

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing Underneath the Simulation Licence

Your Aspen renewal now quotes an AI-assistant uplift. A digital-native process-advisory shop just quoted your largest specialty-chemicals client at forty percent of your fee. Your most senior unit-operations engineer retires next year. A briefing for CTOs, CPOs, and Heads of Digital at mid-market European industrial-process consultancies.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market industrial-process technology and product leaders

DATE

April 2026

Monday 9:30, digital-platform review

Every Head of Digital at a mid-sized European process-engineering consultancy has had the same Monday. Aspen AI assistant adoption at sixty-eight percent across the mid-tier. Feasibility-study drafting time down thirty-two percent. Project margin flat against Q4. Your VP Engineering opens with the dashboard. He is mid-sentence when your phone buzzes. Email from the CTO at your largest specialty-chemicals client, a fourteen-year account: "We're inviting three firms to quote the next plant-expansion feasibility study, including a digital-engineering shop we met atACHEMA. Your proposal is welcome by month end." Your most senior unit-operations engineer mentioned his retirement timeline at the Christmas party. Fourteen months. His reasoning on that client's reactor geometry has never been captured.

You have run digital and engineering technology for this firm for ten years. The pattern used to be simple. License the right simulators. Hold the hazop bar. Ship feasibility studies on time. Meet the board on licence utilisation and billable-hour margin. The business was understandable.

Something changed. Aspen and AVEVA both ship AI on top of their platforms now. What used to take fifty mid-level engineers on a feasibility-study pipeline runs on five seniors and an AI-copilot layer. Ninety-five percent of corporate AI deployments produce no measurable P&L impact, and the gap between licence adoption and project margin is where your seat now lives or loses. A digital-native process-advisory shop at forty percent of your fee just landed a first named quote in your specialty-chemicals book. Your VP Engineering carries correlation-set reasoning that has never been written down, and retirement is fourteen months out.

You are not running one R&D budget. You are running two, and only one is on your scorecard. One funds the Aspen AI assistant seats your engineers already use. The other funds what has to exist by 2028: the reasoning your three most senior unit-operations engineers take with them when they retire, the correlation library no simulator vendor ships, the client-project archive your firm owns the rights to. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist.

This briefing covers what is happening in your sector, what it means for the R&D and digital-platform function you run, and the three questions your team should answer before the next board cycle.

Process advisory just lost the equation that priced the engineering hour

The AI-natives in your category are not experimenting. They are quoting. Small digital-engineering shops built around GPT-5 enterprise, CAPE-Open fluency, and a handful of senior process engineers are pitching specialty-chemicals and refining clients on three-week fixed-fee feasibility studies at forty percent of your rate-card quote. They do not have your fifty-year operator references. They do not need to. Their deliverable lands hazop-defensible enough, fast enough, cheap enough that your client's procurement team is asking your account partner to explain the eight-week timeline in terms their CFO can defend.

The moat your firm ran on for two decades has moved. Mid-tier engineering hours commoditise fast now. A feasibility study that used to require a twelve-person team across three months can be drafted by a three-person team and an AI-copilot layer in three weeks, on modern simulator architecture, priced fixed-fee on an outcome the client can verify. Your engineering-hours edge was never going to hold. What can hold is the layer underneath the simulator: the unit-operations reasoning your seniors carry for specific reactor geometries, the correlation library tuned against twenty-five years of plant data, the client-project archive your firm owns the rights to. Those compound. They take decades to build and they cannot be bought. A digital-native with three people and a CAPE-Open licence cannot ship them because it does not have your project history.

Your feasibility study drafted on Aspen AI assistant is on an eighteen-month productivity curve that every competitor also has. Your captured correlation library against your top ten client reactors is the asset the digital-native cannot replicate without your project archive.

The industry's speed mismatch is the third force. Digital-native process-advisory shops build from scratch with no legacy, no heritage simulator investment, no senior-partner billing model to protect inside their own house. Your firm has all three. You cannot pivot at their fee point on their terms. What you can do is identify what is genuinely defensible in your engineering output, stop defending what is not, and ship the moat they cannot build. The firms already moving are quiet about it. The ones rolling out Aspen AI assistant seats and calling it an AI strategy are the ones the board will replace in 2028.

Your digital-engineering function is feeling it every day

Your build engine. Aspen AI assistant adoption is at sixty-eight percent. Feasibility-study drafting time is down thirty-two percent. Project margin is flat. Your mid-tier engineers are pushing AI-drafted PFDs through review with minimal annotation. The correlation reasoning behind the numbers never gets written down. Twelve months in, your team cannot defend a hazop against a senior challenge because the documentation trail stops at the AI output. The review discipline that worked in 2022 is training the wrong practice on AI-generated volume now.

Your platform and data. Aspen is pushing an AI-assistant uplift on your licence renewal. AVEVA AI Assist is in your mid-tier's inbox. Siemens Gproms is quoting an AI add-on. Every competitor can buy the same seats. What compounds is the layer underneath the licence: the unit-operations reasoning your seniors carry, the correlation library tuned to your top ten client reactors, the client-project archive with hazop history preserved, the data rights renegotiated into the next twenty-four months of master service agreements. A capture-platform built on your project archive is portable across GPT-5, Claude 5, and Gemini 3. The simulator layer commoditises. The capture layer compounds.

Your product thesis. Your board asked for the firm's three-sentence AI thesis last quarter. Your answer was the one from three years ago. Your rate card assumes engineering-hours will price like engineering-hours did in 2020. They will not. Digital-native competitors are quoting fixed-fee feasibility studies against outcomes the client can verify. You quote hours against scope. That is a capital-allocation question about where the moat gets built, not a pricing one. Your CFO and commercial leader need a joint document with you before the digital-native forces the conversation on their terms.

Your R&D bench. You quietly stopped hiring juniors in 2024. The routine simulator work that trained juniors into mid-tier engineers is now running on Aspen AI assistant. No routine work, no training ground, no mid-tier bench in 2030, no internal senior candidate when the current seniors retire. OpenAI and Anthropic both resumed junior hiring in late 2025 specifically to pair juniors with seniors on judgment-intensive work. The process-consultancy that cuts for this year's utilisation number has a cheaper mid-tier in Q4 and an empty senior bench in five years.

Your board-level authority. Thirty percent of European AI-strategy ownership sits with the CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO. You are not on the list. Not because you abandoned the mandate. You have not yet claimed it. Your CEO has a CAIO shortlist. Every month without a visible claim, the thesis drifts further from the seat that can defend it. Your Q1 boundary agreement with the CEO lands before the CAIO shortlist, not after.

The moat that matters in 2028 is not inside the simulation licence

Your sector is repricing from below because digital-native process-advisory shops are quoting fixed-fee feasibility studies at forty percent of your rate card. Your function is compressing because five senior engineers with an AI-copilot layer now produce what twelve did, and the hazop bar that worked for twelve cannot hold the AI-drafted volume five produce. Same force, two altitudes. But there is a third thing happening that most mid-market process-consultancy Heads of Digital have not named yet.

The moat that matters in 2028 is not inside the simulation licence. It is underneath it. Aspen will ship AI on top of its simulator. AVEVA will ship AI on top of its platform. Every competitor buys the same seats. What does not commoditise is the unit-operations reasoning your senior engineers carry, the correlation sets your firm has tuned to specific client reactors over twenty-five years, and the client-project-data estate you own the rights to. A digital-native with three people, GPT-5 enterprise, and CAPE-Open access cannot build that layer because it does not have your project archive. You do. The window closes in roughly eighteen months, which is the time between now and when the digital-native has enough project footprint of its own to start building the same layer, or when your most senior unit-operations engineers start retiring and the capture window closes.

In the late 1980s, process plants ran Programmable Logic Controllers on a point-to-point architecture. Honeywell, Yokogawa, and Foxboro shipped the Distributed Control System. Not a bigger PLC. A different control architecture. Plants that upgraded PLCs piece by piece looked fine for five years. By the mid-1990s they could not integrate with the MES, ERP, and advanced-process-control layer that became the competitive unit. Plants that rebuilt around DCS architecture captured twenty to thirty percent throughput gains and held them for a decade. The engineering leaders who could not make the architecture call lost their operating references to the firms that did.

You can tell which camp a firm is in from a single question. Is the digital-platform budget you defended at the last board meeting a plateau budget, a compounding budget, or one budget pretending to be one when it is two? Most mid-market process consultancies cannot answer. The ones that can have already made the shift.

You have twelve to eighteen months. The Heads of Digital who claim the authorship of this decision in Q1 author the decade. The ones who wait until the CAIO hire arrives inherit a seat that has narrowed to licence utilisation and platform uptime, and a scorecard that stopped mattering the quarter their biggest account opened the quote to a digital-native.

What most firms are doing. What the winners are doing instead.

Every Head of Digital has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH

Roll out Aspen AI assistant, track adoption

Sixty-eight percent seat adoption. Drafting time down. PFDs reviewed lightly. Correlation reasoning never captured. Hazop defensibility erodes quietly.

Run AI budget through the IT line, sign the renewal

CIO owns the roadmap in eighteen months. CAIO arrives in Q3. Your seat narrows to licence utilisation and platform uptime.

Bundle AI-assistant productivity into rate-card hours

Drafting time drops. Fee per project stays. Digital-native quotes fixed-fee against outcomes. Your largest account asks for a benchmark.

Cut the junior layer for the Q1 utilisation win

Cheaper mid-tier this quarter. Empty senior bench by 2031. External senior hire at a premium the firm cannot absorb.

THE WINNERS' PATH

Buy the simulator AI, redesign the review bar

Senior review floor per AI-drafted PFD. Correlation-set documentation mandatory. Plant walk-down checklist against every AI output.

Route digital-platform capital through the M&A hurdle

Two budgets, two hurdle rates, two kill criteria. Plateau delivers margin. Compounding delivers the 2028 capture moat.

Stand one fixed-fee line on protected P&L

Parallel pricing architecture with CFO and commercial. Rate-card holds installed base. Fixed-fee wins net-new. No cannibalisation.

Pair juniors with seniors during the capture workstream

Juniors ride plant walk-downs and hazops for eighteen months. Capture doubles as an apprenticeship. Mid-tier bench in 2029.

Same engineering team. Same twelve months. Opposite outcomes. The difference is whether the Head of Digital authors the capital allocation or signs the licence renewal.

Engineering Throughput. Product Defensibility. Capital Allocation.

Three questions every process-consultancy Head of Digital is tracking. The third is the crux. The first two are how you earn the right to answer it.

01 · ENGINEERING THROUGHPUT

Is our AI-copilot speed shipping hazop-defensible engineering, or drafts that fail plant walk-down?

Aspen AI assistant adoption at sixty-eight percent. Drafting time down thirty-two percent. Project margin flat. The correlation reasoning behind AI-drafted PFDs never gets written down. Buy what the simulator vendors can carry. Own the review discipline that makes the engineering ship past a senior challenge.

02 · PRODUCT DEFENSIBILITY

What does our firm ship that a digital-native at forty percent of our fee cannot copy?

Your engineering-hours edge commoditises every eighteen months as simulator AI adds the next capability. Your correlation library and client-project archive compound. Your unit-operations reasoning sits in three senior engineers' heads, and two retire inside twenty-four months. The moat sits underneath the licence, not above it.

03 · CAPITAL ALLOCATION

Is our digital-platform budget one instrument or two?

Plateau capital buys simulator AI-assistant seats at lower unit cost. Compounding capital builds the 2028 capture moat. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The Head of Digital who walks in with one budget runs the same programme every peer is running. The one who walks in with two, each defended separately, authors the decade.

What the firms winning this play are actually doing

Four concrete moves any process-consultancy Head of Digital can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture the firm builds on.

1

Build engine: absorb AI-drafted volume without the hazop bar slipping

Instrument review-time-per-PFD as a floor, not just drafting cycle time. AI-drafted PFDs above complexity threshold require senior review and documented correlation-set reasoning before release. Plant walk-down checklists mandatory on every AI-drafted hazop input. Evaluation harness equivalent: a sample of AI outputs audited against senior judgment weekly, no audit trail, no ship. Rebalance the senior-to-mid-tier ratio around the composite rather than preserving the 2022 team shape. The dashboard is no longer drafting time. It is review depth against AI-generated volume.

2

Platform and data: build the moat underneath the simulator layer

Three concrete programmes that stack over eighteen months. Structured capture of senior unit-operations reasoning, drawn from active project work rather than retroactive interviews. Proprietary correlation library hardened into a queryable firm asset against your top ten client reactor types. Data-rights audit plus renegotiation across top client master service agreements, tied to the next twenty-four months of renewals, so your firm owns the project-data estate it generates. A capture-platform portable across GPT-5, Claude 5, and Gemini 3. The simulator layer commoditises. This layer compounds.

3

Product thesis: stand one fixed-fee line on protected P&L

Do not migrate the whole book. Keep rate-card hours for the installed base that does not want to switch. Ship one fixed-fee or outcomes-priced feasibility-study line on a separate P&L for net-new and for clients explicitly asking. Write the cannibalisation trigger and kill criteria in Q1: if fixed-fee revenue on an account exceeds rate-card revenue, trigger the migration conversation. The joint Pricing Architecture Document with your CFO and commercial leader is signed before the first fixed-fee deliverable closes.

4

R&D bench: rebuild the junior pathway around senior and agent pairing

Your 2024 decision to stop hiring juniors saves Q1. It empties the 2031 senior bench. Reverse it. Juniors in 2026 do not start on routine simulator runs. The AI-copilot layer absorbed that work. They start on plant walk-downs, hazop audits, and paired work with a named senior during the capture workstream. The career curve compresses because the senior's judgment is now legible in the capture loop and the junior learns it on actual engineering, not template exercises. By 2029 you have a mid-tier bench the competition cannot hire away because the training surface did not exist anywhere else.

The only firm that joins engineering authorship to working capture

The process-consultancy Head of Digital has four places to go. None can do what we do.

The big tech-transformation practices ship a standardised AI-governance framework scaled across a sector book. Their cost structure requires reusability. Your seat is idiosyncratic. A reusable framework on an idiosyncratic seat produces governance theatre you cannot defend to the board.

The simulator-vendor professional-services arms (Aspen, AVEVA, Hexagon pro-serve) price the engagement as enablement for licence expansion. Their incentive structure requires more seats. Telling you to spend freed capacity on a firm-owned capture moat instead of expanded seats is structurally impossible for them to say.

The simulator AI add-on vendors (Aspen AI assistant, AVEVA AI Assist, Bentley AssetWise AI) ship drafting and documentation productivity on a per-seat quote. Their product is the productivity gain. Telling you the gain is worth nothing if you take it to margin rather than to capture is not a conversation they are built to have.

An in-house or fractional CAIO splits the seat. Strategy CAIO is revenue-measured; Platform CAIO is adoption-measured. One person cannot hold both. Two-thirds of newly-appointed CAIOs spend year one compiling inventories instead of setting direction.

The four options you have already looked at. None can run the capital-allocation conversation the seat requires. GRAIL is the fifth.

WHAT YOU GET

Senior advisory that speaks hazop-review, not transformation-deck

Working capture infrastructure your team uses the same week

Custom agents built on your project archive, your correlation sets, your clients

The start of a capture-and-data moat your firm owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried CPO-adjacent seats

Business first, technology second

Three weeks from first meeting to working capture infrastructure

Your team keeps the infrastructure, you own the IP

Every next step earned by what just shipped, never assumed

Three sessions with your digital leadership. Working capture from week one.

We start with your digital and engineering leadership team because the change cannot outrun the 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

Walk through what is happening in your sector and inside your digital-platform function. Build the first piece of infrastructure in the room: a capital-allocation agent that takes your real simulation-licence plus AI-copilot plus capture-platform budget, splits it two ways, and runs the alternatives analysis against your named digital-native entrants. Your leadership team leaves with a shared view and the first working artefact in hand.

WORKSHOP 02

Your capture infrastructure

A half-day build. Your VP Engineering, Head of Digital, Head of Platform, and Head of Data each leave with five to ten pieces of working infrastructure tuned to their function. Hazop-review audit agent. Correlation-set capture scaffolding. Data-rights renegotiation playbook. Capture-platform prototype against your top client project archive. Infrastructure your team is using the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined capture infrastructure based on real use. The board document that splits the digital-platform budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the CEO on who authors the AI thesis. The joint Pricing Architecture Document with your CFO and commercial leader. The eighteen-month capture build plan.

Three weeks. No IT project. No board approval needed to start. The capture infrastructure your digital-platform function is already running on when we finish.

For your next digital-platform review

The questions where your digital leadership cannot agree on an answer are the ones worth an hour on the agenda. Ask these honestly.

Is your digital-platform budget today one instrument or two? If one, what does the board see that splits plateau from compounding, and what is the kill criterion on each?

Name the three digital-native process-advisory shops quoting in your specialty-chemicals and refining books. What fee model are they running, at what percent of your rate card? If you cannot name them, that is the first thing the briefing changes.

If your three most senior unit-operations engineers retired inside twenty-four months, how many months would it take to reconstruct the correlation-set reasoning behind your top ten client projects? When does capture start?

Aspen AI assistant adoption is at sixty-eight percent. Drafting time is down thirty-two percent. Project margin is flat. Where did the freed hours go, and is your hazop bar still the one you want the deliverable trained on?

Heidrick's February 2026 data: thirty percent CIO, twenty-three percent CDAO, twenty-one percent CAIO, one percent COO, zero percent CTO-CPO on the AI-strategy-ownership list. Is your Q1 boundary agreement with the CEO written, or is it waiting until after the CAIO shortlist lands?

If any of these produces a disagreement in your digital leadership, that is the conversation worth having before the board asks.

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

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

► YOU ARE READING

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