

Enterprise AI adoption is often touted as the next frontier in digital transformation, but trust remains the critical bottleneck. As companies like **Garrett Motion**, **PMI**, and **Pandora** explore AI-driven planning assistants, the question arises: should they lean solely on Large Language Models (LLMs) or augment with knowledge-graph-based agents? Understanding the nuances of *LLM-only limitations* and the benefits of auditable, rule-aligned reasoning is paramount to making informed decisions that don't just yield accuracy but build trust.

The Context: Enterprise AI Planning Assistants

Planning assistants powered by AI promise to accelerate complex decisions involving resource allocation, forecasting, and supply chain nuances. This domain demands precision, consistency, and above all, explainability. Technologies in play include:

- **Large Language Models (LLMs):** These models generate fluent, contextually relevant responses by training on vast text corpora.
- **Retrieval-Augmented Generation (RAG):** Enhances LLM outputs by retrieving and grounding responses in dedicated knowledge databases.
- **Knowledge-Graph-Based Agents:** Structured, ontological representation of business logic, entities, and their relationships to underpin reasoning.

Why Trust Is the Bottleneck for Enterprise AI Adoption


Despite the hype, many enterprises wrestle with the “black box excuse”: “The model decided that.” This circular explanation is insufficient when stakes are high. Garrett Motion, a leader in automotive systems, highlights that without clear reasoning, compliance, and audit trails, AI adoption stalls.

Trust requires more than accuracy metrics. It requires:

- **Explainability:** How did the AI arrive at this plan?
- **Consistency:** Will it produce stable recommendations over time?
- **Business-Logic Alignment:** Does the output respect internal rules and domain constraints?

LLM-Only Planning Assistants: Opportunities and Limitations

LLMs excel in language fluency and contextual understanding from massive datasets. Pandora's experiments with LLM-only planning assistants illustrate how natural language interfaces can democratize forecasting insights across teams.



Strengths of LLM-Only Approaches

o9solutions.com

- **Generative Flexibility:** Ability to synthesize complex, human-like narratives explaining recommendations.
- **Contextual Reasoning:** Integrates diverse data points expressed in unstructured text input.
- **Rapid Prototyping:** Easy to deploy without heavy upfront modeling of business logic or ontologies.

Limitations of LLM-Only Assistants

Category LLM-Only Limitation Why It Matters Black Box Excuse Examples Explainability Lack of transparent, auditable reasoning Hard to verify how recommendations comply with business rules "Model decided" or "It's probabilistic" Consistency Output variance on repeated prompts Inconsistent advice undermines user confidence "Randomness is inherent to LLMs" Business Rules Difficulty enforcing strict constraints Plans may violate organizational policies unknowingly "The model learned that implicitly" Hallucination Risks Fabricated facts or false confidence Misleads planners, causing costly decisions "The model is sure, so it must be right"

Hallucinations and the Risk of Fabricated Certainty

LLMs often produce plausible-sounding but inaccurate information, which is risky in planning domains. PMI's AI teams emphasize that even when accuracy improves, hallucinations erode trust because users can't trace outputs back to verified data sources.



This is why retrieval-augmented generation (RAG) is crucial, providing a way to ground LLM outputs in trusted data repositories, partially mitigating hallucinations but not entirely closing the auditability gap.

Knowledge-Graph-Based Agents: Why Adding Structured Logic Matters

Knowledge graphs explicitly encode entities, relationships, and business rules. Garrett Motion leverages knowledge graphs to ensure recommendations respect engineering constraints and regulatory compliance.

Key Advantages

1. **Auditable Reasoning:** Each inference step is transparently linked to data nodes and rules.
2. **Business Rules Alignment:** Constraints can be explicitly enforced, preventing invalid plans.
3. **Consistency:** Deterministic logic reduces arbitrary output variability.
4. **Explainability:** Visualizable paths through the graph explain why a suggestion was made.

Challenges in Knowledge-Graph-Based Agents

- **Complex Setup:** Building and maintaining comprehensive knowledge graphs requires domain expertise and effort.
- **Limited Flexibility:** Rules need to be updated as business evolves; graphs may lag behind dynamic needs.
- **Language Interface:** May require more structured user queries versus free-form text.

LLM + Knowledge Graph Hybrid: Best of Both Worlds

Here's a story that illustrates this perfectly: wished they had known this beforehand.. Pandora and PMI are pioneering hybrid architectures combining RAG-enabled LLMs with knowledge graphs. This approach grounds the generative capabilities of LLMs in auditable, rule-driven knowledge representations.

- LLMs handle natural language understanding and generate human-like explanations.
- Knowledge graphs ensure adherence to business constraints and support traceable logic.

- Retrieval systems keep outputs up-to-date and aligned with certified data.

Such hybrids provide a balance between flexibility and trustworthiness.

Practical Recommendations for Enterprises

1. **Prioritize Audit Trails:** What would I show in an audit? Aim for solutions that capture rationale, data provenance, and decision logic.
2. **Beware Black Box Excuses:** Don't accept vague claims like "the model decided." Demand explanations tied to your business rules.
3. **Test for Consistency:** Run repeated prompts and scenarios. Inconsistent outputs are red flags.
4. **Mitigate Hallucinations:** Use retrieval-augmented techniques and domain-specific knowledge bases.
5. **Invest in Knowledge Graphs:** Though more resource-intensive, they provide aligned, explainable reasoning critical for trust.

Conclusion

Choosing between LLM-only planning assistants and knowledge-graph-based agents boils down to trust, not just accuracy. LLMs are impressive in language generation but suffer from explainability gaps, inconsistency, and hallucination risks that limit enterprise readiness. Knowledge graphs bring auditable reasoning and strict rule alignment but at higher setup costs.

Leaders at Garrett Motion, PMI, and Pandora increasingly adopt hybrid strategies, blending LLM generative power with the transparency and control of knowledge graphs. This fusion addresses the enterprise imperative: AI planning assistants that are not only accurate but demonstrably trustworthy.

Enterprises must ask themselves: **what would I show in an audit?** The answer will guide the choice—favoring solutions that illuminate, not obscure, decision logic.