




DATA GOVERNANCE FOR THE AI ERA

Governing for Answers The Framework

Governing for Answers in practice: principles, practices, maturity levels, metrics, regulatory mapping and 90-day checklists for each pillar.



AUTHOR
DUNCAN BROWN

STATUS
WORKING DRAFT V1.1

DATE
OCTOBER 2026

From the dataset to the answer.

This framework turns the five pillars of *Governing for Answers* into practices you can adopt, assess and evidence. It is a working draft, published to be challenged.

00 How to use this framework

Structure, maturity levels, risk tiers and what changed in v1.1

01 Meaning

Definitions AI can read

02 Scope

Govern what AI actually reads

03 Automation

Governance as code

04 Accountability

An owner for every answer

05 Assurance

Evidence as a by-product

06 Self-assessment scorecard

Score yourself and set targets

07 Feedback and references

How to challenge this draft

Five chapters, one structure.

Each of the five chapters covers one pillar, and every chapter follows the same structure:

- **The problem:** the gap AI exposes.
- **Principles:** what good looks like, in a few statements.
- **Practices:** the controls and patterns that deliver the principles.
- **Maturity:** four levels to place yourself on.
- **Metrics:** how to tell whether it is working.
- **90-day starter checklist:** where to begin.
- **Regulatory mapping:** which obligations the pillar supports.

The same four maturity levels apply across all five pillars:



Level 4 is not the goal for every pillar in every organisation, and it does not mean removing people. Some of the strongest controls are deliberate human checkpoints; what makes them level 4 is that they are built into the work and can't be skipped. The right target depends on your AI risk. A firm using AI only for internal productivity may be well served at level 2 or 3.

WHAT CHANGED IN V1.1

- **Meaning:** every critical metric now needs an unambiguous identity and a stated context of use, so legitimate variants are separate terms, not rival definitions. Owners can be named per family of terms.
- **Scope:** each approved source states what it is authoritative for. Approved does not mean appropriate for the answer.
- **Risk tiers:** a wrong answer now covers wrong content, audience, basis or action, with a rule for sensitive data.
- **Automation:** governance defines the rules and engineers build them in; governance does not have to engineer them itself.
- **Accountability:** what the use-case owner is and is not accountable for is spelled out.
- **Maturity:** level 4 is renamed Embedded, and the scorecard has a scoring rule and guidance on targets.

These changes came from challenges by [Jess Baker](#). See Acknowledgements.

One question sets the tier.

Every AI use case and agent is rated low, medium or high risk, and the tier decides how much governance it needs. Ask one question: what is the worst a wrong answer could do?

An answer can be wrong in four ways. Rate the use case on the worst of them.

- **Wrong content:** the answer is incorrect.
- **Wrong audience:** a correct answer reaches someone not entitled to see it.
- **Wrong basis:** a correct answer is built on data used outside its permitted purpose, or on a source that isn't the right one for the question.
- **Wrong action:** an agent does something consistent with its inputs that it should not have been authorised to do.

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

Then apply four rules:

1. **If the AI acts on its own, move it up a tier.** An agent that sends, changes or approves things without a person checking each action carries more risk than one that drafts for a person to review.
2. **Likely high-risk under the EU AI Act means high.** Treat anything that may fall under Annex III as high, so it is ready before those obligations apply from 2 December 2027. Practices banned under Article 5 are never permitted.
3. **If it can reach restricted or sensitive personal data, it is at least medium.** A correct answer shown to the wrong person is still a failure.
4. **Re-tier when things change:** a new model, new sources, new permissions for an agent, or its output being reused somewhere new.

USING THE TIERS IN THIS FRAMEWORK

Where a practice or control says “high-risk”, it applies to high-tier use cases; “medium or high” applies to both. Low-tier use cases need only the baseline: an entry in the inventory, a named owner and a self-certified checklist.

01 Meaning

Definitions AI can read

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

AI can only be as consistent as the definitions it is given. When one name covers three meanings, an AI tool will pick one silently and present it as fact.

The problem

Most organisations have a glossary. Few have one that is complete, current, linked to the data, and used. Often the meanings are not wrong, just different: several legitimate measures share one name. Analysts cope by asking colleagues which one applies. AI tools cannot ask, so they infer meaning from column names, documents and whatever else they retrieve.

IN PRACTICE

A fund's quarterly investor report shows occupancy at 94%, while the asset management dashboard shows 89% for the same portfolio on the same date. One counts leases signed but not yet started; the other counts only space tenants physically occupy. Neither is wrong. An AI assistant asked “what is the fund's occupancy?” will pick one without saying which. Fixing it means certifying each measure as its own term, here physical and leased occupancy, plus economic occupancy, which weights by rental value. Each has its own identity, calculation and context of use, and the bare word “occupancy” is retired.

PRINCIPLES


1. Every critical metric has an unambiguous identity, one governed meaning and a stated context of use.
2. Every critical term, or family of related terms, has a named owner.
3. Definitions are linked to the physical data that implements them.
4. Definitions are available to machines, not just readable by people.
5. AI tools choose the governed meaning that fits the question, say which they used, and say so when meaning is uncertain.

Practices

- **Critical term register.** Identify the 20 to 50 terms that drive the most important decisions and reports. Give each a unique identifier and short code, so legitimate variants of one concept are distinct terms rather than competing definitions. Certify these first rather than attempting the whole glossary.
- **Context of use.** For each term, record when to use it, when not to, and the terms it is confused with. This is what lets a person, or an AI tool, choose between legitimate variants.
- **Ownership by family.** Group related terms into families, such as occupancy, lease term and rent, and name one owner per family, typically the owner of the data product that serves them. Owners certify; stewards maintain.
- **Definition to data linkage.** Link each certified term to the tables, columns and calculations that implement it in the catalogue, and to its source of record on the approved-sources list.
- **Semantic layer as the source of truth.** Define metrics once, in a governed semantic or metrics layer, and have dashboards and AI tools query through it.

- **Machine-readable context.** Expose certified definitions, context of use, synonyms and caveats to AI tools, for example through the catalogue's API or as retrieval context.
- **Certification status as a signal.** Tag assets as certified, draft or deprecated, so AI tools can prefer certified sources and caveat the rest.
- **Change control.** Version definitions, and notify consumers, including AI use-case owners, when one changes.

Maturity



02 Scope

Govern what AI actually reads

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

Governance must follow the data AI uses, wherever it lives. Today that means documents, prompts, vector stores and AI outputs, not just the warehouse.

The problem

Retrieval-augmented generation (RAG) turns shared drives, wikis and policy libraries into production data sources overnight. Most were never owned, classified or checked for freshness. Prompts and system instructions shape every answer, yet often sit in code or config with no review. AI outputs are pasted into reports and emails, becoming new data with no record of where they came from.

Approval is also not the same as fitness. Five approved, current, well-owned sources can each be authoritative for a different purpose. Governing what AI *may* read is half the job; the other half is what it *should* read for the question in front of it.

IN PRACTICE

A lease question-and-answer assistant is pointed at the asset team's shared drive. Asked when a tenant can break, it quotes a date from "Lease v3 FINAL", an unsigned draft. The executed lease, with a different break date, sits in a subfolder the index never reached. Nobody owned the folder, so nobody had marked which version was authoritative.

PRINCIPLES


1. Every AI use case is known, with its purpose, owner and sources.
2. AI may only retrieve from approved sources.
3. Every approved source has an owner, a sensitivity class, a review date and a statement of what it is authoritative for.
4. Approved does not mean appropriate: for each kind of question there is a source of record, and AI uses it.
5. Any AI answer can be traced back to the sources it used.
6. Prompts are governed as configuration. AI outputs are labelled as such, and are not authoritative until certified.
7. AI retrieves only what the person or agent asking is entitled to see.

Practices

- **AI use-case inventory.** Register each use case and each agent, including AI features switched on inside SaaS tools you already use. Record its owner, purpose, model, risk tier and every source it retrieves from or was tuned on.
- **Approved-sources list.** Govern what AI is allowed to read rather than trying to govern every document. Retrieval connects only to sources on the list.

- **Source metadata and authority.** For each approved source, record an owner, sensitivity class, review date and what it is authoritative for: the subjects, terms and purposes it is the source of record for. Everything else it holds is reference only.
- **Fit-for-the-answer routing.** Bind each certified term to its source of record, and give each AI use case a context pack that says which term and source to use for which kind of question. Have the tool state which it used and why, and test it with questions that have more than one plausible source.
- **Permission-aware retrieval.** Apply the requesting user's or agent's access rights at query time, not only when the index is built. An index built with a service account can hand people content they are not entitled to see.
- **Retrieval lineage.** Log which documents and chunks supported each answer, along with the version of the chunking and embedding pipeline.
- **Prompt management.** Store prompts and system instructions in version control, with an owner and a review before changes go live.
- **Output provenance.** Label AI-generated content that re-enters reports or systems, and treat it as a derived dataset where it is stored. It becomes authoritative only by the same route as any other source: an owner, a review and approval on the approved-sources list.
- **Discovery scans.** Periodically scan for unregistered AI use, such as new connectors to shared drives or AI features switched on in existing tools.
- **Third-party model check.** Before data is sent to an external model provider, confirm how long it is retained, whether it can be used for training, and where it is processed.

Maturity



Metrics

- Share of live AI use cases registered with a named owner (target: 100%).
- Share of approved AI sources with an owner, an in-date review and a stated authority.
- Share of AI answers that can be traced to their sources.
- Number of unregistered AI use cases found by discovery scans.
- Share of retrieval systems that pass a permission test (a user cannot retrieve content they are not entitled to see).

90-DAY STARTER CHECKLIST

- ☐ Launch the AI use-case inventory and ask every team to register
- ☐ Identify the top ten sources feeding AI tools today
- ☐ Name an owner, a review date and what each is authoritative for
- ☐ Move one RAG use case to an approved-sources list
- ☐ Put the prompts for that use case under version control

Regulatory mapping

- **EU AI Act, Article 50:** transparency, including labelling AI-generated content.
- **EU AI Act, high-risk obligations:** an inventory is the precondition for classifying systems.
- **UK GDPR:** purpose limitation, and knowing where personal data flows into AI tools.
- **ISO/IEC 42001:** defining the scope of the AI management system and its AI systems.

03 Automation

Governance as code

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

Rules that run automatically in the flow of data are the only rules that keep pace with AI. Governance requirements have to become executable: governance defines the rule, engineers build it in, and owners set the thresholds.

The problem

Most governance frameworks are enforced through documents, reviews and people remembering to follow them. Quality issues are often found after the fact, in a report or a complaint. An AI agent working at machine speed will have used bad data thousands of times before a quarterly review notices.

IN PRACTICE

A third-party property manager changes its rent roll export, and one building's passing rent arrives monthly instead of annually. The figure flows into the valuation model and the quarterly investor pack. A fund controller spots the drop in rent per square foot six weeks later. A pipeline check with a plausible range for rent per square foot would have stopped the load on day one.

PRINCIPLES

1. Critical rules are expressed as code, versioned and tested like any other code.
2. Checks run where data moves, and critical failures stop bad data rather than just report it.
3. Producers and consumers agree expectations up front in data contracts.
4. Access and masking policies are enforced by the platform, not by trust.
5. Metadata flows to the catalogue automatically.


WHO	THEIR PART IN MAKING A RULE EXECUTABLE
Governance office	Defines the requirement and its severity
Business and data owners	Set the thresholds and accept exceptions
Data and platform engineers	Build the checks, contracts and policies into pipelines and platforms
Security	Enforces identity and access
AI engineers	Own guardrails and retrieval checks

Practices

- **Quality checks as code.** Define checks for completeness, uniqueness, validity, freshness and volume in a declarative checks tool, stored in version control next to the pipeline.

- **Severity-based responses.** Classify each check: critical failures block the load, warnings alert the owner, and informational checks feed trend reporting.
- **Data contracts.** Agree schema, semantics, quality thresholds and freshness between producing and consuming teams. Treat AI applications as consumers with their own contracts.
- **Policy as code.** Use tag-based classification with masking, row-level access and retention rules enforced at query time on the data platform.
- **Metadata automation.** Have pipelines publish lineage, check results and schema changes to the catalogue, so it reflects reality without manual upkeep.
- **Retrieval checks.** For RAG pipelines, check that source documents are in date, approved and correctly classified before they are indexed.
- **Fitness-for-AI checks.** Traditional checks are not enough for AI. Add, by use case: **representativeness** (does the data reflect the situations the AI will meet?), **bias** (are groups over- or under-represented?), **decision freshness** (is the data current enough for this decision?), **authority** (is this the approved, current version, and the right source for this question?) and **leakage** (does it contain anything that should not be exposed?). Set thresholds per use case, not per dataset.
- **Guardrails on prompts and outputs.** Scan what goes into and comes out of AI tools for personal and confidential data, and redact or block according to its classification.
- **Agent permissions as code.** Give every agent its own identity, never a person's or a shared service account, with least-privilege access that expires. Central sign-in standards such as the Model Context Protocol's enterprise-managed authorisation control which systems an agent can connect to, but not what it does once connected, so action limits need their own rules.

Maturity



Metrics

- Share of critical quality rules running automatically in pipelines.
- Share of AI data feeds covered by a data contract.
- Number of incidents stopped by a check before reaching a consumer.
- Share of sensitive columns protected by platform-enforced policies.

90-DAY STARTER CHECKLIST

- ☐ Pick the three data feeds most used by AI tools
- ☐ Write critical checks as code for each and run them in the pipeline
- ☐ Agree a simple contract with the producer of one feed
- ☐ Make one critical check block the load on failure
- ☐ Baseline the share of critical rules that are automated

Regulatory mapping

- **EU AI Act, Article 10:** examining data quality, completeness and bias for high-risk systems.
- **UK GDPR, Articles 5 and 25:** accuracy, and data protection by design.
- **ISO/IEC 42001:** operational controls and monitoring for AI systems.

04 Accountability

An owner for every answer

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

Someone must be accountable for what an AI application tells people, not just for the data it reads. Without that, a wrong answer becomes everyone's problem and nobody's job.

The problem

Traditional governance assigns owners to data domains. AI adds a new layer: a use case combines several sources, a model, prompts and a user interface. When the answer is wrong, the fault could sit in any of them. Data owners say the data was fine; the AI team says the model did what it was asked. The issue goes unresolved.

IN PRACTICE

An AI-drafted investment committee paper overstates a portfolio's weighted average unexpired lease term (WAULT) by 1.4 years. The data team says the lease data is right. The AI team says the model summarised what it retrieved. The deal team says it does not own the tool. The error is caught at committee, but the root cause, an expired lease still marked active, stays unfixed for a month because no one owns the answer.

What “an owner for every answer” means

It does not mean one person checks every output. The use-case owner is accountable for three things: the decision to deploy the use case and what people may rely on it for; the controls being in place and working; and failures being fixed at source. Everyone else stays accountable for their own part.

PART OF THE ANSWER	WHO IS ACCOUNTABLE
The decision to deploy, and what the AI may be relied on for	Use-case owner
How an answer is used in a decision	The person who uses it
The data and sources it reads	Data owners
The meaning of the terms it uses	Term owners
How the model, prompts and retrieval behave	Use-case owner, through a named technical lead
What an agent is allowed to do	Agent owner
Each control	The control's owner
Getting a failure fixed, whoever's part it sits in	Use-case owner

PRINCIPLES

1. Every AI use case has a named owner, accountable for deploying it, for its controls and for getting failures fixed.
2. Data owners remain accountable for the sources the use case reads.
3. People remain accountable for how they use an AI answer in a decision.
4. Every agent has a named person accountable for what it is allowed to do, and what it does.
5. Governance is engaged before go-live, in proportion to risk.
6. Users can report a wrong answer, and every report reaches an owner.
7. Final decisions on disputed issues have a clear escalation route.

Practices

- **AI use-case owner role.** A business owner, not IT, accountable for the use case's purpose, approved sources and prompts, and for getting failures fixed. Pair them with a named technical lead, so the owner does not need to be an engineer and the engineer does not carry business accountability.
- **Agent owner role.** A named person for each agent, accountable for its identity, permissions, limits and behaviour. Automation carries out controls; it never owns the outcome.
- **RACI per use case.** Record who is responsible, accountable, consulted and informed across use-case owner, data owners, engineering, governance and risk.
- **Risk-tiered gate.** Use the three risk tiers. Low-risk use cases self-certify against a short checklist; medium and high-risk use cases get a governance review before go-live, deeper for high.
- **Agent limits and a kill switch.** For every agent rated medium or high risk, write down what it may do, what it must escalate, any spending or action limits, and how to stop it. Test the stop.
- **Human oversight that is real.** Name who reviews what, when, and with what authority to override. Track override rates: almost none can mean rubber-stamping, a great many can mean the AI is not good enough.
- **Safeguards for automated decisions.** Where an AI decision has a legal or similarly significant effect on a person, the UK Data (Use and Access) Act 2025 requires notice, a chance to make representations, meaningful human intervention and a right to contest. Stricter rules still apply if special category data is involved. Design these in, and log each case.
- **Answer feedback loop.** Add a simple way for users to flag a wrong answer. Route each flag to the use-case owner, who triages it as a data, prompt or model issue and passes it to the right owner.
- **Issue management integration.** Log AI answer issues in the same process as other data quality issues, so root causes are fixed at source.
- **Escalation.** Name a forum, such as a data or AI council, with authority to decide disputed issues.

Maturity

LEVEL 1

Ad hoc

Nobody is clearly accountable for AI outputs; issues are handled informally

LEVEL 2

Defined

Use-case owners are named, but there is no gate or feedback route

LEVEL 3

Managed

A risk-tiered gate runs before go-live and user flags reach owners

LEVEL 4

Embedded

Flags route automatically to the right owner, with resolution tracked and reported

Metrics

- Share of AI use cases with a named owner and a completed RACI.
- Share of high-risk use cases that passed the gate before go-live (target: 100%).
- Median time from a flagged answer to a fix at source.
- Share of flagged answers traced to a root cause.
- Share of agents with their own identity and a named owner (target: 100%).
- Human override rate for each use case with human oversight.

90-DAY STARTER CHECKLIST

- ☐ Define the use-case owner role in one page
- ☐ Name owners for every registered AI use case
- ☐ Agree the risk tiers and the gate for each tier
- ☐ Add a “flag this answer” route to one live AI tool
- ☐ List every agent, give each its own identity and a named owner, and write down its limits
- ☐ Name the escalation forum

Regulatory mapping

- **EU AI Act:** deployer obligations for high-risk systems, including human oversight and monitoring.
- **ISO/IEC 42001:** roles, responsibilities and authorities for AI.
- **UK GDPR, as amended by the Data (Use and Access) Act 2025:** the accountability principle, and safeguards for significant automated decisions since 5 February 2026.
- **UK financial services:** FCA and PRA principles, the Consumer Duty and the Senior Managers and Certification Regime, under which senior managers stay accountable for outcomes in their areas, including those AI produces.

05 Assurance

Evidence as a by-product

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

The cheapest evidence is evidence you never had to gather. When the other four pillars run well, most of what an auditor or regulator needs already exists.

The problem

Assurance is often a scramble before an audit: collecting screenshots, chasing approvals and rebuilding lineage from memory. With AI, the questions get harder: which data trained or informed this system, who approved it, how is it monitored? Regulation is also still moving, so a control set built for one regime dates quickly.

IN PRACTICE

A pension fund investor's due diligence questionnaire asks which AI tools touch valuation, ESG and investor reporting data, who approved them, and how they are monitored. Answering takes six weeks of emails across four teams, and the final answer still includes “to be confirmed”. With an inventory, approval records and control results already captured, the same response takes days.

PRINCIPLES

1. Evidence is captured automatically as work happens.
2. Each control is mapped once to every framework it satisfies.
3. Assurance effort is proportionate to each use case's risk tier.
4. The framework is reviewed as regulation and standards change.

Practices

- **Evidence by design.** Store check results, lineage, contract versions, gate approvals and answer logs with timestamps, retained according to policy.
- **Unified control library.** Keep one library of controls, each mapped to the EU AI Act, ISO/IEC 42001, UK GDPR and any sector rules that apply.
- **Risk-tier evidence packs.** Define the evidence each tier must hold. A high-risk use case holds a full pack; a low-risk one a short record.
- **Decision logging.** For high-risk use cases, record with each material decision the model, prompt, sources and data versions used, so it can be reconstructed later.
- **Impact assessment.** Before a high-risk use case goes live, complete a data protection impact assessment and an AI impact assessment. ISO/IEC 42005:2025 gives a method for the second.
- **Continuous control monitoring.** Report control health from system data, such as checks passing, owners assigned and reviews in date, rather than by survey.
- **Regulatory watch.** Assign someone to track changes, such as the 2026 Digital Omnibus amendments to the EU AI Act, and update the control mapping.
- **Independent review.** Have internal audit or second-line risk test a sample of use cases each year.

Maturity



Metrics

- Time to produce a complete evidence pack for one use case on request.
- Share of controls with automatically captured evidence.
- Share of high-risk use cases with a complete, in-date evidence pack.
- Number of audit findings related to AI data governance.

90-DAY STARTER CHECKLIST

- ☐ Draft a control library of 15 to 25 controls across the five pillars
- ☐ Map each control to the EU AI Act, ISO/IEC 42001 and UK GDPR
- ☐ Define evidence packs for each risk tier
- ☐ Test producing one pack for a live use case and time it
- ☐ Name the owner for regulatory watch

Regulatory mapping

- **EU AI Act:** technical documentation, record-keeping and registration for high-risk systems, applying from 2 December 2027 for Annex III systems.
- **ISO/IEC 42001:** internal audit, management review and continual improvement.
- **ISO/IEC 42005 and 42006:** a method for AI system impact assessment, and the rules for bodies that certify against ISO/IEC 42001.
- **UK GDPR:** records of processing and the ability to demonstrate compliance.

Score yourself. Set a target. Start with the gap.

Score each pillar against the maturity tables in its chapter, then set a target based on your AI risk. The largest gaps on your highest-risk use cases are where to start.

PILLAR	KEY QUESTION	CURRENT LEVEL	TARGET LEVEL
Meaning	Do AI tools use certified definitions, choose the right one, and say which?		
Scope	Do we know every AI use case, every source it reads, and what each source is authoritative for?		
Automation	Do critical rules run in pipelines and stop bad data?		
Accountability	Is someone accountable for every AI use case, with a route to fix a wrong answer?		
Assurance	Could we produce evidence for a regulator in days?		

How to score

- **Score what happens today**, with evidence you could show. Plans don't count.
- **A level counts only when it is fully met** for your high-tier use cases and mostly met elsewhere. If you meet some but not all of a level, score the level below.
- **Use the pillar's metrics as the evidence**. If you can't measure it yet, you are probably at level 1 or 2.

How to set a target

- **Target what your risk justifies, not the highest level**. If all your use cases are low-tier, level 2 is usually enough. With medium or high-tier use cases, aim for level 3.
- **Reserve level 4** for controls that must work at machine speed or can't rely on memory, such as agents acting on their own or checks that block a load.
- **Value the step, not the score**. Moving one pillar from 2 to 3 where it changes outcomes is worth more than reaching 4 everywhere.

A gap of two levels or more on any pillar is a signal to act. A gap on Scope usually comes first, because the other pillars depend on knowing what to govern.

Challenges are more useful than agreement.

This framework is published as a draft so practitioners can test it against their own experience. Version 1.1 reflects the first round of challenge.

The questions most worth answering:

1. Is anything missing that belongs in a sixth pillar, or is anything here unnecessary?
2. Do the maturity levels describe real organisations, or are the steps too big or too small?
3. Which practices have you tried, and what happened?
4. Where would this framework fail in your organisation, and why?
5. Should AI outputs and prompts be governed as data, or as something else?
6. Can an AI tool reliably choose the right governed meaning for a question, and explain why?

Comment on the LinkedIn post, or message directly. Feedback will be reflected in the next version, with credit to contributors who want it.

ACKNOWLEDGEMENTS

Version 1.1 was shaped by challenges from [Jess Baker](#). Her questions on metric identity, source authority, who builds the rules and where accountability sits led to the changes summarised under “What changed in v1.1”. Contributors are credited in a personal capacity. The views, and any errors, are the author's.

REFERENCES

1. European Parliament and Council, [Regulation \(EU\) 2026/1744 \(Digital Omnibus on AI\)](#), Official Journal, 24 July 2026
2. White & Case, [EU AI Omnibus enters into force, amending the AI Act](#), 2026
3. UK Parliament, [Data \(Use and Access\) Act 2025](#)
4. ICO, [The Data \(Use and Access\) Act 2025: what does it mean for organisations?](#)
5. Travers Smith, [UK data protection reforms take effect: automated decision-making](#), 2026
6. Covington, [UK financial services regulators' approach to AI in 2026](#), April 2026
7. ISO, [ISO/IEC 42001:2023 AI management systems](#)
8. ISO, [ISO/IEC 42005:2025 AI system impact assessment](#)
9. ISO, [ISO/IEC 42006:2025 Requirements for bodies certifying AI management systems](#)
10. NIST, [AI Agent Standards Initiative](#), February 2026
11. InfoQ, [Model Context Protocol adds centralised auth for enterprise](#), July 2026
12. Gartner, [Lack of AI-ready data puts AI projects at risk](#), February 2025
13. 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.

