AMYGDA

ENTERPRISE AI IN TRANSPORT: WHAT'S WORKING, WHAT ISN'T, AND WHERE TO START

LEADERSHIP BRIEFING

AI is moving fast. Transport is starting to catch up. This briefing cuts through the noise on what is actually delivering results in operations, where the real obstacles sit, and what a sensible first step looks like.

www.amygdalabs.com

The Board Problem

The pattern is identical across every sector: board demands AI, CEO commissions a consultant, a centralised pilot gets stood up, nobody aligns the operations around it, and it quietly dies.

Transport is no exception — in fact it is worse, because transport has some of the most fragmented, legacy-heavy data environments of any industry.

A CEO can launch a project. Only the head of operations, the fleet director, or the regional MD can make it part of how the business actually works. If the business unit owner is not accountable for the outcome — if this sits with a central transformation team or IT — it will not stick.

AI initiatives that are driven from the top without genuine operational buy-in don't work. Period.

Secondly, the technology is not the failure - it is behaving exactly as built.
It is organisations bolting AI onto existing workflows rather than being open to redesigning workflows around AI.

There is a difference, and it matters enormously. And you do not need to pay consultant $500,000+ to find why it matters. Talk to me if in doubt!

Before We Brief - Who are we?

About Amygda

Amygda is the operational intelligence platform for transport industry — the decision layer for trains, aircraft, and airport equipment that cannot afford to fail in peak hours.

We deploy AI that turns the data buried in operations into decisions engineering teams can trust and act on. Backed by the Boeing company, we work with the world's largest transport brands.

Faizan Patankar

Faizan is the founder and CEO of Amygda, the operational intelligence platform for transport. He spent his early career inside Rolls-Royce, working on the data and AI behind some of the most demanding engineering systems on earth - jet engines.

Faizan regularly writes and speaks on what it takes to deploy AI inside mission-critical operations.

What AI in Transport Doesn’t Do

Here is the uncomfortable truth / fact / current state (as of June 2026): AI does not help you integrate anything.

Your TMS does not talk to your depot management system. Your traffic planning tool does not talk to finance. Driver hours data sits somewhere else entirely. An AI agent given dispatcher-level access hits exactly the same walls your human dispatcher does — except the agent will not know to ring Sarah in compliance or ask Gary what the real position is in the yard.

Any transport business older than ten years is a mass of systems waiting to be integrated.

The answer is not to centralise your data or build a data lake. Transport businesses are already swimming in data. The question is not whether you have it. The question is whether an AI agent can read the authoritative version of it — the system that holds the real number, not a downstream copy that was last synced three days ago. That is a targeted access problem.

Lots of transport leaders, in search of ‘MVPs’ and ‘AI POCs’ often pressure teams into creating another source of data just for the project, which then adds to the admin budget of IT teams to look after.

Don’t do that. Tackle the hard problems as part of the AI project - if a human level access is not possible, you cannot expect a human level outcome from AI.

What Actually Works

The instinct to "add AI" by putting a chat interface in front of an existing process is fine as a starting point, but limited in what it changes.

Infact it’s probably a bit late to treat that as frontier AI now. It was a 2023 moment (Nov 2022 is when ChatGPT launched for public use).

And suggesting your employees use Copilot is not the answer - your employees are already using Copilot, they probably tried it, and found how thrash it is. *They probably get more value from ChatGPT than Copilot. Yes, quote me on that.*

More useful is asking: where in our operation does someone currently spend time finding, compiling, or sense-checking information?

That is where AI adds value fastest — because it can do that work without touching your core processes at all. It is also a read-only operation, which means AI won’t write or edite to your current source dataset or databse.

Start with read-only use cases, first:

- An agent that pulls together cross-fleet performance data before the morning briefing
- A summary of a customer fleet account ahead of a account meeting
- Anomaly flagging across utilisation or compliance data that your team would never have bandwidth to spot manually

None of these require changing how you work. They sit alongside existing process. It is currently a *Job to be done (JTBD).* Let AI do that first, before eventing new tasks.

Over time, as confidence builds and the business unit sponsor sees results, workflows naturally shift. Change management becomes part of the conversation at that point — as a consequence of something that is already working, not a condition of starting.

Immediate Shift You Need to Make

Stop thinking of AI as software you deploy - and forget. Start thinking of it as a new type of employee you need to onboard. And no we are not talking about human level competence.

That means:

- Giving agents proper access credentials, not borrowed logins
- Building clean data pipelines so agents read authoritative sources, not stale copies
- Treating AI-generated work the same way you would treat work from a new hire: supervised, reviewed, and checked

The acting AI agents — ones that reschedule runs, amend orders, reallocate drivers — come later, once you have established trust in the data and the controls.

What the Productivity Numbers *Actually* Look Like

Ignore the 10x headlines.

The realistic gain for a well-run transport operation using AI properly in targeted areas is anywhere between 3x to 4x in near term. That is still significant.

But it only materialises when AI agents have access to the right data and there is a person checking the output. This is called a *Human in the loop (HITL) process.*

There are multiple reasons why HITL is needed in transport AI. But mainly because this new field of AI is non-deterministic and prone to non-reproduceability. Which means 2 out of 10 times it is possible the AI gives incorrect answer that needs reviewing.

The bottleneck in making the project sticky is rarely the AI. It is the review process.

You need someone with operational context to sense-check what the agent produces, particularly early on. Budget for that time — it is not overhead, it is how you build the trust that eventually lets you reduce it.

Will AI Kill Jobs?

No. And the record on this is consistent enough to be worth stating plainly. What will change is the nature of the work.

IBM's pitch in 1965 was that computers would eliminate accountants. There are more accountants now than at any point in history — because once you can process financial data at speed, the scope of what you can do with it expands, and you need more people to manage that complexity, not fewer.

The same pattern played out in law. Automation was supposed to hollow out the profession. Every lawyer you encounter today is a computerised lawyer — citations from databases, documents in track changes, briefs typed rather than dictated. The profession is substantially larger than it was thirty years ago.

The pattern in transport will be the same. Planners, coordinators, and operations managers are not going away. What changes is the nature of the work. The best people will shift toward directing and reviewing AI-generated outputs rather than producing everything manually. That is a better job, not the absence of one.

The concern about job losses tends to come from a static view of what work is. The amount of useful work available expands with the capability of the tools available to do it. That has been true of every major technology shift.

There is no compelling reason to think this one is different.

3 Qs before launching AI Project

1. **Do we have a business unit sponsor —** someone with real operational accountability — driving this, or is it sitting with a central team that does not own the outcome?
2. **Can an agent access the authoritative version of the operational data**, or will it be reading stale copies from secondary systems?
3. **When the AI gets something wrong — and it will — is there a person in the loop** who catches it before it becomes a customer or compliance problem?

If you cannot answer yes to all three, that is where to start.

LEADERSHIP BRIEFING

ENTERPRISE
AI IN TRANSPORT: WHAT'S WORKING, WHAT ISN'T, AND WHERE TO START

AI is moving fast. Transport is starting to catch up. This briefing cuts through the noise on what is actually delivering results in operations, where the real obstacles sit, and what a sensible first step looks like.

June 2026.

Faizan Patankar
CEO, Amygda
faizan@amygdalabs.com
www.amygdalabs.com

Add me on LinkedIn
Msg: +447538332580

AMYGDA