



AN INDUSTRY BRIEFING FROM GRAIL

What's Changing Underneath Your Validated Stack

Your digital R&D budget is two instruments pretending to be one. Your platform sits above a model layer that commoditises every eighteen months. Your scorecard measures system uptime, validation status, and submission turnaround, and none of those is the 2026 question. A briefing for CTOs, CPOs, and Heads of Digital R&D at mid-market European life sciences firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market life sciences technology and platform leaders

DATE

April 2026

Monday 9:30, R&D platform review

Every life-sciences CTO at a mid-sized European firm has had the same Monday. Platform uptime green. Last quarter's Part 11 audit observations closed. The IND filing for your lead programme going to the agency Friday. Your Head of Bioinformatics opens the review with the dashboard. She is mid-sentence when your phone buzzes. It is your principal bioinformatician asking for coffee. She had a recruiter call from Genesis Therapeutics over the weekend. Your CSO forwards a press release: Recursion filed an IND in your core therapeutic area using AI-assisted target identification. Your CFO pinged this morning: "Board wants a three-year digital R&D envelope ahead of July, worth discussing." Your CEO has a Chief Digital Officer job description on her desk from a board member who thinks the firm needs one.

You have run this platform for a decade. The pattern used to be simple. Keep the validated stack running. Pass the audits. Ship the submission on time. Meet the board on system uptime and validation status. The business was understandable.

Something changed. What used to take twenty computational chemists, clinical data managers, biostatisticians, regulatory writers, or CMC modellers now runs with agent augmentation on a team of five plus the senior judgment layer. That compression arrived faster than the discipline to absorb it under GxP. Ninety-five percent of corporate AI deployments produce no measurable P&L impact, and in a regulated environment the gap between what the tools can do and what your validated systems absorb is where your seat now lives or loses. An AI-native drug-discovery platform just filed an IND in one of your therapeutic areas. Your principal scientist carries translational reasoning that has never been written down, and two AI-native recruiters called her this month.

You are not running one digital-R&D platform. You are running two, and only one of them is on your validation master plan. One runs today's regulated systems. The other has to exist by 2028: the longitudinal clinical corpus queryable across your historical trial book, the regulatory-submission platform that drafts modules your writers can defend, the CMC data layer that turns a decade of process development into a training corpus. 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 digital R&D platform you run, and the three questions your leadership team should answer before the next board cycle.

Life sciences just lost the assumption that priced the platform

The AI-natives in your sector are not running demos. They are filing. Recursion took preclinical candidates into IND filings in roughly eighteen months in oncology. Insilico Medicine moved a candidate from target identification to Phase I in about thirty months. Iktos, Genesis Therapeutics, and Absci each hold a slice of what used to be exclusively incumbent territory on the discovery side. Saama, Unlearn, and Paradigm4 compress the clinical-data and biostatistics cycle from quarters to weeks. The regulators read the filings. Your therapeutic area's submission queue is repricing around a timeline the 2018 validated stack was never built for.

The moat your platform ran on for a decade has moved. The validated-systems layer still matters, and the audit still has to pass. What used to be defensible in itself, the LIMS-ELN-CDMS-eCTD architecture that took you eight years to validate, now has an AI-native peer that did not have to carry eight years of change control. A capability that used to take a thirty-person informatics team three years to deploy can be rebuilt by a twenty-person AI-native in twelve to eighteen months, on a modern architecture, with validation controls authored native to the platform rather than retrofitted. Your validated-stack edge was never going to hold on its own. What can hold is the layer underneath the model: the longitudinal clinical corpus, the regulatory-interaction history, the CMC development corpus, the data rights the AI-native cannot renegotiate without your customer and CRO relationships.

The feature your team ships on GPT-5 is on an eighteen-month obsolescence curve. The longitudinal clinical corpus you accumulated over ten years is the asset a 2024 AI-native cannot reconstruct at any price.

The speed mismatch is the third force. AI-natives build from scratch with no legacy validation debt, no eighteen-month system change-control cycles, no board audit committee reading every platform change. Your firm has all three. You cannot pivot at their speed on their terms. What you can do is identify what is genuinely defensible underneath your validated stack, stop defending what is not, and ship the moat they cannot build. The firms that are already moving are quiet about it. The ones that are rolling out AI pilots in clinical operations and calling that an AI strategy are the ones the board replaces in 2028.

Your digital R&D platform is feeling it every day

Your build engine. Your bioinformatics, biostatistics, regulatory-writing, and CMC-modeling teams are shipping faster. AI assistance is compressing literature review, primary screening analysis, statistical analysis plan drafting, module authoring, and stability modeling. The dashboard looks good. Twelve months in, your Part 11 audit trail has agent-generated entries that were never mapped to a validation protocol, and your next FDA inspection will ask. The review discipline that worked for human-only workflows is training the wrong practice on agent-augmented volume now.

Your platform and data. Your LIMS was validated in 2016. Your eCTD pipeline was commissioned in 2018. You cannot replatform the validated stack in eighteen months while defending against a Recursion IND in your therapeutic area. You can buy the emerging regulated-AI substrate from vendors where validated options exist and redirect the compounding capital into the layer vendors cannot sell: the longitudinal clinical corpus queryable across your historical trial book, data rights renegotiated across your next eighteen to twenty-four months of CRO and CDMO contract renewals, integration depth published for your top ten regulated systems, a regulatory interaction history structured across rapporteurs and competent authorities. The model layer commoditises. The data layer compounds.

Your product thesis. Your board asked for the digital-R&D three-year roadmap last quarter. Your answer was an extension of the 2023 plan with AI inserted as a horizontal capability. The AI-native in your therapeutic area is running a pipeline, at a fraction of your per-asset unit economics, on a platform validated for its use case from day one. That is a capital-allocation question about where the moat gets built, not a modernisation question. Your CSO and Head of Regulatory need a joint document with you before the next pipeline review, not after.

Your R&D bench. Your principal bioinformaticians, regulatory writers, and CMC authors carry knowledge accumulated over years. Biotech layoffs reached forty-two thousand seven hundred in 2025 and the shortage in regulatory affairs, bioinformatics, and clinical development worsened. AI-native recruiters are calling the exact seniors your firm depends on. The routine work that trained juniors into seniors is running on agents. No routine work, no training ground, no senior bench in 2030, no internal successor for the principal-scientist chair three years from now. The retention conversation with the top five is a Q1 conversation, not a Q4 conversation.

Your board-level authority. Thirty percent of European AI-strategy ownership sits with the CIO, twenty-three percent with the Chief Data Officer or Chief Analytics Officer, twenty-one percent with the Chief AI Officer. The life-sciences CTO or CPO is not on the list. Not because technology leaders abandoned the mandate. Most have not yet claimed it. Your CEO has a CDO job description on her desk. Every month without a visible claim, the thesis drifts further from the seat that can defend it. Your Q1 boundary agreement with the CEO and CSO lands before the CDO search, not after.

The moat that matters in 2028 is not above the model layer

Your sector is repricing from below because AI-native discovery platforms are filing INDs and AI-native clinical-trial platforms are running pivotal studies at a fraction of your per-asset unit economics. Your platform is compressing because five specialists now produce what twenty did, and the validated review bar that worked for twenty cannot hold the volume five produce with agents. Same force, two altitudes. But there is a third thing happening that most mid-market life-sciences CTOs have not named yet.

The moat that matters in 2028 is not above the model layer. It is underneath it, and in life sciences the layer underneath is data under GxP discipline. The feature your team ships this quarter on GPT-5 is on an eighteen-month obsolescence curve. The next foundation model rolls in 2027 and the feature commoditises. What does not commoditise: the longitudinal clinical corpus a ten-year-old firm has that a 2024 AI-native cannot reconstruct, the regulatory interaction history that predicts what the FDA reviewer will object to on your next submission, and the CMC process knowledge an AI-native cannot match without a decade of your scale-up data. A Recursion or Insilico cannot build this layer without a customer base or clinical-history corpus of comparable depth. The window closes when those firms either partner their way into a corpus, acquire a firm that has one, or wait long enough to generate their own.

In 2014, a small software company set out to build a structured oncology data layer across community oncology practices. The EHR incumbents had wired the hospitals. Conventional wisdom said the infrastructure fight was over. Four years later Roche paid one point nine billion dollars for that company. Not for the software. For the longitudinal data across three million cancer patients, captured in a form that matched how researchers and regulators needed to query outcomes. The EHR platforms kept working. The data corpus underneath them became the asset the industry priced.

You can tell which camp a firm is in from a single question. Is the digital R&D 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?

You have twelve to eighteen months. The CTOs who claim authorship of this decision in Q1 write the next decade. The ones who wait until the Chief Digital Officer hire arrives inherit a seat narrowed to platform uptime and validation status.

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

Every life-sciences CTO has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH

Pilot AI in clinical ops, track adoption

AI-assistants shipped without validation wrappers. Part 11 audit trail gaps. Foundation rots under a green dashboard.

Run AI budget through IT, sign the requisition

CIO owns the roadmap within eighteen months. CDO arrives in Q3. CTO seat narrows to platform uptime.

Defer the data-rights renegotiation to the next renewal cycle

CRO and CDMO contracts renew on the old terms. Your corpus stays fragmented. The AI-native partners with a competing CRO in Q3.

Cut the junior layer for the Q1 headcount win

Cheaper function this quarter. Empty principal-scientist bench by 2031. External senior hire at a premium the firm cannot absorb.

THE WINNERS' PATH

Buy the regulated-AI substrate, redesign the review bar

Validation wrappers on every agent-in-workflow. Part 11 controls on audit trail. Eval harness as first-class infrastructure.

Route digital-R&D capital through the pipeline hurdle

Two budgets, two hurdle rates, two kill criteria. Plateau keeps the audit clean. Compounding builds the 2028 corpus moat.

Renegotiate data rights on the next eighteen-month renewal wave

Structured corpus rights codified across CRO and CDMO boundaries. Integration depth published across LIMS, ELN, CDMS, eCTD.

Pair juniors with seniors on compounding work

Principal bioinformaticians, regulatory writers, and CMC authors pair with juniors on judgment-intensive work from day one.

Same platform team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the capital allocation or signs the quality-systems requisition.

Build Velocity. Product Defensibility. R&D Capital Allocation.

Three questions every life-sciences CTO is tracking. The third is the crux. The first two are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our platform shipping output that passes the next audit, or speed that fails inspection?

Bioinformatics and regulatory-writing teams are faster. Audit trails have agent-generated entries nobody mapped to a validation protocol. An AI-assistant that compresses literature review tenfold is a Part 11 finding waiting for inspection. Buy the regulated-AI substrate. Own the discipline that makes it pass.

02 · PRODUCT DEFENSIBILITY

What does our platform do that a Recursion at a fraction of our unit economics cannot copy?

The feature edge commoditises every eighteen months when the next foundation model rolls. The longitudinal clinical corpus, the regulatory interaction history, and the CMC process knowledge compound. Your principal bioinformatician carries reasoning never written down, and two recruiter calls landed this month.

03 · R&D CAPITAL ALLOCATION

Is our digital-R&D budget one instrument or two?

Plateau capital runs the validated stack at lower unit cost. Compounding capital builds the 2028 corpus moat. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The CTO who walks in with two, each defended separately, writes the decade.

What the firms winning this play are actually doing

Four concrete moves any life-sciences CTO can start this quarter. Each addresses one sub-function. Pulled together, they become one architecture the firm builds on.

1 **Build engine: absorb agent-generated velocity without a validation observation**

Instrument review-time-per-output on AI-assisted work as a floor, not just turnaround. Validation wrappers mandatory on every agent-in-the-workflow pattern. Part 11 audit-trail controls on model-assisted output, with every agent action mapped to a named validation protocol. Eval-harness coverage as a first-class metric on outputs destined for regulatory submission. Rebalance the senior-to-junior ratio around the composite operator pattern. The dashboard is no longer turnaround time. It is audit-readiness against output volume.

2 **Platform and data: build the corpus underneath the model layer**

Three programmes that stack over eighteen months. Longitudinal clinical corpus engineered across the historical trial book, queryable in a form that matches how your regulatory writers, translational scientists, and biostatisticians actually work. Data rights renegotiated across your next eighteen to twenty-four months of CRO and CDMO renewals. Integration depth published across LIMS, ELN, CDMS, eCTD, and CTMS, maintained as structured metadata your regulatory and BD teams can walk into a rapporteur meeting or a partnership conversation with.

3 **Product thesis: stand AI-assisted submission on one protected programme**

Do not retrofit the whole stack. Pick one upcoming submission where AI-assisted drafting on a validated substrate runs in parallel to the traditional team. Publish the validation package. Measure cycle time from protocol lock to submission, and the reviewer-comment pattern. Establish the shape the next ten programmes follow. The joint document with your CSO, Head of Regulatory, and CMC lead is signed before the first AI-assisted module ships, not after a competitor's AI-assisted IND forces the conversation on worse terms.

4 **R&D bench: keep the seniors, rebuild the junior pathway around composite pairing**

Retention conversation with the top five principal scientists, bioinformaticians, regulatory writers, and CMC authors in Q1, before the next AI-native recruiter call. Juniors in 2026 do not start on routine work. The agent layer absorbed it. They start on audit, exception-handling, and paired composite work with a named senior for their first eighteen months. The career curve compresses because the senior's judgment is now legible in the agent-and-junior review loop. By 2029 you have a mid-tier bench competitors cannot hire away.

The only firm that joins digital-R&D authorship to regulated work

The life-sciences CTO has four places to go. None can do what we do.

The big life-sciences advisory practices ship a standardised AI-governance framework scaled across a sector book. Their cost structure requires reusability. Your seat is idiosyncratic. A reusable framework produces governance theatre you cannot defend to the audit committee.

The hyperscaler healthcare professional services arms price the engagement as enablement for cloud consumption. Their business model requires platform lock-in. Every build-or-buy decision routes through their pricing for three years.

The life-sciences platform vendors have a quote on your desk. Veeva, IQVIA, Certara, Saama, Medidata, Ennov. Their product is a validated-license intervention measured on seats. Telling you to spend freed capacity on your own corpus instead of expanded modules is structurally impossible for them to say.

An in-house or fractional Chief Digital or AI Officer splits the seat. Strategy CDO is pipeline-measured. Platform CDO is uptime-measured. One person cannot hold both. Two-thirds of newly-appointed CAIOs spend year one compiling inventories.

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 GxP, not transformation-deck

Working infrastructure your team uses the same week

Custom agents built on your validated systems and data

The start of a corpus-and-integration moat your firm owns

No vendor lock-in at any phase of the engagement

WHY IT MATTERS

Founded by operators who have carried digital-R&D seats

Business first, technology second

Three weeks from first meeting to working infrastructure

Your team keeps the infrastructure, you own the IP

Every next step earned by what just shipped, never assumed

Three sessions with your R&D platform leadership. Working infrastructure from week one.

We start with your digital R&D leadership 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 R&D function. Build the first piece of infrastructure in the room: a capital-allocation agent that takes your real digital R&D budget, splits it two ways, and runs the alternatives analysis against the AI-native drug-discovery or clinical-trial platforms moving in your therapeutic area. Your leadership team leaves with a shared view and the first working artefact already in hand.

WORKSHOP 02

Your R&D infrastructure

A half-day build. Your Head of Bioinformatics, Head of Clinical Data, Head of Regulatory Operations, and Head of CMC Data Sciences each leave with five to ten pieces of working infrastructure tuned to their function. Validation-wrapper pattern. Submission-drafting agent. Data-rights renegotiation playbook. Integration depth graph prototype. Infrastructure your team uses the same week, not decks about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The board document that splits the digital R&D budget into two instruments with two hurdle rates and two kill criteria. The boundary agreement with the CEO, CSO, and Head of Regulatory on who authors the AI thesis. The joint pipeline-capital document with your CFO and CSO. The eighteen-month corpus-and-integration build plan.

Three weeks. No IT project. No board approval needed to start. The infrastructure your digital R&D function is already running on when we finish.

For your next R&D leadership review

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

Is your digital R&D 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 AI-native in your therapeutic area or service category. What is their filing cadence or deployment timeline, and at what fraction of your per-asset unit economics are they running? If you cannot name the top three, that is the first thing the briefing changes.

If your principal bioinformatician, lead regulatory writer, and senior CMC author left for Recursion or Insilico tomorrow, how many months would it take to reconstruct the translational reasoning and regulatory history behind your core programmes?

Your bioinformatics and regulatory-writing teams are shipping faster with AI assistance. Where did the freed hours go, and is every agent-generated output mapped to a named validation protocol your next Part 11 audit will accept?

Thirty percent of European AI-strategy ownership sits with the CIO, twenty-three percent CDAO, twenty-one percent CAIO. The life-sciences CTO is not on the list. Is your Q1 boundary agreement with the CEO and CSO written, or is it waiting until after the CDO search lands?

If any of these produces a disagreement in your digital R&D 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