MIDDLEBROOK
DATA & AI GOVERNANCE

DATA GOVERNANCE FOR AI REPORTING

A MASTER CLASS

Building the trusted data foundation that AI-driven reporting depends on — the Why, What, Where, Connect, and How, plus a clear-eyed look at the reporting of the future.

BARRY MIDDLEBROOK

24 years governing enterprise data in regulated industries — data governance, quality, lineage, metadata, master data, and SOX / CFPB-grade regulatory reporting.

middlebrookdataaigovernance.com

Part 1 — WHY: The Stakes Just Changed

For thirty years, a flawed report had a human circuit-breaker. An analyst built it, a manager reviewed it, and a wrong number usually got caught before it reached a decision. Reporting was slow — but it was supervised.

AI removes the circuit-breaker. When AI generates a report — or answers "what was Q3 margin in the Western region?" in plain English — it does so instantly, confidently, and at scale, with no analyst in the loop. That is both a gift and a risk:

- Garbage in, garbage out — now at machine speed. Ungoverned data no longer produces a bad report; it produces a thousand confident, wrong answers before lunch.
- Confident wrongness. An AI will happily invent a metric definition, blend two incompatible sources, or hallucinate a number that looks right. Without governance, nobody can tell.
- The trust ceiling. Leaders only act on reporting they trust. The first time an AI answer is caught being wrong, trust in all of it collapses — and adoption stalls.
- Regulation, arriving fast. The EU AI Act, ISO/IEC 42001, NIST's AI Risk Management Framework, and existing mandates (SOX, sector rules) increasingly require automated outputs to be explainable, traceable, and auditable. "The AI said so" is not a defense.
- Your model-risk program doesn't cover it. Banks already run disciplined model-risk management (SR 11-7), but the OCC's revised guidance explicitly excludes generative and agentic AI — so the most consequential AI in your stack sits outside the controls your second line actually runs. That gap is your exposure.

The new equation: AI reporting is only as trustworthy as the data governance beneath it. AI doesn't reduce the need for governance — it makes governance the foundation everything stands on.

Part 2 — WHAT: Defining Data Governance for AI Reporting

Data governance, classically, is the system of people, policies, processes, and standards that keep data accurate, consistent, secure, discoverable, and trusted — with clear ownership and accountability.

What's new for AI reporting is that you must now govern three surfaces, not one:

1. The data that feeds the AI — the inputs (sources, pipelines, the warehouse/lakehouse). Classic governance, raised stakes. 2. The interpretation layer — the meaning the AI uses: metric definitions, business glossary, the semantic/metrics layer, master data. This is where AI most often goes wrong: not bad math, but the wrong definition of "revenue" or "active customer." 3. The AI's outputs — the generated reports, answers, and narratives themselves, which now need their own quality, provenance, and assurance controls.

Core building blocks (the vocabulary):

Term	What it means	Why AI reporting needs it
Data quality	Accuracy, completeness, consistency, timeliness, validity, uniqueness	AI can't sanity-check; bad data → bad answer, silently
Data lineage	The traceable path of every number from source to output	Makes AI answers explainable & auditable
Metadata & catalog	Data about the data; a searchable inventory	So AI (and people) find the right, governed asset
Business glossary / semantic layer	One canonical definition of each metric and entity	Stops the AI inventing its own "revenue"
Master & reference data (MDM)	The single trusted version of core entities (customer, product, account)	The nouns the AI reasons over must be unambiguous
Data contracts	Enforceable agreements on a dataset's schema, meaning, and SLAs	Catch breakage before it reaches the AI
Access governance	Who — and which AI agents — may touch what	Least privilege; PII protection; AI can't leak what it can't see
Output governance	Validation, provenance, and assurance on AI-generated results	Trust, compliance, and a human check where it matters

The seven forms of AI reporting

"AI reporting" is not one capability — it shows up across the stack in seven forms, each a different way AI touches a number that reaches a decision-maker or a regulator. Every one depends on the same governed foundation:

- Self-service & conversational ("ask your data") — business users query in plain English. Fails when the AI guesses what "revenue" means and answers confidently with no provenance.
- Narrative generation — AI drafts the commentary: MD&A, variance explanations, board-pack write-ups. Fails when the explanation reads well but isn't true.
- Agentic workflows — agents pull data, compute, and assemble the pack, sometimes autonomously. Fails when a chain of unsupervised steps leaves no audit trail — and it sits outside model risk entirely.
- Regulatory & filing assistance — AI prepares or checks filings, disclosures, and SOX evidence. Fails when an error lands in a regulated submission.
- Anomaly & exception reporting — AI flags outliers, control breaks, and fraud signals. Fails when ungoverned inputs cause missed exceptions or drown the team in noise.
- Predictive & forward-looking — forecasts and model outputs feed reports (expected loss, stress-test inputs). Fails when model-driven numbers flow into reports with no lineage.

- Extraction → reporting — AI lifts figures out of contracts, statements, and documents. Fails when a misread number is silently structured and trusted.

You don't govern seven things — you govern one foundation, and all seven get safer at once.

Part 3 — WHERE: The Points of Control

Governance for AI reporting must live at every stage of the data's journey. Map the lifecycle, then place a control at each handoff:

```
SOURCE → INGESTION → STORAGE → MODELING → SEMANTIC/METRICS LAYER → AI ACCESS → OUTPUT → CONSUMPTION
   |         |          |          |              |                    |          |          |
 owner &   data       quality &   lineage &    canonical metrics    least-priv  validation  audit &
 contract  contracts  validation  metadata     & glossary           access      & provenance feedback
```

- Source & ingestion — assign an owner; enforce a data contract (schema + meaning + freshness). Most "AI got it wrong" incidents start with a silent source change here.
- Storage & modeling — quality checks, lineage capture, metadata tagging. Nothing reaches the AI un-profiled.
- The semantic / metrics layer — the single most important control for AI reporting. This is where "revenue," "churn," "active user" are defined once. Point the AI at this layer — never at raw tables — and it can no longer improvise definitions. This is your single source of truth made machine-consumable.
- AI access — the AI (and any agents) get least-privilege access to governed, certified datasets only, with PII masked or excluded by policy.
- Output — every answer carries provenance (which sources, which definitions, as-of when) and passes validation rules; high-stakes outputs route to a human.
- Consumption & feedback — usage is logged for audit; users can flag wrong answers, feeding the quality loop.

Where it breaks without governance: the AI reaches past the semantic layer into raw data, blends two sources that define "customer" differently, uses a stale extract, and produces a confident, un-traceable, wrong number — with no one able to say why. Governance is what closes each of those gaps.

Part 4 — CONNECT: How AI Actually Talks to Your Data

There are a few ways to connect an AI to your data — and the right one, for trustworthy reporting, is exactly where your governance background becomes the value-add. From simplest to most governed:

1. Text-to-SQL. The AI translates a plain-English question ("Q3 margin by region") into a SQL query, runs it against the database, and narrates the result. How: the LLM is given the schema (table/column names + descriptions), it writes SQL, an app executes it, and results go back to the LLM to summarize. Risk: it can write wrong joins, pick the wrong table, or invent a definition. Powerful — but it needs guardrails.

2. RAG (Retrieval-Augmented Generation). For unstructured data (docs, PDFs, tickets). Content is chunked and stored in a vector database; the AI retrieves the most relevant chunks and answers from them. Great for "what does our policy say," weaker for precise numbers.

3. Semantic / metrics layer in front (the governed way). Instead of letting the AI touch raw tables, you put a semantic layer between them (dbt Semantic Layer, Cube, LookML, or a metrics API). The AI queries defined metrics ("revenue," "active customer") — not raw SQL — so it physically cannot invent a definition. This is the architecture you want for trustworthy AI reporting, and it's the core point of this master class.

4. Tool calling / MCP (the modern standard). The AI is given tools (functions) it can call — e.g., a `run_query` tool, or increasingly an MCP server (Model Context Protocol — the open standard for connecting AI to data sources). The AI decides to call the tool; the tool executes the query with all your guardrails baked in and returns results. This is how production AI-to-data connections are built in 2026 (there are MCP servers for Postgres, Snowflake, and more).

5. AI agents. Orchestrate multiple tool calls plus reasoning for multi-step questions ("compare this quarter to last, and explain the drivers").

The reference architecture

How a governed AI-reporting request flows at runtime. Data flows down, results return up, and the teal layers are the governance controls you own — the difference between a demo and a system you can put in front of an auditor.

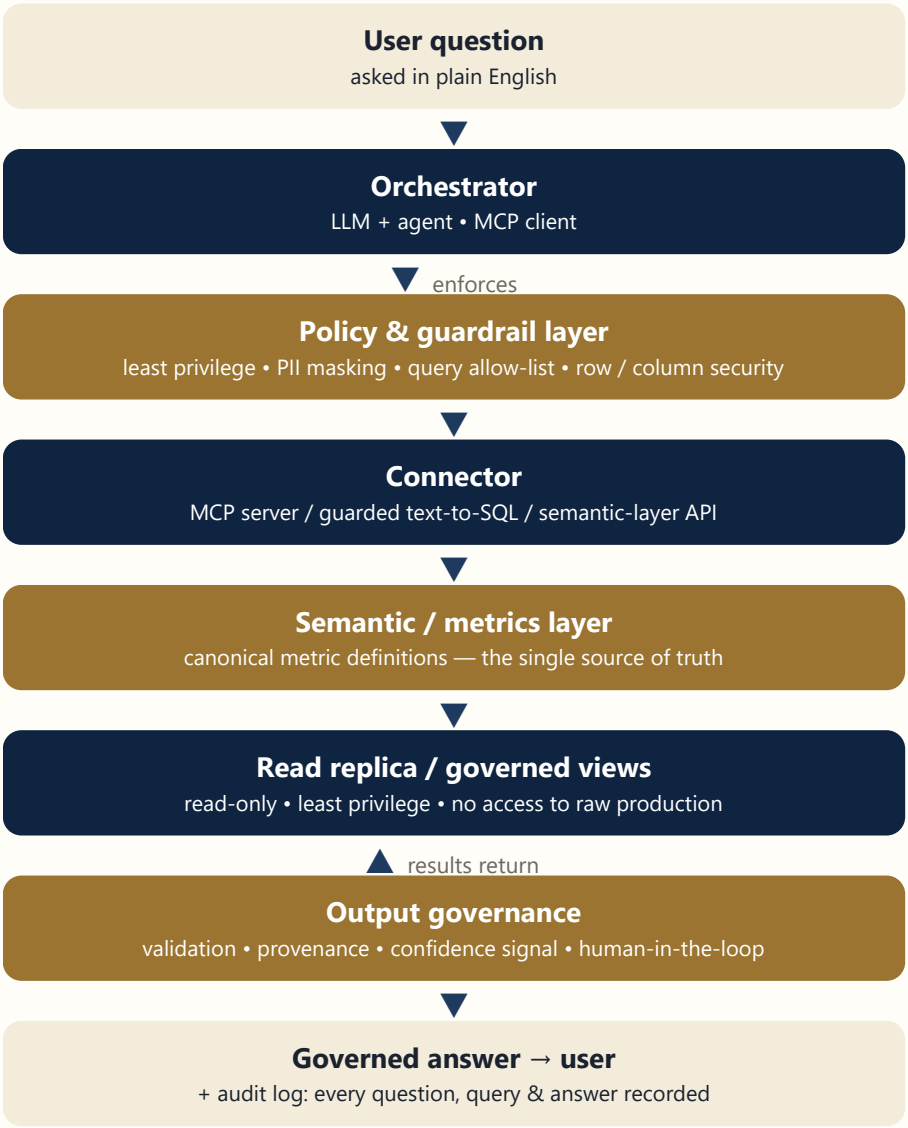

Where you come in — the governance guardrails

Anyone can wire an LLM to a database in an afternoon. The reason 98.5% of organizations can't staff this is that almost nobody knows how to do it safely and correctly. That is the entire value-add:

- Least privilege + read-only. The AI connects through a service account that can `SELECT` from governed views only — never write, never `DROP`, never see raw everything. Point it at a read replica, never prod.
- Semantic layer, not raw tables. So definitions are canonical — no hallucinated "revenue."
- Row/column security + PII masking enforced at the connection, so the AI can't surface what it shouldn't.
- Query guardrails — allow-lists, validation, row/cost limits, timeouts; reject anything that isn't a safe read.

- Provenance + lineage returned with every answer (which source, which definition, as-of when) — explainable and auditable.
- Human-in-the-loop for high-stakes outputs.

That list is "Part 5: How" — the practitioner's framework — made concrete at the exact point where AI meets the database. The connection is easy; the governed connection is the skill.

Part 5 — HOW: The Practitioner's Framework

This is the work. Ten disciplines, in build order.

1. Ownership & stewardship. Name a data owner for every critical domain and data stewards who maintain it day to day. Stand up a lightweight governance council to set policy and adjudicate definitions. Governance fails as a technology project and succeeds as an accountability structure.

2. The canonical metrics / semantic layer. Define each business metric and entity once, in writing, with the formula and the owner. Make that the only thing the AI is allowed to query for reporting. This single move eliminates the most common class of AI-reporting error.

3. Data quality controls. Profile, measure, and monitor against the six dimensions — accuracy, completeness, consistency, timeliness, validity, uniqueness. Publish data-quality scorecards. Wherever practical, enforce data contracts so upstream breakage is caught automatically, not discovered in a board deck.

4. Lineage & metadata. Capture end-to-end lineage and maintain a data catalog + business glossary. This is what lets an AI answer cite its sources — and what lets an auditor trace a number from the report back to the system of record. No lineage, no trustworthy AI reporting.

5. Master & reference data (MDM). Establish the golden record for core entities. The AI must reason over one customer, one product hierarchy, one chart of accounts — not three conflicting versions.

6. Access & security governance. Apply least privilege to people and AI agents. Classify and protect sensitive/PII data; the safest way to keep AI from leaking data is to ensure it can't access what it shouldn't see in the first place.

7. Governing the AI layer. Ground the AI in governed data — retrieval over certified sources and the semantic layer (the disciplined version of "RAG"). Add guardrails: the model may only use approved datasets/definitions, must refuse when data is missing rather than guess, and cannot fabricate metrics.

8. Output governance & assurance. Treat AI outputs as governed artifacts: automated validation against business rules, provenance on every answer (sources + definitions + as-of date), confidence/uncertainty signaling, human-in-the-loop for high-stakes (financial, regulatory, safety) outputs, and an audit trail of every question and answer.

9. Monitoring & observability. Watch for data drift (inputs changing) and model drift (behavior changing). Set quality SLAs, alert on breaches, and run an incident-response process for bad answers —

the same rigor you'd apply to a production outage.

10. Compliance & auditability. Map your controls to the frameworks: NIST AI RMF (Govern/Map/Measure/Manage), ISO/IEC 42001 (AI management system), the EU AI Act (risk tiers, transparency, human oversight), and DORA (AI as ICT / operational-resilience risk) — plus existing mandates: SOX, sector rules, and your model-risk program (SR 11-7), which must now be extended to cover the generative and agentic AI it originally excluded. Build the evidence trail as you operate, not in a fire drill before an audit.

A maturity model

Level	State	AI reporting you can trust?
1 — Ad hoc	No ownership; definitions vary by spreadsheet	No — AI amplifies the chaos
2 — Reactive	Governance only after something breaks	Risky — you find errors after the decision
3 — Defined	Owners, glossary, quality rules documented	Partially — for governed domains
4 — Managed	Quality measured, lineage captured, semantic layer live	Yes — for certified data
5 — AI-trusted	Output governance, monitoring, audit-ready, self-service	Yes — trusted, explainable, at scale

A 90-day starter playbook

- Days 1–30 — Anchor. Pick one high-value reporting domain. Name an owner + steward. Document the top 10 metrics with canonical definitions. Baseline their data quality.
- Days 31–60 — Build the spine. Stand up the semantic/metrics layer for those 10 metrics. Capture lineage to source. Stand up a catalog/glossary entry for each.
- Days 61–90 — Govern the AI. Point the AI only at that governed layer. Add provenance + a validation rule + a human check on the highest-stakes output. Show a leader an AI answer that cites its sources. That demo sells the whole program.

Part 6 — THE FUTURE: AI Reporting, Done Right

Here is the end state this all builds toward.

Reporting becomes conversational and self-serve. A leader asks, in plain English, "How did margin trend by region last quarter, and what drove the change?" — and gets an immediate, correct, narrated answer. No ticket to the BI team. No week of waiting.

And — because it's governed — that answer is trustworthy:

- It's computed from canonical, governed metrics, not improvised math.

- It cites its lineage — which sources, which definitions, as of when — so it can be verified and audited.
- It carries a confidence signal and says "I don't have governed data for that" instead of guessing.
- It leaves an audit trail, satisfying SOX, the EU AI Act, and the auditor's questions by default.
- High-stakes outputs still pass a human check — accountability never disappears, it relocates.

The human role shifts — and rises. People stop hand-building reports. Their job becomes governing the system that generates them: owning the definitions, curating the trusted data, setting the guardrails, monitoring quality, and standing behind the numbers. The data governance professional moves from back-office to mission-critical — because in a world where AI generates the insights, the scarcest, most valuable thing is the person who can guarantee the insights are right.

Trust becomes the product. Two companies will have the same AI models. The one whose AI reporting is governed — accurate, explainable, compliant — will out-decide the one whose AI confidently makes things up. Governance is the moat.

The throughline

AI did not make data governance obsolete. It made it the foundation everything else stands on. The organizations — and the professionals — who master Data Governance for AI Reporting won't just survive the shift to AI. They'll be the ones everyone else has to trust.

Glossary — Tech & Business Terms

Plain-language definitions of every technical and business term used in this master class.

Term	Definition
Access governance	Rules controlling who — and which systems or AI agents — may read or use which data; enforced through least privilege.
Agentic AI	AI that plans and takes multi-step actions autonomously — querying, computing, assembling, sometimes acting — rather than answering once. The most powerful and hardest-to-govern form of AI reporting.
AI agent	An AI system that plans and takes multi-step actions by calling tools (e.g., querying data) to reach a goal, rather than answering once.
Audit trail	A time-stamped, logged record of who asked what, which data was used, and what was produced — required for compliance and investigations.
Business glossary	The agreed, plain-language definitions of business terms and metrics (e.g., what "active customer" means), owned by the business.
CFPB	U.S. Consumer Financial Protection Bureau; its consumer-finance reporting and data rules govern mortgage and lending data.

Term	Definition
Confidence signal	An indication of how certain an AI is in an answer (and when it lacks governed data), so users know when to trust or verify it.
Connector	The component linking the AI to the data — an MCP server, a guarded text-to-SQL engine, or a semantic-layer API.
Data catalog	A searchable inventory of an organization's data assets and their metadata, so people and AI can find the right, governed data.
Data contract	An enforceable agreement defining a dataset's schema, meaning, quality, and delivery — so upstream changes don't silently break downstream use.
Data drift	A change over time in input-data patterns or distribution that can quietly degrade the accuracy of reports and AI outputs.
Data governance	The people, policies, processes, and standards that keep data accurate, consistent, secure, discoverable, and trusted, with clear ownership.
Data lineage	The traceable path of data from source through every transformation to final use; how any number can be traced back to its origin.
Data owner	The person accountable for a data domain — sets policy and standards and answers for its quality and proper use.
Data quality	The fitness of data for use, measured across six dimensions: accuracy, completeness, consistency, timeliness, validity, and uniqueness.
Data steward	The person who maintains a data domain day to day — applying standards, resolving issues, and curating definitions.
DAMA-DMBOK	The Data Management Body of Knowledge from DAMA International; the industry-standard framework of data-management disciplines.
DORA	The EU Digital Operational Resilience Act; brings ICT — and now AI — risk, third-party dependencies, incident reporting, and resilience testing under one operational-resilience regime for financial entities.
ETL / ELT	Extract-Transform-Load (or Extract-Load-Transform); the pipelines that move and reshape data from sources into a warehouse or lakehouse.
EU AI Act	European Union regulation that classifies AI systems by risk and imposes transparency, human-oversight, and accountability obligations.
Golden record	The single, authoritative, de-duplicated version of a core entity (e.g., one true customer record) produced by master data management.
Governance council	A cross-functional group that sets data and AI governance policy, approves definitions, and resolves disputes.

Term	Definition
Governed views	Curated, permission-controlled database views the AI is allowed to read — instead of raw tables.
Hallucination	When an AI produces confident but false or fabricated information — e.g., inventing a metric or a number unsupported by the data.
Human-in-the-loop	A control requiring a person to review or approve an AI output before it is acted on; used for high-stakes decisions.
ISO/IEC 42001	The international management-system standard for governing artificial intelligence responsibly across its lifecycle.
KPI	Key Performance Indicator — a defined metric used to measure progress against a business objective.
Lakehouse / warehouse	Central, governed stores for analytical data; a warehouse is structured, a lakehouse blends that structure with data-lake flexibility.
Least privilege	The security principle of granting the minimum access needed — applied to people and to AI / service accounts alike.
LLM	Large Language Model — the AI (e.g., Claude, GPT) that understands and generates language and powers conversational reporting.
Master data	The core, shared business entities — customers, products, accounts — that must be consistent across every system.
MDM	Master Data Management — the discipline and tooling that create and maintain one trusted golden record per master-data entity.
MCP	Model Context Protocol — an open standard for securely connecting AI models to data sources and tools through governed "servers."
Metadata	Data about data — names, definitions, formats, owners, lineage — that makes data findable and understandable.
Model drift	A decline in an AI model's performance over time as the world or the data changes from what it was built on.
Model risk / SR 11-7	The discipline — and U.S. supervisory guidance — for validating, monitoring, and documenting models. Its revised guidance excludes generative and agentic AI, leaving a gap firms must close by extending model risk to cover them.
NIST AI RMF	The U.S. NIST AI Risk Management Framework; a voluntary structure (Govern, Map, Measure, Manage) for managing AI risk.
Observability	Continuous monitoring of data and AI systems — quality, freshness, drift, errors — so problems are caught early.

Term	Definition
Orchestration	The layer that coordinates the AI's steps and tool calls (the application, agent, or MCP client) between the user and the data.
Output governance	Controls applied to AI-generated results: validation, provenance, confidence signaling, audit logging, and human review.
PII	Personally Identifiable Information — data that can identify an individual; subject to privacy law and strict access controls.
Provenance	The record of where a result came from — which sources, definitions, and as-of date — attached to an output so it can be verified.
RAG	Retrieval-Augmented Generation — a pattern where the AI retrieves relevant data or documents and uses them as grounding context.
Read replica	A read-only copy of a production database used for queries and reporting, so the live system isn't risked or slowed.
Reference data	Standardized lookup values (country codes, status codes, categories) used consistently across systems.
Row/column-level security	Access rules restricting which rows or columns a user or AI can see within a dataset (e.g., only their region; mask SSNs).
Schema	The structure of a database: its tables, columns, data types, and relationships.
Semantic / metrics layer	A governed layer that defines business metrics and entities once, so every query — human or AI — uses the same canonical definitions.
Single source of truth	One authoritative, agreed place for a given piece of data or metric, so everyone (and every AI) gets the same answer.
SLA	Service-Level Agreement — a committed standard (e.g., data freshness or quality threshold) a data product must meet, with alerts on breach.
SOX	Sarbanes-Oxley — U.S. law requiring accurate, controlled, auditable financial reporting; a major driver of strong data controls in finance.
SQL	Structured Query Language — the standard language for querying and manipulating data in relational databases.
Text-to-SQL	A pattern where the AI converts a plain-English question into a SQL query, runs it, and returns or narrates the result.
Tool calling	The mechanism by which an AI invokes a defined function (e.g., "run this query") to act, rather than only generating text.
Vector database	A database that stores data as numeric embeddings for similarity search; the retrieval engine behind RAG.

Frameworks referenced: DAMA-DMBOK · NIST AI Risk Management Framework · ISO/IEC 42001 · EU AI Act · DORA · SOX · SR 11-7.