flow, rather than data wrangling.

- **Abstract data access & semantics**

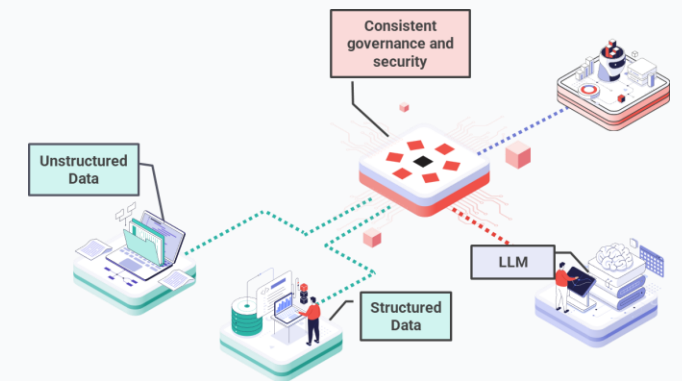
Move policy enforcement and schema complexity out of the AI code and into the platform.

- **Consume governed data products**

AI consumes data as ready-to-use, governed services rather than raw tables.

Key Principle


Treat data as a governed service, not a raw ingredient to be processed inside the model.



Separate the Agent From the Data Layer

Production-ready AI requires clear separation between reasoning and governed data access.





- Agent layer handles planning, memory, and orchestration
- Data layer manages connectivity, semantics, and policy
- Prevents hard-coded data logic in prompts
- Allows models and governance to evolve independently



What This Changes: Speed and Outcomes





By separating AI development from data complexity, organizations deliver faster while scaling value with control and confidence.

For Delivery Speed

-  **No need to interpret schemas or rebuild logic**
Eliminate redundant coding
-  **No custom security per use case**
Unified policy enforcement
-  **No manual performance engineering**
Automated query optimization
-  **Faster iteration and model evolution**
Rapid prototyping cycles

+

For the Business

-  **More teams able to build AI safely**
Democratized innovation
-  **Higher success rate moving to production**
From pilot to enterprise scale
-  **Real-time, context-aware AI decisions**
Better business intelligence
-  **Controlled cost and risk as AI scales**
Sustainable growth



Q&A



denodo
DATA
DAY
ROADSHOW



Thanks!

Come visit us at booth #2

Contact Carlisle Gomes cgomes@denodo.com



www.denodo.com

info@denodo.com

© Copyright Denodo Technologies. All rights reserved

Unless otherwise specified, no part of this PDF file may be reproduced or utilized in any for or by any means, electronic or mechanical, including photocopying and microfilm, without prior the written authorization from Denodo Technologies.