



PLAYBOOK

The AI Meter: A Leader's Playbook for Budgeting AI That Charges by the Task

Six steps for when AI stops being a flat licence and starts billing like a contractor. How to find the meters you already have, price work by the task, give every agent a budget and an owner, set spending guardrails before you switch anything on, measure value honestly, and build the skills that make the spend worthwhile. Written for C-suite and directors in the UAE and Saudi Arabia.

7 October 2026

PEOPLE FOR TOMORROW

What's inside

For most companies, AI has been a flat fee: so many seats, so much a month. That is changing. The most useful AI — agents that do real, multi-step work — is moving to pay-as-you-go. This playbook is about running that meter well: knowing what you're paying for, who's accountable for it, and whether the work is worth it.

-
- 01** Find the meters you already have
 - 02** Price the task, not the seat
 - 03** Give every agent a budget and an owner
 - 04** Set the guardrails before you switch anything on
 - 05** Measure value honestly, and give it time
 - 06** Build the skills that make the spend worth it
-

01

STEP 1 OF 6

Find the meters you already have



Before you plan anything new, find out what's already running. Pay-as-you-go AI arrives quietly: a credit pack on a productivity suite, an API key a team set up for a pilot, a usage add-on inside your CRM or HR system. Each one is a small meter, and most companies don't have a single list of them.

Start with finance and IT together. Pull the last six months of invoices and look for

anything billed by usage, credits or tokens. Then ask each department head one question: what AI tools is your team using that someone pays for? You'll usually find more than you expected.

That list is your baseline. It tells you where spend is today, who started it, and which tools are doing the same job twice.

- ▶ **You can't manage AI spend you can't see. Make one list of every meter first.**

02

STEP 2 OF 6

Price the task, not the seat



Seat pricing asked one question: how many people need it? Usage pricing asks a different one: which pieces of work are worth paying an agent to do? The unit that matters now is the task.

For each workflow you're considering, write down four numbers. How often the task happens. What it costs today in people's time and outside spend. What it will cost with AI, including the human time to check

the output. And what a mistake costs.

Microsoft's own example shows how quickly this moves: a \$30 seat covering everyday use becomes a modelled \$73 once five complex agent tasks are added.

Some tasks will look brilliant on paper.

Others — low volume, high stakes, lots of checking — won't be worth it yet. Both answers are useful.

- ▶ **Every agent proposal should come with a cost per task, today versus tomorrow.**

03

STEP 3 OF 6

Give every agent a budget and an owner



An agent without an owner is a meter nobody's watching. Put one named person against every agent — the business owner of the work, not IT — and give them a monthly budget for it.

This is the same idea the FinOps community calls showback: make each team's AI usage visible to that team, even before you charge it back to them. When people can see what

their agent costs each month, they start asking the right questions on their own. Is this the right model for the job? Is it running tasks nobody needs?

Match the model to the task, too. The newest, most capable models cost the most. Plenty of routine work — sorting, summarising, first-pass drafting — runs perfectly well on a cheaper one.

► **One agent, one owner, one monthly budget they can see.**

04

STEP 4 OF 6

Set the guardrails before you switch anything on



The best time to set a spending limit is before the first task runs. Most enterprise AI platforms now give you the controls: budgets, hard limits, alerts, and restrictions on which teams can use which models. Microsoft, for example, keeps its metered Copilot services switched off until an administrator creates a spending policy.

Use those controls deliberately. Set an alert at 75% of each budget and a hard stop at 100%, at least for the first quarter. Decide in

advance who can approve extra spend and how quickly.

In the Gulf, add one more guardrail: where the work runs. With more local cloud and AI capacity arriving — Microsoft alone has committed more than \$10bn across the UAE, Saudi Arabia, Kuwait and Qatar through 2030 — you can increasingly choose to keep sensitive AI work in-country. Write that choice into the policy, not just the contract.

- ▶ **Budgets, alerts and hard limits go live before the agent does — not after the first invoice.**

05

STEP 5 OF 6

Measure value honestly, and give it time



Usage is easy to measure. Value is harder, and it's the number that matters. For each agent, track two things alongside cost: how much of the task it completes without help, and how often a person has to correct it.

Be realistic about timing. Gartner's research with finance leaders found that routine tasks — pulling data from documents, automating payables and receivables, building reports — typically deliver their

expected value within nine to 10 months. More complex work like forecasting takes longer.

That's a reason for patience, not an excuse. Review each agent every 90 days against the numbers you wrote down in step two. Keep it, change it or stop it. And don't let the quick wins crowd out the slower projects — that's often where the better decisions come from.

► **Cost, completion rate, correction rate. Review all three every 90 days.**

06

Build the skills that make the spend worth it



The same AI tool can be good value in one team and wasted money in the next. The difference is usually the people using it. Gartner's finance research found that low AI literacy is now the biggest barrier to getting value from AI — ahead of finding specialist talent.

So fund the skills alongside the software. Give people real projects to use AI on, a safe sandbox to experiment in, and short on-the-job tasks rather than long courses. Teach

them to check the output, not just to write the prompt.

In Saudi Arabia, SDAIA's new Data and AI Professional Standards Framework, built with the Ministry of Human Resources and Social Development, now defines the roles and competencies behind AI work. It's a useful yardstick for any Gulf employer deciding what "good" looks like. And where you need skills faster than you can build them, bring in specialists on a flexible basis while your own team catches up.

- ▶ **Every AI budget needs a skills line. The meter only pays off when people use it well.**



PEOPLE FOR TOMORROW

None of this needs a new department. It needs one owner per agent, a budget with an alert on it, an honest cost per task, and a team that knows how to use the tools well. Put those four things in place before the meter starts, and AI spend becomes something you manage — not something that surprises you.

Read the playbook → blog.tascoutsourcing.com/hubfs/TASC_Outsourcing/TASC%20AI%20Newsletters/Edition%209/TASC-Report-Edition-09.pdf