Designing Digital Pathology Transformation for the AI Era: Four Challenges, One Systems Problem

David Clark — Clinical Lead for Computational Pathology, East Midlands Pathology Network; Consultant Haematopathologist, Nottingham University Hospitals NHS Trust

DPA Companion Meeting, European Congress on Digital Pathology, Graz, 17 June 2026

The argument in brief

Observation. Digital pathology programmes produce highly variable outcomes despite similar technologies and broadly similar ambitions.

Diagnosis. The technology is only one component of a much larger socio-technical system. Most programmes define the system boundary too narrowly and therefore optimise locally rather than systemically.

Consequence. Problems appear as procurement failures, interoperability failures, governance failures, workforce failures, sustainability failures and AI deployment failures. But these are manifestations of the same underlying design issue.

Solution. Engineering Better Care provides a practical way to understand People, Systems, Design and Risk as interacting perspectives on the same problem.

Call to action. Before specifying technology, define the system you are actually trying to create.

Declaration of interests: I have no commercial or financial links to any non-NHS organisation. I am the founder of www.gembasuite.org a not-for-profit providing free improvement tools for healthcare.

The analysis and opinions in this presentation are my own and not those of the institutions I work within.

I qualified in medicine 45 years ago this summer. When I first walked onto a ward, there were no computers. All the information was on paper. Lab results were pasted into the patient's notes by hand. Now, close to retirement, I work in a completely virtual environment. I still assimilate, correlate and integrate information to make decisions about patient care, but all of it now comes from digital sources: the EPR, the LIMS, digital pathology, flow cytometry plots, molecular genetic data, and for some of us, outputs from AI. Like many of you, all I need is my laptop, a decent screen and wifi and I can do my job anywhere.

However… We've built a digital world, but we pasted it together with a paper-era mindset.

I say that as someone who helped build it. I was clinical lead for digital pathology implementation at Nottingham University Hospitals: two hospitals, one cellular pathology lab, 68,000 cases a year. We started implementation in late 2021. By 2025 we were fully digital: 100% primary reporting, with AI in routine use.

I stepped down when our implementation was completed last year. For the past seven months I've been working on a computational pathology strategy across seven hospital organisations in our regional network. I was asked in because of frustration at the rate of progress — other lab organisations bought similar scanners, similar storage, similar software and worked under similar regulatory constraints. Yet outcomes varied dramatically.

The implication was: if Nottingham can do it, why can't they?

I've reframed the question and removed the implied blame, asking: why does the same investment, and the same technology, produce such different results from one place to the next?

Let's start with the investment strategy.

The NHS didn't make a small bet on digital pathology. It made a system-wide one. Over £50 million since 2019. A real national commitment, to a single technology, across dozens of independent Trusts that together make up the NHS pathology system.

Seen as a whole, this was a single, very large NHS-wide transformational change programme. And its purpose was to help address the long-term mismatch between capacity and demand in cellular pathology services across the NHS.

I want to pause here and look at what happens to large programmes across most domains, from large infrastructure projects to complex IT system delivery.

Bent Flyvbjerg, at Oxford University's Business School, has studied the largest dataset: more than sixteen thousand megaprojects, as he describes them, from multiple domains across the globe. The pattern he observed is so consistent he named it the Iron Law of megaprojects: **over budget, over time, under benefits, over and over again.**

Individual elements may work. But only one in two hundred programmes, taken as a whole, delivers what was promised at the outset. Flyvbjerg describes the systemic causes and biases that contribute to megaproject failure.

There is no guaranteed recipe for success in megaprojects, but his golden rule is: **Think slowly, act fast.**

The NHS made the classic mistake of inverting this: it pushed investment in scanners into an unprepared system.

The parts that worked — for example, Nottingham, and the Leeds-led NPIC network — delivered everything that was asked of them. They transformed the parts of the system they owned. But the whole-system transformation has not occurred.

Diagnosing the underlying causes means addressing **survivorship bias**, which permeates conferences like this.

The archetypal example is the WW2 US air force, which wanted to armour its bombers where returning planes showed most damage. The statistician Abraham Wald stopped them: they were only looking at the planes that came back. The damage that mattered was on the ones that didn't return.

Every success story you'll hear today, mine included, is a plane that came back. Studying the survivors tells us local success is achievable. It doesn't tell us whether we were aiming at the right goal.

For the rest of this presentation I want to focus on what can go wrong in complex system redesign.

Every digital pathology programme, large or small, faces the same four challenges that can derail them. They're not unique to a country, a vendor, or a lab.

First, integration of systems. Getting the scanners, the viewers and the laboratory information system to operate as one system. The most basic level is LIMS integration with the digital pathology system — without this a project cannot pass go — but if we look to the future, much deeper levels of integration.

Second, workflow redesign. Not digitising the old process but rebuilding it: how slides move and are scanned, how cases are allocated, how reports are created. That requires Lean work across the entire end-to-end process, and most NHS labs have little spare management capacity for it, and many lack access to Lean expertise. Rebuilding the reporting process leads straight into the third challenge.

Professional change — people. Pathologists, biomedical scientists and trainees, and importantly the people who manage the service, will need to radically change how they work, day to day. This issue is consistently underestimated. It is the biggest change most pathologists will experience in their entire working lives. Some adapt easily; others struggle. Support and training are essential and this needs to be anchored to a rethinking of the skills that will be required in the future..

Finally, sustaining the investment. Holding the business case together as AI moves from pilot to routine use, once the launch pump-priming funding has gone. The new digital service requires ongoing investment from the local organisation to fund more scanner capacity, archive storage and AI deployment. This is both essential and difficult to achieve in the current financial environment most hospitals are operating within

These barriers to system redesign are interdependent. That's an empirical observation.

So let me share my experience implementing digital pathology, and how a failure to understand the importance of system boundaries has caused problems in delivering system-wide transformation.

First, at the level of the local institution. In Nottingham we did the right things. We implemented an excellent on-premises solution with our partners Indica Labs and Hamamatsu in less than six months. We applied Lean thinking and rebuilt the workflow around it, with voice recognition for pathologists in a virtual reporting environment. We hit our targets. We were asked for 80% digital reporting in three years; we reached 100%. I've now been reporting all my work digitally for over four years.

But at the next level up, the regional pathology network, all the images sat locked behind the Trust firewalls. The key benefit the network needed — importing a difficult case from another

Trust for expert review, and being able to pool reporting capacity — was never designed in. That is a boundary drawn at the edge of the Trust, policed by the firewall and our organisation's cybersecurity policies. It could only be adequately addressed by collective action across the network, and that was not part of the plan.

Now another issue, one level up, at national level. Further contradictions. The Royal College lobbied hard for this national investment in digital pathology. But the training curriculum and FRCPath examination are still not aligned to digital working. So fully digital departments like Nottingham still keep a parallel glass-slide infrastructure running, purely so trainees can practise for an exam that hasn't moved. The NHS sets the national investment strategy, the RCPath owns the training curriculum, and that sits outside the system boundary.

And here is an uncomfortable truth. In 2024 Nottingham won a Royal College Team Achievement Award for its digital transformation. The award is real, and it is deserved for what it measured: we hit every target we were set, and I'm incredibly proud of what our team achieved. Studying success tells you a team hit its targets. It cannot tell you whether those were the right targets. We had real control over local execution, and no control at all over where the system boundary was drawn. That was decided above us, by the funding vehicle at national NHS level, when they allocated funding to local labs with local targets.

NUH successfully optimised the system we were asked to optimise within the system boundaries we were given.

Viewed from the perspective of service transformation digital pathology is not the system. It is one component of the system that delivers an integrated histopathology service. The real system includes people, organisations, workflows, governance, data, infrastructure, funding and risk — a complex system made up of multiple systems.

The challenge is not implementing digital pathology. The challenge is understanding and designing the system digital pathology has to operate within.

The problem was that the system boundary was too small. We hit the targets we were given, but those targets were not sufficient to deliver the wider system change that was intended.

What was needed, and didn't have, was a framework to map the minefields at every level of the system.

The Royal Academy of Engineering's Engineering Better Care framework developed in partnership with the Royal College of Physicians treats the whole system as the unit of analysis. It gives you four perspectives on a system: **People, Systems, Design and Risk.**

Now look back at the four challenges I listed earlier. They are not four problems waiting for four separate solutions. They are one system redesign, seen from four directions.

Four challenges, one system, four perspectives.

People asks: who needs to use the output of this system, and where are they? Asked at national programme level, the answer was across the network, not just inside the Trust firewalls.

Systems asks: what does this system depend on that it doesn't control? Other Trusts' firewalls, and other Trusts' governance decisions. None of them ours to set at local level.

Design asks: at what level is the system boundary drawn, and was that level chosen, or assumed? It was assumed. We inherited it from the funding vehicle. The money came to the Trust, so the system was a Trust system, and we built and optimised it accordingly.

Risk asks: what contradictions exist between this system and the wider structures it sits inside? A national strategy that funded digital programmes, alongside a professional training curriculum whose examination still runs on glass. That contradiction was there the whole time. Nobody owned it.

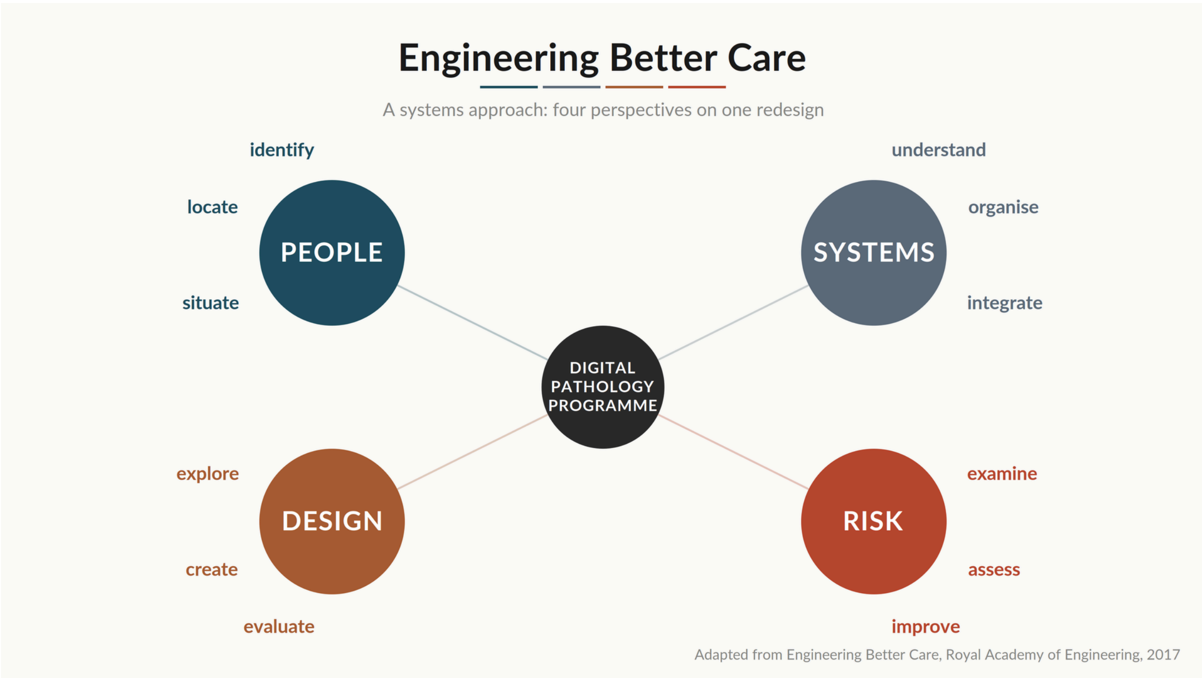

These are not clever questions. They are common sense, which is exactly why it helps to have them written down in a clear framework.

Notice that none of these are technology questions. But none of them are independent of technology either. They are system questions.

If you accept that digital pathology is only one component of a larger system, then the practical question becomes: how do you know whether you have understood that system properly?

If there is one slide to take away, this is it. Eight key questions, two for each of the four perspectives. They are the fundamental questions every change programme needs to ask, at every level of the system, at every stage of the programme.

Eight questions for your programme

PEOPLE

Who needs to use the output of this system, and where are they?

Whose working practice must change for the clinical benefit to be delivered, and who owns that change?

SYSTEMS

What systems outside this programme's control must integrate with it, and who owns their roadmap?

Where does data need to flow across an institutional or governance boundary?

DESIGN

At what level is the system boundary drawn — and was that level carefully chosen, or assumed?

Do your success metrics include the clinical benefit you are building towards, or only the operational outputs you control?

RISK

Which failure modes are already documented in comparable programmes, and have you modelled them in yours?

What contradictions exist between this system and the wider structures it operates inside?

Adapted from Engineering Better Care, Royal Academy of Engineering, 2017

I've included a link to the University of Cambridge website that hosts a comprehensive guide and toolkit for Engineering Better Care.

Let me bring this closer to the present. I work as part of a regional network — six partner Trusts plus my own. Same NHS funding settlement, same national strategy, same clinical need. And, as of this year, three different states.

Three Trusts have been funded since 2022 and have not been able to start implementation, because their aged laboratory information systems will not integrate with the image management system. Two other labs were delayed for three years by the same problem.

What happened next is revealing. The national programme was framed around digital pathology implementation. But the dependency on legacy LIMS platforms sat outside that programme boundary.

Only after local contracts had been signed did many Trusts discover that direct integration wasn't possible.

The solution was specimen-tracking middleware, acting as a bridge between systems. The technology existed. The problem was that the requirement had never been identified as a programme-level risk.

Trusts suddenly found themselves writing additional business cases for infrastructure that wasn't part of the original funding assumptions.

The business case changed after approval.

That's not a procurement failure. That's a boundary-definition failure.

Three years of a programme's momentum, absorbed by the same boundary-definition problem.

Nottingham is the only hospital with a fully digital workflow, because we already had specimen-tracking middleware in our laboratory.

And we are digitising behind individual Trust firewalls. There is no inter-hospital case sharing. The primary network-level benefit of digital pathology — sending a difficult lymphoma case to a specialist like myself working in the regional centre — has not been designed in. Glass slides that have already been scanned are sent and scanned a second time.

Same investment. Same technology. Three outcomes.

That is not an execution problem. Every team in this network has done their job. The decisive variables are system design decisions that were made — or not made — upstream, before anyone bought a scanner: the boundary was drawn at the Trust, the network-level joins were nobody's specification, and the inability of legacy LIMS to integrate with the IMS was not identified as a risk when the national investment strategy was written.

Now add AI. Over the past six months I've been deploying a CE-marked prostate algorithm across the network, with central NHS funding. The evidence package was complete, because we had been running it in Nottingham for over a year with excellent results. We could share all the work done in Nottingham: the validation methodologies, the risk assessments, the GDPR compliance work, the detailed clinical safety case, the vendor support package.

And at every new site the deployment was slow and costly, because each boundary had to be crossed from scratch. No one owns the validation work at network level, so it gets repeated at every site instead of done once and reused. It costs N times what it should, where N is the number of boundaries.

Dr Orly Ardon made the point earlier about financial stewardship. This is what could derail AI in the NHS: every algorithm implementation is a bespoke task in a fragmented system, with variance baked into processes, without pooled expertise.

This will repeat every time we add a new AI application. Same package, same boundaries, crossed from scratch each time.

AI deployment is a stress test of the system you have already built.

If your governance is fragmented, AI inherits it.

If your image access is fragmented, AI inherits it.

If your workforce model is fragmented, AI inherits it.

Let's step back again. Picture the system levels nested inside each other. The lab sits inside the Trust. The Trust sits inside the regional network. The network sits inside the national system — the NHS, the profession and the regulators together. And all of that sits inside the international layer, where the standards actually live.

So the most fundamental question is where the boundaries are drawn, and who owns that decision. A system boundary is a choice, not a fact. But in complex systems, it may not be obvious that an investment or procurement decision is actually a system boundary decision. That's why we need a systems engineering framework.

I want to close by looping back to Flyvbjerg's number: about one project in two hundred lands on budget, on schedule, and on benefit.

So what do the most successful types of megaprojects do differently? Flyvbjerg and Gardner call them *modular natives*: built from repeatable units governed by one set of design rules. They are not bespoke.

Their metaphor is Lego — Flyvbjerg, after all, is Danish. A simple house and a complex castle use the same brick, and the connector is identical on both. The ambition of the model is infinite; the join is constant. That is the source of predictability, of both risk and outcome.

For the NHS, every network, every laboratory, every AI application is a bespoke build. The rescanning problem isn't an imaging failure; it is a connector failure. The images exist. The standard connection required to move them between organisations was never specified.

Bespoke megaprojects regress to the fat tail of unpredictable risk and unlimited overruns.

Find your Lego means: **standardise the connectors, not the components.**

And I mean connectors in the widest sense:

- the language we use to define our systems,
- the protocols that validate safety and quality,
- the APIs between internal layers,
- the standards for image format and calibration,
- Interface every AI tool has to meet.

This is not only a digital-pathology problem. Pathology is where imaging, molecular and clinical data already converge, and AI will soon have to work across all of them.

Which brings us to the hardest joins. I showed you the levels of the hierarchy earlier, and most of these standards sit beyond local control. Our role is to influence upward. The work for every vendor, provider and professional body in this room is to support the work to develop the standards and connectors that will make system-wide redesign possible.

Digital pathology programmes produce variable results because digital pathology is only one component of a larger socio-technical system. Technology, governance, workforce, workflow, interoperability and investment are interdependent. Large-scale programmes succeed or fail according to how well that system is understood and designed.

Engineering Better Care provides a practical framework for doing so. Flyvbjerg's Iron Law predicts what will happen if we are unable to apply it.

With the benefit of hindsight, the best time for the NHS to have applied systems thinking to this transformation was 2019, when the investment strategy was written. The next best time is now.

References and further reading

Engineering Better Care

- Royal Academy of Engineering — Engineering Better Care: https://raeng.org.uk/policy-and-resources/engineering-policy/health/engineering-better-care/
- Engineering Better Care toolkit, University of Cambridge Engineering Design Centre: https://www.iitoolkit.com/improvement/ebc.html
- Clarkson, P.J., Bogle, D., Dean, J., Tooley, M., Trewby, J., Vaughan, L., Adams, E., Dudgeon, P., Platt, N. and Shelton, P. (2017) Engineering Better Care: a systems approach to health and care design and continuous improvement. London: Royal Academy of Engineering. ISBN 978-1-909327-35-1. Full report (PDF): https://www-edc.eng.cam.ac.uk/files/engineering-better-care.pdf

Megaprojects and the Iron Law

- Flyvbjerg, B. and Gardner, D. (2023) How Big Things Get Done: The Surprising Factors That Determine the Fate of Every Project, from Home Renovations to Space Exploration. London: Macmillan. ISBN 978-1-0350-1895-6. Publisher page: https://www.panmacmillan.com/authors/bent-flyvbjerg/how-big-things-get-done/9781035018956