AI-ready or AI-native? A TMS vendor scorecard for bulk and specialty fleets

An **AI-ready TMS** is an existing system opened up so separate AI tools or agents can read its data or operate its screens. An **AI-native TMS** is built so AI sits inside the data model, the decision engine, and the interface where dispatchers work. Both can use the same optimization methods. The difference is what the AI can see and what happens to its plan. Use the five questions below in every vendor demo. Ask to see each one live, on your own data if possible. Score each answer 0, 1, or 2.

Question	Ask the vendor to show you	Scores 0: AI-ready signal	Scores 2: AI-native signal	Vendor A	Vendor B	Vendor C
1. What data does the AI plan on, and how current is it?	Change a trailer's last-contained product or a driver's remaining hours. Show how fast the plan reflects it.	Nightly exports or sync jobs. Wash history, training records, or customer site requirements kept in other systems or spreadsheets.	One live data model covering inventory, yard limits, drivers, trailers, washes, and customer site requirements.			
2. Where do dispatchers do their work?	Accept, edit, and override a recommendation. Show where the override is recorded.	Work spread across the TMS, ELD, spreadsheets, whiteboards, and paper forms.	One integrated workspace. Overrides and notes are captured as structured data.			
3. Does it plan the whole load at once, and replan when things change?	Plan one order end to end: source, trailer, driver, wash, backhaul. Then simulate a driver calling in sick at 2 p.m.	Separate tools for routes, driver assignment, and trailer status. Disruptions are fixed by hand.	Compares complete options across all open orders, shows the trade-offs, and re-solves automatically.			
4. How and when are constraints enforced?	Edit the plan so it breaks a compatibility, hazmat, or hours-of-service rule.	Rules checked only when the suggestion is made. Edits are never rechecked.	Every rule rechecked on live data whenever the plan or the data changes.			
5. Does it learn from outcomes and from dispatchers?	Show how actual wash times, dwell times, or an override change future recommendations.	Improves only with vendor updates. Outcomes stay in the legacy system.	Outcomes and dispatcher decisions feed back into the model continuously.			
Total (out of 10)	Score 1 when the answer is partial, or shown on slides but not live.					

Reading the total. 8 to 10: AI-native. 4 to 7: partly integrated; probe where the gaps are. 0 to 3: AI-ready, with AI added on top of a legacy system.

Is the switch worth it? If you suspect your business is leaving money on the table because decisions aren't optimal, an AI-native platform is worth the implementation effort when planning quality itself is the bottleneck: dispatchers juggling constraints in their heads, wash costs and empty miles you can't explain, institutional knowledge that walks out the door when a planner retires, or growth your current team can't absorb without adding headcount.

One more request for every vendor: a proof of value or a post-implementation customer case, measured against what that customer's dispatchers actually did, such as wash spend, deadhead miles, layovers avoided, or planning time per load.

Published by BeyondTrucks, beyondtrucks.com. Free to share, print, and reproduce with attribution.