

The Four Layers of Trusted Enterprise AI

An executive briefing for governing AI beyond the model

Enterprise AI is often discussed as though trust can be delivered by the model alone. It cannot. A capable model may understand and recommend, but the organization must still determine which information it should use, what it is permitted to do and how the resulting decision can be reconstructed. This briefing separates those responsibilities into four layers.

01 Intelligence

What the AI can interpret, reason about and generate

02 Context

The governed business data and meaning supplied to it

03 Authority

What the AI is permitted to access, recommend or do

04 Evidence

What allows the decision to be understood and defended

The model may provide the intelligence, but the enterprise remains responsible for the context, authority and evidence surrounding every decision.

Cameron Price · Founder and CEO, Data Tiles

02 — THE FOUR LAYERS

Four connected responsibilities. None can substitute for another.

01 — INTELLIGENCE

What can the system understand, infer, recommend or generate?

Intelligence is the model's ability to interpret information, recognize patterns, reason about a task and recommend actions. Capability answers what the system can do. It does not answer what the organization should allow it to do.

02 — CONTEXT

What information should the system use for this particular decision?

Context is the trusted, fit-for-purpose information a decision requires: approved definitions, relationships, quality expectations, business rules and the purpose for which the information may be used. More accessible data is not always better context. Trusted context is selected, understood and governed for a purpose.

03 — AUTHORITY

What may the agent access, change, communicate or initiate?

Authority separates the ability to propose an action from permission to carry it out. It is enforced through enterprise identity, permissions, workflow, approvals and runtime controls, not through instructions the model could reinterpret.

04 — EVIDENCE

Can the organization reconstruct what happened and why?

Evidence connects what informed the decision with the policy, approval and action that produced the outcome. A persuasive explanation generated afterward is not the same as an independent record.

Intelligence without context can be confidently wrong. Context without authority cannot safely become action. Authority without evidence cannot be defended.

The governed data product is the central business asset.

Within this framework, the governed, decision-ready data product created in Latttice is the central business asset. It supplies the trusted context and data-product evidence that people, applications, workflows and AI can reuse, bringing together business meaning, governed enterprise data, lineage, quality expectations, policy and access conditions for a defined purpose.

WHAT A DECISION-READY DATA PRODUCT PROVIDES

Approved definitions	Current, fit-for-purpose information
Business rules and policies	Ownership and access conditions
Quality, lineage and provenance	Evidence supporting the decision

Lenz is one environment for designing AI agents around a defined business decision and the trusted data products they need. The lasting enterprise asset is the governed data product. It can be reused across people, BI, applications, workflows and AI without depending on one model or agent builder.

Latttice does not replace existing data platforms, catalogs, identity systems or infrastructure controls. Data Tiles does not replace the enterprise control environment. Latttice helps AI operate from governed business context and preserves data-product evidence that can be combined with enterprise approvals, actions and logs.

The model should contribute intelligence. The enterprise must provide the context, retain authority and preserve the evidence.

Five questions every leadership team should ask.

- 01 What important business decision is the AI supporting?
- 02 Which governed data product supplies its business context?
- 03 Who determines what the AI is authorized to access or do?
- 04 What evidence will be retained about the recommendation or action?
- 05 Could the organization explain and defend the resulting decision?

If a leadership team cannot answer all five questions, the issue is not necessarily that the AI model is inadequate. One of the surrounding enterprise layers may be missing.

Build the trusted business context your AI needs.



Read the full framework.

data-tiles.com/four-layers-trusted-enterprise-ai

Explore Lattice

letstalkdata@data-tiles.io