

AN INDUSTRY BRIEFING FROM GRAIL

What's Changing in the Discovery Moat You Encode

Your firm holds twenty years of programme-execution intelligence in three or four senior heads. Your AI-native competitor has shiny generative chemistry. Your scorecard is system uptime and validation cycle time, and neither is the 2026 question. A briefing for CTOs and Heads of Digital R&D at mid-market European specialty biotechs, CDMOs, and diagnostics firms.

PREPARED BY

Johan Grönstedt · GRAIL

FOR

European mid-market biotech
technology and digital R&D
leaders

DATE

April 2026

Monday 9:15, R&D tech review

Every CTO at a mid-sized European specialty biotech has had the same Monday. Validation cycle time on the last eCTD release flat against last quarter. Bioinformatics agent adoption at sixty-four percent. Your Head of Bioinformatics opens the review with the dashboard. He is mid-sentence when your phone buzzes. It is your Head of HR. Two senior regulatory architects had recruiter calls from Recursion last week. Your CEO forwarded a partnership pitch from Insilico Medicine on your second indication. Your CFO sent the Shopify R&D-percentage line from 2025 with a one-line note: "What's our equivalent for digital R&D?" Your Head of CSV walked into your office Friday with a question about whether your AI-augmented eCTD authoring needs a new validation strategy before the next IND in Q3.

You have run this function for ten years. The pattern used to be simple. Hire strong bioinformaticians. Validate every system. Hold the GxP review bar. Ship the eCTD package on time. Meet the board on platform uptime and validation cycle time. The R&D function was understandable.

Something changed. What used to take a thirty-person bioinformatics, eCTD authoring, and lab informatics build team now runs on six. That compression arrived faster than the discipline to absorb it inside a 21 CFR Part 11 environment. Ninety-five percent of corporate AI deployments produced no measurable P&L impact, and the gap between what the tools can do and what your validated stack absorbs is where your seat now lives or loses. An AI-native discovery firm just signed a partnership with a top-five pharma in your indication. Your senior regulatory architects carry reasoning behind every IND your firm has filed for fifteen years that has never been written down, and two of them had recruiter calls last month.

You are not running one digital R&D function. You are running two, and only one is on your scorecard. One funds what you already validate and ship. The other funds what has to exist by 2028: the encoded programme-execution intelligence that lives in three senior heads, the structured retrievable graph of fifteen years of regulatory submissions, the validation pathway for AI-augmented work that your CSV team has not built yet. 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 function you run, and the three questions your team should answer before the next board cycle.

Specialty biotech just lost the discovery moat that priced the category

The AI-native discovery firms in your sector are not experimenting. They are building. Recursion runs phenomic screening at a scale Insilico Medicine does not. Insilico ships generative chemistry pipelines Iktos does not. Isomorphic Labs carries DeepMind's structural-prediction lineage. BenchSci and Causaly are rebuilding the literature-and-target-rationale layer. Each one is encoding a discovery-data moat that compounds with every screen, every model run, every published partnership. Their economics let them sign top-five-pharma partnerships at hundreds of millions upfront on platforms still pre-clinical. Your board is reading the same headlines.

The moat your firm ran on for three decades has moved. The discovery science that priced your platform a decade ago is being rebuilt on foundation models. A capability that used to take a 50-person discovery and translational team three years can be assembled by a 30-person AI-native in 12 to 18 months on a modern model stack. Your scientific edge was never going to hold against compute and data scale. What can hold is the layer underneath the model: the twenty years of programme-execution intelligence your firm holds and they do not. Every regulatory-meeting outcome that shaped your IND strategy. Every CMC tech-transfer failure that shaped your manufacturing playbook. Every clinical-development judgment that shaped which indications you abandoned. Those compound. They take fifteen years to generate. An AI-native discovery firm cannot ship them because they were not in the room when the FDA pre-IND went sideways in 2014.

Your discovery science was always going to commoditise. Your accumulated programme-execution intelligence is the asset the AI-native cannot replicate without your fifteen years of submissions, tech-transfers, and regulatory interactions.

The sector is also repricing underneath everyone at the same time. EU HTA Joint Clinical Assessments went live for oncology and ATMPs in 2025 and expand further in 2026. The EU AI Act puts clinical decision support in high-risk scope. FDA AI/ML guidance and the EMA Reflection Paper on AI are setting validation expectations no Big Four pharma deck has fully translated yet. Specialist-fund distributions sit at decade lows, opening a BD window for disciplined buyers and punishing firms without a sharpened digital R&D story. The mid-market sits where these shifts converge. The firms treating them as separate compliance projects rather than one data architecture are walking into a JPM conversation they will not win.

Your digital R&D function is feeling it every day

Your build engine. Foundation-model assist on bioinformatics scripting and first-pass eCTD authoring is at sixty percent adoption. Validation cycle time has not moved. Your senior architects are pushing through agent-generated drafts your IT-Quality-Assurance team has not figured out how to validate. The auditor walks in expecting GAMP 5 traceability the new pipeline does not produce.

Your platform and data. Pfizer, Novartis, AstraZeneca built data-strategy investments from 2008 to 2015 and have a seven-year head start on AI-augmented discovery and submission. You cannot replicate seven years in eighteen months. You can buy the build-side compression at the platform layer and redirect the compounding capital into the layer the vendors cannot sell: a structured retrievable graph of fifteen years of submissions, tech-transfer documentation, and BD packages. Sponsor and CRO data rights renegotiated. Senior-judgment encoding before your architects accept offers. The model layer commoditises. This layer compounds.

Your product thesis. Your CEO forwarded the Insilico partnership pitch this morning. The default answer is yes: outsource the discovery layer, focus on clinical execution and commercialisation. The structural problem is the AI-native captures the moat the partnership transfers. The right answer for most mid-market specialty biotechs is build the encoded-judgment moat first and partner selectively where the moat does not reach.

Your R&D bench. Two senior regulatory architects accepting Recursion offers is the catastrophic event. Each carries fifteen years of judgment about how the FDA reads your indication and how the EMA negotiates your scientific advice. None of it is encoded. Your junior bioinformaticians run on agent-generated work that used to be apprenticeship material. The pathway from junior to senior has silently broken. The firm that cuts the junior layer for this year's number has a cheaper function in Q4 and an empty bench in 2031.

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. Your CEO has a Chief Digital Officer shortlist next to the JPM deck. Every month without a visible digital R&D thesis, the authority drifts further from the seat that can defend it. Your Q1 boundary agreement with the CEO and CSO lands before the CDO shortlist.

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

Your sector is repricing from the discovery layer because AI-native firms are shipping target identification, generative chemistry, and translational modelling at compute-and-data scale your platform was never going to match. Your function is compressing because six bioinformaticians and regulatory architects now produce what thirty did, and the validation discipline that worked for thirty cannot hold the volume six produce with agents. Same force, two altitudes. There is a third thing happening that most mid-market specialty biotech CTOs have not named yet.

The discovery moat that matters in 2028 is not above the model layer. It is in the encoded judgment underneath it. AI-native firms ship better target identification and better generative chemistry every six months. What does not commoditise is the twenty years of your firm's programme-execution intelligence. The regulatory-meeting outcomes that shaped your IND strategy. The CMC tech-transfer failures that shaped your manufacturing playbook. The clinical-development judgment that shaped which indications you abandoned and why. A two-person AI-native cannot reconstruct it because they were not in the room. You can encode it. The window closes in roughly eighteen months, which is the time before two senior regulatory architects accept offers from Recursion or Insilico, after which the encoding cost rises by an order of magnitude.

In 2002, lab-notebook digitisation arrived in pharma R&D. Symyx, IDBS, LabArchives sold to Pfizer, Novartis, AstraZeneca as bench-efficiency tools. Make the scientist faster, replace paper, save filing time. Most firms bought them as efficiency tools. They got modest gains and a software bill. Three firms recognised them as the structured-data substrate that would underpin AI-augmented discovery a decade later. Pfizer, Novartis, and AstraZeneca made data-strategy investments from 2008 to 2015 that produced a seven-year head start when foundation models arrived in 2024. The 2002 scorecard never turned red on the firms that bought ELN as efficiency. The unit of judgment had moved.

You can tell which camp a firm is in from a single question. When your senior regulatory architect retires in 2027, can your next IND submission be drafted by walking the same programme-history graph she would have walked? Most mid-market biotechs cannot. The ones that can have already started encoding.

You have twelve to eighteen months. The CTOs who claim the digital R&D thesis in Q1 author the decade. The ones who wait until the CDO shortlist arrives inherit a seat that has narrowed to platform uptime and validation cycle time, and a scorecard that stopped mattering the quarter their senior bioinformatician took an Insilico offer.

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

Every digital R&D leader has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones.

THE DEFAULT PATH	THE WINNERS' PATH
Roll out AI assist on eCTD and bioinformatics Sixty-percent adoption. Validation cycle time flat. CSV finds traceability gaps the pipeline never produced.	Buy the platform, lead the validation group IT-Quality-Assurance partnership on AI/ML lifecycle. First validated AI-augmented IND becomes a peer reference.
Run the budget through IT, sign the requisition CIO owns the roadmap. CDO arrives in Q3. CTO seat narrows to platform uptime.	Author the thesis before the CDO shortlist Two budgets, two hurdle rates, two kill criteria. Plateau funds the stack. Compounding funds the moat.
Sign the AI-discovery partnership offer Recursion or Insilico takes the discovery layer. The AI-native captures the moat.	Build the encoded moat, partner selectively Encode fifteen years of programme intelligence. Own the layer no AI-native can replicate.
Cut the junior layer for the Q1 headcount win Empty senior bench by 2031. External hire at a premium when your architect retires.	Pair juniors with seniors and agents from day one Apprenticeship compresses from five years to eighteen months. Senior judgment becomes legible.

Same digital R&D team. Same twelve months. Opposite outcomes. The difference is whether the CTO authors the encoded moat or signs the partnership offer.

THREE QUESTIONS

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

Three questions every specialty biotech CTO is tracking. The second is the crux. The other two are how you earn the right to answer it.

01 · BUILD VELOCITY

Is our digital R&D speed shipping validated work, or notebooks that fail audit?

Foundation-model assist on eCTD and bioinformatics is at sixty percent adoption. Validation cycle time flat. The auditor finds traceability gaps the new pipeline never produced. Pfizer, Novartis, and AstraZeneca spent seven years on data-strategy foundations that let agents ship inside GxP. You cannot replicate seven years in eighteen months. Buy what the platform layer can carry. Lead the validation working group.

02 · PRODUCT DEFENSIBILITY

What does our firm encode that an AI-native discovery startup cannot copy?

Recursion has phenomic data Insilico does not. Insilico has generative chemistry Iktos does not. None of them have your firm's twenty years of programme-execution intelligence. Regulatory-meeting outcomes, CMC tech-transfer failures, clinical-development judgment, BD diligence patterns. It sits in three senior heads. Two of them had recruiter calls last month. The window to encode it is eighteen months.

03 · R&D CAPITAL ALLOCATION

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

Plateau capital runs the existing clinical-data platform, eCTD tooling, and lab informatics at lower unit cost. Compounding capital builds the encoded-judgment moat that has to exist by 2028. On one hurdle rate the first wins every quarter. On one scorecard the second does not exist. The CTO who walks into the board with one budget runs the same programme every peer is running.

What the firms winning this play are actually doing

Four concrete moves any specialty biotech 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 work without breaking validation

Lead the IT-Quality-Assurance working group on validation strategy for AI-augmented eCTD authoring, bioinformatics scripting, and clinical-data quality. GAMP 5 traceability gates at every merge: no traceability, no submission. CSV pathway documented for every AI/ML lifecycle including retraining. The first IND filed under the new framework becomes a reference your peers do not have.

2

Platform and data: encode the moat underneath the model layer

Three programmes that stack over eighteen months. Senior-judgment encoding: pair every senior regulatory, clinical, and CMC architect with a structured-interview agent transcribed into a programme-execution-intelligence graph. Programme-data restructuring: convert fifteen years of submissions and tech-transfer documentation into structured retrievable form. Sponsor and CRO data rights renegotiated across renewals.

3

Product thesis: own the encoded moat, partner selectively

The default response to the Recursion or Insilico partnership pitch is yes. That answer transfers the moat to the AI-native. The right answer for most mid-market specialty biotechs is build the encoded-judgment infrastructure first, then partner where the moat does not reach. Write the thesis with your CEO and CSO before the partnership offer sets the frame.

4

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

The 2024 decision to thin the bioinformatics and regulatory-tooling junior layer saves Q1 headcount cost. It empties the 2031 senior bench. Reverse it. Juniors start on audit, exception-handling, and paired composite work with a named senior architect for eighteen months. The career curve compresses from five years to months because the senior's judgment is now legible in the agent-and-junior review loop.

The only firm that builds your encoded moat with you, not for you

The specialty biotech CTO has four places to go. None can do what we do.

The big pharma-consulting practices ship standardised AI-governance frameworks 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 audit committee.

The pharma-tech platform vendors have a quote on your desk: Veeva Vault AI, Benchling AI, MasterControl AI, Saama. They bill on platform penetration. Telling you to redirect freed capacity into your own encoded-judgment moat is structurally impossible for them to say.

The AI drug-discovery partnership offers are landing on your CEO's desk: Recursion, Insilico Medicine, Isomorphic Labs. Their economics require they keep the discovery moat. The partnership transfers your accumulated programme intelligence to them.

An in-house or fractional Chief Digital Officer splits the seat. Strategy CDO is revenue-measured. Platform CDO is uptime-measured. One person cannot hold both. Two-thirds of newly-appointed CDOs spend year one compiling system inventories instead of authoring the thesis.

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

WHAT YOU GET

Senior advisory that speaks regulatory and engineering, not vendor-deck

Working agents your team uses the same week

Custom encoding tools built on your submissions, your senior architects, your stack

The start of an encoded-judgment 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 agents

Your team keeps the infrastructure, you own the IP

Every next step earned by what just shipped, never assumed

Three sessions with your digital R&D leadership. Working agents from week one.

We start with your digital R&D 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 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 your named AI-native discovery competitors and your active partnership offers.

WORKSHOP 02

Your encoded R&D infrastructure

A half-day build. Your VP Engineering, Head of Bioinformatics, Head of Regulatory Tech, and Head of CSV each leave with five to ten pieces of working infrastructure tuned to their function. Senior-judgment encoding agent. Programme-execution-intelligence graph prototype. Validation-pathway scaffolding for AI-augmented work. Sponsor data-rights audit playbook. Working agents your team uses the same week, not decks about agents.

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 CFO on who authors the digital R&D thesis. The joint Pricing Architecture Document with your CFO and BD desk. The eighteen-month encoded-moat build plan.

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

For your next R&D tech 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 discovery firm in your indication. What does your firm encode that they cannot? If you cannot name the encoded asset in three sentences, that is the first thing the briefing changes.

If your senior regulatory architect or two senior bioinformaticians left for Recursion or Insilico tomorrow, how many months would it take to reconstruct the reasoning behind your last three IND submissions?

Foundation-model assist adoption is at sixty percent. Validation cycle time has not moved. What did your CSV team learn from the last AI-augmented submission, and is your IT-Quality-Assurance pathway ready for the next IND?

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 and CSO written, or is it waiting until after the CDO shortlist 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