



BoardAcumen
BOARD INTELLIGENCE

BOARD REPORT · DIGITAL COMPETENCY ASSESSMENT

The starting point

Digital competency assessment of the Board of Vitalis Health Group.
What the first measurement says — and what is built from it.



52/100

AGGREGATE
COMPETENCY

6/6

DOMAINS
ASSESSED

II

DIRECTORS
IN SCOPE

49

SCORED
OBSERVATIONS

CYCLE 1 · PERIOD 2026-Q3 · ANALYSIS DATE 29 JULY 2026

ANDIARA PETTERLE, PHD · BOARDACUMEN · ILLUSTRATIVE SAMPLE — NOT A REAL CLIENT

SAMPLE DOCUMENT · FICTIONAL ORGANISATION · ILLUSTRATIVE DATA

How this is organised

This report has four parts. The first establishes the starting point, with the rigour with which it was measured. The second places that starting point in the sector this board operates in. The third — the longest — is the development agenda that follows from it. The fourth proposes three structural decisions that outlive the cycle.

I The starting point

· what was measured, and how to read it

Opening note on the assessment	03
Executive summary — seven conclusions	04
Method, bands and the limits of the measurement	05
The aggregate portrait	06
The six domains, in order	07
The shape of the board	08
Individual dispersion — the real finding	09
Proficiency bands and coverage	10
The Board Competency Index	11

II The context in which this number exists

· sector, peers, regulation

Healthcare in 2026 — the numbers that frame the board's job	12
What comparable boards have already done	13
The coincidence that sets priority: exposure × competency	14
The regulatory calendar 2025–2029	15
The asset this board already has — and is not yet using	16

III The development agenda

· twelve months, six domains, a second measurement

The principle: board competency is not technical competency	17
Cybersecurity and Data Governance — the diagnosis and the curriculum	18
AI and Advanced Analytics — the diagnosis and the curriculum	19
The ten questions that go into the minutes	20
Change and ethics — the domains that decide whether the rest works	21
Foresight and leadership — the levers, not the problems	22
The twelve-month programme	23
How we will know it worked	24
Close, resolutions and sources	25

IV Structural recommendations

· three decisions that exceed the cycle

I · Literacy as a permanent function, not a project	26
II · Board composition and a technology advisory committee	27
III · A strategic plan that reaches the people	28

No director is identified in this document. All results are presented in aggregate, in ranges and distributions — a condition of method, not courtesy: the assessment exists to develop a collegiate body, not to rank people.

Reading the assessment

I will start with the conclusion, because it is the only honest way to open a document like this: **the result of this assessment is good enough to be useful, and low enough to be urgent.** Both things at once.

The board's aggregate is **52 out of 100** — the band the methodology designates as *developing*. It is not a poor result. It is, in fact, a result consistent with self-declaration surveys of boards doing this exercise for the first time, and considerably better than the position of a board that has not yet measured itself — because that board has no number at all. A board that measures itself is, by definition, one cycle ahead of a board that presumes.

But there is a second reading, and it is what justifies this report running to twenty-eight pages rather than two. **The two lowest domains in the measurement — Cybersecurity and Data Governance (44) and AI and Advanced Analytics (46) — are precisely the two domains where this group carries the sharpest verifiable exposure.** The first sits under the strictest special-category data regime in law, under NIS2 obligations that name the management body personally, and in a sector that in the first half of 2026 absorbed an average of 2.3 ransomware attacks per day. The second faces a high-risk classification calendar that reaches triage, clinical decision support and health-insurance pricing from December 2027. This is not unfortunate coincidence: it is the predictable pattern of any body that has followed the digital agenda as a management topic and not as a supervision topic. It is corrected with a programme, not with alarm.

This document does not assess people. It assesses the capacity of a collegiate body to ask the right questions in time — and proposes how to increase it.

Three method decisions I took in preparing this report, and which I ask the board to validate:

- 01 **No names.**
Not in an annex, not in an appendix, not in conversation. The individual assessment belongs to each director and to their own dashboard. What reaches the board is the aggregate, the range and the distribution. A development instrument that converts into an exposure instrument stops producing honest answers in the following cycle — and the second measurement is worth more than the first.
- 02 **Less diagnosis, more curriculum.**
Part III is the largest section of this document by deliberate choice. An assessment report that devotes more space to the problem than to the path is a badly calibrated report.
- 03 **The starting point exists to be surpassed, not defended.**
The usefulness of this 52 is not in its absolute value. It is in the existence of a baseline against which twelve months of work can be measured. Without it, any future progress would be an opinion.

There is one final point, and it is the most important. Vitalis does not enter this agenda as an organisation in difficulty trying to catch up with its sector. It enters as a group holding the two assets that clinical artificial intelligence cannot manufacture: **a thirty-year longitudinal clinical record across nine hospitals and a diagnostics network, and the trust of a patient population and 4,200 affiliated clinicians.** Those assets are worth more than at any point in the group's history, and they are — for the first time — technically transferable, contractually licensable, and irreversibly losable in a single governance failure. The competency this report measures is the competency to capture that value before it is spent. That is why the development agenda has a deadline.

Andiara Petterle, PhD

External adviser responsible for the assessment · BoardAcumen

29 July 2026

Seven conclusions

01 The board sits in the «developing» band, consistently.

The aggregate is 52/100. The six domains spread across a narrow interval, between 44 and 61 — seventeen points. There is no domain of excellence compensating for a critical domain: there is a common plateau. That is operationally easier to treat than an unbalanced profile, because a single programme serves the whole board.

02 One domain of six crosses the proficiency threshold — and it is the one that matters least on its own.

The methodology's proficiency threshold is 60. Only Ethical Governance of Digital Initiatives (61) crosses it. In a healthcare group that is unsurprising and genuinely creditable: this is a board accustomed to clinical ethics committees, informed consent and equity of access. But ethics without cyber and AI competency produces a board that can articulate the principle and cannot audit the system.

03 The dispersion between directors is the most relevant finding — more than the average.

Among directors with a complete assessment, individual results run from **34 to 71**. A 37-point range in a body that decides by collegiate deliberation means the quality of the discussion depends on who speaks — and that there exists, inside the board, capacity sufficient to teach the rest of the board. It is simultaneously the risk and the resource.

04 The two lowest domains are the two of greatest verifiable exposure.

Cybersecurity and Data Governance: 44. AI and Advanced Analytics: 46. Under NIS2 the management body of an essential entity approves cyber-risk measures, oversees their implementation, may be held liable for failure, and is required to follow training. From 2 December 2027, AI used in emergency triage, in clinical diagnosis and treatment support, and in health-insurance risk assessment and pricing is classified high-risk under Annex III of the AI Act. Both are matters of record, not of opinion.

05 Four of eleven directors have no complete assessment on record.

Of the 11 directors, **seven** have a complete profile across the six domains, **two** have a partial assessment and **two** have no record. Completing coverage is the intervention with the highest immediate return — not because it improves the result, but because it makes the result representative of the body that actually deliberates.

06 The composite index is depressed by incompleteness, not only by competency.

The Board Competency Index stands at **43/100** — the «digital awareness» band. Two of the index's six components measure participation and declared composition, not knowledge; and the momentum component is excluded from the calculation by design until a second measurement exists. The index rises with administrative work before it rises with training.

07 The recommendation is a twelve-month programme with a second measurement in 2027-Q3.

Four quarterly cycles, anchored to the regulatory calendar and to the group's own three strategic scenarios; a set of standing questions that pass into the minutes; and a re-measurement that converts this portrait into a movement. **Proposed target: no domain below 55, at least four domains above the threshold of 60, with coverage at 11 of 11.**

Results are directional, human-interpreted and referenced to a version of the instrument. They do not constitute an automated verdict on any person, nor an assessment of individual performance.

What this number is — and what it is not

Before any reading, the board must be able to trust the instrument. This page describes what was measured, on what scale, and where the known limits of the measurement lie.

The six domains

The assessment uses the **BoardAcumen Board Digital Competency Instrument v1.0** and covers six domains. All six were assessed — **6 of 6**. Each domain aggregates a set of sub-competencies, assessed through situational-response items rather than simple self-declaration of knowledge.

D1
Leadership in Digital Strategy

D2
Cybersecurity and Data Governance

D3
AI and Advanced Analytics

D4
Ethical Governance of Digital Initiatives

D5
Technology Foresight and Trend Analysis

D6
Change Management and Digital Culture

The scale and the bands

BAND	INTERVAL	BOARD READING
CRITICAL	0 – 50	The topic is not interrogated with autonomy. The board depends on what it is brought.
DEVELOPING	51 – 70	The topic is recognised and discussed. The capacity to challenge is still irregular.
PROFICIENT / LEADING	71 – 100	The board leads the topic, sets the standard and anticipates the agenda.

Within this scale the methodology fixes **60 as the proficiency threshold** (0 = no competency demonstrated; 100 = leadership practice). It is against this 60 – and not against the midpoint of the scale – that domain results should be read.

Three limits I assume explicitly

- A single measurement window.** Only the 2026-Q3 period exists. Without a second measurement there is no trend — and the momentum component of the composite index is, by design, excluded from the calculation until two periods exist. This report describes a position, not a trajectory.
- Partial coverage.** Nine of eleven directors have an assessment on record; seven have a complete profile. Domain averages are calculated over those who responded — between seven and nine observations per domain. They are indicative of the body, not exhaustive.
- External comparison references.** No scored population of boards exists for direct comparison with this instrument. Every peer comparison here uses identified external sources — cited on page 25 — and they are self-declaration surveys, not equivalent scores.

RECONCILIATION NOTE, FOR THE RECORD

The aggregate of **52** is calculated over the **49 individual domain scores** collected, weighted by observations per domain — not over the simple arithmetic mean of the six domain values, which is 52.5. Both are correct; they answer different questions. Where this report says «the board average», it means the observation-weighted 52. Individual results are as calculated and displayed by each director's own dashboard.

Where the board stands

52/100

AGGREGATE DIGITAL
COMPETENCY44⁻⁶¹

RANGE ACROSS DOMAINS

34⁻⁷¹RANGE ACROSS
DIRECTORS

I/6

DOMAINS ABOVE THE
THRESHOLD OF 60

The average, read precisely

The aggregate value of **52** places the board in the first third of the «developing» band (51–70). It is one point above the top of the critical band. It is a statistically fragile position — small variations in coverage move it — and a clear position in the practical sense: **this board has established digital awareness and does not yet have installed proficiency.**

The distinction is not semantic. Awareness is knowing that a topic exists and that it matters. Proficiency is being able to look at a proposal to deploy an AI triage tool in the emergency department, or at an incident affecting the patient record, and formulate the question the presentation did not bring. It is that second thing which, from December 2027, the law presumes a management body can do.

What holds it up and what pulls it down

Three domains sit above the aggregate: Ethical Governance (61), Leadership in Digital Strategy (58) and Technology Foresight (55). These are domains where top-management experience transfers well — perceiving trends, formulating strategy, weighing reputation and patient equity are competencies this board already exercises in other subjects.

Three domains sit below: Change Management and Digital Culture (51), AI and Advanced Analytics (46) and Cybersecurity and Data Governance (44). These are domains where management experience *does not* transfer on its own, because they demand specific technical vocabulary. **It is precisely at this frontier that training produces the largest gain per hour invested.**

The reading I propose to the board

WHAT IS SOUND

The board measured itself, across all six domains, with an external instrument and a fixed scale. That places it ahead of the two thirds of boards that report limited or non-existent AI knowledge — and therefore do not know where they stand.

WHAT IS CONCERNING

The two lowest domains coincide with the two of greatest exposure: patient data under a regime that names the management body personally, and clinical AI entering a high-risk classification. The window between measuring and being required is quarters, not years.

WHAT IS DECIDED

The difference between this report being useful or decorative rests on a single decision: approve the twelve-month programme and schedule the second measurement. Without a second measurement, this 52 is a curiosity.


COMPARATIVE CONTEXT

According to Deloitte's Global Boardroom Program, **66% of boards report limited or non-existent AI knowledge or experience** — an improvement on the 79% of the previous survey, but still two thirds of the global population surveyed. In a PwC survey of executives, only **32% consider their board to have the right combination of skills and knowledge**. Sector-specifically, an estimated **10–15% of large health-system boards had formal AI oversight structures in place in early 2026**, against roughly half of surveyed healthcare organisations already running generative AI in production. No population of boards scored with this instrument exists — the comparisons on this page are with self-declaration surveys, not with equivalent scores. The pattern is nonetheless consistent: the starting point of Vitalis is not atypical.

Deloitte Global Boardroom Program • Governance of AI • PwC, Annual Corporate Directors Survey • Censinet, AI governance in health systems, 2026

In order, with the threshold visible

The cyan line marks the **proficiency threshold (60)**. One domain reaches it. The correct reading is not «the board fails in five domains» — it is «the board sits at a short and tractable distance from the standard, with one flank already there».




DOMAIN	SCORE	DIST. TO THRESHOLD	OBS.	BAND
Ethical Governance of Digital Initiatives	61	+1	7	DEVELOPING
Leadership in Digital Strategy	58	-2	8	DEVELOPING
Technology Foresight and Trend Analysis	55	-5	8	DEVELOPING
Change Management and Digital Culture	51	-9	8	DEVELOPING
AI and Advanced Analytics	46	-14	9	CRITICAL
Cybersecurity and Data Governance	44	-16	9	CRITICAL
Board average	52	-8	49	DEVELOPING

Obs. = number of individual scores per domain. The board average is calculated over the 49 individual scores and not over the mean of the six domains — hence the difference against the simple arithmetic mean of the column (52.5). Source: BoardAcumen · Board Overview and Director × Domain Competency Heatmap · 29 July 2026.

The structural observation. The range between the highest and the lowest domain is seventeen points. In a board with an unbalanced profile — say a 78 in strategy and a 35 in cybersecurity — the problem would be one of composition and would be solved by recruiting. Here the problem is one of **common plateau**, and it is solved with a single programme for the whole body. That is cheaper, faster, and does not require touching the composition of the board.

The shape of the board



■ BOARD AVERAGE □ THRESHOLD 60

Rings: 25 • 50 • 75 • 100. 0 = no competency demonstrated • 100 = leadership practice.

A hexagon that is almost regular

The shape matters as much as the area. A profile with peaks and troughs indicates a board with isolated specialists — strong where somebody knows, blind where nobody does. The Vitalis profile is **close to regular**: all vertices fall within a narrow band, and only one crosses the dashed threshold line.

This has three practical implications.

- **There is no dependency on a single person.** No domain is being propped up by a single director with exceptional knowledge. The key-person risk in digital oversight is low.
- **There is also nobody pulling the board up.** The absence of peaks means there is no director today to whom the others default on AI or on data. That role is vacant — and it is assignable.
- **One common curriculum serves everyone.** With seventeen points of range across domains, the same programme produces gain across the whole board. There is no need to design six distinct tracks.

WHERE THE SHAPE DISTORTS

The only relevant asymmetry lies on the D4–D2 axis: Ethical Governance (61) sits seventeen points above Cybersecurity and Data Governance (44).

The board articulates the principle better than it audits the system. For a healthcare group whose ethical commitments are entirely dependent on the integrity and confidentiality of the clinical record, that asymmetry is exactly the one worth inverting. A commitment to patient dignity that cannot be verified in the infrastructure that holds the patient's data is a statement, not a control.

IF THE SHAPE STAYS REGULAR

A regular hexagon that rises uniformly is the best outcome of a common-core programme: the whole body crosses the threshold together and no domain becomes dependent on one person. That is the target profile for 2027-Q3.

IF ONE VERTEX SPIKES

A single sharp peak in the second measurement would mean the programme trained the interested rather than the board. It would raise the average and widen the 37-point range — the opposite of the objective.

IF D2 STAYS WHERE IT IS


Cybersecurity and Data Governance is the vertex that carries a legal obligation attached to the management body itself. It is the one place where an unchanged score is not a matter of pace — it is an exposure the board has knowingly left open.

Directional results, human-interpreted, referenced to a version of the instrument. They do not constitute a verdict on any director.

The finding the average hides

The board average is 52. Individual results among directors with a complete assessment run from **34 to 71**. It is this range — not the average — that determines the quality of a collegiate deliberation.

DENSITY OF INDIVIDUAL RESULTS • SEVEN COMPLETE PROFILES • NO IDENTIFICATION, NO UNIT VALUES



Illustrative density curve, with no unit values plotted: three of the seven complete profiles fall in the critical band (≤ 50), three in the developing band (51–70) and one reaches the proficient band (71–100). Two directors with a partial assessment (four and five domains) have no overall result and do not appear in this distribution. Source: BoardAcumen • Score Distribution • 29 July 2026.

Why the range matters more

A board does not decide by the average of its members. It decides by the discussion it can sustain — and that discussion is bounded, on each topic, by who has the vocabulary to lead it and by who does not have the vocabulary to contest it.

With **37 points of range**, digital oversight at Vitalis is today unequal by construction: there are directors in a position to interrogate a clinical-AI deployment proposal with technical autonomy, and there are directors who depend more on management's confidence. Neither position is censurable individually — it is the normal portrait of a body that had never measured this competency. **It is precisely that inequality which the common core of the Part III programme was designed to close.**

The resource the range reveals

The inverse reading is the more useful one. **There exists, inside the board, capacity sufficient to raise the board.** One complete profile sits above the proficiency threshold, at 71; the median sits roughly nine points below the threshold.

This opens an option that requires no external budget: **designate a director as owner of the digital and AI agenda** — not as technical specialist, but as the person accountable for ensuring the question is asked and recorded. It is the practice NACD recommends explicitly: attribute clear ownership of AI oversight rather than leaving it diffuse across the body.

NACD • Board AI Governance Guide • The Conference Board, Governing AI in the Corporation


THE IMPLICATION FOR THE PROGRAMME

A programme designed for the average would fail at both ends: trivial for the top and inaccessible for the base. The Part III proposal is therefore built on **two storeys** — a mandatory common core for the whole board, which installs vocabulary and standing questions, and a deeper module for whoever takes ownership of the AI and data agenda.

Who was measured, and in which band they landed

Distribution by proficiency band

Applying the scale fixed on page 05, of the seven directors with a complete profile, **three sit in the critical band (≤ 50), three in the developing band (51–70) and one reaches the proficient band (71–100).**



Against the fixed three-band scale of page 05. Source: BoardAcumen · Score Distribution, recalculated against the threshold of 60 · 29 July 2026.

Coverage of the measurement

The representativeness of the portrait depends on how many members of the body appear in it. Today, four of eleven directors do not have a complete profile.

ASSESSMENT STATUS	N	%
Complete profile · six domains	7	64%
Partial assessment · four and five domains	2	18%
No assessment on record	2	18%
Board	11	100%



HIGHEST IMMEDIATE RETURN

Complete coverage to **11 of 11** before the end of 2026-Q4. It does not improve the average — it may even lower it — but it converts a partial portrait into a portrait of the body, and it moves two of the six components of the composite index (page 11), which measure participation and composition rather than knowledge.

READING 01

One director reached the proficient band. In a healthcare group of this scale that is a genuine asset and an atypical one — and it is the single most exploitable fact in this report. Ownership of the agenda has a natural candidate.

READING 02

Three of seven complete profiles sit in the critical band. That is what makes the common core mandatory rather than optional: a programme that only the interested attend widens the range instead of closing it.

READING 03

Two directors with no record is a governance signal, not a personal one. It indicates that digital onboarding of directors does not exist as a process — a correction of procedure, not of people.

The Board Competency Index: 43

The composite index stands at **43/100** — the «digital awareness» band. It is the lowest number in this report, and also the easiest to misread. It is worth taking apart, because **a quarter of what it measures is not knowledge** — and a further 15% is, by design, outside the calculation while only one measurement exists.

The six components

COMPONENT	WEIGHT	MEASURES
Competency coverage	30%	Mean of directors' overall results
Competency distribution	20%	Dispersion of coverage by domain
Participation rate	15%	Assessed over total members
Momentum	15%	EXCLUDED requires 2nd measurement
Critical gap	10%	Severity of domains below 50
Composition	10%	Diversity of declared experience and roles

With the momentum component excluded, its 15% weight is redistributed proportionally across the remaining five. BoardAcumen methodology · BCI v1.

The index bands

INTERVAL	BAND
85 – 100	Digital leader
70 – 84	Digital maturity
55 – 69	Developing digital capability
40 – 54	Digital awareness VITALIS · 43
0 – 39	Digital readiness deficit

WHAT THIS MEANS PRECISELY

Vitalis sits in the **lower half of the «digital awareness» band** — twelve points below the next band. But the components pulling it down are not all competency: participation and composition together account for **25% of the index**, and both depend on administrative records — how many directors completed the assessment, and how many filled in their experience profile on the platform.

Three reasons this index rises before the board learns anything

- 01

Coverage closes.

Participation is a simple ratio of assessed directors to total members. Taking the board view from 9/11 to 11/11 is an estimated gain of roughly five points — not a full correction.
- 02

The experience profile fills in.

Composition assesses the diversity of experience and roles declared. Unfilled profiles count as absence of experience — understating a board that does hold relevant clinical, regulatory and operational experience, simply not on record.
- 03

Momentum starts to exist.

While only one measurement exists, 15% of the index sits outside the calculation. The second measurement does not merely measure progress: **it unlocks an entire component.**

PERIMETER NOTE — FOR THE BOARD RECORD

The index is calculated over the platform's organisational register — **26 director and chair accounts** across the group's subsidiary boards, of which 14 have a completed assessment — while this report refers to the parent board view of **11 directors and 49 scored observations**. The perimeters do not coincide. The index is therefore presented as *directional context*, not as a measure of the body, and reconciling the register is an administrative item in cycle C1. All domain averages and the aggregate of 52 refer, without ambiguity, to the perimeter of 11.

Indices and scores are directional, human-interpreted and referenced to a version of the instrument. They do not replace the board's judgement — they inform it.

The sector in which this 52 has to work

A board's competency is not assessed in the abstract. It is assessed against the difficulty of the decisions that board has in front of it. These are the numbers that frame the job.

2.3/day

RANSOMWARE ATTACKS ON
HEALTHCARE, H1 2026

410

ATTACKS IN H1 2026
▲14% VS H2 2025

50%

HEALTHCARE ORGS
RUNNING GENERATIVE AI

10–15%

HEALTH-SYSTEM BOARDS
WITH FORMAL AI
OVERSIGHT

The asset is appreciating and being spent at the same time

Two things happened to healthcare data between 2024 and 2026. It became the scarcest input in clinical AI — models are abundant, longitudinal outcome-linked records are not — and it became the most actively targeted asset in the ransomware economy. **In the first half of 2026 the sector absorbed 410 ransomware attacks, an increase of roughly 14% on the previous half-year, of which 247 hit hospitals, clinics and direct care providers.**

The link to this assessment is direct. The domain that protects that asset — Cybersecurity and Data Governance — is the lowest of the six, at 44, sixteen points below the threshold. It is not a coincidence to be noted. It is the finding.

The gap has moved from adoption to oversight

Roughly half of surveyed healthcare organisations had generative AI in production use by April 2026 — in clinical documentation, coding, triage support, imaging and patient communication. In the same window, an estimated **10–15% of large health-system boards had a formal AI oversight structure in place.**

That is the widest gap in this report, and it is a governance gap rather than a technology gap. The organisation is deploying faster than the body supervising it can follow. The consequence is not theoretical: **AI-related malpractice claims rose 14% between 2022 and 2024**, before the current deployment wave had matured.

Where the risk actually enters: the third party

74%

of hospital and health-system security leaders name EHR, AI and cloud vendors as their top emerging cyber risk.

63%

reported a vendor-linked security incident in the preceding 24 months.

THE BOARD READING

67% experienced clinical disruption from a vendor outage or cyber event. In a healthcare group, third-party risk is not a procurement matter. It is a continuity-of-care matter — and continuity of care is a board matter.

There is also new revenue here, and it is board business. Longitudinal clinical data — de-identified, consented, governed — has become a licensable asset for research partnerships, real-world-evidence programmes and algorithm validation. The European Health Data Space Regulation, in force since March 2025, builds the legal architecture for exactly that secondary use. **Vitalis holds thirty years of it.** Whoever negotiates those arrangements needs a board capable of assessing them: what is being licensed, on what lawful basis, with what re-identification risk, for what price, and with what right to withdraw. *That is not a compliance question. It is a capital-allocation question with a compliance floor.*

Comparitech • Healthcare Ransomware Roundup, H1 2026 • Censinet • AI governance in health systems, 2026 • Regulation (EU) 2025/327 (European Health Data Space). Figures for Vitalis itself are illustrative and fictional; sector figures are as published by the sources named.

What the boards that have already decided are doing

The legitimate question from any director on seeing a 52 is: *and everyone else?* No population of boards scored with this instrument exists for direct comparison — but sector self-declaration surveys show a consistent pattern. On declared competency, everyone else is also poorly prepared. On **governance structure**, the market moved very fast in the last year — and it is there that Vitalis is genuinely behind.

WHERE THE MARKET IS ALSO WEAK

66% of boards report AI knowledge or experience that is limited or non-existent — an improvement on the 79% of the previous survey.

32% is the proportion of executives who consider their board to have the right combination of skills.

2.7% of S&P 500 directors have identified AI competency — against 1.5% in 2021. Four years, not one.

Deloitte Global Boardroom Program • PwC, Annual Corporate Directors Survey • EY, AI governance disclosure analysis

WHERE THE MARKET HAS ALREADY ORGANISED

83% of S&P 500 companies disclose AI as a risk factor — 12% in 2023.

20% of S&P 500 companies assigned AI oversight to a dedicated board committee — 6% the year before.

62% of directors surveyed already reserve agenda time to discuss AI.

EY • The Conference Board, Governing AI in the Corporation • NACD, Director Essentials: Implementing AI Governance

The conclusion that matters

The international standard has stopped being *having AI competency on the board* — almost nobody has that. The standard has become **having ownership, agenda and method**: somebody accountable, time reserved, and questions that appear in the minutes. Three things that do not require hiring a specialist, and that are within this board's reach inside a quarter.

OWNERSHIP

NACD recommends that AI oversight be **integrated into existing risk governance**, with ownership explicitly attributed, rather than treated as a separate function. The dominant trend is not to create new committees — it is to clarify who answers.

COMPETENCY ASSESSMENT

The same guidance expressly recommends conducting **AI competency assessments by director**. Vitalis has already complied with that recommendation — this report is the result. Few boards, in any geography, can say the same.

CONTINUING EDUCATION

Deloitte recommends integrating AI training into new-director onboarding and maintaining it as a continuing programme, with the same seriousness as financial or legal literacy. Certifications dedicated to AI oversight for directors already exist in the market.

THE CREDIT THIS BOARD ALREADY HAS

It is worth stating explicitly, because an assessment report tends always to sound like a deficit: **measuring is the advanced practice**. While two thirds of boards report having no AI competency — and therefore do not know where they stand — this board has six scored domains, a known range and identified coverage. It has moved from opinion to data. It is from that point, and only from that point, that an auditable programme can be built.

The coincidence that determines priority

If all six domains were equally exposed, the order of work would be a matter of indifference. They are not. Crossing each domain's score with the group's concrete exposure in that domain, priority stops being a choice and becomes a reading.

DOMAIN	SCORE	CONCRETE EXPOSURE AT VITALIS	PRIORITY
Cybersecurity and Data Governance	44	Special-category health data for a 1.2M-life plan and thirty years of clinical records across nine hospitals; NIS2 obligations as an essential entity, in which the management body approves cyber-risk measures, oversees implementation and may be held liable; a sector under 2.3 ransomware attacks a day.	CRITICAL
AI and Advanced Analytics	46	Systems already in production — imaging triage support, clinical documentation, no-show prediction, plan risk scoring. From 2 December 2027, emergency triage, clinical diagnosis and treatment support, and health-insurance risk assessment and pricing are Annex III high-risk.	CRITICAL
Change Management and Digital Culture	51	18,000 employees and 4,200 affiliated clinicians. Clinical adoption is not a training problem — it is a problem of workflow redesign and of professional autonomy. AI value does not come from the tool; it comes from the redesign, and the redesign fails on culture before it fails on technology.	HIGH
Ethical Governance of Digital Initiatives	61	Patient trust is the group's central asset and the most fragile. Algorithmic bias in triage or in plan pricing is an equity-of-access event with direct regulatory and reputational effect. The board's strongest domain — and the one with the least margin for error.	HIGH
Leadership in Digital Strategy	58	The approved capital plan concentrates digital investment in the plan and in diagnostics while the hospital estate carries the maintenance burden. That is a capital-allocation decision the board supervises.	MEDIUM
Technology Foresight and Trend Analysis	55	The three scenarios already identified in-house — care migrating out of the hospital, the payer-provider boundary being redrawn by risk models, and clinical data as a licensable asset — require anticipation, not reaction.	MEDIUM

THE SENTENCE THAT SUMMARISES THIS PAGE

The two domains in which this board scores lowest are exactly the two in which, within eighteen months, it acquires verifiable obligations. The order of the programme is not a preference — it is the calendar.

WHY THIS IS NOT AN ALARM

The obligation does not require technical directors. It requires **informed human oversight, auditable records and documentation** — and that the organisation supports the development of AI literacy among those who operate these systems. These are duties of means, dischargeable with a programme.

WHY IT IS ALSO NOT OPTIONAL

Under NIS2, training of the management body is not a recommendation — it is an express requirement placed on members of that body. The question is not *whether* someone asks. It is what the board can show it did, when they do.

What is already in force and what enters now

DATE	MILESTONE	RELEVANCE TO THE BOARD
2 Feb 2025	AI Act prohibitions enter into force, together with the AI literacy duty (Art. 4) for providers and deployers.	Already applicable. Covers every person operating AI systems in the group.
26 Mar 2025	The European Health Data Space Regulation (EU) 2025/327 enters into force, staged to 2035.	Legal architecture for primary and secondary use of the group's clinical data.
2 Aug 2025	Obligations for general-purpose AI models apply.	Relevant to contracted providers, not to the group directly.
In force	NIS2 — healthcare providers as essential entities. The management body approves cyber-risk management measures, oversees their implementation, may be held liable for failure, and its members are required to follow training .	The single strongest legal anchor for this report's recommendation.
2 Aug 2026	Article 50 transparency obligations apply. Persons interacting with an AI system must be informed; synthetic content must be machine-readably marked and, where it informs the public on matters of public interest, disclosed. Directly reaches patient-facing chat and triage interfaces and AI-assisted clinical correspondence.	DIRECT IMPACT
26 Mar 2027	EHDS general date of application.	Reference date for the data-access and secondary-use framework.
2 Dec 2027	Annex III stand-alone high-risk systems apply. Includes AI for emergency-call evaluation and dispatch, triage of emergency patients , assistance to clinical professionals in diagnostic and treatment decisions , and risk assessment and pricing in life and health insurance . Date per the 2026 simplification package.	DIRECT IMPACT
2 Aug 2028	Annex I: high-risk AI embedded in products already regulated under sectoral legislation — including medical device software under MDR and IVDR .	Planning horizon for the imaging and diagnostics estate.
26 Mar 2029	EHDS: exchange of the first priority categories of health data — patient summaries and ePrescriptions — becomes required. Imaging, laboratory results and discharge reports follow by March 2031.	Interoperability horizon; shapes the data architecture decided now.

THE SANCTIONING REGIME, IN THREE TIERS

Up to **€35M or 7%** of worldwide annual turnover for prohibited practices; up to **€15M or 3%** for the generality of the remaining obligations, including high-risk; up to **€7.5M or 1%** for supplying incorrect or misleading information to authorities. NIS2 sanctions and personal accountability of the management body run in parallel, on their own basis.

THE MINUTE QUESTION FOR 2 AUGUST 2026

«Which of our patient-facing interfaces are AI systems, who is told what and when, which clinical correspondence is AI-assisted, and where is that recorded?» **Without a written answer there is no compliance** — and the answer is not obtained in a meeting. It is obtained from an inventory.

Regulation (EU) 2024/1689 (AI Act), Articles 4, 6(2), 50, 99 and Annex III · Directive (EU) 2022/2555 (NIS2), Article 20 · Regulation (EU) 2025/327 (European Health Data Space) · Regulation (EU) 2017/745 (MDR) and 2017/746 (IVDR). Dates reflect the position at the analysis date of this sample and should be re-verified against current sources before use.

What this board already has — and is not yet using

This assessment is not an isolated exercise. Three of the four anchors of the strategic agenda the executive team has already opened correspond directly to pairs of measured domains; the fourth cuts across all of them.

ANCHOR 01 • THE TRUST ASSET

Patient trust and clinical reputation are what allow a group to hold sensitive data at all. They are governed by **Ethical Governance (61)** and **Cybersecurity and Data Governance (44)** — the board's highest and lowest domains. The group holds the asset and does not yet hold oversight to match it.

ANCHOR 02 • THE LONGITUDINAL RECORD

Thirty years of outcome-linked clinical records, across nine hospitals and 140 diagnostic units, is the input clinical AI cannot manufacture. Deciding what is licensed, to whom, on what basis and at what price sits with **Leadership in Digital Strategy (58)** and **AI and Advanced Analytics (46)**.

ANCHOR 03 • REDESIGN, NOT AUTOMATION

AI value does not come from the tool; it comes from redesigning the work — in clinical documentation, in the revenue cycle, in scheduling, in prior authorisation. That is the territory of **Change Management and Digital Culture (51)**. It is the work that has not yet started.

ANCHOR 04 • GOVERN AS ADVANTAGE

The group does not start from zero: traditional analytics are in production. The deficit is in generative and agentic AI and in its **governance**. That is a decision-architecture problem, not a courage problem — and it is solved with method.

The artefact that already exists, and what it lacks

The executive team has produced a **Clinical AI Deployment Standard** — written in-house, covering validation, clinical sign-off and post-deployment monitoring. That document resolves the executive side: how it is decided, what is asked, what is not delegated below the medical directorate. **What this assessment shows is that the board side is missing** — the set of questions the supervising body asks systematically, independently of who is presenting, and which is recorded in the minutes. That is the object of Part III.

WHY THE CLINICAL-ETHICS TRADITION IS THE LEVER, NOT THE OBSTACLE

This board has, for decades, run clinical ethics committees, informed-consent frameworks, adverse-event review and equity-of-access commitments. Every one of those is structurally the same instrument the AI Act now requires under a different name: a documented risk classification, a named accountable person, human oversight of a consequential decision, and an auditable record of why it was permitted. **Vitalis does not need to learn a new governance grammar. It needs to apply the grammar it already speaks to a class of system it has not yet treated as clinical.** That is a shorter distance than 44 and 46 suggest — and it is the fastest path in this report.

CONTINUITY, IN ONE LINE

The executive team has written how it decides on clinical AI. This report measures the board's capacity to supervise that decision — and proposes the twelve months that close the difference between the two.

Board competency is not technical competency

The whole programme that follows rests on a distinction I ask the board to adopt as doctrine: **a director is not asked to know how to build an 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 months, not years.

WHAT THE PROGRAMME IS NOT FOR

- Training directors in data science or programming.
- Producing technical opinions on architectures or models.
- Substituting for the competency of 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**, in any circumstance.
- Knowing **what has to be written down**, and where.

The four capacities, applied

CAPACITY	WHAT IT MEANS IN PRACTICE	WHERE IT IS VERIFIED
Knowing how to ask	Faced with a clinical-AI proposal, the board identifies what is missing: which data feed it, which patient cohorts are represented and which are not, who answers for an error, what happens when it fails, how it is switched off.	Standing questions in the minutes
Knowing how to judge the answer	Distinguishing an adoption metric («the tool was used in 4,000 encounters») from an impact metric («time to diagnosis fell, and readmissions did not rise»). It is the difference between adopting and redesigning.	Decision dossier
Knowing the non-delegable	Some decisions do not descend: what is published in the group's name, what is done with patient and member data, what is licensed and to whom, and where a clinician's judgement remains sovereign over a model's output.	Map of reserved decisions
Knowing how to record	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.	Minutes and AI inventory

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**: to support the development of a sufficient level of literacy in those who operate or use these systems, proportionate to context and risk. It does not require diplomas. It requires a programme — adequate, documented and actually executed. NIS2, in turn, **requires members of the management body to follow training** on cyber risk and encourages the same for employees. *A board development plan, approved in the minutes and measured at the end, is precisely the proof of both.*

DOMAIN 2 · CRITICAL PRIORITY · LOWEST IN THE MEASUREMENT

Cybersecurity and Data Governance

44

-16 TO THRESHOLD · 9 OBS.

This is the board's lowest domain and, at the same time, the one that protects the asset on which everything else in this group depends: **the confidentiality and integrity of the clinical record**. The combination is the reason this domain cannot wait for year two of the programme.

Working hypotheses to confirm in the programme

The score of 44 measures the board's *supervision* competency — not the state of the organisation's security. The hypotheses below are what the C3 cycle (p. 23) was designed to test.

- **Risk of security surfacing only as an incident report.** When that happens, the body sits as retrospective auditor. Effective oversight is anterior: on resilience, on critical dependencies, on what happens to care delivery if a system falls.
- **Ownership inventory to confirm.** Nine hospitals, a diagnostics network and a health plan accumulate data in silos under different lawful bases. Session S2 builds the data inventory with lawful basis, with the executive team — producing it if no consolidated version exists.
- **The link between data and AI is what the programme joins.** Every AI system is, before anything else, a consumer of data. That is why this domain precedes D3 in the programme sequence (p. 22).

The curriculum · three sessions and a standing exercise

- S1 **What a board supervises in cybersecurity.** Moving from the incident list to resilience: continuity of care, critical dependencies, tested recovery, third parties.
90 min · plenary
- S2 **Data governance as an asset.** Inventory, ownership, lawful basis, retention — and which datasets are licensable under EHDS and which are not.
90 min
- S3 **Crisis simulation.** Tabletop on a ransomware event that simultaneously affects the patient record and elective scheduling across two hospitals. The board decides under time pressure, with incomplete information.
120 min
- Ex **Quarterly data and security panel** — five fixed indicators, always the same, so that variation is legible over time.
permanent

The five fixed questions of the quarterly panel

01

Which systems are indispensable to delivering care tomorrow morning, and what is the tested recovery time of each?

02

What personal and special-category data do we process, on what lawful basis, and how many data-subject requests did we receive and meet in time?

03

Which third parties have access to our data or systems, and when was that access last reviewed?

04

What incidents occurred, how long did detection take, and what changed in the process because of them?

05

Which datasets could be lawfully licensed or contributed to research, and what is missing for that to be possible?

NOTE

The fifth question is deliberately built in attack mode, not only defence. Data governance that only serves to avoid fines leaves half the value on the table.

The calibration point: **the objective of the first year is not to reach excellence in cybersecurity**. It is to move from 44 to above 55 — which means, concretely, that any director can read a security panel, identify what is missing in it, and know what the next question is. That is achievable in three sessions and one standing quarterly panel.

DOMAIN 3 · CRITICAL PRIORITY

AI and Advanced Analytics

46

-14 TO THRESHOLD · 9 OBS.

This is the central domain of the programme. Not by fashion: by arithmetic. It is the second lowest in the measurement and the one of greatest regulatory density from December 2027 — and the group has **systems already in production** whose formal governance this assessment, by design, does not measure directly.

Working hypotheses to confirm

- **Possibly asymmetric vocabulary.** A competency of 46 is consistent with a board that discusses AI mainly in the vocabulary of automation — efficiency, throughput, staffing substitution — and less in the vocabulary of decision under uncertainty: training data, cohort representativeness, calibration drift, hallucination, human oversight. Session S1 confirms or dismisses this.
- **Risk of conflating the layers.** The analytics the group already runs — imaging triage support, no-show prediction, plan risk scoring — is different in nature from generative and agentic AI. Treating them as one topic tends to produce misaligned decisions: underestimating the risk of the first and overestimating the novelty of the second.
- **Inventory to confirm.** Session S3 builds, with the executive team, the inventory of AI systems in use — with purpose, accountable owner, clinical sign-off status and risk classification — producing it in formal terms if it does not yet exist.

The curriculum · four sessions

- S1 The three layers of AI and what changes in each.** Predictive, generative, agentic: what each does, where it fails, and what question each demands.
90 min · plenary
- S2 Seeing the machine answer.** Live demonstration on the group's own terrain: what a model does with an ambiguous presenting complaint, where it is confident and wrong, and how that appears — or fails to appear — to a clinician.
90 min · plenary with exercise
- S3 From inventory to risk classification.** Building, in session, the inventory of AI systems in the group and classifying each against Annex III. Produces an artefact, not a summary.
120 min · group work
- S4 Deliberating a real proposal.** Simulation of an AI investment decision with an incomplete dossier — in the format in which proposals actually reach the board.
120 min · simulation

Design rule: no session ends without an artefact that stays in the board's file. 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.

THE OPPORTUNITY THIS DOMAIN OPENS — AND IT IS NOT DEFENSIVE

It is worth recording that work in this domain does not serve only to comply. It serves to decide on **the secondary use of the clinical record** — research partnerships, real-world-evidence programmes, algorithm validation — a market the EHDS is explicitly building the legal architecture for; and on **risk models in the health plan**, where the boundary between actuarial precision and unlawful discrimination is exactly the line the board is accountable for holding. Neither of those is bought off a shelf. The decisions about what to do with them are the most valuable this board will take in the next three years — and they demand precisely the competency this domain measures.

The ten questions that go into the minutes

A board's competency materialises in a set of questions that are asked **always** — independently of who presents, of how convincing the presentation is, and of how far advanced the project already is. I propose that these ten enter the board's standing procedure for any initiative involving AI.

- 01 **What decision does this system take, and what decision remains human?**
If the answer does not distinguish the two, the system is not governed.
- 02 **What data feed it, whose are they, and on what lawful basis do we use them?**
Applies to patient records, plan-member data, imaging and administrative claims alike.
- 03 **Which patient populations are represented in the training data — and which are not?**
In healthcare this is the equity question, and it is the one that turns a technical failure into a clinical and regulatory event.
- 04 **When this fails, how do we know — and in how long?**
Detection before correction. A system with no failure signal is a system without oversight.
- 05 **Who is the person — a name, not a function — accountable for this system?**
Diffuse ownership is absence of ownership.
- 06 **Where does a clinician's judgement remain sovereign over the model's output, and is that written down?**
The non-delegable, stated as a list rather than as a principle.
- 07 **What is communicated to the patient, and how do they know?**
Obligation from 2 August 2026 for AI interaction and for synthetic content.
- 08 **How is it switched off, who has authority to switch it off, and has that been tested?**
The existence of the switch is not enough: it has to have been exercised.
- 09 **What impact metric — not adoption metric — will we see within two quarters?**
«The tool was used» is not a result. It distinguishes adopting from redesigning.
- 010 **Where is the answer to the previous nine questions recorded?**
The tenth is the one that makes the other nine auditable.

THE MINIMUM DOSSIER

No AI initiative reaches the board without: 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. Nine fields.** If management cannot fill it in, the decision is not mature.

THE INTENDED SIDE EFFECT

These questions do not discipline only the board — they discipline the organisation upstream. When teams know a proposal will have to answer nine fields, **the quality of the proposal changes before it reaches the room.** It is the cheapest mechanism available for raising a group's AI maturity: change what the board requires.

Change and ethics — the domains that decide whether the rest works

DOMAIN 6 • HIGH PRIORITY

Change Management and Digital Culture

This is the domain that determines whether everything else produces effect. AI value does not come from the tool — it comes from redesigning the work. And redesign fails on culture long before it fails on technology. In healthcare that failure has a specific name: **a clinical workforce that experiences the system as surveillance rather than as support, and routes around it.**

What the board supervises here

- Whether transformation reaches **all four fronts** — clinical, revenue cycle, plan operations and support functions — or only the visible ones.
- Whether there is an explicit distinction between work that is automated, work that is augmented, and work that remains integrally human.
- Whether clinician resistance is treated as an information signal or as an obstacle to route around. The first reading produces better decisions.
- Whether new competencies are being built or bought — and what that implies for retention in a market with structural clinical shortage.

Dedicated session: a redesign clinic on one real front — not the clinical front, to force the discussion out of the comfortable place. I propose the **revenue cycle**: it is measurable, it is large, and it is where effect on the result appears fastest.

5I

DOMAIN 4 • HIGH PRIORITY • HIGHEST IN THE MEASUREMENT

Ethical Governance in Digital Initiatives

In a healthcare group, digital ethics is not a chapter of social responsibility — it is **management of the principal asset**. The trust patients and clinicians place in this group is the accumulated result of decades of clinical decisions. Algorithmic bias in triage or in plan pricing can destroy part of that capital in a single episode, at scale, and in minutes.

What the board supervises here

- Whether a **public policy on clinical AI use** exists — and whether the group would be comfortable publishing it.
- How the use of AI is disclosed to patients, and whether that disclosure meets Article 50 from 2 August 2026.
- Which clinical decisions are **non-delegable to a system**, written as a list rather than as a principle.
- Who answers publicly when an AI-assisted clinical process fails — decided before it fails.

ARTEFACT

A one-page clinical-AI ethics policy, publishable, approved by the board. Publishable is the test: a policy the group would not be comfortable making public is not a policy — it is an intention.

6I

WHY THESE TWO DOMAINS TRAVEL TOGETHER

Change without ethics produces speed without direction — and in a business built on clinical trust, speed in the wrong direction is the worst available outcome. Ethics without change produces the opposite: a group that is unimpeachable and irrelevant. **Vitalis can afford neither**, and that is why these two domains are worked in the same cycle of the programme, with the same closing session.

Foresight and leadership — the levers, not the problems

These are two of the board's stronger domains. The temptation is to leave them out of the programme. That would be a mistake: **it is precisely because they are higher that they serve as levers for the other four.** They are worked not to correct them, but to use them.

DOMAIN 5

55

Technology Foresight and Trend Analysis

Five points from the threshold, and the domain the board can cross first — and the first crossing matters, because it demonstrates to the body that the programme produces results.

How it is used

- Anchor foresight in the group's **three in-house scenarios** rather than in generic trends. A board that debates its own scenarios develops competency faster than a board that watches market presentations.
- Institute a **fixed horizon slot** in every meeting: fifteen minutes, one external signal, one implication for the group. Fifteen minutes, twelve times a year, is three hours of structured foresight — more than most boards manage.
- Use the scenario of **care migrating out of the hospital** as the worked case: it is the one where the group most explicitly has to choose a side, and therefore the one that most demands anticipation.

THE SEQUENCE I PROPOSE

Begin with the strong domains in the first cycle — foresight and leadership — so that the board **crosses the threshold in at least one domain before the end of 2026**. The morale effect in a twelve-month programme is not a detail: it is what sustains attendance through the hard cycles.

THEN THE CRITICAL ONES

Second and third cycles on data and cyber, then on AI — the two domains of regulatory exposure. This is where the greatest effort sits, and the greatest risk avoided.

FINALLY, CONSOLIDATION

Fourth cycle on change and ethics, integrating everything prior into a single artefact — the clinical-AI policy and the map of non-delegable decisions — immediately before the second measurement.

DOMAIN 1 • GREATEST CARRY-OVER EFFECT

58

Leadership in Digital Strategy

With eight observations, this is one of the best-covered domains — and the one with the greatest downstream effect on the others. A well-supervised capital allocation corrects, downstream, what no technical training corrects.

How it is used

- Put to the board the question the capital plan itself raises: **is the concentration of digital investment in the plan and in diagnostics, while the hospital estate carries maintenance, a deliberate choice or a consequence of allocation?** A competent board requires that answer in writing.
- Distinguish, in every initiative, **digital as a channel** from **digital as a care model**. Conflating the two is the most expensive strategic error in the sector.
- Require each digital bet to declare which of the **three scarce assets** — clinical trust, the longitudinal record, the care network — it strengthens. A bet that strengthens none of the three is cost, not strategy.

Twelve months, four cycles, a second measurement

Four quarterly cycles. Each with plenary sessions, mandatory artefacts and a minute item. Total load per director: **19 hours of sessions over twelve months** — a little over an hour and a half a month — plus roughly 45 minutes for the second individual assessment.

CYCLE	DOMAINS	CONTENT AND ARTEFACT	MINUTE ITEM
C1 2026-Q4 LAUNCH	D5 Foresight D1 Leadership + register administration	Foundation and first crossing. Horizon session anchored in the three in-house scenarios; review of digital capital allocation against the approved plan. 2 sessions × 90 min · 3h00 Artefacts: fixed horizon slot instituted · assessment coverage taken to 11/11 · experience profiles completed on the platform · ownership of the AI agenda attributed to a named director.	Approval of the programme and designation of the owning director
C2 2027-Q1 CRITICAL	D2 Cybersecurity and Data Governance	From incident to resilience. What a board supervises; data as an asset; crisis simulation with simultaneous impact on the patient record and on elective scheduling. 90+90+120 min · 5h00 Artefacts: quarterly panel of five fixed indicators · data inventory with lawful basis · list of datasets with licensing potential and what is missing for them to be lawful.	Approval of the permanent quarterly panel
C3 2027-Q2 CRITICAL	D3 AI and Advanced Analytics	The four AI sessions. Three layers; seeing the machine answer on clinical terrain; inventory and risk classification against Annex III; deliberation of a real proposal. 90+90+120+120 min · 7h00 Artefacts: inventory of the group's AI systems with owner and classification · the ten standing questions adopted into procedure · nine-field minimum dossier in force.	Adoption of the standing questions and of the minimum dossier
C4 2027-Q3 CONSOLIDATION	D6 Change Management D4 Ethical Governance	Integration and commitment. Redesign clinic on the revenue cycle; publishable clinical-AI ethics policy; close with the second measurement. 120+90 min · 3h30 Artefacts: one-page clinical-AI policy approved by the board · map of non-delegable decisions · second measurement completed and progress report.	Approval of the policy and reading of the 2nd measurement cycle

19^{h00}

SESSIONS PER DIRECTOR OVER 12 MONTHS · 3H00 + 5H00 + 7H00 + 3H30

13

ARTEFACTS PRODUCED AND FILED IN THE BOARD'S RECORDS

2

MEASUREMENTS — THE FIRST DONE, THE SECOND IN 2027-Q3

TWO STOREYS, PER PAGE 09

The whole programme above is a **mandatory common core** for all eleven directors. A **second storey** is added for the director who owns the AI agenda and for anyone who wishes to follow: deeper work in cycles C2 and C3, plus access to an external certification dedicated to AI oversight for directors, which already exists in the market. It is the answer designed for the 37 points of range — the programme cannot be calibrated to the average.

How we will know it worked

A development programme with no targets declared before it starts is not assessable at the end. I propose the following targets to the board for the second measurement, in 2027-Q3. They are targets for the body — never for a person.

INDICATOR	CYCLE 1 2026 - Q3	TARGET 2027 - Q3	HOW IT IS REACHED
Assessment coverage	9 / 11	11 / 11	Administrative. Complete in cycle C1.
Aggregate average	52	≥ 59	Consistent with four domains ≥60 and the lowest ≥55, below.
Domains above 60	1 / 6	≥ 4 / 6	D4 already there; D5 and D1 in C1; D2 in C2; D3 in C3.
Lowest domain	44	≥ 55	No domain in the critical band. This is the most important target.
Range across directors	37 pts	≤ 20 pts	Body indicator — the expected effect of the mandatory common core, which raises the base without any individual record.
Composite index (administrative)	43	≥ 58	Directional context, not a governance target (see p. 11) — participation, composition and momentum approximated by the 2nd measurement.
Register reconciled	No	Yes	Platform register aligned with the effective composition of the board.

The instruments that stay installed

Independently of the score, the programme leaves five permanent instruments. It is these — not the number — that prove diligent oversight.

- 01 **The ten standing questions**, in the board's procedure, for any AI initiative.
- 02 **The nine-field minimum dossier**, as a condition of admissibility for any AI proposal reaching the board.
- 03 **The inventory of AI systems** in the group, with purpose, named owner and risk classification, reviewed semi-annually.
- 04 **The quarterly data and security panel**, with five fixed indicators.
- 05 **The clinical-AI ethics policy**, publishable, and the map of non-delegable decisions.

THE TARGET THAT IS NOT IN THE TABLE

None of the numbers above is the real objective. The objective is that, in 2027, when a clinical-AI proposal 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

If the second measurement arrives with an expectation that everybody must improve, it will stop producing honest answers. I ask the board to fix now, in the minutes, that **the second measurement carries no individual consequences of any kind**. It is the condition on which the instrument stays worth anything in cycles three and four.

All scores are directional, human-interpreted and referenced to a version of the instrument. Comparisons between cycles require the same instrument version and the same perimeter — both recorded in each cycle's report.

What I ask of the board

This report opened with a sentence and closes with the same one: the result is good enough to be useful and low enough to be urgent. What makes it one or the other is decided in this meeting.

#	PROPOSED RESOLUTION	DEADLINE
01	Approve the twelve-month programme and the targets on page 24 , across four cycles — approximately 19 hours of sessions per director plus a 45-minute second individual assessment — and schedule the second measurement for 2027-Q3. Accountable: chair of the board, with BoardAcumen.	This meeting
02	Designate a director as owner of the digital and AI agenda — not as technical specialist, but as the person accountable for ensuring the question is asked and recorded. Accountable: the board, by resolution at this meeting.	This meeting
03	Complete assessment coverage to 11 of 11 and reconcile the platform register perimeter. Accountable: company secretary, with BoardAcumen.	By 2026-Q4
04	Instruct the executive team to produce the inventory of AI systems in use , with purpose, named owner and Annex III risk classification — and a written answer to the Article 50 question: which patient-facing interfaces are AI systems, what is disclosed, and where it is recorded. Accountable: chief executive, with the CMIO.	PRIORITY Preliminary inventory by 31 August; complete by 30 September 2026
05	Adopt the ten standing questions and the nine-field minimum dossier as board procedure for any AI initiative. Accountable: the director owning the AI agenda (resolution 02).	By 2027-Q1
06	Record in the minutes that the assessment carries no individual consequences , in any cycle, and that individual results remain exclusively with each director. Accountable: the board.	This meeting

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. The competency this report measures is, exactly, the competency to capture that value before it is spent.

Sources

ASSESSMENT DATA

BoardAcumen · Board Overview, Director × Domain Competency Heatmap, Score Distribution, Critical Gaps and Board Competency Index · Vitalis perimeter · extraction of 29 July 2026. Board Digital Competency Instrument v1.0 · BCI v1, six weighted components. **Vitalis Health Group is a fictional organisation and all of its figures are illustrative.**

GOVERNANCE AND BOARDS

Deloitte Global Boardroom Program · PwC Annual Corporate Directors Survey · EY, AI governance disclosure in the S&P 500 · The Conference Board, Governing AI in the Corporation · NACD, Board AI Governance Guide · World Economic Forum, Future of Jobs Report 2025.

SECTOR AND REGULATION

Comparitech · Healthcare Ransomware Roundup, H1 2026 · Censinet, AI governance in health systems, 2026. Regulation (EU) 2024/1689 (AI Act), Articles 4, 6(2), 50, 99 and Annex III · Directive (EU) 2022/2555 (NIS2), Article 20 · Regulation (EU) 2025/327 (European Health Data Space) · Regulations (EU) 2017/745 and 2017/746 (MDR / IVDR).

Andiara Petterle, PhD

External adviser responsible for the assessment · BoardAcumen
29 July 2026 · Cycle 1 · Period 2026-Q3



RECOMMENDATION I OF III

Literacy as a permanent function, not a project

The twelve-month programme of Part III resolves one cycle. This recommendation asks for something else: that training of the board and the executive team stop being a project and become a **permanent function of the organisation** — with an annual plan, a budget line, a named owner and a compliance record, as already happens with financial and legal literacy.

Why it has to be permanent

- **Because the subject matter changes faster than a mandate.** A single training fixes the board at the state of the technology on the date of the session. Between the first measurement and the second, the object of oversight will already have changed.
- **Because the law asks for continued effort, not an event.** The AI literacy duty is an obligation of means: supporting a sufficient level of literacy in whoever operates or uses these systems. An annual plan, approved and executed, is the proof; an isolated session is not. NIS2 goes further, placing the training requirement on members of the management body directly.
- **Because rotation dilutes the investment.** Without an integration module, each new director enters at a prior plateau — which is what the measurement suggests has already happened.
- **Because the board cannot require of the organisation what it does not require of itself.** It is the condition of credibility for everything in Part III.

Three audiences, with distinct objectives

AUDIENCE	OBJECTIVE	LOAD / YEAR
Board 11 directors	Oversight competency. Knowing how to ask, how to judge the answer, what is non-delegable, and how to record.	19h00
Executive team and clinical leadership chief executive, CMO, CMIO, CIO, direct reports	Decision competency. Designing, contracting, governing and switching off systems — and answering to the board with a complete dossier.	24h
Organisation clinical, revenue cycle, plan operations, support	Use literacy. What can and cannot be done with these tools, with which data, and what must be disclosed to the patient.	to define

The board's load is that of Part III. The remainder are sizing proposals, to be validated with the executive team.

The six mechanisms that make this verifiable

01 · ANNUAL PLAN

Board training plan approved in the minutes in the first quarter of each year, with domains, sessions and dates.

02 · BUDGET

Its own identifiable budget line. What has no budget has no priority when the quarter tightens.

03 · OWNERSHIP

A named owner of the plan — the same director who owns the digital and AI agenda proposed in Part III.

04 · INTEGRATION

Mandatory module for a new director within the first 90 days, with an initial assessment. This is Deloitte's express recommendation.

05 · RECORD

Attendance and artefacts filed in the board's records. Training with no record is not opposable to a regulator.

06 · MEASUREMENT

Re-measurement each cycle, with the same instrument and the same perimeter. Without it, the plan is activity rather than result.

The World Economic Forum estimates that **39% of workers' core skills will change by 2030**, and identifies AI and big data, cybersecurity and technological literacy as the three fastest-growing skills — exactly three of the six domains of this assessment. Among employers surveyed, **85% place reskilling among their priorities**. *Vitalis is not deciding whether it invests in literacy: it is deciding whether it does so before or after needing it.*

World Economic Forum · Future of Jobs Report 2025 · Deloitte Global Boardroom Program · NACD Board AI Governance Guide · Regulation (EU) 2024/1689, Art. 4 · Directive (EU) 2022/2555, Art. 20.

RECOMMENDATION II OF III

Board composition and a technology advisory committee

There are competencies that are built and competencies that are brought in. Part III deals with the first. This recommendation deals with the second — and proposes the route that does not require touching board composition now: **an advisory committee with independent members who already hold the competency.**

What the market has already done

Among S&P 500 companies, the proportion assigning AI oversight to a dedicated board committee moved from **6% to 20% in a year**, and disclosure of AI as a risk factor rose from 12% to **83%** — while AI competency among directors remained at 2.7%. *The committees are the response to the scarcity of competency, not a consequence of its abundance.*

PATTERN	STRUCTURE ADOPTED
Standing committee, external	A permanent digital technology committee composed exclusively of independent directors with verified competency.
Interdisciplinary advisory group	Directors, management, internal specialists and a senior external adviser on cybersecurity — advisory rather than deliberative.
Digital advisory board	Directors, the executive team and external technology leaders convened on a fixed cadence around a defined agenda.
Clinical-technical hybrid	In healthcare specifically: pairing the existing clinical governance or quality committee with data, AI and cyber competency, rather than creating a parallel structure.

EY • The Conference Board • Russell Reynolds Associates, on advisory boards as a route to supplying scarce technological competency.

The healthcare-specific argument

This group already has the structure. A healthcare board of this scale runs a **clinical governance or quality and safety committee**, with a tradition of reviewing adverse events, protocols and equity of access. That committee already does — for clinical processes — precisely what the AI Act asks for AI systems: documented classification, named accountability, human oversight of a consequential decision, an auditable record.

The proposal is therefore not to create a parallel structure. It is to extend the mandate of the committee that already exists and to seat in it the competency it does not have.

THE PROPOSAL, IN TWO STEPS

Step 1 • now, without changing composition. Constitute a **Technology, Data and AI Advisory Committee**, without deliberative power: two directors, two to three independent external members with verified competency, and the CIO and CMIO as standing invitees. Written mandate, quarterly report to the board.

Step 2 • at the renewal cycle. Approve a **board skills matrix** that includes AI, data and cyber as a declared and verified competency, and use natural vacancies to supply it — rather than treating each replacement as an isolated problem.

The three independent profiles that are missing — and why

PROFILE A

Clinical data and applied AI

Someone who has taken a clinical AI system from validation into production and lived with what followed. Covers the critical domain (46) and the decisions on secondary use of the record, where the board will assess contracts with no internal precedent.

PROFILE B

Cyber resilience in critical infrastructure

Someone accountable for continuity where downtime has clinical consequences. Covers the lowest domain (44) and the asset the group cannot afford to lose — the profile the hybrid pattern above seats with its own place at the table.

PROFILE C

Payer-side risk models and equity

Someone who has worked on risk assessment and pricing in health insurance under regulatory scrutiny — the domain entering Annex III in December 2027, and the one least tolerant of learning by trial.

THE LIMITS, SO THE COMMITTEE DOES NOT BECOME SOMETHING ELSE

The committee **advises and does not decide**: it does not receive the board's non-delegable competencies, does not substitute for the executive team, and does not interpose itself in the reporting line. External members sign confidentiality and a declaration of independence and conflicts, verified annually — with particular attention to relationships with clinical technology vendors and the payer market. And the mandate is written and dated: an advisory committee without one converts, over time, into influence without accountability.

RECOMMENDATION III OF III

A strategic plan that reaches the people

The group has scenarios and it has a budget. What the assessment suggests is missing is the **bridge between the two**: a plan that declares, for each long-horizon strategic decision, which competencies are required to execute it — and by which route they arrive.

Why the bridge is missing

It is not a failure of rigour. It is the consequence of planning and workforce management living in different documents with different horizons — strategy at three years, recruitment at the quarter — so each digital bet is approved without it being written down *who* will execute it.

The sector evidence is direct. Roughly half of healthcare organisations already run generative AI in production, while an estimated 10–15% of large health-system boards have formal AI oversight — and **74% of hospital security leaders name their own technology vendors as the top emerging risk**. None of that is a question of tools. It is a question of who is there, what they can run, and who answers when it fails.

The corollary is what Part III says in another form: **AI value does not come from the tool, it comes from the redesign** — and the redesign is not done with the competency mix that designed the previous process.

What we propose should exist

A **competency annex to the strategic plan**, approved alongside it and reviewed with it. Four elements:

- 01 **Competency map by scenario.**
For each of the three scenarios: which competencies become critical — and, hardest to write, which cease to be.
- 02 **Build / buy / advise, per front.**
Per critical competency and front — clinical, revenue cycle, plan operations, technology. Recommendation II is «advise»; Recommendation I is «build».
- 03 **Link to the budget.**
Each initiative declares the competency it requires and the route by which it arrives. A bet with no competency attributed is a bet with no executor.
- 04 **Quarterly indicators for the board.**
Not for people management: coverage of critical competencies, time to fill critical roles, retention in new roles, and share of the workforce with AI literacy training completed.

HORIZON	OPEN STRATEGIC DECISION	COMPETENCY IT REQUIRES	ROUTE
Now 0–6 months	Disclosure of AI use in patient-facing channels and clinical correspondence; the answer to Article 50.	Clinical AI governance · compliance	BUILD
Short 6–18 months	What to do with thirty years of clinical records: license for research, build a product, or neither.	Clinical data and applied AI · secondary-use agreements	ADVISE
Medium 12–36 months	Position on care migrating out of the hospital — the scenario where the group must choose a side.	Digital care models · distributed clinical operations	HIRE
Medium 12–36 months	Redesign of the revenue cycle and plan operations — the largest addressable administrative cost base.	Process redesign with AI · change management	BUILD
Structural permanent	Resilience of clinical operations and governance of the patient record.	Cyber resilience · data governance	ADVISE

Illustrative table, to structure the board's discussion — it does not replace the executive team's strategic plan. Sources: World Economic Forum, Future of Jobs Report 2025 · Censinet, 2026 · Comparitech, H1 2026.

THE THREE RECOMMENDATIONS, IN ONE SENTENCE

Train the people already at the table, bring in those who know what the table does not, and write the strategy and the people who will execute it into one document. None of the three is a difficult decision. Together they are what turns a result of 52 into a starting point.

Part IV · Andiará Petterle, PhD · BoardAcumen · 29 July 2026. Vitalis Health Group is a fictional organisation created to demonstrate the BoardAcumen method; all figures relating to it are illustrative.