

In today's enterprise landscape, Artificial Intelligence (AI) promises transformative potential, especially in forecasting business-critical metrics such as sales, demand, and inventory levels. Yet, despite advances in Large Language Models (LLMs) and retrieval-augmented generation (RAG) techniques, a fundamental bottleneck remains: **trust**. Organizations like Garrett Motion, PMI, and Pandora face this challenge head-on. The question is no longer about whether AI can generate accurate forecasts but whether stakeholders can trust those numbers.

Why Trust is the Bottleneck for Enterprise AI Adoption

Many enterprises have sprinted to deploy AI-powered forecast tools. However, even when models improve accuracy marginally, executives hesitate to rely on the outputs. Why? Because improved accuracy alone does not create trust. Instead, trust hinges on **explainability**, **consistency**, and alignment with **business logic**.

Take Pandora, the well-known jewelry brand, which leverages AI to forecast demand fluctuations worldwide. Despite improvements in predictive performance, supply chain managers demand transparency before changing order quantities that impact millions in inventory. Similarly, Garrett Motion, a global automotive supplier, requires rigorous traceability of how a forecast transforms raw sensor and market data into actionable decisions.

LLMs and RAG have unlocked fresh capabilities, but they often induce hallucinations—fabricated facts with confident-sounding language—which heightens the risk of misleading stakeholders. With so much on the line, enterprises need more than just predictive accuracy; they need a reliable audit trail that exposes the *why* and *how* behind every forecast number.

The Anatomy of an AI Audit Trail for Forecast Numbers

What exactly does an AI audit trail for a forecast number look like? At a minimum, it requires three foundational pillars:

1. **Data Lineage:** Documenting the origin and transformation history of every data element that feeds the forecast
2. **Business Rules Trace:** Mapping how domain-specific logic and constraints shape intermediate calculations and final outputs
3. **Decision Justification:** Summarizing the rationale behind the selected forecast figure, ideally linking model inferences to source data and business context

1. Data Lineage: Tracking the Roots

Forecast models ingest vast amounts of data—historical sales, seasonality patterns, macroeconomic indicators, supplier lead times, and external signals harvested via RAG pipelines. An effective audit trail logs:

- The precise dataset versions (e.g., which snapshot of PMI's procurement data)
- Preprocessing steps such as cleansing, normalization, and feature extraction
- Model input parameters and environmental conditions (e.g., date, timezone)

For Garrett Motion, even minimal discrepancies in sensor data calibration can cascade into significantly different forecasts. Thus, audit systems must enable stakeholders to trace each input value back to its source reliably. This [AI hallucinations in business](#) reduces the risk of "black box excuses" like "the model decided," which do not satisfy governance or audit standards.

2. Business Rules Trace: Making the Logic Transparent

Forecasts do not emerge solely from statistical patterns; they embody embedded business logic. For example, Pandora might impose constraints on minimum order quantities and promotional periods that modulate forecast outputs. The audit trail must show:

- Which business rules applied to the forecast calculation
- How rules interact or override pure model predictions
- Changes or exceptions logged with timestamps and responsible users

This business rules trace ensures that forecasts comply with established policies and aids in identifying misalignments between AI inferences and real-world priorities.



3. Decision Justification: Explaining the Output

Finally, the forecast number needs a clear narrative tying data and rules together. Techniques leveraging LLMs and RAG can generate natural language summaries explaining:

- Key drivers influencing the forecast (e.g., a recent surge in PMI's raw material costs)
- Confidence intervals and scenario analyses
- Any hallucinated or uncertain elements requiring human review

This transparency reduces the risk of fabricated certainty, a prevalent issue when AI systems produce plausible but incorrect explanations. By integrating decision justification into the audit trail, organizations empower both technical and business users to question, validate, or challenge forecast metrics.

How Garrett Motion, PMI, and Pandora Demonstrate Effective AI Audit Trails

Leading enterprises set benchmarks in implementing audit trails that address trust bottlenecks.

Company Practice Key Impact Garrett Motion Detailed sensor data lineage plus dynamic business rules applied in real-time Reduced forecasting errors from unforeseen hardware malfunctions by 18% PMI (Philip Morris International) Embedded compliance rules with explainable LLM-generated decision summaries Faster regulatory audits and reduced risk from black box algorithms Pandora Cross-functional traceability dashboards exposing

model inputs, rule overrides, and forecast justifications Greater confidence in inventory planning, minimizing excess stock by 12%

Addressing Key Challenges: Hallucinations and Output Variance

Audit trails must also expose known shortcomings of AI technologies:



- **Hallucinations:** LLMs sometimes fabricate details that sound plausible but lack basis in data. Audit logs should flag unverified statements and track confidence scores.
- **Random Output Variance:** Non-deterministic elements in RAG pipelines may cause inconsistent explanations for identical inputs. Version-controlled audit records help isolate cause and effect.

Without this transparency, decision-makers might be misled into false precision or dismiss genuine insights, undermining trust.

Conclusion: Building Trust Beyond Accuracy

In sum, an AI audit trail for forecast numbers is not just about technical logging; it embodies a cultural and procedural commitment to transparency and accountability. Amidst the fast rise of complex models and generative AI, enterprises like Garrett Motion, PMI, and Pandora demonstrate that only through rigorous **data lineage**, **business rules trace**, and **decision justification** can trust be earned and sustained.

Without these pillars, forecasts risk becoming inscrutable black boxes whose outputs invite skepticism rather than strategic action. Incorporating thoughtfully designed audit trails enables organizations not only to meet regulatory and ethical standards but also to transform AI-driven forecasts into true business assets.

Next time you see a forecast generated by an AI— *what would you show in an audit?*