

Touchless planning — automating complex business decisions without human intervention — is the holy grail of enterprise AI. Yet, achieving high touchless rates remains elusive for many companies. Recently, a leading brewer announced a 20% increase in their touchless planning rates, sparking questions: How did they crack the code? What can other enterprises learn from this success?

While companies like **Garrett Motion**, **PMI**, and **Pandora** have made significant strides with AI-driven planning, the brewer's breakthrough highlights critical themes often overlooked in AI adoption:

- **Trust** as the fundamental bottleneck
- The crucial roles of **explainability, consistency, and business-logic alignment**
- Why **accuracy gains alone don't build trust**
- The hazards of **hallucinations** and fabricated certainty from Large Language Models (LLMs)

We'll unpack these themes and demonstrate how tools like **retrieval-augmented generation (RAG)** and structured feedback loops support continuous improvement in planning accuracy and trustworthiness.

Why Trust Is the True Bottleneck in Enterprise AI Adoption

Across industries, from manufacturing to entertainment, AI hype has promised cost savings and optimization gains. Yet actual adoption lags behind. Why? Surveys from enterprise users reveal a common answer: *lack of trust*.

Garrett Motion, a leader in automotive technology, reported that their AI deployment stalled when planners doubted black-box recommendations. Similarly, **Pandora** found their music operations teams hesitant to rely on opaque AI forecasts absent clear explanations of how decisions aligned with their unique business rules.

Simply put, no matter how accurate or fast the model's output, operations teams will not fully trust or adopt AI unless they understand the "why" behind it. This trust gap manifests starkly in touchless planning volumes—AI outputs deemed trustworthy proceed automatically, while suspicious outputs trigger manual review, creating a variance review cycle that slows ROI.

The Brewer's Approach: Beyond Accuracy to Explainability

The brewer faced a similar challenge. Early trials of LLM-based demand and supply plans produced promising accuracy improvements. Yet touchless planning rates plateaued. Why? Planners lacked actionable explanations about how recommendations matched dynamic business rules.

To overcome this, the brewer integrated transparent models incorporating:

- **Rule adherence checks** surfaced alongside AI suggestions, confirming alignment or exposing conflicts
- **Traceability links** showing data provenance and rationale paths back to source facts
- A **feedback loop planning** mechanism allowing planners to flag, correct, and refine AI decisions continuously

This shift moved planner perception from AI as a black box to a collaborative assistant. Trust grew, enabling a notable 20% rise in touchless planning rates and reducing the costly variance review cycle.

Explainability, Consistency, and Business-Logic Alignment Are Non-Negotiable

Why do explainability and business-logic alignment matter so much?

Because enterprise planning embodies complex, evolving rules shaped by regulatory mandates, customer contracts, and internal policies. These rules are often manually embedded within spreadsheets, siloed databases, and tribal knowledge. AI models that disregard or distort them invite fatal distrust.

Key Theme Why It Matters How to Address Explainability Users understand why an AI recommendation is made
Show decision paths, rule matches, exceptions, and confidence levels Consistency Repeated prompts yield stable, predictable outputs
Standardize input formats, use deterministic AI components, and log outputs Business-Logic Alignment AI respects domain rules and policies
Embed rules in model constraints or augment with RAG retrieval of validated documents

At **PMI**, for example, investments in rule-aligned AI planning tools significantly reduced post-decision exceptions, smoothing downstream execution. Pandora likewise enhanced planning adoption by embedding domain constraints directly into RAG frameworks that surfaced explanations with knowledge citations.

Tools Like RAG Make the Difference

Retrieval-Augmented Generation (RAG) combines generative AI (LLMs) with external trusted knowledge bases to ground outputs in verified facts. Instead of hallucinating plausible but baseless plans, the model “retrieves” relevant rules, historic patterns, and policy documents to reference in its output explanations.

This architectural pattern is key to minimizing hallucinations—the infamous AI pathology where confident yet fabricated answers destroy trust.

Why Accuracy Gains Don’t Automatically Create Trust

AI developers often tout improved accuracy metrics as proxies for success. However, as our brewer discovered and industry leaders confirm, accuracy gains without explainability and auditability are insufficient.



Why not?

1. Operators can’t verify or debug opaque predictions
2. Minor output variance isolates models as unpredictable black boxes
3. Random “hallucinatory” errors erode confidence faster than slow iterative improvement
4. Failing to tie decisions back to business rules prompts manual overrides and variance reviews

Trust requires consistent, transparent, and aligned decisions traceable to known data, not just higher hit rates or machine learning scores.

Feedback Loop Planning Drives Continuous Improvement

Successful enterprises embed human-in-the-loop feedback cycles to close trust gaps and foster continuous model evolution. The brewer institutionalized a **feedback loop planning** process:

- Planners review automated outputs flagged as low confidence or misaligned
- Corrections and rationales feed back into the AI training and rule definition layers
- Periodic variance review cycles expose systemic bottlenecks and rule evolution needs
- Results feed back into model tuning, expanding the valid operating envelope

This dynamic process boosts both touchless planning rates and confidence, turning variance cycles into engines of continuous improvement.

Hallucinations and the Risk of Fabricated Certainty

LLMs possess stunning language generation capabilities but inherently risk hallucinating—generating false but fluent outputs.

In mission-critical enterprise planning, fabricated certainty is toxic. Stakeholders who encounter plausible but incorrect AI explanations distrust the system outright. This can result in:



- Suspended AI adoption
- Deterioration of model governance
- Escalated variance reviews and manual decision bottlenecks

Mitigating hallucinations requires:

- Integrating RAG to ground output in vetted knowledge bases
- Enforcing strict post-processing checks aligning model proposals with business rules
- Building audit trails for each AI decision to backtrace its rationale

- Educating users on AI limitations and creating interfaces that encourage healthy skepticism and collaborative correction

Lessons from the Brewer, Garrett Motion, PMI, and Pandora

What common threads emerge from successful enterprise AI planning deployments?

1. **Trust beats standalone accuracy:** Transparency, explainability, and auditability are essential for broad adoption.
2. **Align AI with business logic:** Ground AI in existing domain knowledge through techniques like RAG to prevent drift and hallucinations.
3. **Close the feedback loop:** Use variance review cycles not to reject AI but to iteratively improve outputs and expand touchless planning reach.
4. **Monitor output consistency:** Reduce random variation on identical prompts to build operator confidence and reduce noise in variance reviews.
5. **Educate and empower users:** Build human-AI collaboration, not black-box automation, to harness the unique strengths of planners and machines.

Conclusion: Building Trustworthy Touchless Planning Systems

The brewer's 20% improvement in touchless planning rates is not a fluke but a blueprint for enterprise AI success. Trust—a blend of explainability, consistent alignment with business rules, and continuous improvement—is the gatekeeper to unlocking AI's full potential on the factory floor, in supply chains, and beyond.

Enterprises must move beyond chasing pure accuracy metrics and invest in transparent architectures leveraging LLMs enhanced by RAG, rigorous audit trails, and feedback loop [Look at this website](#) planning processes. Only then can AI evolve from a distrusted black box to a dependable co-pilot driving efficient, scalable, and resilient business planning.

What would I show in an audit? A complete trace of how every plan recommendation links back to data, rules, and user feedback—the antidotes to excuses like “the model decided.” Without this clarity, all gains remain fragile and provisional.