

Intelligent at Scale: How Maya Powers Multi-Site WMS Operations

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

Executive Summary

Modern supply chains are no longer centralized. Enterprises operate across regions, time zones, and warehouse models — each with unique workflows, customizations, and support challenges. As operations scale, the complexity of support scales with it.

This whitepaper explores how Maya, our AI-powered WMS expert, delivers intelligent, real-time support across multi-site environments. Whether your operations span five sites or fifty, Maya brings consistent resolution, deep contextual awareness, and expert-level support at scale — all without increasing headcount.

The Multi-Site Challenge

Global WMS operations are inherently complex:

- Different sites often have different rulebooks
- Operational priorities vary by region
- Language, time zones, and team capabilities differ
- Escalation paths become fragmented and inconsistent

Traditional support models struggle here. You end up with:

- High reliance on local experts
- Inconsistent support quality
- Slow cross-site knowledge transfer
- Escalations that stall due to lack of context

Maya was built to change that.

Maya's Architecture: A Deep Dive

Maya is architected to scale — not just in volume, but in intelligence:

- **Central Knowledge, Local Context:** Maya operates on a shared rules framework while adapting to each site's unique rulebooks, workflows, and exceptions.
- **Language-Optimized Query Engine:** Supports 30+ languages with built-in input optimization to handle regional language nuances.
- **Decentralized Execution, Central Oversight:** Maya processes queries independently at each site, while enabling centralized monitoring, dashboards, and governance.
- **No-Code Customization:** Rule updates and new scenarios can be onboarded in under a week, without requiring WMS application changes.
- **Token-Based Pricing:** Scales seamlessly from small pilot rollouts to full enterprise deployment — with cost predictability.

Site-Level Adaptability Without Rework

Each site may have unique configurations — different carriers, order types, pick strategies, exception flows. Maya adapts without needing custom code or duplicated efforts:

- Custom rulebooks per site
- Localized data validation logic
- Conditional process routing
- Context-aware responses that reflect local exceptions

Support teams no longer need to reinvent the wheel for each new site. Maya carries forward organizational intelligence and applies it smartly based on context.

Case-in-Point: Scaling Without Scaling Teams

In deployments across global warehouse networks, Maya has shown:

- Up to 10X workload capacity per support analyst
- Consistent 97%+ resolution accuracy across sites
- 2–3X reduction in support costs over three years

A typical enterprise that spends \$1M annually on WMS support can expect Maya to maintain performance while significantly reducing cost year-over-year — with predictable scale.

Unified Visibility With Role-Based Controls

Maya supports both local and central roles:

- Site-Level Users: Raise and resolve issues within their operational scope
- Central Admins: Monitor cross-site activity, review exceptions, and manage usage
- Support Leadership: Access analytics to guide staffing, process improvement, and automation opportunities

This model ensures alignment without micromanagement.

Seamless Integration With Enterprise Architecture

Maya is designed to fit — not disrupt:

- Runs on AWS and Google Cloud
- Interfaces with WMS databases (like MS SQL Server) in read-only mode
- Integrates with support tools like ServiceNow
- Mobile-ready and responsive
- SSO and RBAC supported

Whether it's a central command center or a mobile user in a remote site, Maya delivers a consistent, intelligent experience.

What This Means for Enterprise Operations

By replacing traditional support with Maya, organizations unlock:

- **Faster decision-making** at the floor level
- **Lower dependency** on specialized consultants
- **Consistency** in how issues are diagnosed and resolved
- **Operational resilience** across shifts and sites

For CIOs and COOs, this means predictable costs, higher customer satisfaction, and the confidence to scale without friction. For WMS teams, it means fewer fire drills and more time to focus on strategic improvements..

Looking Ahead: Toward Proactive, Autonomous Support

Maya today handles support queries intelligently and independently. But we're not stopping there.

The next evolution is a brain outside the brain — a model where Maya separates reasoning from memory. It will:

- Learn continuously from every issue it resolves
- Store context in knowledge graphs and scenario libraries
- Enable onboarding of new rules and exceptions in under a week

This shift will make support not just intelligent, but proactive. Maya will begin to anticipate issues, suggest optimizations, and reinforce best practices — all while integrating seamlessly into your existing WMS environment.

Conclusion

Support doesn't have to be slow, manual, or expensive. With Maya, it becomes a strategic advantage.

We built Maya not just as a tool, but as a solution grounded in years of experience supporting real-world warehouse operations. It combines domain expertise, AI precision, and operational scalability in a way that's never been done before.

The future of WMS support is autonomous. It's fast. It's cost-efficient. And it's already here

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