

AI in Engineering Consulting

Engineering Consulting


What changes for the engineering consultancy as AI enters its industry

“The firm that adapts to AI by getting more efficient is solving the wrong problem with extraordinary discipline.”

01 What is the price floor of an hour you sell?

02 Who captures the efficiency, you or the client?

03 What is your earnings mix worth on the multiple?



What changes for the engineering consultancy as AI enters its industry

For forty years engineering consulting survived every technology wave by selling more hours. CAD, BIM, digital twins. Each new tool came with a new reason for clients to pay for more time. AI is the first wave that compresses the work without expanding the scope. The pricing mechanism that absorbed every previous wave has nothing to say about it.

For the CEO of a mid-market engineering consultancy, this is not a technology question. It is a revenue question, a profit question, and a valuation question. Revenue, because the hours your firm sells are the asset whose price floor is now under structural pressure. Profit, because the margin headroom available to firms that reposition early is large enough to change the income statement. Valuation, because the same EBITDA earned through outcome-based and recurring revenue is worth a meaningfully higher multiple than EBITDA earned through pure billable hours.

The strategic move is not faster hourly delivery. It is shifting the firm's revenue mix so AI-driven efficiency is captured as margin instead of passed to clients as discounts, and using the freed-up senior capacity to build the recurring advisory layer the leverage model never produced.

This briefing lays out the shift, the three CEO-level questions it raises, the strategic choice in front of the leadership team, and the four levers that compound when used together. It closes with how GRAIL works with engineering consultancy management teams to act on it.

This briefing is about what that move looks like, and how a leadership team gets it done before the next board cycle.

The first technology wave that compresses execution instead of expanding it

Every previous technology in engineering consulting created more work. CAD made drawings cheaper to produce, so clients ordered more of them. BIM made models richer, so projects produced more deliverables. Digital twins extended billable engagement across the full asset lifecycle. Each wave came with a new reason for clients to pay for more time. The pricing model survived four decades of technology change because every new tool gave it a new reason to keep working.

AI does not behave the same way. The client still needs the same bridge. The same compliance check. The same MEP design. The scope of work is unchanged. But the routine production work that absorbed most of the hours can now be done in a fraction of the time. Under hourly billing, that is not a productivity gain. It is a revenue problem. Faster delivery times the same hourly rate equals less revenue at the same scope of work.

The firm that adapts to AI by getting more efficient is solving the wrong problem with extraordinary discipline.

The instinctive response is to absorb AI the way every previous wave was absorbed: roll out the tools, update the templates, train the staff, and let efficiency arrive as a quiet productivity gain. That response is the trap. It lowers the cost of delivery without changing the price floor of the work. It hands the savings to procurement departments that are already estimating projects with their own AI. And it leaves the firm with the same revenue model running in a market where the underlying unit of trade is depreciating in real time.

Revenue. Profit. Valuation.

Under the old model these three questions had three different answers. AI ties them together. Move the revenue mix and you move all three at the same time.

01 · REVENUE

What is the price floor of an hour you sell?

The hours a mid-market engineering firm sells are now under structural pressure from a counterparty that did not exist three years ago: a client procurement function with its own AI benchmark of what comparable scope should cost. The price floor is sliding. The question is not whether to defend it. It is what to sell instead.

02 · PROFIT

Who captures the efficiency, you or the client?

AI removes hours from the same scope of work. Whoever owns the pricing mechanism captures the difference. Under hourly billing, the client takes it as a discount. Under outcome-based pricing, the firm takes it as margin. In a single-digit margin business, that distinction is the difference between operating profit collapsing and operating profit nearly doubling.

03 · VALUATION

What is your earnings mix worth on the multiple?

Public engineering consultancies trade on a multiple anchored to the project-based billable-hour model. The same EBITDA earned through outcome-based delivery and recurring senior advisory is worth meaningfully more on the same multiple. Repositioning the mix is not just a margin story. It is an enterprise value story the board has not yet been asked to consider.

Two columns. The whole industry sorts itself between them.

The pattern is consistent across every function in the firm. Routine production work compresses; human judgment work appreciates in value. The bottom-left of every matrix is where firms currently deploy most of their junior and mid-level staff. AI does not replace the senior engineer who walks the site. It replaces the work that three junior engineers do to support that senior. The leverage model, the recruitment pipeline, and the cost structure all change at the same time.

DEPRECIATING

Standard calculations and code compliance checks

Routine drawing generation and BIM model population

Proposal boilerplate and historical cost benchmarking

Progress reports, schedules, photo logs

Tendering and qualification responses

APPRECIATING

Complex structural judgment under unusual constraints

Reading a site, a subcontractor, a room of decision-makers

Pricing creativity and outcome-based deal structuring

Trust, relationship, the politics inside a client organization

Senior advisory: paid judgment, not paid time

Whichever column the firm trains its people into over the next three years is the firm it becomes. The decision is not whether AI affects the work. It is which side of the line the leadership team deliberately moves toward.

Two paths. Same technology. Opposite outcomes.

The financial logic of the choice in front of the leadership team is unforgiving in either direction. The same AI technology, deployed into the same firm, produces opposite outcomes depending on which pricing mechanism it is married to.

THE DEFAULT PATH

Adopt AI as a productivity tool

Roll out the tools. Train the staff. Update the templates. Let efficiency arrive as a quiet productivity gain absorbed inside the existing model.

Keep selling hours

The unit is unchanged. Each project takes fewer hours. Revenue compresses at the same scope. The firm becomes more efficient and less valuable at the same time.

Hand the savings to procurement

Clients with their own AI benchmarks already know what AI-efficient delivery should cost. They take the difference as a discount. Margins compress toward the cost of delivery.

Defend the existing valuation multiple

The firm stays inside the same earnings basket every other engineering consultancy is in. Multiple stays anchored to the project-based hourly model.

THE REPOSITIONING PATH

Treat AI as a pricing question

Use AI to redesign how the firm charges, not just how it delivers. Capture efficiency as margin, not as discounts.

Shift the revenue mix

Repackage repeatable scopes as fixed-price outcomes. Productize the senior premium as recurring advisory. Open markets the unit economics used to close.

Capture the savings as margin

Margin headroom is large enough to materially change the income statement. Operating profit can nearly double on a 30% mix shift before any growth assumptions.

Earn a new valuation multiple

The same EBITDA earned through outcome-based and recurring advisory revenue is worth more on the multiple. The mix shift compounds the margin shift.

Same technology. Opposite outcomes. The difference is the pricing decision, not the AI decision.

The growth levers that compound when used together

Repositioning is not a single decision. It is four levers, each available to the firm right now, that compound when they are pulled together. Used in isolation, any one produces a modest gain. Used together, they describe a different business by 2028.

1

Package the bread and butter

Find the 95%-repeatable service the firm delivers as custom work today. Decompose it by role and task: what is automatable, what benefits from augmentation, what stays purely human. Rebuild it as a fixed-price packaged offer with agentic delivery. Take it to market proactively as cheaper, faster, and higher quality than the industry baseline.

2

Segment the pricing

For the hourly holdouts, a proactive offer: deliver in fewer hours at higher per-hour rates. Net-zero on the specific engagement. But the price floor rises across all hourly work, which is the setup for the rest of the decade.

3

Productize the human premium

As AI handles the production layer, the human layer becomes scarce and valuable. Sell it explicitly as a distinct premium offer: a senior advisory tier on top of augmented delivery, priced as a retainer. Recurring revenue, the type engineering consultancies do not have today and the type the valuation multiple rewards most.

4

Open markets you closed before

The business the firm turns away today because unit economics do not work. Small clients where overhead ate the fee. High-volume low-margin work that diluted averages. Adjacent geographies too small to justify setup cost. Automation changes the math on all of them. The question: what did you say no to in 2025 that becomes viable in 2027?

Business first. Technology second.

GRAIL is an AI-native management consultancy. We help European mid-market leadership teams figure out the one or two things that fundamentally change about their business as AI enters their industry, and then build the capability to act on it before competitors do. We are not an engineering software vendor: Autodesk and Bentley sell better design tools, and they always will. We are not a strategy house: the big consulting firms cannot give you outcome-based AI advice without confronting their own hourly model. We sit in the gap between the two: business-model strategy delivered as working AI systems your leadership team uses in their own work.

The market offers two things: vendor lock-in dressed as transformation, and consulting reports dressed as strategy. GRAIL delivers a third option: working AI capability that you own.

WHAT YOU GET

Senior advisory direct to the CEO and GMT

Working AI capability you own from day one

Custom agents built around your business logic

Outcomes in weeks, not quarters

No lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators, not technologists

Business-first strategy, technology second

Context architecture, change management, workflow design

Results before you commit to scale

Each next step earned, never assumed

Three sessions with your management team. No IT involvement. Working tools from week one.

GRAIL starts with the management team because organizational transformation cannot outrun the leadership team's personal understanding of what is possible. Three workshops. Each delivers standalone value. Each earns the right to the next.

WORKSHOP 01

How to think

The strategic foundation. Where AI creates real value in engineering consulting and where it is noise. The augmentation framework. The pricing transformation thesis. Your leadership team leaves with a shared frame for the rest of the work, and with the first working agent built live in the room.

WORKSHOP 02

Your AI team

A half-day activation. Each member of the management team receives a personal suite of five to ten custom agents trained on their workflow, their decisions, their competitive landscape. Board preparation. Tender analysis. Pricing simulation. Working tools from week one, not slides about tools.

WORKSHOP 03

Competitive edge

Refinement from real usage. Agents upgraded based on how the team actually used them. The pricing transformation roadmap. The cross-country AI infrastructure plan. The recommendation for what comes next, earned by what just shipped.

Three weeks. No IT involvement. No board approval needed to start. The sharpest strategic conversation your leadership team will have this year, backed by working tools they are already using.

What to ask in your next leadership meeting

The questions where the room cannot agree on an answer are the ones worth a longer conversation. Ask these honestly. The point is not the polished response. The point is whether the answer would survive contact with a major client three months from now.

If a client asks next quarter why they should pay for the same scope at the same price when AI helps you deliver it in significantly fewer hours, what is your answer?

Which 30% of your current revenue could be repackaged as fixed-price, outcome-based delivery in the next twelve months, and which would be the first to break?

What recurring revenue does your firm have today that does not depend on the next project? If the answer is none, what would you sell to create it?

If you shifted your earnings mix from pure project to part-outcome and part-recurring, what does the board's view of your enterprise value look like a year later?

When the production layer stops eating the week, where do the freed-up senior hours go: more client work, internal innovation, proactive business development, or back into the same hourly pool?

If any of these questions land somewhere uncomfortable, that is the conversation worth having.

IN THIS SERIES

There's a version of this briefing for every seat at the table.

GRAIL has produced industry briefings calibrated to the specific reality of each C-suite role.

CEO

▶ YOU ARE READING

Revenue. Margin. Valuation.

Three numbers your board reads. AI is rewriting all three in your industry.

CFO

Capital. Forecast. Story.

What you allocate this year is the story your board reads next year.

COO

Cheaper. Better. Faster.

The triple threat that used to be a tradeoff is now table stakes.

CHRO

Capacity. Capability. Promise.

The biggest talent transition of a career. Compete on capacity, not headcount.

CPTO

Moat. Speed. Allocation.

The platform moat that survives 2028 is being chosen this year.

CCO

Margin. Pace. Position.

The whole GTM playbook is being rewritten in real time, across every commercial team at once.

Visit grail.works for the full series.

GRAIL

An AI-native management consultancy helping European mid-market leadership teams figure out what fundamentally changes about their business as AI enters their industry.

www.grail.works