

1 AI A DAY · @JJWONG.AI

Local or Cloud: The Honest Maths

1 AI A Day · Day 17 · @jjwong.ai

JJ Wong · jjwong.ai

How to decide whether running your own AI model is worth it for your business — a four-question test, a cost worksheet that prices your own time in, and what I actually run, pay for and cancelled.

The four-question test

Work down the list. Stop at the first yes.

- 1. Must the data stay on your premises?** Because a contract or a regulation says so — not because it would be nice. If yes, local is probably right.
- 2. Is the task simple and very high-volume?** Tagging, extracting fields, classifying thousands of rows a month. If yes, local may well pay for itself.
- 3. Is your connectivity genuinely unreliable** where the work happens? A warehouse, a factory floor. If yes, local keeps working when the line drops.
- 4. Do you enjoy running it?** A real reason. Just be honest that it's the reason.

Four noes: pay the subscription. Spend the weekend on your business instead.

Notice "it's cheaper" isn't on the list. For most small businesses it isn't, once your time is priced at anything at all.

The cost worksheet

Fill in your own numbers. Be honest about the hours.

LOCAL	YOUR NUMBER
Hardware (machine, upgrades) ÷ 24 months	RM ____ / month
Electricity for the hours it runs	RM ____ / month
Setup time, spread over 12 months	____ hours ÷ 12
Maintenance time — updates, things breaking	____ hours / month
Your hourly value	RM ____ / hour
Total monthly cost = hardware + power + (setup hours + maintenance hours) × hourly value	RM ____

CLOUD	YOUR NUMBER
Subscription or API spend	RM ____ / month
Time spent managing it	____ hours / month × hourly value
Total monthly cost	RM ____

For most people doing general thinking work — writing, planning, analysis — the cloud side wins before the hourly value is even filled in.

The maths flips at volume. Running one simple task ten thousand times a month is exactly where local starts to make sense, because API costs scale with usage and your hardware doesn't.

Where the quality gap shows

Local models are genuinely good at:

- Summarising and extracting
- Tagging and classifying against clear rules
- Reformatting and cleaning data

They're noticeably weaker at:

- Long reasoning across many steps
- Careful writing that has to sound like you
- Following a long, complicated set of rules without drifting
- Anything where a subtle mistake is expensive

Match the model to the task, not to a principle.

The privacy trap

Plenty of people run a model locally "for privacy" — then paste the same data into a cloud tool at the next step of the workflow.

Privacy is a property of the whole pipeline, not one node in it. If any step sends the data off your machine, the local step bought you nothing. Map every step first. Then decide.

What I actually run, pay for, and cancelled

Running locally

TOOL / MODEL	FOR WHAT	WHY LOCAL
[JJ: fill in]		

Paying for

TOOL	FOR WHAT	ROUGHLY PER MONTH	WHY IT'S WORTH IT
[JJ: fill in]			

Cancelled

TOOL	WHY I CANCELLED
[JJ: fill in]	

The one thing to remember

The real comparison isn't free versus paid. It's your weekend versus a subscription. Answer the four questions, fill in the worksheet with your own hourly value, and let the number decide.

Questions? Reply to this email — I read them all. More at [@jjwong.ai](#) — one AI play for SEA sellers, daily.