Consulting’s self-inflicted wound: firing the experts who make AI work

Marceau Partners



© 2026 Marceau Partners Corporation. Share with credit and a link to the original source.

GenAI is reshaping the consulting value proposition in just two years. Models can draft analyses in minutes, summarize complex documents, and even propose strategies. But the more machines “produce,” the more vital human judgment becomes.

Senior experts and partners (people who have seen cycles, failures, and edge-cases) are the only ones truly positioned to review, challenge, and validate both inputs and outputs.

Regulation, client risk tolerance, and the hard data all point to the same conclusion: **this is a human-in-the-centre era, not a human-optional one.**

Why senior judgment now matters more; not less

- **Regulation mandates oversight.** The EU AI Act (entered into force on August 1, 2024; fully applicable by August 2, 2026, with some earlier obligations) explicitly requires *effective human oversight* for high-risk AI systems. This isn’t a suggestion. It’s compliance-critical design. Firms that strip out senior review are building processes that will fail regulatory muster and client audits.
- **GenAI still hallucinates, badly in complex domains.** Recent legal-domain research finds general-purpose LLMs hallucinate on legal queries **58%–82% of the time. It is a** sobering reminder that fluent text is not the same as truth. Senior reviewers are the circuit-breakers who catch these confident errors before they hit a client deck or a courtroom
- **Productivity upside doesn’t remove accountability.** McKinsey estimates GenAI could add **$2.6–$4.4 trillion** in annual value and lift labor productivity by **0.1–0.6 percentage points** per year through 2040 (*if* organizations redeploy human time wisely). Translation for consulting: use machines to draft faster, then devote senior time to framing the problem, stress-testing assumptions, and making the call.

The commercial irony: firing the very people clients pay to trust

Here’s the paradox: even as oversight becomes indispensable, many firms have tightened partner tracks and reduced senior ranks. This is mostly explained through utilization math or short-term profit protection.

- **Layoffs and “partner-light” models.** Accenture has repeatedly cut roles while reshaping its workforce for AI, including **11,000+** roles in the past months and earlier rounds totalling **~19,000** (about 2.5% of staff) since 2023. Other firms have slowed or cut partner promotions (e.g., McKinsey’s partner class reportedly down ~20% year-on-year), and UK Big Four units have trimmed or reshaped partner pools to protect margins.

© 2026 Marceau Partners Corporation. Share with credit and a link to the original source.

- **Utilization and margin pressure are real, but misleading.** Industry benchmarks show billable utilization has sagged: SPI’s 2025 PS Maturity data pegs average utilization at **~68.9%** with EBITDA margins at a decade-low **~9.8%**; 2024 professional-services benchmarks also flagged drops in pipeline, utilization, and project margins. Responding by cutting senior oversight may lift near-term utilization optics, but it degrades quality control and risk management; the very elements clients pay premiums for.
- **Profit optics often rely on cutting people, not growing trust.** Multiple reports across the Big Four show partner pay and promotions whipsawing with job cuts and cost controls. This protects this year’s P&L, but erodes institutional judgment and mentorship; the assets that make AI outputs credible to boards, regulators, and procurement.

What senior experts uniquely do in an AI-first workflow

1. **Frame the problem correctly.** Most AI errors start with the wrong question or the wrong dataset. Senior partners understand context (industry structure, regulatory landmines, political economy), and they know which signals actually matter. That’s not prompt engineering. This is experience.
2. **Interrogate and validate the inputs.** “Ground truth” data is noisy, partial, and biased. Senior reviewers spot survivorship bias, spurious correlations, and stale reference classes that younger team members (or the model) miss.
3. **Enrich and validate the outputs.** When models hallucinate or confabulate sources, senior judgment decides what holds up in the boardroom, the market, and the press release. In high-risk domains, that’s now a legal expectation; not a stylistic preference.
4. **Translate into action.** Clients don’t buy slideware anymore: they buy outcomes. Seasoned operators weigh feasibility, stakeholder dynamics, and sequencing. That is precisely where AI is weakest and where trust premiums accrue.

Rethinking the business model: from “billable hours” to “assured outcomes”

If firms cling to legacy utilization math, senior expertise will always look “uneconomic.” But the market is making the opposite bet:

- **Clients want transformation with accountability.** Surveys and market commentary show firms that integrate external expertise effectively are materially more likely to hit strategic goals. In practice, that means putting experienced humans at the centre and using GenAI as an amplifier—not a substitute.
- **Utilization targets need modernization.** Averages of ~65–70% may be “normal,” but they’re the wrong North Star for AI-heavy delivery, where automation compresses junior hours and elevates the marginal value of senior review. Chasing 80%+ across all levels creates perverse incentives to cut non-

© 2026 Marceau Partners Corporation. Share with credit and a link to the original source.

billable activities like QA, coaching, and model governance; the exact work regulators and clients now expect.
- **Governance is the new differentiator.** With the EU AI Act's staged application schedule already underway, clients will ask not just "Can you do it?" but "How do you *govern* it?" Firms that can point to named senior reviewers, documented human-oversight controls, and escalation paths will win.

A pragmatic blueprint

- **Institutionalize "Senior Sign-Off."** Make an identifiable senior reviewer accountable for the integrity of every AI-assisted work product in regulated or high-impact domains. Map this to Article 14 human-oversight controls and client audit trails.
- **Rebalance pyramids.** Shrink low-value manual analysis layers (where GenAI excels) and reinvest in senior talent, QA, and data governance.
- **Change what you measure.** Track *decision quality, rework rates, and client risk incidents* alongside utilization and margin. Tie partner compensation to quality and governance, not only revenue.
- **Sell assurance, not slides.** Productize "AI-assured consulting" with transparent oversight workflows, model cards, and issue logs that clients can inspect.

Bottom line

GenAI is not eliminating the need for senior experts; it is **raising the cost of not having them**. The firms that will thrive aren't those that automate judgment away, but those that **codify** it and making human oversight visible, auditable, and central to the value proposition. In a market where models can draft anything, **wisdom** is the scarcest, and most defensible asset.

© 2026 Marceau Partners Corporation. Share with credit and a link to the original source.