


DATA GOVERNANCE FOR THE AI ERA

Governing for Answers The Proposal

A proposal for data governance in the AI era: five pillars
that move governance from the dataset to the answer.



AUTHOR
DUNCAN BROWN

STATUS
WORKING DRAFT V1.1

DATE
OCTOBER 2026

Govern the answer, not just the dataset.

When AI turns data into fluent answers, every weakness in meaning, ownership and quality reaches people as something they believe. This paper proposes a five-pillar model that moves governance from the dataset to the answer, without abandoning the fundamentals that still do the heavy lifting.

01 Meaning Business definitions machines can read, not just people.	02 Scope Governance extended to documents, prompts, vector stores and AI outputs.	03 Automation Rules enforced in pipelines as code and contracts.	04 Accountability A named owner for every AI use case and agent.	05 Assurance Evidence mapped to regulation, produced as a by-product.
---	---	--	--	---

What this paper argues

- AI has removed the human checkpoint that used to catch bad data.
- Governance must follow data into documents, prompts and outputs.
- Rules that live in PDFs cannot keep pace; rules in pipelines can.
- Every AI use case needs an owner, and every control needs evidence.

What it gives you

- A model you can explain in one diagram (sheet 3).
- A one-question risk tier to keep effort proportionate.
- A 12-month roadmap with two clear gates.
- Seven measures that track trust, not paperwork.

60%

of AI projects unsupported by AI-ready data will be abandoned through 2026, Gartner predicts. The same research found 63% of organisations lack, or are unsure they have, the right data management practices for AI. · GARTNER, FEBRUARY 2025

A dashboard gets questioned. An AI answer gets believed.

For twenty years, governance has had a silent partner: the analyst who spots a strange number, asks which definition applies, and checks before a figure reaches a board. AI removes that checkpoint.

SHIFT 1

The consumer has changed

Data is increasingly read by models and agents, not analysts. A model cannot ask which of three occupancy measures you meant.

SHIFT 2

The estate has changed

Retrieval turns shared drives, wikis and policy libraries into production data sources overnight. Most were never owned, classified or checked.

SHIFT 3

The pace has changed

An agent can make thousands of decisions an hour. Quarterly reviews and steering committees were designed for a human pace of change.

Today's typical governance model struggles on all three. It is organised around structured data in a warehouse, enforced through documents and reviews, and brought in after a use case has been built. Each of those habits leaves a gap that AI exposes, and exposes at scale.

That is also the opportunity. The question the business asks is no longer “why does governance matter?” It is “can we trust what the AI just told us?” Only governance can answer that, and answering it well puts data governance at the centre of every AI programme rather than at its edge.

The fundamentals don't change. What changes is where they are applied, how fast, and by whom.

Every AI answer has a supply chain.

An AI answer is assembled from sources, shaped by definitions, filtered by retrieval, generated by a model following a prompt, then acted on by a person or an agent. A failure at any link reaches the user looking exactly like a correct answer.




FIG. 1 – THE ANSWER SUPPLY CHAIN, WITH THE FIVE PILLARS MAPPED ONTO IT

LINK	WHAT GOES WRONG	PILLAR THAT CLOSSES THE GAP
Sources	An unowned, out-of-date document is indexed alongside the current one	Scope
Meaning	Three legitimate measures share one name; the model picks one silently	Meaning
Data in motion	A feed changes format and bad values flow through for weeks	Automation
Retrieval	The index returns content the person asking should not see, or an approved source that isn't the right one for the question	Scope
Model + prompt	A prompt is edited in production with no review or version	Scope
Answer	It is wrong, and no one knows whose job it is to fix it	Accountability
Decision	An agent acts on it, and nobody can later show why	Accountability · Assurance

This is why the proposal is framed around answers. Each pillar governs a link in the chain; together they make the whole chain trustworthy and provable.

Five pillars on the foundations you already have.

None of the pillars replaces ownership, classification, quality or lineage. Each one takes those fundamentals somewhere AI has made them newly necessary, and each ends with a single test.




FIG. 2 – GOVERNING FOR ANSWERS: FIVE PILLARS ON THE FUNDAMENTALS

01 Meaning: definitions AI can read

Business definitions usually live in glossaries that people rarely open. AI tools need them as working context, because a model will infer meaning from column names and whatever documents it retrieves.

- Treat the glossary and semantic layer as production infrastructure: versioned, owned and tested.
- Give every critical metric an unambiguous identity, one governed meaning and a stated context of use, starting with the 20–50 terms that drive decisions, and link each to the columns that implement it.
- Expose certified definitions to AI tools through the catalogue or semantic layer, so answers use agreed meaning and say which one.
- Mark uncertified fields so AI can caveat an answer or decline to give one.

IN PRACTICE

A fund's investor report shows occupancy at 94%; the asset dashboard shows 89% for the same portfolio and date. One counts leases signed but not started; the other only occupied space. Both are legitimate. An AI assistant will pick one and not say which.

THE TEST Can your AI assistant tell a user which definition it used, and why?

02 Scope: govern what AI actually reads

The AI estate is wider than the warehouse. Governance has to follow the data wherever AI picks it up: documents, prompts, vector stores and the outputs that get pasted back into reports.

- Keep one inventory of AI use cases and agents, including AI switched on inside existing SaaS tools, with every source each one reads.
- Govern what AI is allowed to retrieve through an approved-sources list, rather than trying to govern every document.

- Give each approved source an owner, a review date and a statement of what it is authoritative for. Approved does not mean appropriate for every question.
- Apply the user's own access rights at query time, so AI never returns content the person asking cannot see.
- Manage prompts as configuration, and label AI outputs that re-enter systems. They are not authoritative until certified.

IN PRACTICE

A lease assistant is pointed at a shared drive. Asked about a break date, it quotes “Lease_v3_FINAL”, an unsigned draft. The executed lease sits in a subfolder the index never reached. Nobody owned the folder.

THE TEST For any AI answer, can you say which documents it drew on, who owns them, and why they were the right ones?

03 Automation: governance as code

Policies written in documents cannot keep pace with AI. Rules have to run where the data moves, and critical failures have to stop bad data rather than describe it afterwards. Governance defines the rule; engineers make it executable.

- Agree data contracts between producers and consumers, and count AI applications as consumers.
- Run quality checks as code in the pipeline; critical failures block the load.
- Add AI-specific checks where they matter: representativeness, bias, freshness for the decision, authority of sources and leakage.
- Enforce classification, masking and access at query time on the platform.
- Give every agent its own identity and least-privilege access that expires.

IN PRACTICE

A property manager changes its rent roll export and one building's rent arrives monthly instead of annually. It flows into the valuation model and investor pack for six weeks. A range check on rent per square foot would have stopped it on day one.

THE TEST What share of your critical quality rules run without a person starting them?

04 Accountability: an owner for every answer

Data owners are accountable for inputs. Someone must also be accountable for what the AI does with them, or a wrong answer becomes everyone's problem and nobody's job.

- Name a business owner for every AI use case, alongside the data owners of its sources. They are accountable for the decision to deploy, for the controls, and for getting failures fixed, not for checking every output.
- Name an owner for every agent, with written limits on what it may do and a tested way to stop it.
- Add a governance checkpoint before go-live, scaled to the use case's risk tier.
- Give users a “flag this answer” route, and treat each flag as a data issue with a root cause.

IN PRACTICE

An AI-drafted investment committee paper overstates a portfolio's WAULT by 1.4 years. The data team says the data is right, the AI team says the model summarised what it found, and the deal team says it doesn't own the tool.

THE TEST **When an AI answer is wrong, does anyone know whose job it is to fix it?**

05 Assurance: evidence as a by-product

Regulators, auditors and investors will ask for proof. The cheapest proof is the kind the other four pillars create automatically, as the work happens.

- Capture lineage, check results, approvals and answer logs with timestamps as the work happens.
- Keep one control library, each control mapped once to every framework it satisfies.
- Scale evidence to the risk tier: a full pack for high risk, a short record for low.
- Log the model, prompt and sources behind each material high-risk decision.

IN PRACTICE

An investor's due diligence questionnaire asks which AI tools touch valuation and reporting data, who approved them and how they are monitored. Answering takes six weeks and four teams. With evidence captured by default, it takes days.

THE TEST **Could you produce evidence for a regulator in days rather than months?**

One question sets the risk tier.

What is the worst a wrong answer could do? Governance effort should follow the answer. A team summarising meeting notes should not wait weeks for a review; a model supporting credit or hiring decisions should never skip one.

“Wrong” covers four things: wrong content, a correct answer shown to the wrong person, a correct answer built on data used out of context or on the wrong source, and an action an agent should not have been allowed to take.

TIER	WHAT A WRONG ANSWER COULD DO	EXAMPLE	GATE BEFORE GO-LIVE
Low	Cause internal inconvenience that a person would spot	Summarising meeting notes	Self-certify against a short checklist
Medium	Mislead a business decision	Portfolio analytics assistant	Review by the domain's governance lead
High	Affect customers, investors, employees, financial reporting or a regulated decision	Tenant screening, investor reporting, valuations	Full review, impact assessments and senior sign-off

FOUR RULES

If the AI acts on its own, move it up a tier. An agent that sends, changes or approves things without a person checking carries more risk than one that drafts for review. **Likely high-risk under the EU AI Act means high. If it can reach restricted or sensitive personal data, it is at least medium. Re-tier when the model, the sources, an agent's permissions or the use of its output change.**

Familiar roles, new responsibilities, and two new ones.

Governance should be designed into AI use cases at the start, not reviewed at the end.

ROLE	ACCOUNTABLE FOR	WHAT CHANGES IN THE AI ERA
Data owner	Quality, access and meaning of a data domain	Includes documents and knowledge sources, not just tables
Data steward	Definitions, classification and issues	Maintains the definitions AI will use; triages flagged answers
AI use-case owner · new	An AI use case: deploying it, its controls and getting failures fixed	Owens the purpose, the prompts and the approved sources; makes sure the right owner fixes a wrong answer
Agent owner · new	What an agent is allowed to do, and what it does	Sets its permissions and limits, and can switch it off
Data and platform engineers	Pipelines, contracts and platform controls	Build checks, contracts and policies as code
Governance office	Framework, standards and assurance	Defines the rules; runs the gate, the inventory and the evidence
Risk, legal and compliance	Regulatory interpretation	Confirms tiers; maps controls to regulation

Every new AI use case follows the same light path:



FIG. 3 — THE PATH FOR EVERY NEW AI USE CASE

The deadline moved. The work didn't.

The EU's Digital Omnibus on AI (Regulation (EU) 2026/1744) came into force on 27 July 2026. It deferred the heaviest high-risk obligations and left others on their original dates.




FIG. 4 – KEY EU AI ACT DATES AS AMENDED, WITH THE UK'S AUTOMATED-DECISION REFORMS

Transparency duties under Article 50 apply by function rather than risk, so chatbots and content generators are in scope now. Annex III covers areas such as employment, credit and insurance, where many firms already use AI. Treating the delay as a reason to pause inventory and classification would be a mistake: that work is slow, and it is squarely governance's job.

UK No AI statute UK GDPR, as amended by the Data (Use and Access) Act 2025, governs personal data in AI, including new safeguards for automated decisions since February 2026.	STANDARD ISO/IEC 42001 The AI management system standard gives Assurance a certifiable structure. ISO/IEC 42005 adds a method for AI impact assessment.	BODY OF KNOWLEDGE DAMA-DMBOK 3.0 A community update that will bring AI into the canon. This paper aims to complement it with practical patterns.
---	--	---

This section summarises public sources and is not legal advice. Firms should confirm their position with their own legal and compliance teams.

Twelve months from inventory to evidence.

The first 90 days are about visibility: you cannot govern AI you haven't found. Controls and assurance follow, each unlocked by a simple gate.



FIG. 5 — THREE PHASES, TWO GATES

Gate one: every live AI use case and agent is in the inventory with a named owner. **Gate two:** no high-risk use case can reach production without passing the review.

Measure trust, not paperwork.

Each pillar gets one or two measures, baselined in the first 90 days. Set targets locally: a number borrowed from another organisation says little about your own starting point.

PILLAR	MEASURE	DIRECTION
Meaning	Critical metrics with one governed meaning linked to source columns	Up
Scope	Live AI use cases and agents in the inventory with a named owner	Up, to 100%
Scope	AI knowledge sources with an owner, an in-date review and a stated authority	Up
Automation	Critical quality rules running automatically in pipelines	Up
Automation	Incidents caught by a check before they reach an AI answer	Up
Accountability	Median time from a flagged AI answer to a fix at source	Down
Assurance	Time to produce evidence for one use case on request	Down

A proposal is only credible if it names its weak points.

“It’s just good governance with AI in the title.”

Partly true, and deliberate. The fundamentals don't change. Scope, speed and the consumer do, and they change what good looks like in practice.

Governing documents may not scale.

Firms hold millions of files. The answer is to govern only what AI is allowed to retrieve, through an approved-sources list.

Governance as code needs engineers.

Many governance teams are staffed for policy, not pipelines. Governance defines the rules and engineers build them, so the model may need a shared team with data engineering, or a slower path for smaller firms.

New owners could blur accountability.

If a use-case owner and a data owner disagree about a wrong answer, someone must have the final say. This needs testing in real organisations.

Open questions for the community

1. Should AI outputs that re-enter systems be governed like any other data, or as a separate class?
2. Who should own prompts and system instructions: the use-case owner, IT, or governance?
3. Is a single semantic layer realistic, or will each AI platform keep its own?
4. Can an AI tool reliably choose the right governed meaning for a question, and explain why?

ACKNOWLEDGEMENTS

Version 1.1 was shaped by challenges from [Jess Baker](#), whose questions on metric identity, source authority, who builds the rules and where accountability sits changed the Meaning, Scope, Automation and Accountability pillars. Contributors are credited in a personal capacity. The views, and any errors, are the author's.

10 · REFERENCES

1. Gartner, [Lack of AI-ready data puts AI projects at risk](#), February 2025.
2. European Parliament and Council, [Regulation \(EU\) 2026/1744 \(Digital Omnibus on AI\)](#), Official Journal, 24 July 2026.
3. Cloud Security Alliance, [EU AI Act's high-risk deadline: deferred, not cancelled](#), August 2026.
4. UK Parliament, [Data \(Use and Access\) Act 2025](#); ICO, [What the DUAA means for organisations](#).
5. ISO, [ISO/IEC 42001:2023](#) and [ISO/IEC 42005:2025](#).
6. DAMA International, [DAMA-DMBOK 3.0 Project](#).

“In practice” examples are illustrative scenarios drawn from commercial real estate and investment management, not accounts of specific organisations.

This is a working draft. Challenge it.

The most useful feedback is disagreement: where a pillar is missing something, where the maturity levels don't match reality, and where this model would fail in your organisation. Comment on the LinkedIn series or message me directly. Contributors will be credited in the next version if they wish.

