



DoorDash Didn't Buy an AI Platform. It Built One.

DoorDash launched an internal AI research lab and its own benchmark, tuned on its own codebase and incident history. The lesson isn't about DoorDash's scale. It's about the question every business should be asking.

By **Mike Walker** · Founder, SL8 Agency

DoorDash just launched an internal AI research lab. Not a customer-facing product. Not a tool to sell. A research arm, with its own benchmark, built to answer one question: how do we run AI on our own business, well.

The detail worth sitting with is the architecture, not the announcement.

What they actually built

Instead of buying an off-the-shelf AI tool, DoorDash built a model-agnostic routing system. Simple tasks go to cheap open-source models. Hard tasks go to frontier models. The router decides which is which, and the whole system is tuned on DoorDash's own codebase and incident history, not generic training data.

Two things fall out of that design.

Control. Code and business data never leave the building. No vendor sees the codebase to route a support ticket.

Cost. Frontier model pricing only gets paid for work that actually needs a frontier model. Everything else runs on cheaper open models, at a fraction of the cost, with no drop in output for the task at hand.

Why "full context" is the real feature

A generic AI tool knows how AI tools work in general. It doesn't know your incident history, your codebase's specific failure patterns, or the shorthand your team uses in a ticket. DoorDash's system does, because it was tuned on exactly that.

This is the part that generalises past DoorDash's scale. The value isn't the model. It's the business context wrapped around the model: the data, the history, the specific patterns of how *this* company actually operates. That context is what turns a general-purpose model into something that makes fewer wrong calls on your specific business.

The question this raises for smaller businesses

You don't need an internal AI research lab to apply the same logic. Most mid-market businesses don't need frontier-model pricing for 90% of what they'd use AI for, and most don't need a single vendor holding their operational data to get value from AI either.

The build-versus-buy question isn't really about engineering capacity. It's about which parts of your AI stack you're comfortable renting, and which parts need to run on your own data, under your own control, priced for the actual difficulty of the task.

DoorDash answered that question at their scale. Every business running AI in production, at any size, is answering a smaller version of the same question, whether they've named it or not.

Mike Walker

Founder, SL8 Agency

Ready to stop experimenting and start building?

One conversation. A straight read on where you actually stand.

BOOK A CALL



If the button doesn't open: `calendar.app.google/7J2ThQMgXULCqeZe6`