ORTENT WHITEPAPER

Selling software and AI into the NHS

What actually decides whether a product gets in.

By Andrew Wyatt, Founder, Ortent Advisory Ltd

July 2026 · For founders, boards and investors in health software and AI

Contents

Executive summary

Most software and AI companies approach the NHS as a sales problem. Sharper positioning, a better demo, a more persuasive deck. Then the pipeline fills with enthusiastic clinicians and nothing closes, and two quarters disappear.

The problem is rarely the product. The buyer is not deciding whether your product is good. They are deciding whether adopting it is safe for them personally and affordable for their organisation this year. That is a different problem, and it does not yield to a better deck.

> *The NHS does not buy technology. It buys trust. Everything in this paper serves that one sentence.*

Three consequences follow, and most companies learn them the expensive way. Your first customer is not the trust, it is Information Governance. Your real competitor is not another vendor, it is the status quo and the cost of change. And a pilot without an agreed scaling budget is not a sales strategy, it is funded market research.

This paper is about selling into that system. It covers, in order, why companies fail and how buying actually works; who holds power; the regulatory assurance stack that makes you deployable; the evidence and business case that make you sellable; the routes to market and the competition you are really against; the pilot-to-scale problem and a five-level readiness model; three worked go-to-market examples; a before-and-after case; and the reform moving the whole map. England-focused, verified July 2026.

The instrument sits alongside it. The companion diagnostic at ortent.co/tools/nhs-readiness scores your own product across twenty markers and places it on the five-level model. Read the paper, then run it.

Why most companies fail

Companies do not usually fail in the NHS because the product is weak or the market is closed. They fail because they run a commercial motion designed for a different kind of buyer. In ordinary enterprise software, an enthusiastic user with a budget can buy. In the NHS, the enthusiastic user rarely has a budget, the budget holder is not the user, and both sit behind a wall of information-governance and clinical-safety sign-off that has nothing to do with how good the demo was.

The five own-goals

- Pitching features and technology to clinicians who care about problems and workflow, not architecture.
- Selling to an enthusiastic clinician with no budget and no executive sponsor, which produces endless pilots and no purchase.
- Leaving information governance, the Caldicott Guardian and the data protection officer until late, then discovering a blocker that stops everything.
- Bringing no credible economic case, or a case whose savings land in a budget other than the one being asked to pay.
- Treating a framework listing as demand generation. A framework makes you buyable. It does not make anyone want to buy.

Each of these is a symptom of the same root error: treating the NHS as one account with one buyer. It is not. The rest of this paper is about the structure that error ignores, and how to work with it rather than against it.

How the NHS actually buys

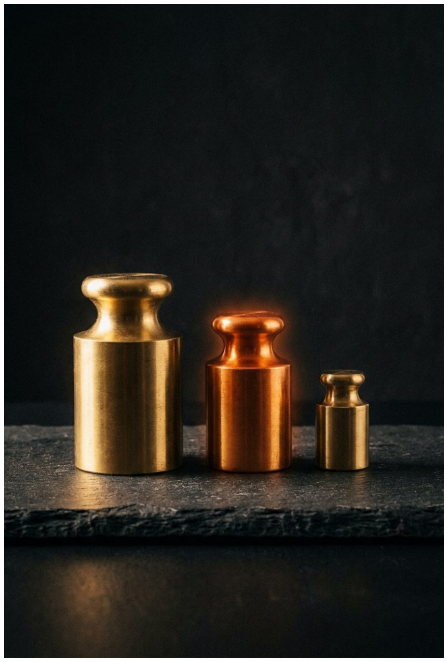

The single most useful mental model is this. Standards and credentials are set nationally. Money and contracts sit regionally and locally. Navigation and evidence come from a separate innovation layer. No one part can buy on its own, and a company that treats the NHS as one account loses.

The three layers

Nationally, bodies such as the MHRA, NICE and the data and clinical-safety standards set the rules of entry. They rarely sign a cheque, but they decide whether you are allowed through the door. Regionally, the money sits with 42 Integrated Care Systems and their statutory Integrated Care Boards, which are consolidating through mergers toward the high twenties. The board commissions and pays; its partnership arm sets strategy and holds no chequebook. Do not confuse the two. Locally, around 200 provider trusts deliver care and buy the clinical systems, and primary care runs on an independent-contractor model in which most GP practices are small businesses whose IT is centrally funded and dominated by two clinical systems.

The two facts about money that decide deals

Capital versus revenue. The NHS separates capital from revenue, each capped and hard to move between. Software is usually a revenue cost, but implementation and some licences can be capital. A business case has to be affordable within both envelopes and account for the revenue consequences of capital, including depreciation and central charges.

Who pays versus who benefits. The saving from a product often lands in a different budget from the one that pays for it. A tool that reduces admissions saves the commissioner but costs the provider to run. A tool that saves nursing time does not cash out unless posts are actually removed. This mismatch is the single most common reason a strong return-on-investment product never gets funded, and it is invisible on a features roadmap.

> *Name the budget that pays and the budget that saves, and get both parties in the room. If you cannot, you do not have a business case. You have a benefits narrative.*

Who holds power

A purchase needs a budget holder, a clinical sponsor, information-governance and safety sign-off, and procurement. Any one of them can stall it. Map every deal against the people you will actually meet, and what each is really deciding.

The most expensive misread is treating information governance as a late-stage formality. In practice the data protection officer, the Caldicott Guardian and the senior information risk owner can each stop a deal on confidentiality or lawful-basis grounds, and they will, if the paperwork arrives after the clinical enthusiasm rather than before it.

Role	What they decide	Champion or blocker
CFO / Director of Finance	Affordability, capital versus revenue, cashable savings this year. Holds the budget.	The critical gatekeeper
CIO / Director of Digital	Architecture, cyber, interoperability, the EPR roadmap and integration cost.	Kills what does not fit
CCIO / CNIO	Clinical usefulness, clinician adoption and the safety of the workflow.	Key clinical bridge
Caldicott / DPO / SIRO	Confidentiality, lawful basis and information risk.	Pure gatekeepers
Clinical champion	Whether it solves a real problem on their ward or pathway.	Most important ally
Procurement	A compliant route, the right framework, value for money.	Blocker if engaged late

The clinical champion

The champion is the ally you cannot manufacture. A respected frontline clinician who advocates because your product solves a problem they feel daily. They provide credibility no company can buy, they defend the business case to peers and the board, and they drive the adoption that is the hardest part of the whole exercise. Without one, digital deals stall at pilot. The corollary risk is real: if your champion leaves the trust, the deal often dies with them, so never stay single-threaded.

> *Your first customer is not the trust. It is Information Governance. Bring it the answers before it has to ask the questions.*

The assurance stack

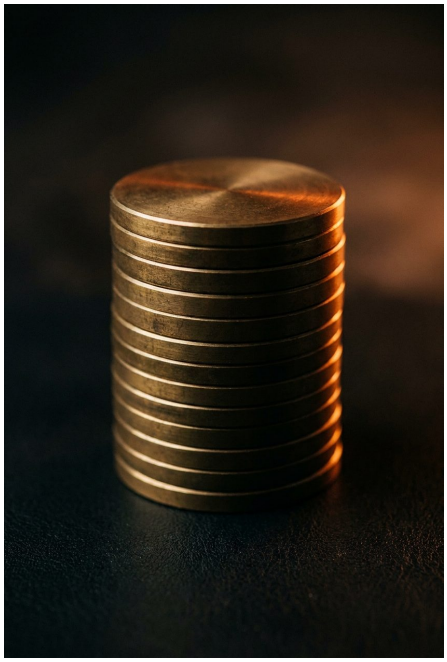

Two categories sit between you and a live deployment. Legally mandatory items, and framework or procurement gates that become contractually mandatory in practice. The single front door that bundles most of it is the DTAC, the Digital Technology Assessment Criteria. Get the DTAC pack right, with real evidence behind it, and you have covered most of the rest.

The legally mandatory core

- Medical device status. If your software has a medical purpose, diagnosis, prevention, monitoring, prediction or treatment, it is a medical device and must be UKCA or CE marked and MHRA-registered. Most clinical decision-support and diagnostic AI lands here. It is the most expensive path, and buyers ask about it early.
- Clinical safety, DCB0129. A named Clinical Safety Officer who is a registered clinician, a clinical safety case and a hazard log. Following the 2025 change to information standards this is now a legal must-comply duty, not advisory.
- Data protection. Most companies are processors acting for the NHS controller. The lawful basis, the DPIA and the data processing agreement all have to be right before data moves.

The AI model-training trap

This one catches AI companies repeatedly, and it is worth stating plainly. Acting as a processor is fine while you only process NHS data to deliver the service the trust commissioned. The moment you use identifiable patient data to train, fine-tune or improve your own proprietary model, you are deciding a new purpose of your own. That makes you a controller or joint controller, triggers your own lawful basis and transparency duties, and in practice gets blocked on sight by most data protection officers and Caldicott Guardians. The workable positions are two: train only on properly anonymised data, where pseudonymised does not count, or secure an explicit separate lawful basis and consent with its own DPIA. Settle the answer in the data processing agreement up front. Discovering it mid-procurement kills deals.

The framework and procurement gates

- DSPT and Cyber Essentials. The annual Data Security and Protection Toolkit at Standards Met, plus Cyber Essentials, and often Cyber Essentials Plus where you handle personal data.
- DTAC. The national baseline bundling clinical safety, data protection, technical security, interoperability, and usability. A new form has been in force since February 2026, and buyers now re-check the underlying artefacts rather than trusting the self-declaration.
- Interoperability. UK Core FHIR and the NHS Number via the Personal Demographics Service. Saying you do FHIR is not enough. It must be UK Core, and it is frequently the deciding factor against the incumbent record.

- Net zero and social value. Every NHS procurement carries a minimum ten per cent weighting for net zero and social value, and larger contracts require a published Carbon Reduction Plan. This is scored whether or not your product is clinical.

For AI specifically, the rules on automated decisions changed in February 2026 when the Data (Use and Access) Act 2025 came into force. Where your product makes or materially informs a significant decision about a person, the safeguards are product requirements, not terms and conditions. The interface has to surface that a decision was automated, offer a genuine route to human review, and let the person contest it. Design these into the user experience and the audit trail, not the small print.

Evidence and the business case

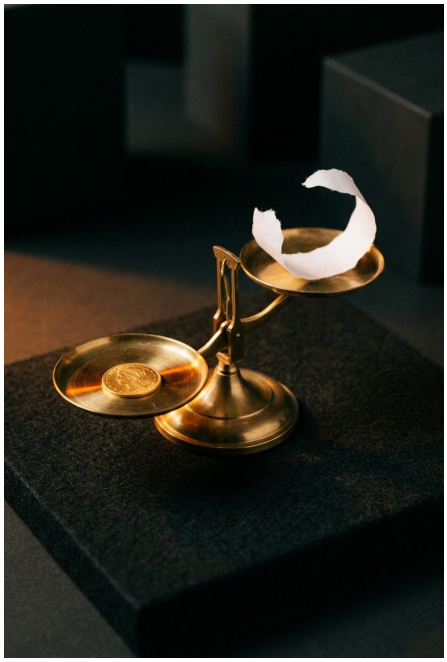

Clearing the assurance stack means you can go live. It gives no one a reason to buy. That comes from two things: evidence of benefit at the right level, and a business case a finance director will sign.

Evidence

NICE's Evidence Standards Framework for digital health technologies tiers the evidence you need by the risk and claim of your product. Use it as your evidence roadmap from day one. Building your study to the right tier is far cheaper than retrofitting evidence after a failed pilot, and commissioners increasingly expect you to self-classify against it. Engagement is not the same as outcomes, and a slide showing high usage is not a slide showing benefit.

The business case

In a system with no spare cash, the deliverable that actually sells is a finance-grade case: which national priority it serves, the cashable saving quantified, which budget pays and which benefits, and the benefits you will report back on. Bring the model. Do not expect an overstretched trust to build it for you, and do not confuse a benefits narrative with a business case. A benefit that cannot be cashed does not free a budget.

The strongest hook is the cost-improvement programme every trust runs each year. If your product can be positioned as a cashable line against those savings targets, by removing agency spend, reducing length of stay, avoiding a hire or cutting a paper process, it moves from a nice-to-have to something that helps a director hit a number they have to hit anyway.

> *Marketing spend before you have evidence is wasted, because the constraint is evidence, not awareness. Diagnose the gate in front of you, and spend on that.*

The route to market and the competition

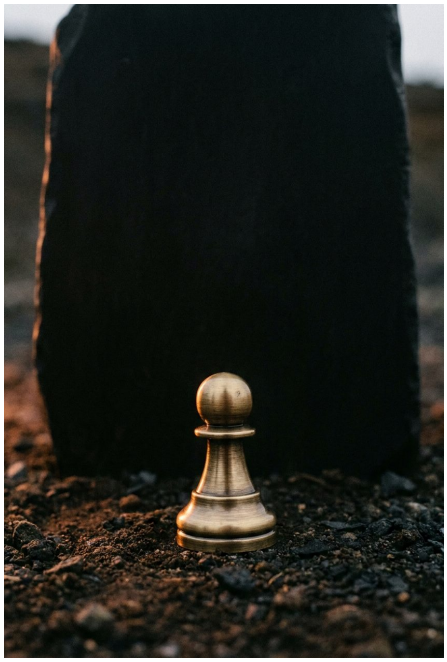

Framework presence is the centre of gravity of NHS go-to-market, because being on a framework or a dynamic purchasing system lets a trust buy you without running a full tender. For software and AI the primary channels are G-Cloud and the Spark dynamic purchasing system. The Spark route is always open, so you can join at any time rather than waiting for a round. Register as a supplier on Atamis, the platform where most NHS procurement and the pipeline of upcoming opportunities now live.

Two regimes run in parallel. The Procurement Act 2023, live since February 2025, governs software, IT and AI. The Provider Selection Regime governs clinical services. Almost always a software company sits under the Procurement Act. A practical warning from the people who run these processes: most framework applications are rejected for missing documentation or wrong formatting, not weak product. Treat the paperwork as the deliverable.

The competition you are actually against

Founders benchmark against other startups. NHS buyers do not. The real competitive set is four things. The incumbent electronic patient record, from Epic, Oracle Health, System C or Meditech, owns the workflow and increasingly ships the adjacent module you are selling, bundled, so you compete with good enough and already paid for. The hyperscalers and ambient AI, from Microsoft, Google and Amazon, are moving into documentation, triage and analytics with brands procurement already trusts. The national platform play, the Federated Data Platform and the Single Patient Record, concentrates data and analytics centrally, so a data or analytics layer has to know whether it is complementary or redundant. And build versus buy, plus plain inertia, because every buyer can choose to do nothing, and the thing that already works is usually a spreadsheet.

> *In the NHS, best product loses to lowest-risk change more often than founders can accept. The incumbent EPR is not just a competitor. It is the gatekeeper. Treat interoperability as commercial strategy, not a technical detail.*

Pilotitis and the readiness model

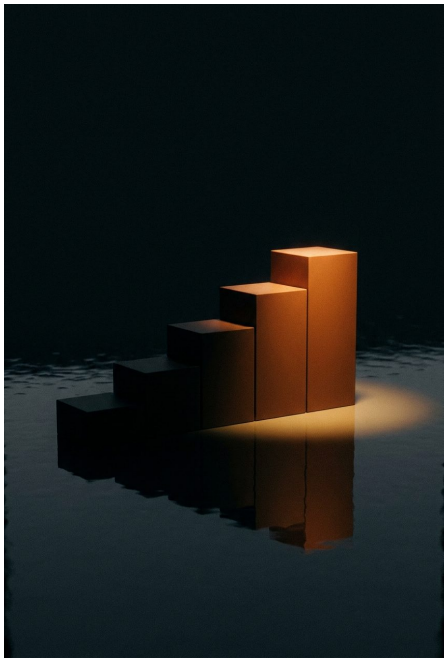

Pilotitis is the defining failure pattern: promising products stuck in repeated small, short, locally funded pilots that never scale. Companies spend two to three hundred thousand pounds per trust navigating repeat clearances, then hit a cliff edge when pilot funding ends. The way out is to know exactly where you sit, and to spend only on the gate directly in front of you.

The trap is spending on the wrong level. Marketing spend at Level 2 is wasted, because the constraint is evidence. Sales headcount at Level 3 burns cash, because the constraint is repeatability. Reference sites are currency: one well-evidenced deployment with a named champion and a quantified benefit is worth more than any amount of marketing.

Level	State	The gate to clear next
1	Not deployable	Close the assurance stack. Nothing else counts yet.
2	Deployable but unsellable	Evidence of benefit and a CFO-grade business case. Not marketing.
3	Sellable but unscalable	A framework and a business case others can lift. Not more sales effort.
4	Repeatable regional sales	A national pull lever.
5	National market access	Hold it. Evidence travels; adoption no longer depends on one champion.

A pilot breaks the cycle only when it is designed for the business case, not the demo: agree the metric, the decision a good result triggers, and whose budget funds scale, before you start.

> *A pilot without an agreed scaling budget and a named owner is not a sales strategy. It is funded market research.*

Three worked examples

The model is easier to trust against real go-to-market shapes. Three composite cases, each drawn from the pattern of companies selling the kind of product named. None is a specific company.

One. An AI imaging tool into radiology

A diagnostic AI that flags suspected findings on chest imaging. It is unambiguously a medical device, so the first gate is MHRA. The company arrives at trusts with a strong clinical study and a keen consultant radiologist as champion, and stalls, because it has treated the sale as clinically-led only. It sits at Level 3: real evidence, a champion, one reference site, but each deal is bespoke and slow. The gate in front of it is not more clinical evidence. It is a business case a finance director will sign, framed as radiologist reporting capacity released and backlog cleared against the elective-recovery target, plus a framework route so trusts can buy without a tender. Fix those two and it moves to Level 4, repeatable across the region, with the national pull lever, a NICE route, as the next gate.

Two. A virtual-ward and remote-monitoring platform

A remote-monitoring platform for virtual wards. Lower device-classification risk, and squarely aligned to the hospital-to-community shift, which is a strong national-alignment answer. Its trap is different: the benefit, admissions avoided, lands in the commissioner's budget while the running cost falls on the provider. Classic who-pays-versus-who-benefits. It can be technically deployable and clinically loved and still fail to get funded. The gate is not the product. It is naming the payer and beneficiary budgets and getting both the ICB and the provider in the same room, with a case that shows where the saving lands and how it is shared. Dynamic purchasing routes suit it well, and a Health Innovation Network can generate the real-world evaluation the business case needs.

Three. An ambient-scribing tool for clinicians

An ambient voice tool that drafts the clinical note from the consultation. It sits directly on a first-wave national investment priority, which helps, but it faces the sharpest version of two problems. First, the model-training trap: if it improves its own models on identifiable consultation audio, it becomes a controller and is blocked. It has to lock that down in the data processing agreement. Second, the competition: the hyperscalers are in this exact space with trusted brands, and the incumbent EPR may bundle the feature. Its defensible position is workflow fit and interoperability with the installed record, not a feature comparison. Where it wins, it wins by reducing the cost and risk of change, and by proving cashable clinician time released, not by out-demoing Microsoft.

Before and after: the trust that said not yet

A composite that shows the model used honestly. A growth-stage company with a genuinely good product runs the readiness diagnostic before a board meeting where it plans to raise on the strength of its NHS pipeline.

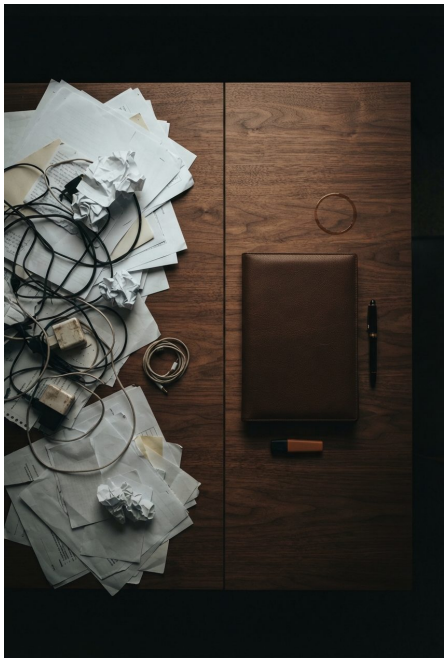

Before

The deck says twelve active NHS opportunities and two live pilots. The founder believes the company is at national scale-up. The diagnostic tells a different story. One mandatory gate is open, because the product is arguably a medical device and classification has not been resolved. That alone caps it at Level 1, not deployable, whatever the pipeline looks like. Underneath, the commercial side is genuinely strong. But there is no CFO-grade business case, the two pilots have no agreed scaling budget, and the company is on no framework.

After

The board changes the plan. Instead of raising on pipeline, it spends one quarter closing the device-classification question and the DTAC pack, which moves the company off Level 1. It builds one finance-grade business case with a named payer budget, and converts its strongest pilot into an evidenced reference by pre-agreeing what a good result triggers and from whose budget. It applies to G-Cloud and Spark. Nine months later the same pipeline converts, because the company is now buyable without a tender and has a case a finance director can defend. The pipeline did not need to be bigger. The company needed to clear the gate in front of it.

> *Twelve promising opportunities on top of an open mandatory gate is not a pipeline. It is a queue behind a locked door. Find the lock first.*

Where this does not apply

The guide is deliberately scoped. Three boundaries are worth naming.

- England, not the four nations. Scotland, Wales and Northern Ireland run different structures, standards and procurement. NHS-approved in England carries limited automatic weight in Cardiff, Edinburgh or Belfast. Treat each nation as a separate market.
- Clinical and data products, not pure back-office. A genuinely non-clinical administrative tool that touches no patient data faces a lighter version of the assurance stack. The buying dynamics still apply; the regulatory gates largely do not.
- Growth-stage, not pre-revenue research. If you have no product and no deployments, the maturity model is your roadmap, not your scorecard. Start with the assurance stack and one evidenced use case.

Two things are also genuinely in flux at the time of writing and should be re-checked before use in a live proposal: the exact number of Integrated Care Boards during the mergers, and the forthcoming AI-in-healthcare regulatory framework expected to follow the national commission's recommendations.

Next steps

Diagnose the level. Fix the gate. Then move. The companion diagnostic turns this paper on your own product: twenty questions across the assurance stack and the buying reality, a read on each axis, legally deployable and commercially viable, and the single gate in front of you. It takes ten minutes, captures no email, and runs in your browser. There is also a Claude prompt that runs the same assessment conversationally and pressure-tests you on the two real gatekeepers, cashable savings and interoperability with the incumbent record.

- Run the diagnostic at ortent.co/tools/nhs-readiness.
- Have your clinical lead, commercial lead and finance lead answer it independently. The markers where they score lower than you are where your view of the state and the state itself have diverged.
- If it surfaces a gate you want a second pair of eyes on, or you want a board-ready read on where a portfolio company sits, book a working session at ortent.co/contact.

ABOUT THE AUTHOR

Andrew Wyatt runs Ortent Advisory, working with growth-stage software and AI companies in life sciences, digital health and AI as a Non-Executive Director, board advisor and fractional commercial leader. He has four exits across his career (Lotus to IBM, Paragon to Phone.com, Apertio to Nokia at $240M, Clearswift to Lyceum) and was CGO at Sapio Sciences and COO at Lumeon. He is not a lawyer, scientist or clinician; this paper is a commercial field guide, not regulatory or legal advice.

England-focused. Verified July 2026. · ortent.co · Clarity when it counts.