

# Six questions worth an afternoon

None requires a consultant, a platform or a budget line. All can be answered internally — and the answers are more useful than any statistic published this year.

Each question is paired with what the 2026 research found when someone asked it at scale — ten studies, nine organisations, roughly 45,000 respondents between them. If your answer is materially better than the figure beside it, you are in the leading fifth. If you cannot answer at all, that is the finding.

|   |                                                                                                                               |                                                                                                                                                                                              |
|---|-------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | Which stage did we actually accelerate — and was it the constraint?                                                           | +442%<br>median code-review time, while task throughput rose 34% and <b>delivery metrics did not move</b><br>Faros AI · 22,000 developers · corroborated by DORA                             |
| 2 | Where did the saved hours go, and can anyone show it?                                                                         | 30%<br>of leaders can say where the freed time went. <b>Seven in ten cannot</b> — the saving has been assumed, never observed<br>McKinsey · 10,018 senior executives                         |
| 3 | How long does one real unit of work take, from request to delivered, including the waiting?                                   | ≈ 20%<br>get measurable value from AI — <b>19, 20 and 21%</b> , three firms, three different questions, no coordination<br>McKinsey 10,018 · Deloitte 3,235 · ServiceNow 4,500               |
| 4 | What share of the data our AI needs can it actually reach?                                                                    | 79%<br>cannot access all the data their initiatives need; <b>18% have fully governed data</b> — against 85% who call their strategy clearly defined<br>Cloudera · 1,270 IT leaders           |
| 5 | Which of our AI systems would a regulator consider high-risk — and could we show six months of logs and explain one decision? | 3.5 : 1<br><b>74% expect agentic AI within two years; 21%</b> have a mature governance model for it — same survey, inside the EU AI Act window<br>Deloitte · 3,235 respondents, 24 countries |
| 6 | Of the steps AI now performs faster, which ones should exist at all?                                                          | 5%<br>are redesigning work rather than automating what already existed — <b>the only figure here that is a decision, not an outcome</b><br>ServiceNow / ThoughtLab · 4,500 executives        |

If these came back uncomfortable, that is the normal result — it is what the data predicts. The organisations in the leading fifth are not the ones with better answers to question one. They are the ones who asked question six first.

The full reasoning, all nine categories and every figure with its source: [unlockedconsulting.ai](https://unlockedconsulting.ai) — Enterprise AI in 2026, read across ten studies.