A SECOND ESSAY ON HEALTHCARE LEADERSHIP · JULY 2026

The Pharmaceutical Edition

AI Is Not a Tool.
It's Infrastructure.

What changes for the pharmaceutical industry when it stops buying AI to move faster - and starts building it to earn trust.

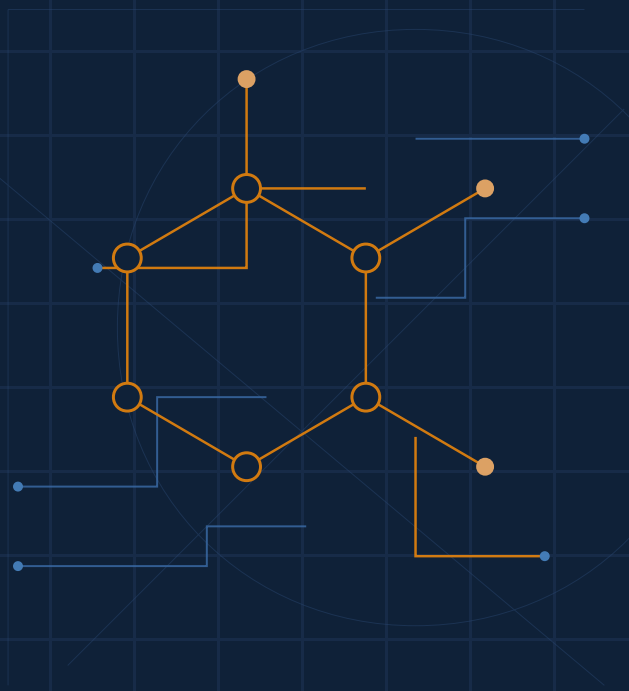

Lucien Engelen
CEO Transform.Health · Healthcare Transformation Architect

FOR CLINICIANS, BOARDS,
REGULATORS & INDUSTRY LEADERS

Primum non nocere.

First, do no harm - the oldest promise medicine makes, and the one AI must now be built to keep.

In my first essay on this subject, I argued that healthcare organisations are still asking the wrong question about artificial intelligence - treating it as a product to procure and deploy, when it has quietly become infrastructure, the substrate on which the whole enterprise runs. That essay was written for hospitals, boards, and policymakers. This one is written for the industry that sits upstream of all of them: the companies that discover, test, manufacture, and steward the medicines themselves.

Pharma has not been idle on AI. Quite the opposite. Few sectors have invested more, or more visibly - in machine learning for target discovery, in generative chemistry, in trial-site selection, in manufacturing quality, in the automation of safety case processing. And yet almost all of that investment has been pointed in one direction: *inward*. It has been spent to make the company faster and cheaper at what it already does. Measured on its own terms, much of it has worked. Measured on the terms that will decide the next decade, it has largely missed the point.

> *The pharmaceutical companies that lead the next decade will not be the ones with the fastest pipelines. They will be the ones that became the trust infrastructure of medicine itself.*

Because the deeper truth is the same one I reached looking at hospitals: this is not, at heart, a technology problem. It is an organisational and a relational one. Pharma's defining challenge has never been the speed of discovery. It is the state of its relationships - with the regulators who govern its medicines, the doctors who prescribe them, the nurses who administer them, and the patients who take them. On every one of those relationships, trust is thin and getting thinner. Deploying AI as a private efficiency engine does nothing to repair that. Done carelessly, it makes the black box darker. Done as infrastructure - transparent, continuous, shared - it becomes the single most powerful trust-building instrument the industry has ever held.

The tool trap

When AI is treated as a tool, it lives where tools live: inside a function, owned by a team, justified by a business case, measured on cycle time. Discovery buys a tool. Clinical operations buys another. Manufacturing buys a third. Pharmacovigilance buys a fourth. Each is a genuine improvement, and each is invisible to everyone outside the wall it sits behind. The regulator never sees the reasoning. The physician never sees the model. The patient never knows it exists. Efficiency compounds internally while opacity compounds externally.

This is the same mistake I watched hospitals make when they filed AI under “facilities” or “research” and governed it through an ad-hoc committee. You would not run a building's electricity through the facilities department as a side project; yet that is precisely how most organisations still manage the current that now runs through everything. In pharma the equivalent is filing AI under R&D or digital, and treating each deployment as a discrete procurement rather than as a layer that spans the entire life of a medicine.

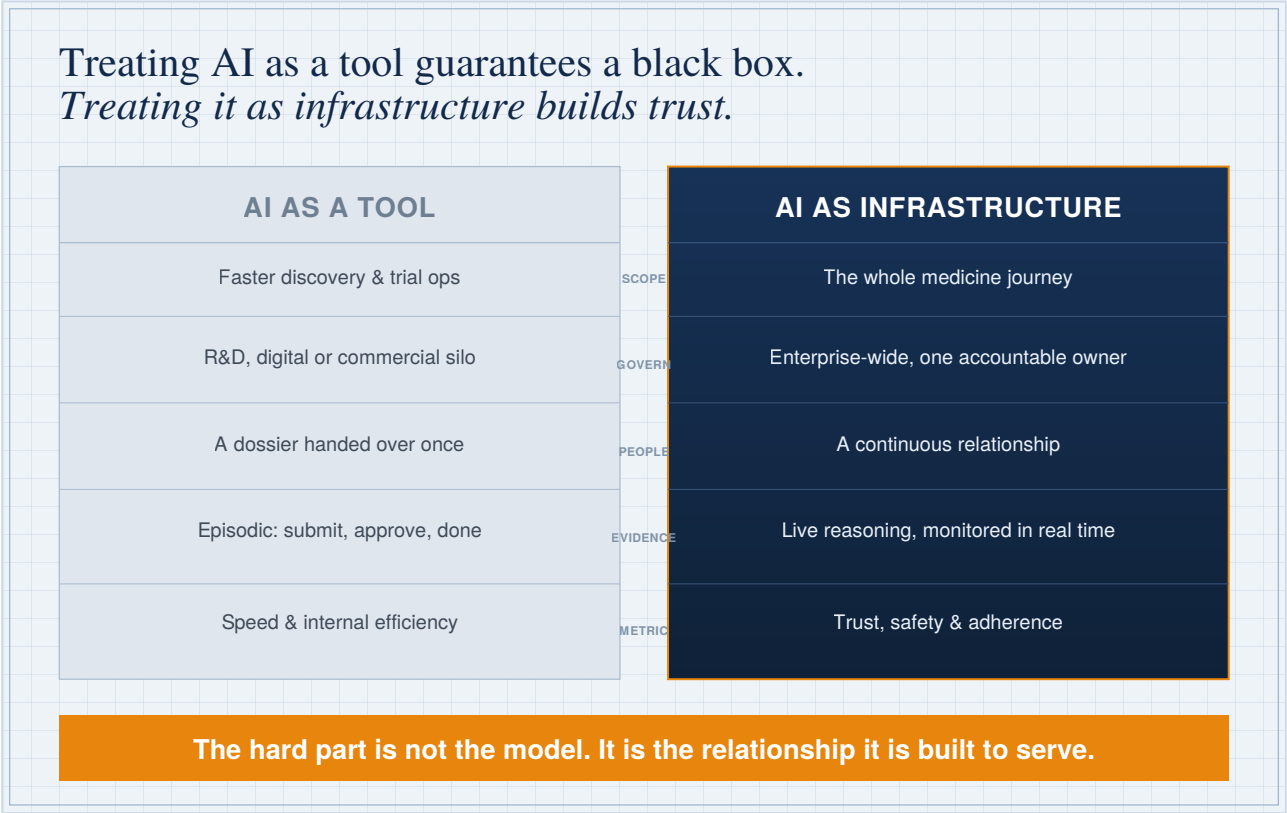

Figure 1. The shift is not a bigger model. It is a change in what AI is *for* - from an internal efficiency play owned by a silo, to a shared, continuously monitored layer that runs the length of the medicine journey.

The distinction matters because infrastructure changes the questions you ask. You stop asking *which use case gives the fastest return* and start asking *how do we build the enterprise around this*. You stop hunting departmental quick wins and start redesigning the operating model. And crucially for pharma, you stop treating each stakeholder interaction as a transaction - a dossier, a detail aid, a package insert, an adverse-event form - and start treating it as a continuous, two-way relationship that the infrastructure keeps alive.

Four frictions that efficiency cannot fix

Pharma enters this moment carrying four frictions at once, and no amount of internal speed resolves any of them. Regulators are demanding more evidence, more often, and are rightly wary of systems they cannot inspect. Doctors are sceptical of recommendations that arrive as outputs without reasoning - and their scepticism is professional judgement, not obstinacy. Patients are more informed than ever and, in the same breath, more exposed than ever to confident misinformation. And nurses, at the very front line, are asked to execute protocols whose logic was decided far from the bedside.

Notice what these have in common. Every one is a failure of *shared understanding*, not of throughput. A faster trial does not make a regulator trust the model that designed it. A cheaper synthesis does not make a physician believe the therapy recommendation. Efficiency is orthogonal to trust; you can have all of one and none of the other. That is why the industry's AI spend has produced impressive internal metrics and almost no movement on the relationships that actually gate its licence to operate.

Beneath the four frictions sit three structural obstacles that pharma must name honestly. The first is the **operating model**: the entire commercial and regulatory machine is built around episodic contact - the submission, the launch, the campaign, the periodic safety update. Infrastructure is continuous by nature, and a continuous

capability sits badly inside an episodic organisation. The second is **governance**: as long as AI is scattered across R&D, digital, quality, and commercial, no single owner is accountable for how it serves anyone outside the company. The third is the **commercial reflex** - the deep habit of treating information as promotion, which is directly at odds with the radical transparency infrastructure demands. These are not reasons to wait. They are the design brief.

What the rest of healthcare already knows

Pharma does not have to invent the infrastructure playbook. It can read it off the parts of healthcare that have already begun to build. The clearest lessons are sitting in hospitals, clinics, and homes right now.

From sepsis models, the lesson of continuous, locally validated monitoring. When a widely deployed early-warning model for sepsis came under scrutiny, the problem was not that prediction is impossible; it was that a model shipped once and monitored rarely drifts away from the population it serves. The health systems that got it right treated the algorithm as infrastructure to be watched continuously, not a feature to be switched on. For pharma, the read-across is exact: post-market safety cannot be a periodic report. It must be a live system that is monitored for drift, bias, and signal in real time - what some are now calling algorithmovigilance.

From radiology, the lesson of explainability at the point of decision. The imaging AI that clinicians actually adopted did not ask them to trust a verdict; it showed the region of interest, the comparable priors, the confidence, the reasoning - and let the radiologist remain the decision-maker. Adoption followed transparency, not accuracy alone. Pharma's therapy recommendations, dosing suggestions, and trial-eligibility calls should arrive the same way: with the reasoning visible, so the prescriber validates rather than defers.

From diabetes care, the lesson of the patient as a trusted sensor. Continuous glucose monitors and closed-loop insulin systems worked because patients understood they were inside a safety system, not under surveillance. They could see how their own data improved their own outcomes. That consent-through-comprehension is exactly what pharma needs from wearables and patient-reported outcomes - monitoring that patients trust because they understand it.

From antimicrobial stewardship, the lesson of decision support embedded in the workflow. Stewardship programmes succeeded by putting guidance where the decision is made, owned by a named team, measured on outcomes, and paired with the humans who carry the accountability. That is the model for medical affairs in an infrastructure era: not a call centre answering queries, but stewardship embedded in the clinical workflow. **And from molecular tumour boards, the lesson of shared decision-making** - AI surfaces the evidence; a room of humans decides. In each case the technology earned its place by making the people around it better, and by being legible to them. That legibility is the whole game.

THE THROUGH-LINE

Every healthcare AI that earned trust did the same two things: it made its reasoning visible to the humans it served, and it was monitored continuously rather than validated once. Neither is a modelling problem. Both are infrastructure choices - and pharma is free to make them today.

Rebuilding the four relationships

If infrastructure is the answer, the test of it is relational: does each stakeholder end up understanding more, trusting more, and participating more than before? This view reframes every relationship the company holds. One current runs through all four - transparent reasoning, continuously monitored - but it lands differently for each.

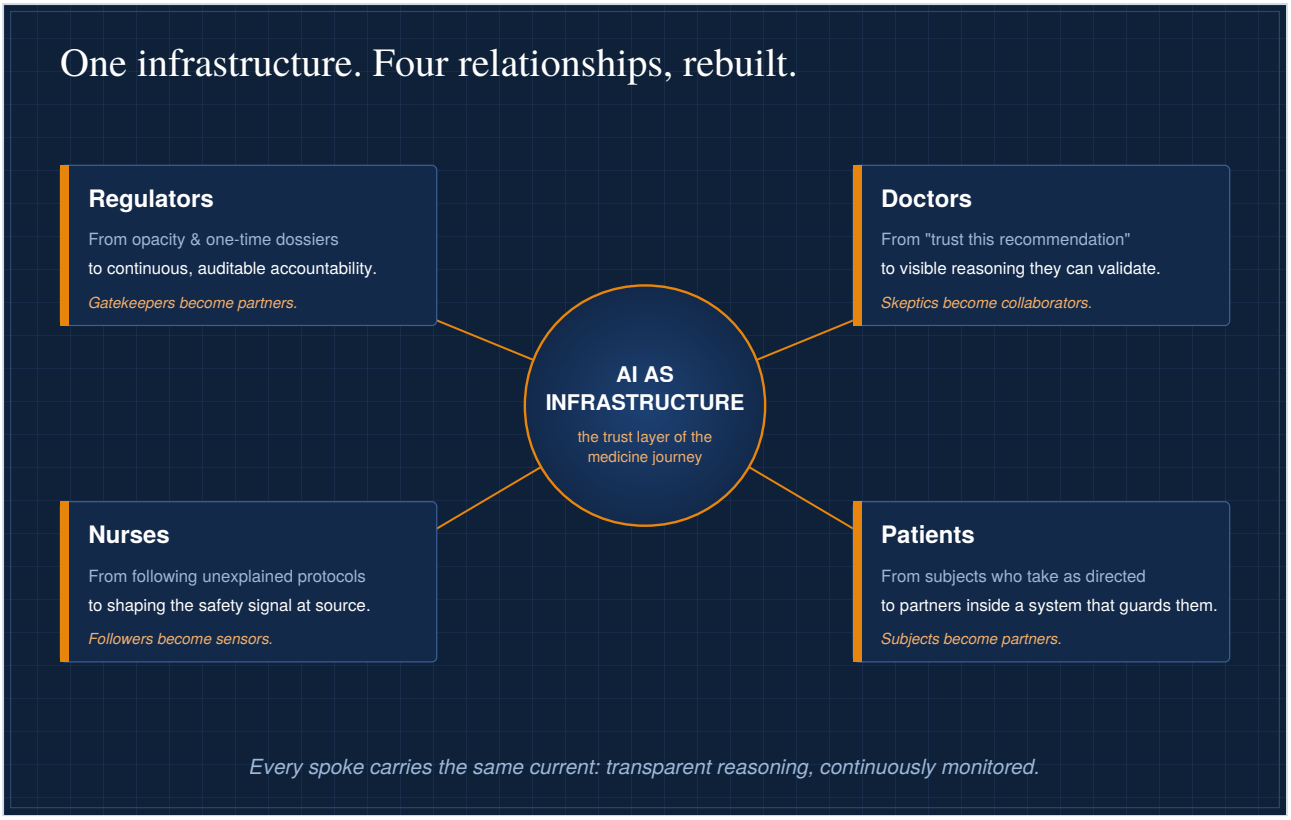

Figure 2. One infrastructure, four relationships rebuilt. The hub is not a product; it is the trust layer of the medicine journey, and every spoke carries the same commitment to visible, monitored reasoning.

Regulators: from opacity to continuous accountability

The traditional model is a one-time handover: run the trial, generate the dataset, submit the dossier, receive the decision. The infrastructure model replaces the handover with a standing dialogue. When an AI proposes a trial cohort, the regulator sees the reasoning, not merely the output. When a manufacturing system flags a quality deviation, they can follow the logical chain. Post-market surveillance moves from periodic to continuous, so that adverse-signal patterns surface weeks or months earlier. And as models evolve, that change is visible - ending the era of “we deployed a new version and did not tell you.” This is not a slower path to approval; it is a more intelligent one, because the regulator understands the nature of what they are approving. The payoff is a change in role: gatekeepers become partners, and credibility deepens well beyond any single decision.

Doctors: from recommendations to shared decisions

In the old model the AI hands the physician a recommendation to accept or override, which isolates them behind a black box. In the infrastructure model the AI amplifies clinical judgement instead of replacing it. The recommendation arrives with its reasoning attached - this therapy, for this patient, given this genetic profile, these comorbidities, this trial evidence, these real-world outcomes we are monitoring. The system keeps learning from practice, and the physician can see it learning. Proactive safety signals reach them for patients matching a given

profile before they could have spotted the pattern alone. Scepticism turns into partnership: the tool stops being something to be trusted on faith and becomes something that visibly makes the clinician better.

Nurses: from protocol-followers to safety sensors

A new drug traditionally arrives with protocols written by people far from the bedside, followed diligently but without full sight of the rationale. Infrastructure gives nurses an active role. When the AI suggests a monitoring interval or a dose adjustment, it explains why - "this patient's profile and response pattern suggest checking levels at 48 hours rather than 72" - language a nurse can act on and pass to the patient. Their observations feed the safety mesh in real time: a reported symptom is correlated, escalated, and recognised as the frontline signal it is. And the system teaches as it works, deepening understanding of drug behaviour and patient response. Followers become sensors - and retention, morale, and safety all improve together.

Patients: from subjects to partners

The old model is one-way: take the drug, follow the instructions, and if something goes wrong, report it and hope. Infrastructure makes patients active participants in their own safety. Information becomes personalised reasoning they can understand - "given your genetics and health profile, expect benefit in four to six weeks, watch for these specific effects, and here is why." Monitoring through wearables or patient-reported outcomes is something they consent to because they grasp how and why, and they can see how their data protects others. When a relevant safety signal appears, they hear it directly and early, not six months later in a letter. Subjects become partners; adherence rises, trust deepens, and word of mouth becomes the most effective marketing the company owns.

From promotion to provision: marketing and information

Nowhere will infrastructure thinking feel more uncomfortable - or prove more valuable - than in how pharma communicates. For a century the commercial model has treated information as promotion: a message to push, a share of voice to win, a leaflet to comply with. Infrastructure inverts this. When the reasoning behind a medicine is transparent, personalised, and continuously updated, the reasoning itself becomes the most persuasive thing the company can offer. The job shifts from promotion to *provision* - providing the legible, trustworthy explanation that a patient or clinician actually needs to make a decision.

Figure 3. The narrative shift for commercial and medical information. Static, one-way, share-of-voice messaging gives way to living, explainable information - and to share of trust.

This is not a softening of the commercial mission; it is a sharper version of it. Consider what changes in practice. The static package insert becomes a living, personalised explanation that adapts to the individual's profile. The undifferentiated campaign gives way to reasoning a specific patient can act on. Medical information stops being a reactive query desk and becomes a proactive stewardship layer. And the metric changes with the method: the old contest for share of voice becomes a contest for *share of trust* - which is both harder to buy and far harder for a competitor to take from you.

There is a regulatory dividend hiding in here, too. The same transparency that earns patient trust is transparency a regulator can read every word of. When your information is grounded in visible, validated reasoning rather than persuasion, the historic tension between commercial communication and compliance narrows. You are no longer trying to say as much as the rules allow; you are trying to explain as clearly as the evidence permits. Those are very different companies to be.

The toolkit that builds the trust layer

Infrastructure is not an attitude; it is a set of things you build. If a pharmaceutical company wanted to make this real, the tools follow directly from the relationships they serve - four stakeholder-facing surfaces sitting on one shared foundation. The sequence matters as much as the list: build the base first.

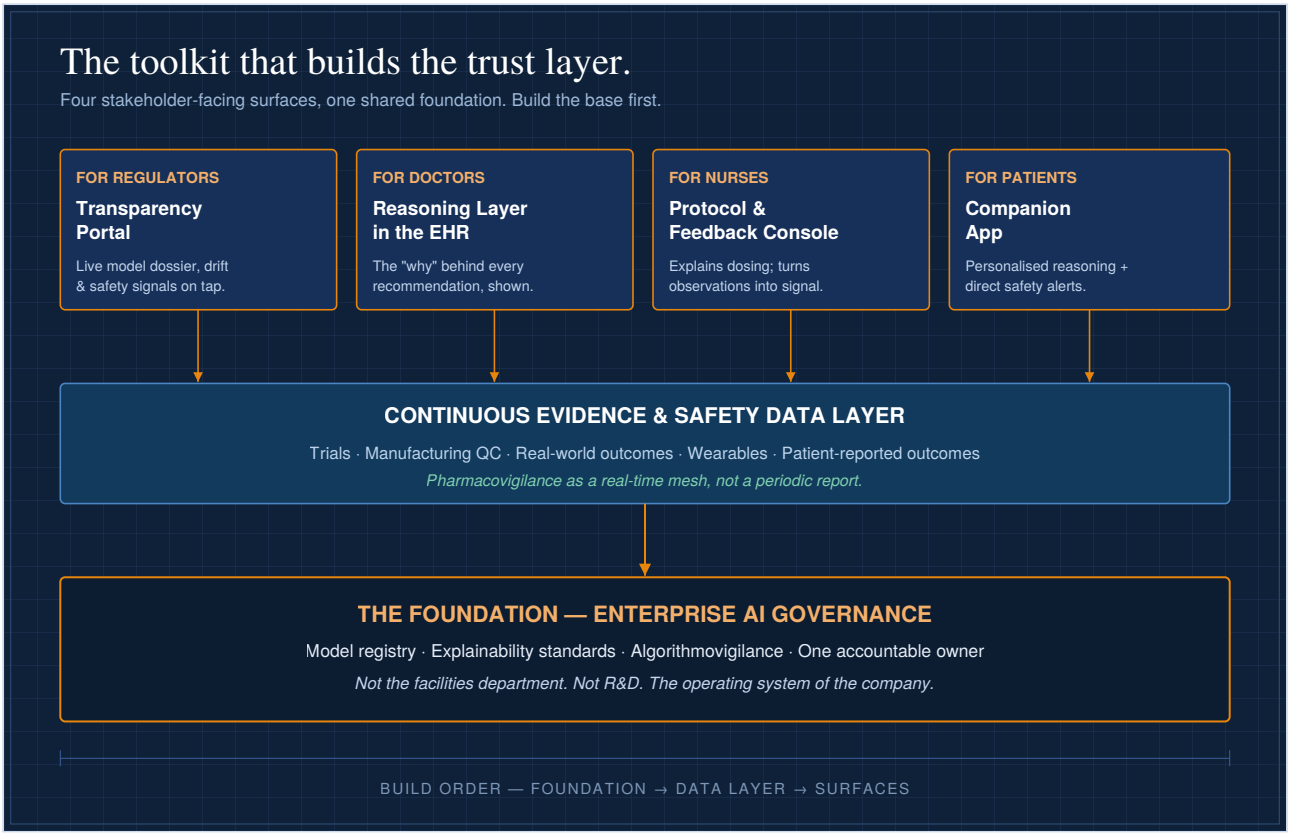

Figure 4. A build order, not a shopping list. Enterprise governance and a continuous evidence layer come first; the stakeholder-facing surfaces are only as trustworthy as the foundation beneath them.

At the base sits **enterprise AI governance**: a model registry, shared explainability standards, algorithmovigilance for drift and bias, and - the single most important design decision - one accountable owner. Not the facilities department, not R&D, but the operating system of the company. Above it runs a **continuous evidence and safety data layer** that unifies trials, manufacturing quality, real-world outcomes, wearables, and patient-reported outcomes into pharmacovigilance as a live mesh rather than a periodic report.

On that foundation, four surfaces. A **regulator transparency portal** - a live model dossier exposing reasoning, versioning, drift, and safety signals on demand. A **clinician reasoning layer** embedded in the electronic record, so the "why" travels with every recommendation. A **nurse protocol-and-feedback console** that explains dosing decisions and turns bedside observations into signal. And a **patient companion** that delivers personalised, understandable reasoning and direct, early safety alerts. Each is worthless without the base beneath it - which is exactly why so many pilots, built surface-first, never earn trust.

Bricks, bytes, and behaviour

My late colleague René Bleeker used to insist that any real transformation rests on three things: bricks, bytes, and behaviour. Most organisations fund one or two and wonder why nothing scales. Pharma has poured resources into the bytes. The bricks - the governance, the operating model, the accountable ownership - and the behaviour - the culture, the literacy, the willingness to be seen - have lagged. Infrastructure demands all three at once.

The evidence from the wider enterprise is consistent and, by now, hard to ignore. Organisations that cross roughly a five-percent-of-budget threshold on AI are markedly more likely to report positive returns - not because of the spend itself, but because crossing it tends to mean they have reorganised around the capability rather than bolting it on. Hybrid governance - enterprise strategy paired with distributed implementation - outperforms both

purely centralised and purely decentralised models. And the widest gap in every sector is not between those who have adopted AI and those who have not; it is between adoption and *scale*. Adoption is a pilot in one function. Scale is the moment governance, workflows, and data practices are integrated across the whole company. Bridging that gap is an act of organisational design, not a technology programme.

> *Healthcare's caution is not a weakness to overcome. Honoured thoughtfully, the regulation, the safety expectations, and the scepticism are the design requirements that produce infrastructure other sectors will envy.*

Six mandates for the board

For the leaders who will actually have to do this, the evidence converges on six commitments. None is about a model. All are about the enterprise that surrounds it.

Figure 5. Six mandates for the pharmaceutical board. Read together, they describe a company that has stopped buying AI to move faster and started building it to earn trust.

Take AI out of the silo and govern it enterprise-wide, with one owner accountable to every stakeholder. Build the foundation before the surfaces - registry and standards first, dashboards second. Make explainability a product, not a disclosure, and ship the reasoning to doctors, patients, and regulators alike. Turn pharmacovigilance into a live mesh of continuous monitoring, wearables, and patient-reported outcomes that produces early signals and direct alerts. Rewire the commercial engine around trust rather than volume - from share of voice to share of trust, from promotion to personalised reasoning. And bring stakeholders into the design room from the start: regulators, clinicians, nurses, and patients shaping the system, not receiving it at launch.

The end of the beginning

We have moved past the phase in which AI in pharma is a novelty or a competitive garnish. It is becoming the fabric of how a medicine is discovered, proven, made, monitored, and explained - and it will define not only which companies succeed, but which remain relevant. The industry that treats it as a departmental efficiency play will be outpaced by the one that treats it as the means of rebuilding trust with everyone who touches its medicines.

As I wrote in *Augmented Health(care): The End of the Beginning*, changes in tooling and infrastructure are not going to take our jobs. Done right, they help us carry the load - and in pharma's case, they let a whole industry do the thing it has always claimed to want and rarely known how to do: earn, and keep, the trust of the people it serves. That trust was always the real product. AI, built as infrastructure, is finally the way to manufacture it at scale.

> *The choice is the same one it always is. Either we build this infrastructure thoughtfully, with patient safety as its foundation - or we let market pressure and competitors build it for us. The choice is ours. And the time is now.*

. . .

A NOTE ON THIS COMPANION ESSAY

This is a second essay, written to extend *AI Is Not a Product. It's Infrastructure.* (Engelen, 2026) to the pharmaceutical industry, and to translate its argument into practical stakeholder guidance. Adoption, scaling, budget-threshold, and hybrid-governance figures are drawn, as in the first essay, from Deloitte's *2026 State of AI in the Enterprise* and related 2025–2026 enterprise and healthcare AI research. Cross-healthcare examples - sepsis early-warning models, explainable imaging, continuous glucose monitoring and closed-loop systems, antimicrobial stewardship, and molecular tumour boards - are used illustratively to show infrastructure principles already at work in care delivery.