


Your competitors are shipping AI features. What's stopping you?


Getting to a working AI feature isn't mostly an AI problem. It's an infrastructure problem, and most companies find that out later than they'd like.


DOES THIS SOUND LIKE YOU?


AI keeps coming up in board meetings, in competitor announcements, in every industry newsletter you've read for the past year. At some point, you sat down with your CTO and asked a straightforward question: how long would it take to ship something? The answer was longer than you expected, and the reason wasn't the AI part. Before you can connect a model to anything useful, you need clean data, a proper API layer, and an architecture that can handle async workloads. None of which you currently have. Technology isn't the obstacle. The fifteen years of accumulated infrastructure decisions underneath it are.

SOUND FAMILIAR?

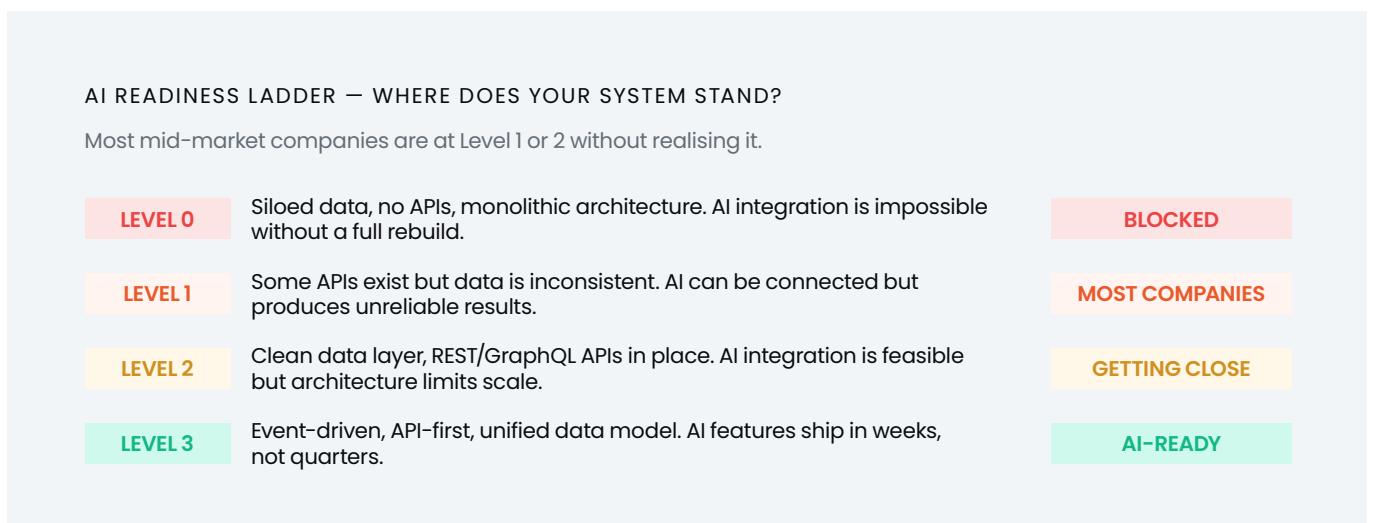
- 

"We have an AI strategy"
It's on slide 14 of the board deck. Nothing has shipped yet.
- 

"How long would it take to connect a model to our data?"
You asked six months ago. The answer is still "it's complicated."
- 

"Our competitors just launched another AI feature"
And you're still figuring out why your data isn't ready for it.
- 

"The ambition is there, the infrastructure isn't"
Every AI conversation eventually hits the same wall.



How we get you to AI-ready.

- 1 Find out what's actually standing in the way | 2 weeks**

Most teams know they're not ready for AI. What they don't know is exactly why, and what it would take to get there. We look at your system across four dimensions: data quality, API surface, architecture patterns, and infrastructure flexibility. You leave with a written report that gives you a straight answer on where you stand, what's blocking you, and what the path forward looks like. Something concrete to bring back to your board instead of another estimate.
- 2 Foundation modernization | months 1-4**

We build what AI actually needs: a unified data model that consolidates your siloed systems, a clean API layer that models can query reliably, and an event-driven architecture that handles async AI workloads. This is the work most companies skip, and the reason most AI projects fail before they start. Getting this right means every AI feature after it takes weeks, not months.
- 3 First production AI feature, months | months 4-8**

We don't stop at infrastructure. We deliver one working, production-grade AI integration: an LLM-powered feature, a recommendation engine, or a predictive analytics module. You go to your next board meeting with something live, not a roadmap slide. And your competitors who've been shipping AI features? You're now one of them.

Ready to talk through it?

Most AI conversations we have start the same way: the strategy is there, the board is asking about it, but nobody has been able to give a straight answer on what it would actually take to ship something. That's a good place to start. Walk us through where you are and what you're trying to build, and we'll tell you honestly what we think is standing in the way.

[Book a 30-min call](#)

WHO WE ARE

We are a Berlin-based software engineering company with 100+ engineers and 15 years of experience delivering complex projects for mid-market companies across the Netherlands, the UK, Scandinavia and others. Our engineering teams operate across Serbia, Bosnia & Herzegovina, and Portugal. We've spent 15 years inside the kind of systems that need modernizing. We know exactly how they break, and how to fix them. We offer both managed modernization projects and dedicated engineering teams, depending on what works best for your organization.

