

Governing what we already run

A three-day intensive for the Board of Vitalis Health Group.

Built from the Cycle 1 assessment — six domains, one plateau, two critical gaps.



3

DAYS • IN PERSON

14

MODULES

19h

CONTACT HOURS

6/6

DOMAINS COVERED

FOR THE BOARD AND THE EXECUTIVE COMMITTEE • IN PERSON • EN
FACILITATION • ANDIARA PETTERLE, PHD • BOARDACUMEN • ILLUSTRATIVE SAMPLE

SAMPLE DOCUMENT • FICTIONAL ORGANISATION • ILLUSTRATIVE DATA

WHAT THIS IS – AND WHAT IT IS NOT

A programme, not a presentation

The objective is that this board leaves able to supervise clinical artificial intelligence with the same fluency with which it already supervises capital and clinical risk — without depending on anyone to ask the right questions for it.

This is not a keynote, it is not a consulting diagnosis, and it is not a sales meeting. Each block **teaches a concept, applies it to this group, and hands the decision back to the table.** Each day is useful on its own.

I am not going to teach you your business. This board has governed nine hospitals through two decades of change. I am going to give you the concepts, the frameworks and the questions to lead clinical AI — instead of absorbing it.

THERE IS NOTHING SOLD AT THE END

The product of the three days is **yours**: the Board AI Manual — your principles, the standing questions you will require, the red lines, and the risks you will oversee. Whoever facilitates appears in the footer, not on the stage.

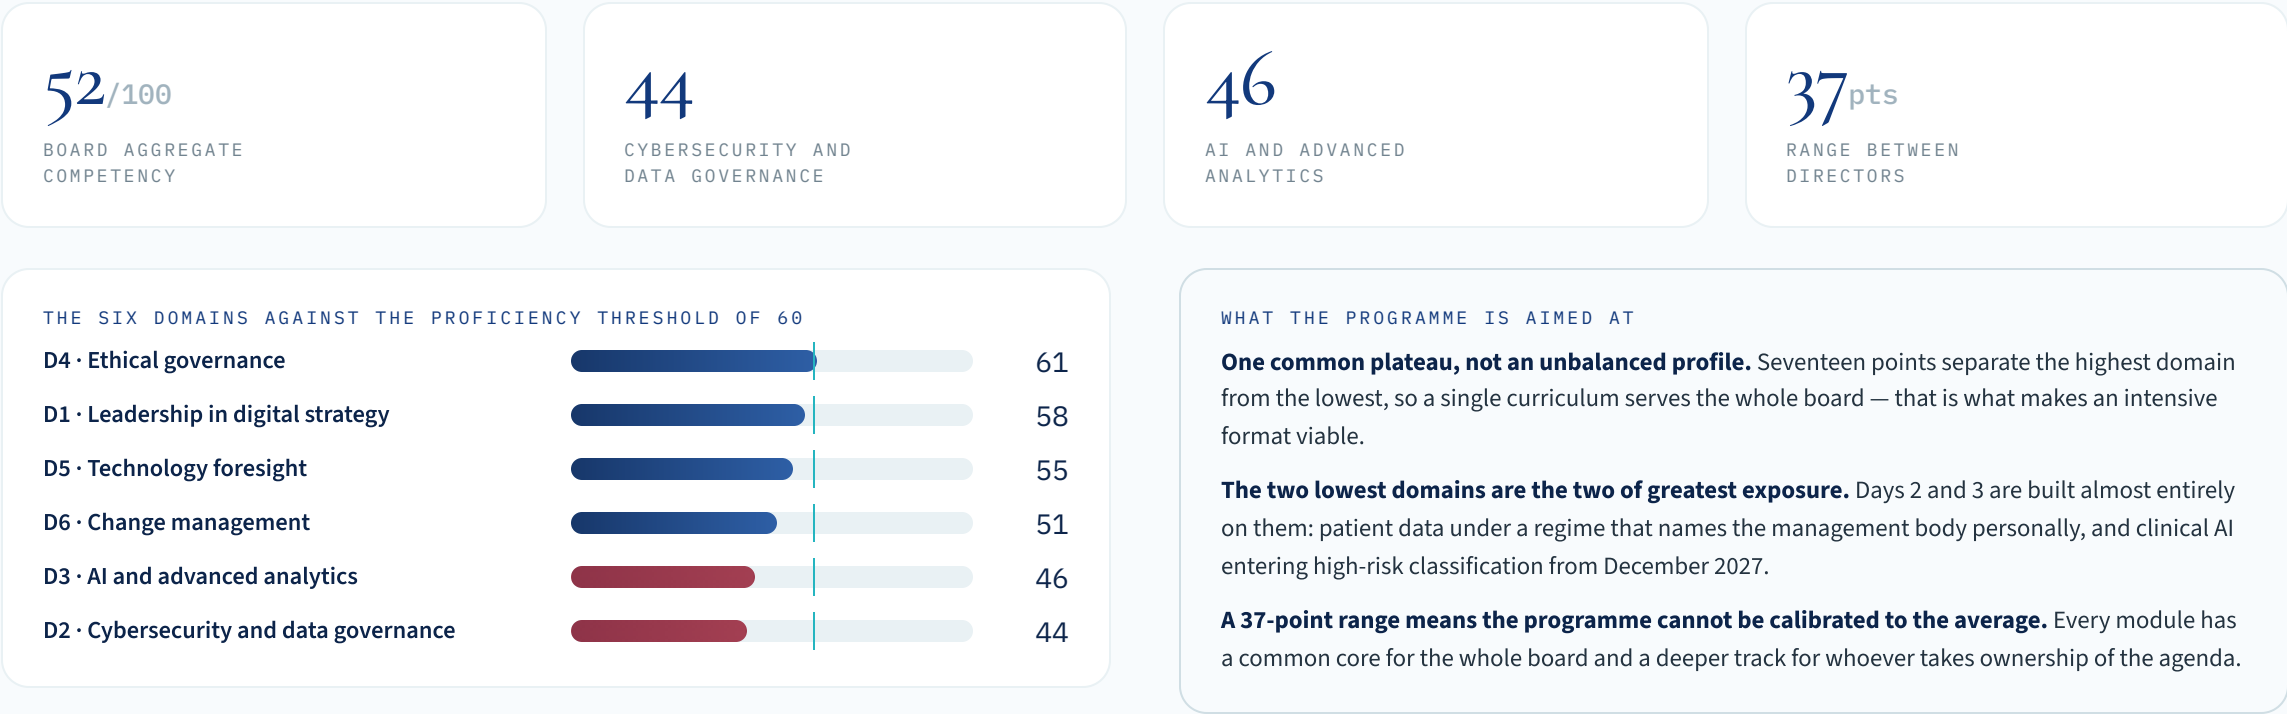
THE SHIFT

From «AI as a diffuse threat and a black box» to «AI as a system we have mapped, classified and can govern — and a decision that is ours».

WHERE THIS PROGRAMME COMES FROM

Built on a measurement, not on an assumption

This is not a generic curriculum. Every module below exists because the Cycle 1 assessment identified a specific gap in this board — and the order of the three days is the order of the exposure.



THE DESIGN PRINCIPLE

Board competency is not technical competency

A director is not asked to know how to build a clinical AI system. A director is asked to know how to recognise a badly governed one. These are different competencies, and the second is learned in days, not years.

WHAT THE PROGRAMME IS NOT FOR

- Training directors in data science or clinical informatics.
- Producing technical opinions on models or architectures.
- Substituting for the executive team, the CIO or the chief medical information officer.
- Turning the board meeting into a technical project review.

WHAT THE PROGRAMME IS FOR

- Knowing **which question has not yet been asked**.
- Recognising an **insufficient answer** when it is given.
- Knowing **what is never delegated** to a system — in any circumstance.
- Knowing **what has to be written down**, and where.

WHY THIS FRAMING IS ALSO THE LEGAL FRAMING

Two obligations converge on the same artefact. The AI literacy duty of Article 4 of the AI Act is an **obligation of means** — supporting a sufficient level of literacy in whoever operates these systems, proportionate to context and risk. NIS2 goes further and **requires members of the management body to follow training** on cyber risk. A board development plan, approved in the minutes and measured at the end, is precisely the proof of both.

THE ARC

Three days, one movement

Each day ends with an artefact and opens the following one. The tension is deliberate: it starts with a number that hurts, escalates into the two domains where the board is weakest, and resolves in a document the board itself writes.

DAY 1 · UNDERSTAND

The machine and the economics

Three layers of AI · what became scarce in healthcare · the group's three scenarios, stress-tested.

M0 – M3 · 5H30

DAY 2 · INTERROGATE

Data, resilience and clinical AI

What a board supervises in cyber · data as an asset · crisis simulation · risk classification against Annex III.

M4 – M8 · 7H30

DAY 3 · GOVERN

Decide, record, commit

Ethics and equity · redesign, not automation · the ten standing questions · a real deliberation · the Board AI Manual.

M9 – M13 · 6H00

Facilitation rule: no expository block runs longer than 18 minutes without interaction. This is training, not a conference.

AGENDA

Reference agenda · Day 1

TIME	BLOCK	WHAT HAPPENS	FORMAT
09h00 – 09h20	M0 · Opening and learning contract	The question that frames the three days	Plenary
09h20 – 10h20	M1 · How the machine actually works · Lab 1	Three layers of AI · model vs agent · seeing the machine answer on clinical questions	Plenary + lab
10h20 – 10h35	Break		—
10h35 – 11h45	M2 · What became scarce in healthcare	«If intelligence became cheap, what became expensive?» — trust, the longitudinal record, the care network	Plenary + World Café
11h45 – 13h00	M3 · The group's three scenarios, stress-tested	Care out of the hospital · the payer–provider boundary · clinical data as a licensable asset	Sub-groups
13h00 – 14h00	Lunch		—
14h00 – 14h30	M3 · Close of Day 1 · scenario sheets	Early-warning signals and no-regret moves — they enter the Manual	Plenary

AGENDA

Reference agenda · Day 2

TIME	BLOCK	WHAT HAPPENS	FORMAT
09h00 – 10h30	M4 · What a board supervises in cybersecurity	From the incident list to resilience: continuity of care, critical dependencies, tested recovery, third parties	Plenary
10h30 – 10h45	Break		—
10h45 – 12h15	M5 · Data as an asset · Lab 2	Building the data inventory with lawful basis, live, with the executive team	Workshop
12h15 – 13h15	Lunch		—
13h15 – 15h15	M6 · Crisis simulation	Tabletop: ransomware affecting the patient record and elective scheduling across two hospitals	Simulation
15h15 – 15h30	Break		—
15h30 – 16h30	M7 · Clinical AI: the three layers inside the hospital	Imaging triage, documentation, no-show prediction, plan risk scoring — what each one actually is	Plenary + lab
16h30 – 17h30	M8 · Risk classification workshop	Classifying the group's own AI systems against Annex III — produces an artefact, not a summary	Group work

AGENDA

Reference agenda · Day 3

TIME	BLOCK	WHAT HAPPENS	FORMAT
09h00 – 10h15	M9 · Ethics, equity and the patient	Algorithmic bias as an equity-of-access event · Article 50 disclosure · what is non-delegable	Plenary
10h15 – 10h30	Break		—
10h30 – 11h45	M10 · Redesign, not automation	80/15/5 · where the clinician rises in the chain · the revenue-cycle clinic	Plenary + clinic
11h45 – 12h30	M11 · The ten questions that go into the minutes	Drafting the board's standing procedure for any AI initiative	Plenary drafting
12h30 – 13h30	Lunch		—
13h30 – 15h00	M12 · Deliberation simulation · «Project Aurora»	A real-format €4.1M clinical-AI proposal with an incomplete dossier. The board decides	Simulation
15h00 – 16h00	M13 · The Board AI Manual + close and commitment	Gallery walk, canvas, dot-voting on priorities · the three questions to take home	Plenary

DAY ONE

OI

Understand the machine and the economics

Before governing artificial intelligence, the board has to see it work — and see where the value in healthcare moved.

Mo · Opening and learning contract

09H00 – 09H20 • 20 MIN

LEARNING OBJECTIVE

Create psychological safety and a clear contract — this is training, everyone is allowed not to know, and the decision stays with the table.

THE CONCEPT TAUGHT

- None yet. The framing comes first.

This board took nine hospitals through two decades of change and never lost a year to a technology. But there is one question yesterday's discipline does not answer: if a model already tells a patient what their symptoms probably mean, and the answer is confident, and it is wrong, and nobody in this room can say who checked it — what exactly is this board supervising? Then leave the silence.

THE ACTIVITY

Round-robin check-in — one word each: «my biggest fear, or my biggest bet, on AI at Vitalis». Recorded on a flip.

THE OUTPUT

The emotional map of the room — fears and bets — which we return to at the close of Day 3.

WHY IT IS HERE

The programme opens with the group's own asset and a Socratic question, not with an agenda. It sets the tone: a board that does not decide on AI has already decided — by omission.

M1 · How the machine actually works

09H20 - 10H20 · 60 MIN · LAB 1

LEARNING OBJECTIVE

Give real literacy — by the end, every director distinguishes predictive, generative and agentic AI, knows what a model does and does not do, and understands why a confident answer can be wrong.

THE CONCEPT TAUGHT

- **The three cognitive layers.** Predictive AI is a very sophisticated calculator (it forecasts, classifies, optimises — no-show prediction, risk scoring). Generative AI is a brilliant intern who never sleeps (it writes, summarises, translates — clinical documentation). Agentic AI is an autonomous colleague that acts, decides and coordinates end to end.
- **Model versus agent.** An agent that acts creates new accountability: who answers when the agent errs?
- **How a model «learns».** It tokenises, abstracts statistically and generates probabilistic syntheses — it does not store and does not retrieve a record. This is exactly why a model can be fluent and wrong, and why «the system said so» is never a defence.
- **Intelligence became a commodity.** The question stops being «do we buy AI?» and becomes «where is our AI defensible?»

A model that is 94% accurate is wrong once in every seventeen patients. In a network that sees 40,000 encounters a month, that is not a rounding error. It is a clinical governance question — and it has a name, a threshold and an owner, or it has none of the three.

THE ACTIVITY

Lab 1 — «seeing the machine answer». Live, we put three real clinical questions to a public assistant: an ambiguous presenting complaint, a drug interaction, and a question about a condition the group's own service is known for. The room watches three things: does it answer? does it hedge? would a tired clinician at 03h00 notice the error? Fallback: prepared screenshots. In pairs, 3 min: «what does this mean, concretely, in our emergency department?»

THE OUTPUT

The visceral experience that fluency and correctness are different properties — and the bridge into what the board has to require.

WHY IT IS HERE

You cannot discuss the economics or the governance of AI without first understanding the machine. The lab turns an abstract concept into something that happened in front of them.

M2 · What became scarce in healthcare

10H35 - 11H45 · 70 MIN

LEARNING OBJECTIVE

Understand where value sits in healthcare once intelligence is nearly free — and what therefore became scarce, and so valuable.

THE CONCEPT TAUGHT

- «If intelligence became cheap, what became expensive?» The answer, and the central thesis of the three days: **clinical trust · the longitudinal record · the care network**. The three defensible assets when anyone can generate a plausible clinical paragraph for cents.
- **The before/after table**. Before: scale of clinicians = advantage, software supports people, compliance is an obligation. After: intelligence is abundant, software performs the work, proprietary outcome-linked data plus trust is the barrier, and governance becomes a commercial asset.
- **Disintermediation**. The first place a patient takes a symptom is no longer a clinician or a search engine — it is an assistant that does not cite, does not refer and does not carry liability.
- **Authority is migrating**. From institutions to interfaces. Trust becomes relational, not institutional — which attacks the central asset of a hospital group directly.

Thirty years of records of what actually happened to real patients is the one input clinical AI cannot manufacture. Every model in the world is available to your competitors by Thursday. Your archive is not. The question is whether this board decides what to do with it — or discovers what was done with it.

THE ACTIVITY

World Café · 30 min. Tables with cards for the group's service lines — hospitals, diagnostics, the health plan, outpatient. Each table places each line on two axes: which defensible asset does it hold today (trust / record / network — or none), and what does the «after» column look like for it. Rotation between tables.

THE OUTPUT

A shared map of the group's service lines by defensible asset — the raw material for Day 2 and Day 3.

WHY IT IS HERE

This is the economics, and it comes after the fundamentals because the migration of value is only visible once you understand the machine.

SUB-GROUPS

M3 · The three scenarios, stress-tested

11H45 - 14H30 · 105 MIN

LEARNING OBJECTIVE

Move from the concept to this group's real exposure — which futures are plausible, and what the current strategy assumes without admitting it.

THE CONCEPT TAUGHT

- Scenario thinking applied to technological disruption — probability × impact × horizon, and the difference between deciding in time and reacting late.

The question is not whether the future arrives. It is which of the three has already begun — and whether the strategy we approved last quarter survives it.

THE ACTIVITY

Stress test in sub-groups · 60 min + 45 min plenary. Three sub-groups, one per real scenario: **(1) care migrates out of the hospital** — probability HIGH, 12–36 months: diagnostics, monitoring and follow-up leave the building; the estate becomes a cost base. **(2) The payer–provider boundary is redrawn by risk models** — MEDIUM-HIGH, 18–36 months: whoever prices risk best captures the margin, and from December 2027 that pricing is high-risk AI. **(3) Clinical data becomes a licensable asset** — HIGH, 12–24 months: research partnerships and algorithm validation become a revenue line, or a liability. Each group answers: which early-warning signals would we see first · what becomes irreplaceable at Vitalis · the no-regret move.

THE OUTPUT

Three scenario sheets with early-warning signals and no-regret moves — they enter the Manual on Day 3.

WHY IT IS HERE

It closes Day 1 with the economics applied to the group, creating the urgency Day 2 channels into data and clinical AI.

DAY TWO

02

Data, resilience and clinical AI

The two domains where this board scores lowest, and the two where obligations become verifiable first.

M4 • What a board supervises in cybersecurity

09H00 – 10H30 • 90 MIN

LEARNING OBJECTIVE

Move the board's view of security from a retrospective incident report to an anterior view: resilience, dependencies, tested recovery and third parties.

THE CONCEPT TAUGHT

- **The board is not the auditor.** When security reaches the board only as an incident history, the body sits as retrospective auditor. Effective oversight is anterior.
- **Four questions replace the incident list.** Which systems are indispensable to delivering care tomorrow · what is the *tested* recovery time of each · which third parties hold access, and when was it last reviewed · what changed in the process because of the last incident.
- **Third-party risk is a continuity-of-care matter.** 74% of hospital security leaders name EHR, AI and cloud vendors as their top emerging risk; 67% have experienced clinical disruption from a vendor event. That is not procurement. That is the board's.
- **NIS2 names you.** The management body of an essential entity approves cyber-risk measures, oversees implementation, may be held liable for failure — and its members are required to follow training.

In the first half of 2026 the healthcare sector absorbed an average of 2.3 ransomware attacks a day. The relevant question for this board is not whether we are a target. It is: if the patient record is unavailable at 07h00 on a Monday, who decides which theatres run — and has that person ever had to decide it before?

THE ACTIVITY

Guided reading of a real security panel, redacted — the board identifies what is missing in it. Then, in plenary, drafting the five fixed indicators the board will require every quarter.

THE OUTPUT

The quarterly data and security panel — five fixed indicators, always the same, so variation is legible over time.

WHY IT IS HERE

It opens Day 2 with the board's lowest domain, and with the one obligation in this programme that names directors personally.

M5 · Data as an asset · the inventory

10H45 – 12H15 · 90 MIN · LAB 2

LEARNING OBJECTIVE

Stop treating data governance as a defence against fines and start treating it as management of the group's second scarcest asset.

THE CONCEPT TAUGHT

- **Inventory, ownership, lawful basis, retention.** Nine hospitals, 140 diagnostic units and a 1.2M-life plan accumulate data in silos under different lawful bases. What a board needs is not the detail — it is the confirmation that the map exists and has an owner.
- **Special-category data changes the arithmetic.** Health data carries the strictest regime in law; consent, purpose limitation and re-identification risk are board-level constraints on any secondary use.
- **The European Health Data Space builds the architecture for secondary use.** In force since March 2025, applying from March 2027. It is the legal frame in which «license the archive» becomes a real option rather than an idea.
- **Governance that only avoids fines leaves half the value on the table.** The fifth panel question is deliberately built in attack mode.

Ask the room: which of our datasets could be lawfully licensed or contributed to research tomorrow? If the answer takes more than a week to produce, the asset is not governed — it is merely held.

THE ACTIVITY

Lab 2 — building the inventory live. With the executive team in the room: the group's principal data domains, purpose, lawful basis, owner, and licensing potential. Where the answer is unknown, that is the finding — and it becomes an action with a name and a date.

THE OUTPUT

The data inventory with lawful basis, plus a list of datasets with licensing potential and precisely what is missing for that to be lawful.

WHY IT IS HERE

Every AI system is, before anything else, a consumer of data. This module has to precede clinical AI, or the AI discussion floats.

SIMULATION

M6 · Crisis simulation

13H15 - 15H15 · 120 MIN

LEARNING OBJECTIVE

Exercise the board's decision-making under time pressure and incomplete information — the only conditions under which resilience is actually tested.

THE CONCEPT TAUGHT

- Crisis governance: what the board decides, what it does not decide, and what it must not learn from the press.

Everyone has a plan until the patient record is encrypted and the theatre list for Monday is inside it.

THE ACTIVITY

Tabletop · 120 min. A ransomware event simultaneously affecting the patient record and elective scheduling across two hospitals. Roles assigned: chair, clinical director, CIO, communications, and a director playing the regulator. Injects every 20 minutes — a journalist calls; a clinician reports a workaround using personal devices; the vendor says restore will take 48 hours not 4; a subset of exfiltrated records appears. The board must decide: divert or continue · notify when and whom · what is said to patients · who signs.

THE OUTPUT

A one-page decision log from the simulation, and the list of decisions the board discovered it had never assigned an owner to.

WHY IT IS HERE

It is the climax of Day 2. A board that has rehearsed once decides differently — and the gaps a rehearsal surfaces are the cheapest gaps this group will ever find.

M7 · Clinical AI: the three layers inside the hospital

15H30 – 16H30 • 60 MIN

LEARNING OBJECTIVE

Map what the group already runs, layer by layer — and see why treating it as one topic produces misaligned decisions.

THE CONCEPT TAUGHT

- **What is already in production.** Imaging triage support, clinical documentation, no-show prediction, plan risk scoring. Predictive, mostly — and the oldest of them is the least examined.
- **Conflating the layers is the expensive error.** It tends to underestimate the risk of what is already running and overestimate the novelty of what is new.
- **Where the clinician's judgement stays sovereign.** Written as a list, not as a principle. A model output that no one is accountable for overriding is a decision the board has delegated without saying so.
- **Failure signal and deactivation.** How do we know it broke, in how long, who can switch it off — and has that ever been exercised?

The imaging triage tool has been in production for three years. Nobody in this room has ever seen its false-negative rate. That is not a technology gap. That is an oversight gap, and it predates generative AI entirely.

THE ACTIVITY

Lab 3. The four systems in production are put on the wall. For each, in pairs: what decision does it take, what decision remains human, whose data feed it, and how would we know if it degraded. Gaps are marked in red — and the red marks are the agenda.

THE OUTPUT

A one-page map of the systems in production with owner, human-oversight point and failure signal — feeds directly into M8.

WHY IT IS HERE

It anchors clinical AI in what the group already does, rather than in what the market is excited about.

M8 · Risk classification workshop

16H30 – 17H30 · 60 MIN

LEARNING OBJECTIVE

Classify the group's own AI systems against the legal framework — and understand the calendar that makes it non-optional.

THE CONCEPT TAUGHT

- **The healthcare-specific high-risk list.** From 2 December 2027, Annex III of the AI Act covers AI for emergency-call evaluation and dispatch, **triage of emergency patients**, assistance to clinicians in **diagnostic and treatment decisions**, and **risk assessment and pricing in life and health insurance**. Three of those four describe systems this group runs or plans.
- **Article 50 comes first.** From 2 August 2026: persons interacting with an AI system must be informed; synthetic content must be marked. That reaches patient-facing chat, triage interfaces and AI-assisted clinical correspondence.
- **Medical device software follows a different track.** Annex I — AI embedded in products already regulated under MDR and IVDR — applies from 2 August 2028. The imaging estate sits here.
- **The sanctioning regime, in three tiers.** Up to €35M or 7% for prohibited practices; €15M or 3% for most obligations including high-risk; €7.5M or 1% for supplying incorrect information.

Regulation as a moat, not a wall. A group that can show a classified inventory, a named owner per system and an auditable oversight record has built something informal competitors cannot cross in a quarter.

THE ACTIVITY

Group work. Building the classification table live: system · purpose · does it touch triage, diagnosis or insurance pricing · classification · owner · what is missing. Where a system cannot be classified with the information in the room, that is recorded as the gap.

THE OUTPUT

The inventory of the group's AI systems with owner and Annex III classification — the single most useful artefact of the three days.

WHY IT IS HERE

It converts Day 2 from understanding into an auditable document. Training that produces only comprehension is not evidence of a programme.

DAY THREE

03

Govern: decide, record, commit

Everything from the first two days converges into one document, written by the board itself.

M9 · Ethics, equity and the patient

09H00 – 10H15 · 75 MIN

LEARNING OBJECTIVE

Use the board's strongest domain as the lever it is — and convert a tradition of clinical ethics into a working instrument for AI.

THE CONCEPT TAUGHT

- **You already speak this grammar.** Clinical ethics committees, informed consent, adverse-event review and equity-of-access commitments are structurally what the AI Act now requires under different names: documented classification, named accountability, human oversight of a consequential decision, auditable record.
- **Algorithmic bias is an equity-of-access event.** Which patient populations are represented in the training data — and which are not — is the healthcare question that turns a technical failure into a clinical, regulatory and reputational one.
- **Disclosure to the patient.** What is communicated, when, and by whom — the Article 50 obligation, met in a way a patient can actually understand.
- **Publishable is the test.** A clinical-AI policy the group would not be comfortable publishing is not a policy. It is an intention.

A commitment to patient dignity that cannot be verified in the infrastructure that holds the patient's data is a statement, not a control. This board articulates the principle better than it audits the system — and that is the asymmetry worth inverting today.

THE ACTIVITY

Drafting session: the one-page clinical-AI ethics policy. Five headings, written in the room, sentence by sentence, with the chair holding the pen.

THE OUTPUT

A publishable one-page clinical-AI policy, ready to be approved in the minutes.

WHY IT IS HERE

Day 3 opens on the board's strongest ground, so the hardest work of the day — the deliberation simulation — lands on confidence rather than on deficit.

Mio · Redesign, not automation

10H30 – 11H45 · 75 MIN

LEARNING OBJECTIVE

Understand that AI value does not come from adding tools but from redesigning the work — and decide where the human stays.

THE CONCEPT TAUGHT

- «AI on old processes produces little. AI on redesigned processes produces transformation.» The adoption funnel: most adopt, few redesign, fewer still capture measurable impact.
- **The 80/15/5 operating model.** 80% automated, 15% human-in-the-loop, 5% senior escalation. The clinician rises in the chain: from operator to curator, investigator and validator.
- **The hybrid organisation chart.** From pyramid to layers — strategic humans, humans plus AI, agents — with emergent roles: agent supervisor, care-experience orchestrator.
- **Adoption metric versus impact metric.** «The tool was used in 4,000 encounters» is not a result. «Time to diagnosis fell and readmissions did not rise» is.

It is not putting AI into the old process. It is accepting that the old process no longer exists. Cutting is not strategy — redesigning is. And redesigning includes deciding where the clinician stays, on purpose, and writing it down.

THE ACTIVITY

Redesign clinic · 45 min. Deliberately not the clinical front, to force the discussion out of the comfortable place: the **revenue cycle** — coding, prior authorisation, denials, collections. Sub-groups classify tasks into the three 80/15/5 bands and name one decision that must never leave a human.

THE OUTPUT

An 80/15/5 map for the revenue cycle plus the list of non-delegable decisions — both enter the Manual.

WHY IT IS HERE

Because redesign fails on culture long before it fails on technology, and in healthcare that failure has a specific name: a clinical workforce that experiences the system as surveillance and routes around it.

MII · The ten questions that go into the minutes

11H45 – 12H30 · 45 MIN

LEARNING OBJECTIVE

Convert three days of concepts into a standing procedure — a set of questions asked always, independently of who presents.

THE CONCEPT TAUGHT

- **Competency materialises as procedure, not as memory.** Ten questions in the board's procedure for any AI initiative.
- **The nine-field minimum dossier.** Purpose in one sentence · risk classification · data used and lawful basis · population coverage of the training data · named accountable owner · human oversight plan · failure signal and deactivation procedure · impact metric with a deadline · vendor dependencies. One page. If management cannot fill it in, the decision is not mature.
- **The tenth question is the one that matters.** «Where is the answer to the previous nine recorded?» It is what makes the other nine auditable.
- **The intended side effect.** These questions discipline the organisation upstream. When teams know a proposal must answer nine fields, the quality of the proposal changes before it reaches the room.

WHY IT IS HERE

It is the mechanism that survives the three days. Everything else is understanding; this is the part that changes what the board requires.

Oversight that leaves no trace does not exist for evidentiary purposes. The record is what distinguishes a diligent board from a board that believes it was diligent.

THE ACTIVITY

Plenary drafting: the board writes its own ten questions, in its own words, on the wall. Then the minimum dossier is agreed as a condition of admissibility.

THE OUTPUT

The ten standing questions and the nine-field minimum dossier, ready for adoption into procedure.

SIMULATION CLIMAX

M12 · Deliberation simulation · «Project Aurora»

13H30 - 15H00 · 90 MIN

LEARNING OBJECTIVE

Exercise the whole competency at once: a real proposal, an incomplete dossier, roles assigned, and a decision the board actually has to take.

THE CONCEPT TAUGHT

- Everything from the first two days, applied under the conditions in which proposals genuinely arrive at a board.

No AI initiative should reach this board without nine fields answered. The purpose of this simulation is to discover how it feels to be asked for them — and how quickly a confident presentation stops being convincing.

THE ACTIVITY

Simulation · 90 min. «Project Aurora» — a fictional but realistic €4.1M proposal: an AI triage and clinical decision-support layer across emergency departments in four hospitals, with a vendor model, a two-year exclusivity clause, and a business case built on length-of-stay reduction. The dossier is deliberately incomplete: no population breakdown of the training data, an owner given as a job title rather than a name, and an impact metric that is an adoption metric. Roles: proposer · value sceptic · risk and governance · clinical lead · patient advocate. Decision rubric provided. Debrief consolidates what the board would now require.

THE OUTPUT

A decision on Aurora, and the refined version of the ten questions after they have been used in anger.

WHY IT IS HERE

It is the climax of the programme. The decision returns to the table — converted into competency and into a document.

ARTEFACT CLOSE

M13 · The Board AI Manual · close and commitment

15H00 – 16H00 · 60 MIN

LEARNING OBJECTIVE

Consolidate everything into one artefact belonging to the board — not a report received, but a manual built — and turn learning into dated action.

THE CONCEPT TAUGHT

- The Manual has six sections: (1) the principles that are not negotiated · (2) the ten standing questions and the minimum dossier · (3) the clinical red lines and non-delegable decisions · (4) the risks to oversee and the quarterly panel · (5) the no-regret moves from the three scenarios · (6) what stays human.

If you forget every number from these three days, keep one sentence: when intelligence becomes free, what stays expensive is trust, the record only we hold, and the direct relationship with the patient. That is what we are here to protect — and to amplify.

THE ACTIVITY

Gallery walk and canvas · 40 min, then **close · 20 min**: the three questions to take home, the urgency calendar (one thing NOW, one this MONTH, one this QUARTER), and commitment cards signed by each director.

THE OUTPUT

The Board AI Manual — the product of the three days, written by the board, ready to be tabled and approved. Plus the twelve-month follow-through: four quarterly cycles and a second measurement in 2027-Q3.

WHY IT IS HERE

It closes the circle opened at the check-in on Day 1 — fears and bets — and hands back agency rather than dependency.

THE METHOD

IV

What is taught, what is produced, what is measured

The frameworks, the three live labs, the artefacts, and how the effect of the programme is measured.

THE TOOLBOX

The frameworks that are taught

All of them part of the teaching method, not borrowed from a commercial proposal. Each one is introduced where it is needed and used again later in the programme.

FRAMEWORK	WHAT IT IS	WHERE IT ENTERS
The three cognitive layers	Predictive (calculator) · generative (intern) · agentic (colleague)	M1
Model versus agent	A tool that answers versus a system that acts and decides	M1
Fluent and wrong	Why probabilistic fluency is not correctness — and what that means clinically	M1 · M7
«What became expensive?»	Clinical trust · the longitudinal record · the care network	M2
The before/after table	Where value in healthcare moved after intelligence became cheap	M2
Disintermediation of the first question	Where a patient now takes a symptom first	M2
Scenario stress test	Probability × impact × horizon · early-warning signals · no-regret moves	M3
The four anterior questions	What a board asks about resilience instead of about incidents	M4

Global cases enter only as proof that something is possible; the substance is always healthcare and always this group.

FRAMEWORK	WHAT IT IS	WHERE IT ENTERS
Data inventory with lawful basis	Ownership, purpose, retention, licensing potential	M5
Adoption funnel	Most adopt, few redesign, fewer capture measurable impact	M10
80/15/5	The AI-native operating model and where the clinician rises	M10
Hybrid organisation chart	Strategic humans / humans + AI / agents	M10
Annex III classification	Triage · clinical decision support · insurance pricing · Article 50	M8
The ten questions + nine-field dossier	The board's standing procedure for any AI initiative	M11
«What AI will not do»	Accountability, meaning, trust, holding an unpopular position	M12 · M13
The urgency calendar	One thing now · one this month · one this quarter	M13

THE THREE LIVE LABS

Where the abstract becomes visceral

LAB 1 • M1 • SEEING THE MACHINE ANSWER

Three real clinical questions put to a public assistant, live. Teaches fluent-and-wrong in the only way it is ever really learned. Fallback: prepared screenshots.

LAB 2 • M5 • BUILDING THE INVENTORY

The data inventory with lawful basis, built in the room with the executive team. Where the answer is unknown, that is the finding.

LAB 3 • M7 • THE SYSTEMS ALREADY RUNNING

The group's four production systems on the wall. For each: what it decides, what stays human, and how we would know it degraded.

All three are learning moments — the board observes, interprets and draws its own conclusion. There is no «and this is what we sell».

AND TWO SIMULATIONS

M6 • Crisis — ransomware on the patient record and the theatre list, with injects every twenty minutes. **M12 • «Project Aurora»** — a €4.1M clinical-AI proposal with a deliberately incomplete dossier, five assigned roles and a decision rubric. Simulations are where competency is revealed rather than declared.

WHAT IS PRODUCED

Nine artefacts that stay in the board's records

Design rule: no module ends without an artefact. Training that produces only unverifiable comprehension is not evidence of a programme — and Article 4 asks for evidence of a programme, not memory of a session.

D1 · M3	Three scenario sheets with early-warning signals and no-regret moves
D2 · M4	The quarterly data and security panel · five fixed indicators
D2 · M5	The data inventory with lawful basis, and the licensable-dataset list
D2 · M6	The crisis decision log and the list of unassigned decisions
D2 · M7	The map of systems in production, with owner and failure signal
D2 · M8	The AI inventory classified against Annex III
D3 · M9	The one-page publishable clinical-AI ethics policy
D3 · M10	The 80/15/5 map and the non-delegable decisions
D3 · M11–13	The ten standing questions, the minimum dossier, and the Board AI Manual

HOW THE EFFECT IS MEASURED

A programme with no declared targets is not assessable

The three days sit inside the twelve-month agenda. The intensive installs the common core; the four quarterly cycles carry it, and the second measurement in 2027-Q3 closes it.

INDICATOR	CYCLE 1 • 2026-Q3	TARGET • 2027-Q3
Assessment coverage	9 / 11	11 / 11
Aggregate average	52	≥ 59
Domains above the threshold of 60	1 / 6	≥ 4 / 6
Lowest domain	44	≥ 55
Range between directors	37 pts	≤ 20 pts

THE TARGET THAT IS NOT IN THE TABLE

None of the numbers above is the real objective. The objective is that when a clinical-AI proposal next comes to this board, **the discussion is led by the board and not by the presentation**. That is not measured on a scale of 100. It is measured in the minutes — in the questions recorded, and in the answers management arrives already having prepared.

A METHODOLOGICAL CAUTION

The second measurement must carry **no individual consequences of any kind**, recorded in the minutes before it is taken. It is the condition on which the instrument stays worth anything in the third cycle and the fourth.

FACILITATION AND PRODUCTION

How the room has to be set up

FACILITATION RULES

- No expository block runs longer than **18 minutes** without interaction.
- Every module ends with an artefact on the wall, not with a slide.
- The chair holds the pen in every drafting session. The facilitator does not write the board's words.
- Nobody is asked to display ignorance. Questions are collected on cards as well as out loud.
- The executive team is in the room for M5, M7 and M8 — and out of it for M6 and M12.

ROOM AND MATERIALS

- **Room:** islands of 3–4 people, free walls for the gallery walk, no fixed rows.
- **Technology:** projection · sound · own hotspot · AI accounts logged in and rehearsed · recorded fallbacks for all three labs.
- **Printed:** service-line cards · the three scenario sheets · the Aurora dossier and rubric · the Manual canvas (A1 plus pocket version) · commitment cards.
- **Before the day:** the preparation reader is sent separately. The programme opens with the check-in in the room — there is no further prior work.
- **Confidentiality:** nothing said in the room is attributed outside it, and no individual assessment result is referenced at any point.

FACILITATION

Who runs the room

Andiara Petterle, PhD

EXTERNAL ADVISER · BOARDACUMEN

Board-level digital competency and governance advisory. The programme is facilitated directly, with the executive team present only in the modules where their answers are the material.

The role in the room is deliberately narrow: teach the concept, apply it to this group, and hand the decision back. Whoever facilitates appears in the footer, not on the stage.

WHAT IS AGREED BEFORE THE FIRST DAY

- The three scenarios are the group's own — confirmed with the executive team, not invented.
- The systems named in M7 and M8 are the systems actually in production.
- The Aurora dossier is calibrated to the format in which proposals genuinely reach this board.
- The assessment results are used in aggregate only. No individual result enters the room.

CLOSE

What the board leaves with

- A shared vocabulary for predictive, generative and agentic AI — and for why fluent is not correct.
- A classified inventory of what the group already runs, with an owner for each system.
- Ten questions that go into the minutes, and a dossier no proposal can skip.
- A publishable ethics policy and a written list of what is never delegated.
- A rehearsed crisis, and the list of decisions that had no owner until Tuesday afternoon.

A hospital group holds the two things clinical artificial intelligence cannot manufacture: the trust of the people who use it, and thirty years of records of what actually happened to them. Neither can be bought. Both can be lost in a single governance failure. These three days exist to make sure that failure is not this board's.

AND THEN

The intensive installs the common core. The twelve-month agenda carries it: four quarterly cycles anchored to the regulatory calendar, and a second measurement in 2027-Q3 that converts this portrait into a movement.