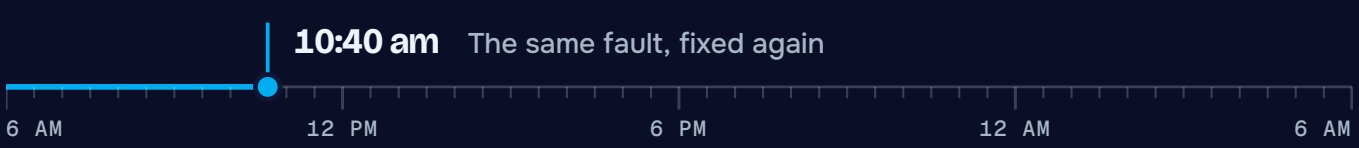
Web Vijayi

SERVICE BROCHURE

LLM fine-tuning on your own data

A model that follows your terms and format, tested before you say yes.





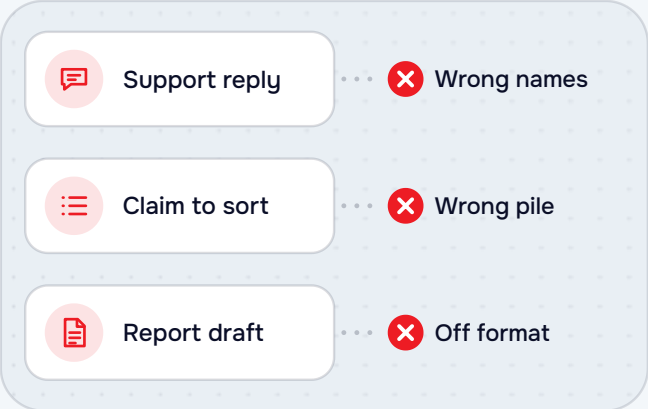
Who it is for

Medium and large companies that use a language model at volume and keep seeing the same faults, and product teams that want a smaller model that's cheaper to run.

- CUSTOMER SUPPORT GROUPS
- CLAIMS AND TICKET TEAMS
- SOFTWARE PRODUCT TEAMS
- DATA-SENSITIVE FIRMS

The same mistake, put right again

- A support assistant gets your product names wrong, or a drafting tool ignores the house format.
- A model that sorts claims or tickets puts the same awkward cases in the wrong pile every week.
- At high volume, every repeated fault means staff putting the same mistake right by hand.



BEFORE: THE SAME FAULTS, FIXED BY HAND

What we set up

First we test whether better prompts fix the faults. When they don't, we train an existing language model further on examples from your business, so it follows your terms, format and tone. Then you see it scored against the original on cases it has never seen.

Prompts tried first

We test better prompts on your real cases first. If they reach the target, we stop there.

Data you approve

We clean your examples, take out personal details the task doesn't need and record each source.

Tested on unseen cases

The tuned model, the original and the prompt-only version run the same held-back tests, safety included.

How it works

YOU DO THIS · WE DO THIS · TOGETHER

1 · TOGETHER

We agree what "better" means

We write success measures as tests and build a test set from real cases. Part of it is kept back and never used for training.

2 · US

We try prompts and record the score

That score is the baseline every later step has to beat. If prompts reach the target, we stop there.

3 · TOGETHER

We prepare the data and tune

You choose the sources and approve the dataset before training. We tune the model, then test it against the original, safety included.

4 · YOU

You decide on go-live

You see the results side by side and say yes or no. The previous version stays ready, so you can switch back.

What you get

- ✓ A written test plan with success measures and a held-back test set you keep
- ✓ Tested prompts with a version history and the score of each version
- ✓ A training dataset you approved, with a record of where each part came from
- ✓ The tuned model or adapter files, stored where you choose
- ✓ A results report comparing the tuned model with the original, safety tests included
- ✓ Training records for every run, and a way back to the previous version

What we need from you

- Someone who knows the work, with time to judge and rank answers
- Real examples, such as past tickets or reports, cleared for training
- Your rules: what the model must never say or use
- Where the model must run, and any security review you need
- Your legal adviser's check on the data and the model's licence

Timeline We agree the timeline with you before we start.

You stay in control

Check it, pause it or stop it at any time.

Nothing goes live without your yes

Nothing is trained on data you haven't approved, and the tuned model replaces nothing until you've seen its results.

You see the test results

The results report sets the tuned model beside the original on the same tests, with records of every training run.

Unusual cases go to your team

Where we also connect the model to your systems, it passes unusual cases to your team and keeps a record.

Pause it or switch back

Check it, change it, pause it or go back to the previous version whenever you like.

Questions people ask

Should we fine-tune or use our documents?

Fine-tuning changes how a model works: your terms, format and tone. Facts that change, like prices, stay in your documents for the model to look up.

How much data does it need?

It depends on the task, and quality matters more than volume. We size it once we've seen your examples and how often your experts agree.

How do you know it worked?

Part of the test set is never used for training. The original, the prompt-only version and the tuned model are all scored on it, side by side.

Who owns the tuned model?

We agree ownership in writing before any work starts. If you need to hold the model yourself, we tune an open model whose licence suits your use.

Tell us what your model gets wrong

In a free consultation we'll hear which answers go wrong today, then tell you whether better prompts, your documents or fine-tuning is the fix.

We quote in writing after a free consultation.

webvijayi.com

care@webvijayi.com

Book a free consultation:
webvijayi.com/contact