

GPT-5.6 isn't one model — it's three, and you're on the wrong tier.

What OpenAI's new three-tier pricing means for your business — and 3 tasks you're overpaying to run on the wrong model.

Source: Engadget · OpenAI rolls out GPT-5.6

— THE NEWS, IN ONE BREATH

This week OpenAI rolled out GPT-5.6 — and it isn't one model, it's three. Luna runs at \$1 per million tokens in, \$6 out. Terra sits in the middle at \$2.50/\$15. Sol, the heavy one, costs \$12.50/\$75 — more than twelve times the price of Luna. Here's the part nobody says out loud: most of the AI work a business actually does — drafting emails, summarizing threads, sorting leads — does not need the \$12.50 model. If your tools default everything to the top tier, you're quietly paying Sol prices for Luna work.

— SO WHAT — IF YOU OWN THE BUSINESS

3 WAYS TO USE IT THIS WEEK

1 *Right-size your everyday AI tasks*

~15 MIN

COST AUDIT

NO CODE

Do it: Drafting, summarizing, and sorting are Luna-tier work. Point the boring, high-volume tasks at the \$1 model instead of letting a tool default them to the flagship.

Where DIY stops: Flipping one tool's model in settings is a click. Making every tool auto-route each task to the cheapest tier that can handle it — that's the wiring.

2 *A cheap first-pass filter on everything inbound*

~20 MIN

SALES / OPS WORKFLOW

NEEDS WIRING

Do it: Let the \$1 tier read and sort every inbound lead, email, or ticket, and only escalate the 10% that genuinely need a human or a smarter model.

Where DIY stops: The sorting logic is easy to describe. Wiring it to fire the instant a lead or email lands, every time, without you touching it — that's the engine.

3 *The tool you shelved because AI was 'too expensive'*

VARIES

CUSTOM BUILD

NEEDS WIRING

Do it: That quote-drafter, report-generator, or intake tool that didn't pencil out at flagship prices now runs for cents per task on the \$1 tier. The math just changed.

Where DIY stops: Cheaper tokens lower the cost of the thinking. Turning it into a tool your team actually opens every day is the part I build.

Here's the move: picking a model tier isn't an engineering decision — it's a pricing decision, and you can make it today. Almost every operator runs one expensive model for everything because nobody told them there were cheaper tiers built for exactly the repetitive tasks. Match the task to the tier and the same work costs a fraction of what you're paying now.

*That's why ProofWright exists. **I'm not an engineer either** — I built 10 of these for my own five companies, zero engineers hired. I build the same engines for a few operators at a time — the ones that run in the background while you run the business.*

Want to know which tier your tools are actually calling?

Tell me what your business does — reply right here — and I'll map which of these three is cheapest to wire up for you now. No pitch, no deck.

Ready to build? Book a Teardown → [proofwright.ai/teardown](#) — 90 min, \$500, credited toward any build in the next 30 days.

→ Reply with your industry. Worst case, you get a straight answer for free.