BUSINESSGPS

THE CAPABILITY BRIEF

The Capability Brief

On 12 November Capgemini reported that only a third of executives can confidently estimate the indirect costs of AI. Before an AI action goes live, put its full running cost on the rule sheet

Week of November 9, 2026

PUBLISHED

NOVEMBER 16, 2026

The two-minute version: price the checking, not only the AI

The scene. An AI action is coming to your programme. Before it goes live, write down what it will cost to run each month. Include the hours of the people who check and correct its work. Then get that figure signed off by name.

An AI action is one task an AI system does on its own: it sorts an email, drafts a reply or matches a payment. A rule sheet is a single page you write for one such action before it goes live. It records the rule the action replaces, what the action does and who owns it. It also records the evidence that will show it works, and how to return to the old way if it does not. This edition adds one row to that page: cost.

In this edition. On 12 November 2026 the Capgemini Research Institute published findings on what large organisations spend on AI. One headline in its press release reads: “Only a third of executives can confidently estimate indirect AI costs, including human oversight, rework and governance”.

What changed on the rule sheet:

- New. A cost row: the running cost of the action, the hours of the people who check it, and who signed that figure off.
- Escalating. The evidence row. A claim that an action saves time is hard to test when nobody has written down how long its work takes to check.
- Holds. The other rows: the rule replaced, the action, the owner and the way back.

Bottom line: count the people who check the AI’s work as well as the AI’s own charge, and get the total signed off before go-live.

Only a third of executives in Capgemini’s survey can confidently estimate AI’s indirect costs

Capgemini’s release of 12 November 2026 reports these findings:

- Only a third of executives can confidently estimate indirect AI costs, “including human oversight, rework and governance”.
- In the past 12 months, 59% of executives say their organisation’s AI-related expenditure was “significantly higher than expected or budgeted”.
- Among the organisations affected, spending was 24% above expectations or budget on average. Around one in five organisations (21%) overshot their AI budget by more than 30%.

Capgemini says its survey asked 500 “CXO-level senior executives” from large organisations, in September 2026. CXO-level means chief officers: the chief executive, the chief financial officer and their peers. Capgemini sells AI and consulting services, so it has an interest in these findings. The figures are what executives report about their own organisations. The release names no organisation’s result.

The 3 costs it names can each be counted on your own programme, in hours, before go-live.

The indirect costs are paid in hours, by people who already have a job

Capgemini's 3 indirect costs each have a plain meaning on a programme.

- Human oversight is the person who reads what the AI produced before it goes anywhere. A credit controller reads each payment reminder the AI drafted before it is sent.
- Rework is the time spent putting right what the AI got wrong. The same controller rewrites the reminder that quoted the wrong invoice.
- Governance is the work that keeps the action under control: a monthly look at a sample, a check of the log, the meeting that approves a change to the action.

None of these appears on the AI supplier's invoice. Each is paid in hours. Those hours often belong to a team whose time is already planned for something else.

So the cost row has 3 parts:

```
Running cost:   the AI's charge per month, at the volume the plan expects
People:         hours per month to check, correct and review, by role
Signed off by:  name, role and date, before go-live
```

Take the AI's charge from the contract, at the volume the plan expects

Ask whoever holds the supplier contract how this action is charged. It may be per use, per user or a flat licence. Turn it into a monthly figure at the volume the plan expects.

Some charges rise with use and some do not. A per-use charge doubles when the volume doubles. A flat licence stays the same until you need more of it. Write which kind this is beside the figure, because it decides how the row moves as more people use the action.

If the plan gives no volume, that is the first gap to fill. Ask the business owner how many items the action will handle in a normal month.

Done when: the row shows a monthly charge, the volume it assumes and the document the charge came from.

Time the checking on 20 real items before go-live

In testing, ask the person who will check the work to time themselves on 20 items. Note how many items needed correcting, and how long each correction took.

Time it as the checker will really work: at their own desk, with their other work going on. A check done in a quiet test room runs faster than one done between phone calls.

Multiply both figures up to the monthly volume. A timed sample of 20 is rough. It is still better than a guess, and it is your own. If the 20 show fewer than 3 corrections, time 20 more before you trust the correction hours.

If the action is already live, time 20 items from this week instead.

Done when: the row shows checking hours and correction hours per month, from a timed sample, with the checker's role.

Add the hours that keep the action under control

List the regular work that keeps this action safe to run. It might be a monthly review of a sample, a look at the log or a short meeting before any change. Put a role and a number of hours a month against each.

If nobody is named for this work yet, write that on the sheet. An action with no one reviewing it still has a governance cost. It is paid later, when something goes wrong.

Done when: the governance work is listed, each line with a role and hours a month.

Get the total signed off by whoever pays for the hours

Add the AI's charge and the people's hours, priced at your organisation's own hourly cost. Finance can give you that rate. If it cannot, use a stated rate and write it on the sheet. Take the total to the manager whose team supplies the hours. Ask them to sign it, with the date.

The checker often sits in a team that did not ask for the AI. Their manager may not know the hours are coming. Signing is how they find out before go-live, not after.

If you work for the integrator, draft the row with the client's own figures. The client's manager signs it. Check with your delivery lead what your firm has agreed to provide after go-live. Those support hours belong on this row too, so the client's manager sees them before signing.

A signed rough figure is worth more than an exact one nobody owns. The signature is what makes the hours real to the team that will supply them.

If no one will sign, write that on the sheet. Tell whoever owns the go-live decision.

Done when: a named person has signed the total, with a date, before go-live.

Compare the estimate with the first full month

Set a date one month after go-live. On that date, put the actual hours and the actual charge beside the estimate.

Write the difference on the sheet, whichever way it goes. If the hours fell as the people learned the work, record that. If they rose, find out why before the volume grows. While the volume is still growing, repeat the comparison at 3 months.

Done when: the review date is in the diary of the row's owner, and the estimate to compare against is on the sheet.

In an invented example, the people's hours cost more than 20 times the AI's own charge

Illustrative, and invented: an AI action sorts incoming customer emails and drafts a reply for a person to send. It handles 2,000 emails a month.

> Running cost: £100 in usage charges each month, from the supplier's price list. People: a service agent reads every draft, 1 minute each; about 1 draft in 10 is rewritten, 5 more minutes each; the team leader reviews a sample and the log, 4 hours a month. Signed off by: the customer service manager, before go-live.

The checking is 2,000 minutes. The rewriting is 200 drafts at 5 minutes, another 1,000 minutes. Together that is 50 hours a month. With the 4 hours of review, it is 54.

At an invented £40 an hour, the people come to £2,160 each month. The AI's own charge is £100.

Use your own figures. The sum shows which part of the cost to count first. It does not show that this action is poor value. A reply drafted in seconds may still save far more than it costs to check.

What to watch

The next Dallas Fed Banking Conditions Survey that repeats the question on customers' AI agents will show more than 11.4% of responding bankers saying they or their service provider have restricted that activity. This applies only if that survey is published by 31 March 2027.

Sources

1. Capgemini Research Institute, press release on the economics of AI, Paris, 12 November 2026, read on 15 November 2026. https://www.capgemini.com/news/press-releases/most-organizations-plan-to-adopt-a-multi-model-ai-strategy-as-nearly-six-in-ten-faced-significantly-higher-ai-costs-than-expected/
2. Federal Reserve Bank of Dallas, Banking Conditions Survey, special questions, November 2026, read on 15 November 2026 (for What to watch). https://www.dallasfed.org/research/surveys/bcs/specquest/2026/bcs2607sq

BusinessGPS publishes capability.ai.