

AI Sovereignty Starts with Data Sovereignty

How to Gain Control of the AI Stack


AN EXECUTIVE WHITEPAPER FOR AI, DATA, TECHNOLOGY, RISK,
AND BUSINESS LEADERS

The provocative
question for every
executive sponsor:
Can you control the
data, context, policies,
workloads, models, and
agent actions behind
every AI outcome?



Contents

I 01	Executive Brief	4
I 02	What is AI Sovereignty?	6
I 03	Data Sovereignty Is the Cornerstone	7
I 04	The Global Regulatory and Best-Practice Landscape	8
I 05	Active Context as the Sovereignty Control Layer	9
I 06	How Denodo Maps to AI Sovereignty Requirements	10
I 07	Popular Approaches to AI Sovereignty and the Gaps Denodo Addresses	11
I 08	AI Sovereignty Self-Assessment	13
I 09	Practical Guidance by Score	14



01 Executive Brief

AI sovereignty does not simply rely on an AI model that operates in an environment you control. It is a control strategy across the entire AI stack, including where data is sourced, stored, and processed.

AI sovereignty is often discussed as a national policy issue: where models are built, whose infrastructure runs them, and which jurisdictions control the data they use. For enterprise leaders, however, AI sovereignty asks whether organizations can control, govern, explain, and trust the data that feeds the AI and whether the organization can, when necessary, control the context, workloads, models, policies, and agentic actions that shape AI outcomes.

Core thesis

AI sovereignty begins with data sovereignty and extends to data portability. If an organization cannot control where data resides, how it is accessed, how it is interpreted and contextualized by AI, where processing occurs, and how quickly data or workloads can move when conditions change, it cannot claim meaningful sovereignty over AI.

The recent AI Trust Gap Report, based on research conducted by Arlington Research, found several gaps between what trustworthy AI requires and what most organizations can deliver today: 66% of organizations say AI data must be real-time or sub-minute to be trustworthy; 63% struggle to identify trustworthy data or prepare and integrate it for AI; 67% struggle with AI data security and access controls; and enterprise AI initiatives draw from an average of more than 400 data sources. [1]

66%

say AI data must be real-time or no more than one minute old to be trustworthy

63%

struggle to identify trustworthy data or prepare and integrate it for AI initiatives

67%

struggle with AI data security and access controls

400+

average original data sources feeding enterprise AI initiatives

The Denodo Platform addresses these gaps by delivering the necessary AI data layer: a governed, semantic, real-time data access layer between distributed enterprise data and the AI systems that consume it. The Denodo AI data layer achieves this through active context, a logical tier that brings together an abstracted view of shared business semantics, live operational data access, lineage, runtime governance, and trusted integration logic across distributed enterprise systems. This abstraction also improves portability: data and workloads can move across cloud, on-premises, sovereign cloud, or edge environments without forcing every AI application, dashboard, API, or agent to be rebuilt. [2][3]

WHAT THE EXECUTIVE SPONSOR SHOULD TAKE AWAY:



AI sovereignty is not achieved by locking all supporting data and metadata into one platform; it is achieved by putting the organization in deliberate control of data, context, policy, workload placement, portability, and AI actions, wherever they reside.



Data sovereignty is the cornerstone of AI sovereignty because AI agents are downstream of data access, data meaning, data movement, and data governance.



Denodo is additive to sovereign cloud systems, private LLMs, AI gateways, lakehouses, catalogs, and governance platforms, because it solves the distributed enterprise data problem underneath them.

02 What is AI Sovereignty?

A practical definition for governments and corporations

AI sovereignty is the ability of an organization, government, or ecosystem to control the AI systems it depends on, including the data, models, infrastructure, policies, workloads, decisions, and actions that shape AI outcomes. It is not achieved simply by owning a model or hosting data locally. Those disciplines may be important, but they are not sufficient. Sovereignty also requires portability: the ability to move data, workloads, and services across providers, regions, cloud systems, on-premises environments, or sovereign infrastructure without breaking any downstream consumer processes.



For governments, AI sovereignty is a matter of national resilience, public trust, economic competitiveness, security, and control over strategic digital infrastructure. For corporations, it is about risk management, IP protection, vendor leverage, regulatory compliance, portability, cost control, operational accountability, and even business continuity.

The AI Leader's Challenge

If your AI depends on data you cannot locate, semantics you cannot reconcile, policies you cannot enforce, workloads you cannot deliberately place, and infrastructure from which there is no easy exit, then your AI is not sovereign. It is dependent on other actors by default.

03 Data Sovereignty Is the Cornerstone

A practical definition for governments and corporations

Data sovereignty means that data is subject to the laws, policies, governance requirements, and operational controls of the jurisdiction or organization responsible for it. In practical enterprise terms, it means knowing and controlling where data is stored, where it is processed, who can access it, what policies apply, how it is copied or moved, how portable it is, how it is interpreted, and how its use is monitored and audited.

The traditional data sovereignty question:

Where is the database, file, application, or data lake located?

Data sovereignty in the context of AI:

Where does data become prompts, vectors, retrieval context, agent inputs, logs, decisions, and actions?

AI makes data sovereignty more difficult because data is not used as it has traditionally been used: for reporting, analytics, and traditional business applications. In that traditional world, it was sufficient to focus on where the data itself was stored and processed. But with AI, it may also be embedded into prompts, converted into vectors, cached in retrieval systems, passed to model APIs, stored in logs, or exposed to autonomous agents at inference time. Each step in the AI agent's lifecycle can create a new sovereignty boundary.

What changes when AI becomes agentic:

The AI Trust Gap Report frames this as the shift from answering to doing: agents perceive conditions, decide what to do next, and execute actions by calling tools or interacting with systems of record. Agents break trust when they lack live situational awareness, act on the wrong data because of missing context or inconsistent semantics, or operate outside compliance and business guardrails. [1]

Design Principle

Data should not move merely because AI needs context. Data should move only when the organization deliberately decides that movement is the right architectural, regulatory, economic, or sovereignty choice. And when it does move, consumers should not be interrupted in their tasks.

04 The global Regulatory and Best-Practice Landscape

AI sovereignty does not simply rely on an AI model that operates in an environment you control. It is a control strategy across the entire AI stack, including where data is sourced, stored, and processed.

AI sovereignty is often a matter of national policy: Establishing where models are created, which hardware supports them, and what legal authorities regulate the data they process. But for enterprise leaders, AI sovereignty has a different meaning: It requires an understanding of whether or not organizations can control, govern, explain, and trust the data that supports the AI, and whether or not organizations can control the AI agent’s context, workloads, models, policies, and actions.

JURISDICTION / FRAMEWORK	WHAT IT EMPHASIZES	WHY IT MATTERS FOR SOVEREIGNTY
EU AI Act + GDPR + Data Act	Risk-based AI rules, personal data transfer controls, data access rights, cloud switching, and interoperability	AI governance, data transfer rules, and portability become architecture-level requirements. The Data Act applies since 12 September 2025; Chapter VI addresses switching between data processing services, and Article 29 phases out switching charges, with a full ban from 12 January 2027. [6][7][15][16]
United States: NIST AI RMF	Voluntary lifecycle framework for trustworthy AI	Provides enterprises with a common language for AI risk, trustworthiness, and controls. [8]
OECD AI Principles	Trustworthy, human-centered AI aligned with human rights and democratic values	Promotes globally interoperable AI policy principles. [9]
Singapore	Practical AI governance via AI Verify, GenAI, and agentic AI frameworks	Focuses on testing, transparency, agent boundaries, and practical assurance. [10]
United Kingdom	Pro-innovation, principles-based governance through existing regulators	Encourages sector-specific accountability without a single AI law. [11]
China	Generative AI service rules, security assessment, algorithm filing, content controls	Shows sovereignty as a matter of national security, public interest, and platform accountability. [12]
India DPDP Act	Digital personal data processing, consent, fiduciary obligations, and data rights	Extends data protection obligations into AI data pipelines and services. [13]
ISO/IEC 42001	AI management system standard	Provides a management-system approach for governing AI risks and opportunities. [14]

Implication

Regulatory compliance cannot be retrofitted at the model layer alone. The control plane must reach into the data and context layer, because that is where AI systems retrieve facts, interpret meaning, apply policies, and increasingly depend on portable infrastructure choices.

05 Active Context as the Sovereignty Control Layer

The Denodo AI data layer gives organizations agency over distributed data, semantics, governance, workload placement, and portability.

Denodo positions the AI data layer as the infrastructure between an enterprise's distributed data sources and the AI, analytics, applications, and business users that need to consume them. Denodo describes active context as the operational state created when enterprise context is live, governed, semantically unified, and available across distributed systems. Because consumers connect to logical data products, views, APIs, and semantics rather than directly to physical infrastructure, the underlying data location can change with far less disruption. [2] [3]

From Data Sovereignty to AI Agency

Denodo creates a control loop between distributed data and the AI systems that act on it



The four pillars of active context:

UNIVERSAL CONNECTIVITY

All sources. One access point.

ZERO-COPY DELIVERY

Live data. No unnecessary movement.

GOVERNED ACCESS

Security that travels with data.

SEMANTIC TRUST

One truth, for humans and AI.

These pillars are directly relevant to sovereignty. Universal connectivity gives organizations visibility across the full data estate. Zero-copy delivery reduces unnecessary replication, latency, residency exposure, and cost. Governed access applies policies at runtime and can enforce access constraints consistently. Semantic trust provides AI systems with consistent business definitions, source authority, and lineage. Together, they also enable portability: data can be repatriated, moved, or re-platformed while consumers continue to use a governed, logical layer. [2][3][4][5]

Agency, not Centralization

The goal is not to force every dataset, workload, or AI use case into one platform. The goal is to give the organization full agency: the ability to decide where data stays, where processing occurs, which policies apply, which sources are authoritative, and how quickly the organization can shift infrastructure without disrupting AI consumers.

06 How Denodo Maps to AI Sovereignty Requirements

A control and portability matrix for executive sponsors and architecture teams

AI SOVEREIGNTY REQUIREMENT	DENODO CONTRIBUTION	WHY IT MATTERS
Control where data comes from.	Connects across cloud, on-premises, and SaaS systems, APIs, warehouses, lakehouses, streams, and files.	AI is not limited to what has been copied into a single platform.
Control where data is maintained.	Provides live access without requiring unnecessary data movement or replication.	Data can remain in the right system, region, repository, or sovereign environment.
Control where processing occurs.	Logical query planning, pushdown, caching, and workload routing help place processing where it makes sense.	Organizations can balance residency, performance, cost, and operational requirements.
Control which AI or human consumers can access what data.	Centralized row-level security, masking, identity-aware access, and runtime enforcement.	AI agents inherit governed access rather than bypassing it through downstream copies.
Control business meaning and AI context.	Universal semantic layer, lineage, relationships, and authoritative business definitions.	Universal semantic layer, lineage, relationships, and authoritative business definitions.
Control reuse and scale.	Governed enterprise data products engineered once and consumed by AI, analytics, apps, and users.	Teams stop rebuilding integration, governance, and context for every AI initiative.
Control portability and exit.	Abstracts consumers behind stable logical views, semantics, APIs, and enterprise data products while the underlying data location or platform can change.	Organizations can repatriate data from cloud to on-premises, move between cloud systems, or shift processing without breaking AI agent processes, apps, dashboards, or APIs.

Denodo's active context enables practices such as data discovery, access management, security, privacy, data integration, preparation, monitoring, FinOps, and sharing, through a logical or virtual representation, instead of directly on each physical source. It also unifies data from diverse sources into a consistent, business-friendly framework with definitions, relationships, lineage, tags, and governance simplification. This logical abstraction is also a portability mechanism: it decouples AI and analytics consumers from the physical location of data and workloads. [3][4][5]

Portability Outcome

The goal is not to be independent of every cloud, model, lakehouse, or application. Each of these may each play critical roles in your data estate. Instead the goal is to avoid lock-in, lighten data gravity, and put portability in your control, not theirs. When the underlying data platform changes, Denodo offers a governed, live, semantic access layer that doesn't change, so consumers keep working while the infrastructure beneath them evolves.

07 Popular Approaches to AI Sovereignty and the Gaps Denodo Addresses

The right architecture is additive, not either-or.

APPROACH	HOW IT HELPS	GAP DENODO CAN ADDRESS
Sovereign cloud	Residency, jurisdiction, infrastructure control	Does not automatically solve multi-cloud, SaaS, on-premises access, semantics, or runtime governance.
Private or on-premises LLMs	Model control, prompt/data exposure reduction	Still need live, governed, semantically trusted enterprise context.
Data localization / replication	Regional control, regulatory alignment	Can create stale copies, unclear lineage, duplicate governance, and cost.
AI gateways	Model routing, observability, access to model APIs	Does not solve the data and context problem underneath the model.
Catalogs and governance tools	Discovery, inventory, stewardship, policy documentation	Metadata becomes active only when it governs access and meaning at runtime.
RAG and vector databases	Grounding AI in enterprise content	Can create uncontrolled copies or retrieval context without source authority and consistent policy.
Cloud repatriation / provider switching	Vendor leverage, cost control, operational resilience, and regulatory portability	Moving data solves only the infrastructure change. Denodo decouples consumers from the physical movement, reducing refactoring, downtime, and downstream disruption.

Sovereign cloud, private models, AI gateways, catalogs, data governance platforms, RAG, and cloud portability programs all have roles to play. But each becomes stronger when it is connected to a governed AI data layer that can reach across the full enterprise data estate without requiring unnecessary data movement or forcing consumers to be rewired when infrastructure changes.

Architectural Principle

Sovereign AI does not require all data to move into one sovereign platform. Instead it requires a sovereign control layer across the data estate: one that can connect, govern, interpret, deliver, and port context as infrastructure choices change.

A Wake-Up Call

The first wave of enterprise AI was about access to models. The next wave will be about control of context, and the freedom to move that context when strategy, regulation, cost, or risk requires it. Organizations that control their AI context will be able to innovate faster, comply with greater confidence, switch technologies with less friction, and trust AI agents to operate closer to the business. Organizations that do not will find that their AI ambitions are constrained by stale data, fragmented governance, semantic confusion, rising costs, platform dependency, and limited portability.

AI sovereignty starts with data sovereignty. Data sovereignty now also requires portability. And Denodo helps deliver both through a powerful AI data layer.



08 AI Sovereignty Self-Assessment


A campaign-ready diagnostic for executive sponsors

Score each question from 1 to 5: 1 = not true today; 3 = true for some major AI use cases; 5 = consistently true and auditable across the AI stack. Add the six scores, divide by 30, then multiply by 100.

QUESTION	WHAT A HIGH SCORE MEANS
1. Data control	You know where AI data resides, where it is copied, where it is processed, and which jurisdictions and policies apply.
2. Live context	AI can access current operational data from authoritative sources without waiting for batch pipelines or creating unnecessary copies.
3. Semantic trust	AI uses consistent business definitions, source-of-truth rules, lineage, and context across domains and systems.
4. Runtime governance	Access controls, masking, entitlements, residency rules, and usage policies are enforced at the point of AI access.
5. Workload placement	You can deliberately decide where AI-related data processing occurs: source, cache, cloud, on-premises, edge, or sovereign environment.
6. Portability and architecture choice	You can change models, cloud systems, lakehouses, vector stores, AI tools, and you can even repatriate selected cloud data/workloads on-premises without rebuilding the governed data foundation or disrupting consumers.

AI Sovereignty Self-Assessment Score

Add six 1-5 scores, divide by 30, then multiply by 100



09 Practical Guidance by Score

Use the assessment to prioritize architecture moves, not just governance paperwork.

SCORE	GUIDANCE
80-100 Sovereign by Design	You have strong control across data, context, governance, architecture, and portability. Focus next on scaling reusable, governed AI data products and continuously validating model, agent, data, and workload behavior.
60-79 Controlled, but Fragmented	You likely have strong controls in some platforms or use cases, but not across the full enterprise. Prioritize a unified AI data layer that extends governance, semantics, live access, and portability across sources.
40-59 Exposed by Architecture	Your AI strategy may depend on copied data, inconsistent semantics, manual controls, or platform-specific governance. Reduce unnecessary movement, centralize runtime policy enforcement, and decouple consumers from physical infrastructure.
Below 40 Dependent by Default	AI initiatives may be moving faster than your control model. Establish governed, live, semantically trusted access to the highest-risk data sources and workflows before scaling agents, and identify where vendor lock-in could block future exit strategies or repatriation.

Citations

[1] Denodo, [The AI Trust Gap Report](#), 2026.

[2] Denodo, [Active Context Layer: Definition, Importance, and Key Components](#).

[3] Denodo Platform / AI Data Layer.

[4] Denodo, [What is Logical Data Management?](#)

[5] Denodo, [Universal Semantic Layer](#).

[6] European Commission, [AI Act overview](#).

[7] European Data Protection Board, [International data transfers](#).

[8] NIST, [AI Risk Management Framework](#).

[9] OECD [AI Principles](#).

[10] AI Verify Foundation / IMDA, [Model AI Governance Frameworks](#).

[11] UK Government, [A pro-innovation approach to AI regulation](#).

[12] Cyberspace Administration of China, [Interim Measures for Generative AI Services](#).

[13] Government of India, [Digital Personal Data Protection Act, 2023](#).

[14] ISO/IEC 42001:2023, [AI management systems](#).

[15] European Commission, [Data Act explained](#).

[16] Regulation (EU) 2023/2854, [Article 29, gradual withdrawal of switching charges](#).

Citations are public-facing references. Regulatory references are summarized for business and architecture context and should not be treated as legal advice.

NEXT STEP:

Contact us for a free consultation
on how Denodo can help to
close the AI trust gap in your
organization.



Visit www.denodo.com | Email info@denodo.com | Discover community.denodo.com

