

Inside Maya: How AI, Rules, and Real-Time Data Drive Expert-Level Resolutions

By Rajesh Madhavan, CEO – eSoftLabs | Head – R&D, AI

Executive Summary

In complex, high-volume warehouse environments, precision and speed are critical. When something goes wrong — a task gets stuck, a rule doesn't fire, or data doesn't align — delays ripple across the entire supply chain. The ability to resolve these issues accurately, consistently, and fast is no longer a support function — it's a competitive advantage.

Maya, our AI-powered WMS support engine, brings expert-level resolution into every corner of your warehouse — without the wait, without the bottlenecks, and without the need for deep internal expertise. This whitepaper takes you behind the scenes to show how Maya operates: from understanding natural language queries to executing real-time backend validations using a combination of fine-tuned models, structured rules, and intelligent agents.

From Query to Resolution: What Happens When You Ask Maya a Question

When a user submits a query — whether it's "Why is my picking task stuck?" or "Why did this PO line fail validation?" — Maya initiates a multi-layered process:

1. **Language Understanding:** The query is interpreted using a fine-tuned natural language model designed specifically for WMS environments. Maya understands both the user's intent and the operational context.
2. **Query Classification:** Maya determines whether the issue is WMS-relevant, and identifies the correct operational domain (Receiving, Picking, Task Interleaving, Returns, etc.).
3. **Rule Mapping:** It refers to a customer-specific rulebook created by WMS architects. These rulebooks mirror the logic of your actual implementation — no guesswork, no assumptions.
4. **Data Retrieval & Validation:** Maya dynamically generates SQL queries and validation logic, connects to the WMS database in read-only mode, and retrieves data in real time.
5. **Decision Execution:** Based on validation results, Maya assembles a coherent, context aware response in plain English. Where needed, it suggests resolution paths, next steps, or exceptions to be raised.

This full cycle typically completes in seconds.

Maya's Architecture: A Deep Dive

At the heart of Maya is a multi-model AI framework, supported by:

- **Three fine-tuned models, each with a specific role:**
 - **Intent Interpreter:** Understands user input and refines vague queries.
 - **Rule Matcher:** Applies logic from the rulebook to classify and validate issues.
 - **Response Generator:** Converts backend results into clear, actionable outputs.
- **Structured Rulebooks:**
 - Authored and validated by seasoned WMS experts.
 - Continuously updated during implementation.
 - Stored in a structured format to support automated lookups and rule validation.
- **Task-Focused Execution Agents:**
 - Each agent is aligned to a WMS process area.
 - They validate rules, fetch data, apply conditional logic, and decide resolution pathways.
- **Real-Time Database Access:**
 - Read-only connection to the customer's WMS database.
 - Queries executed securely, with data privacy handled via Google DLP.

This layered approach delivers speed, consistency, and traceability — while allowing Maya to operate non-intrusively within your existing WMS environment.

Accuracy Without Hallucinations

Maya is not a general-purpose LLM. It does not guess. Unlike typical generative AI tools, every response from Maya is grounded in:

- **Verified business logic**
- **Context-aware rule interpretation**
- **Real-time system data**

This makes Maya especially suitable for enterprise use cases where factual correctness and traceability are non-negotiable.

Speed and Cost Efficiency: Real Metrics, Real Results

Maya delivers what most human-led support models can't:

- Issue resolution in under 3 minutes for most complex scenarios
- Response generation in seconds for simple issues
- Resolution accuracy exceeding 97%
- Support costs that are 2–3X lower over a three-year period compared to traditional support teams

Each issue processed by Maya brings architect-level reasoning without the wait. Enterprises see direct operational gains and substantial ROI.

Maya Is More Than a Support Tool — It's Embedded Expertise

Maya acts like a full-time, always-on WMS analyst — one that doesn't take breaks, doesn't need onboarding, and never forgets a rule.

It supports high-scale environments with precision and consistency, enabling:

- Faster onboarding of new warehouse sites
- Reduced reliance on external consultants
- Continuous improvement via real-world feedback loops

Conclusion: The Future Is Already Here

Maya shows what's possible when domain expertise meets engineering precision. It's not just a tool — it's how support should work in the era of real-time, AI-driven operations.

With Maya, you don't just automate responses. You codify your expertise. You scale your operations without scaling your teams. And you deliver expert-level support to every issue — instantly.

Rajesh Madhavan
CEO, eSoftLabs
Head – R&D, AI