

AI literacy is *already* a legal duty. Here is the whole programme.

A single working instrument for an international school: what the law requires, where you stand today, a six-module staff course, a full student programme mapped to the OECD–EC framework, twelve classroom activities, and the policy set that closes the loop.

Providers and deployers of AI systems must take measures to ensure a sufficient level of AI literacy among staff and other people operating AI on their behalf — proportionate to their technical knowledge, experience, education and the context of use, and to the people the systems are used on.

Article 4, EU AI Act – applicable since 2 February 2025.
Supervised and enforced by national market surveillance authorities from 2 August 2026.

SCHOOL	ACCOUNTABLE OWNER	DATE COMPLETED	NEXT REVIEW
Legal name	Name & role	DD / MM / YYYY	DD / MM / YYYY

Where the law stands right now

The Act arrives in waves. Two of those waves have already broken over schools. A third lands in days. The one everybody worries about — high-risk systems — has been pushed back, and that is exactly why the near-term duties get forgotten.

2 Feb 2025

IN FORCE

Prohibited practices, and the Article 4 literacy duty

Both applied from this date. If your staff have been using AI tools without training since then, the duty was already unmet — it did not begin in 2026.

2 Aug 2025

IN FORCE

General-purpose AI model obligations

Applies to model providers, not to you as a school. Relevant only when you assess a vendor's paperwork.

2 Aug 2026

THIS IS THE ONE

Enforcement of Article 4 begins · Article 50 transparency applies

National market surveillance authorities start supervising and enforcing the AI literacy duty. Separately, transparency and disclosure obligations for AI-generated content and AI interactions become enforceable — these were left untouched by the 2026 Digital Omnibus.

2 Dec 2027

DEFERRED

Annex III high-risk obligations — including education

Political agreement reached on 7 May 2026 moved these from August 2026 to December 2027. Confirm the final adopted text before relying on the later date. The deferral does not touch Article 4 or Article 5.

2028

DEFERRED

Annex I high-risk — AI embedded in regulated products

Rarely relevant to a school unless you deploy regulated medical or safety equipment with AI components.

THE TRAP

Headlines in mid-2026 said "EU delays AI rules." Leadership teams read that and stood their projects down. What was delayed was the high-risk regime. The literacy duty, the prohibitions, and the transparency duties were not. A school that paused on the strength of those headlines is now exposed on the exact obligation that is easiest to evidence and cheapest to meet.

PART 2 / 11

Is your school in scope, and as what

Almost every school is a **deployer**. A small number are also **providers**. The label determines which duties attach, so settle it before anything else.

DEPLOYER – NEARLY CERTAINLY YOU

You use AI systems in your professional activity

Staff use a chatbot for planning. Your MIS has AI features. You run a writing-feedback tool, a translation tool, a safeguarding filter, or a timetabling optimiser. Any one of these makes you a deployer.

Article 4 applies to you in full.

PROVIDER – CHECK CAREFULLY

You develop an AI system, or put your name on one

You built a custom GPT, agent, or marking assistant and made it available to others under the school's name. You white-labelled a vendor tool. You substantially modified a high-risk system.

Provider duties are heavier. Get advice if this is you.

Which school uses count as high-risk

ANNEX III, POINT 3 – EDUCATION AND VOCATIONAL TRAINING

Four education uses are listed as high-risk. Their obligations are deferred, but you should still know which of your systems will land here, because the inventory you build now is the same inventory you will need then.

LISTED USE	WHAT IT LOOKS LIKE IN A SCHOOL	STATUS
Admission and assignment	AI ranking or shortlisting applicants; automated setting or streaming on entry. Includes scholarship triage.	HIGH-RISK · DEFERRED
Evaluating learning outcomes	Automated marking that contributes to a reported grade; AI systems used to steer the learning process. Formative-only tools sit lower, but document the distinction.	HIGH-RISK · DEFERRED
Assessing appropriate level of education	AI-assisted setting decisions, pathway recommendation, predicted-grade generation used for placement.	HIGH-RISK · DEFERRED
Monitoring behaviour during tests	Remote proctoring and cheating-detection during examinations. Check the vendor's emotion-inference claims against Part 3.	HIGH-RISK · DEFERRED
Everything else	Lesson planning, resource generation, admin drafting, translation, differentiation, staff productivity, chatbots for parents.	LIMITED / MINIMAL RISK

THE PRACTICAL READING

Most of what a school actually does with AI is minimal or limited risk. That is good news, and it is also why the compliance job is small enough to finish this term. Your obligations are: don't cross the red lines, know what you run, train your people, disclose where required, and keep the paperwork.

Red lines you must not cross

ARTICLE 5 – PROHIBITED PRACTICES · APPLICABLE SINCE 2 FEBRUARY 2025

These are not risk-managed. They are banned. Penalties reach the highest tier in the Act — up to €35 million or 7% of worldwide annual turnover, with

proportionate caps for smaller organisations. One of the eight prohibitions names education explicitly.

ARTICLE 5(1)(F) – THE ONE AIMED AT YOU

Using AI to infer the emotions of a person in an education institution is prohibited

The exception is narrow: systems intended for medical or safety reasons only, and it is read restrictively. The Commission's guidelines confirm that both "emotion recognition" and "emotion inference" are caught, tied to inference from biometric data. It covers public and private institutions, all levels, in person or online — and it reaches admissions too.

Where this bites in real schools: engagement-tracking in video lessons, webcam attention or "focus" scoring, sentiment analysis of student faces or voices, proctoring tools that flag stress or suspicion from expression, wellbeing dashboards that infer mood from biometrics.

SCREENING QUESTIONS FOR EVERY TOOL YOU RUN

-
- ☐ **No tool infers emotion, mood, engagement or attention from faces, voices or body signals**
Ask each vendor in writing. "Engagement analytics" is the phrase to interrogate — check whether it is clickstream (fine) or biometric (not).
-
- ☐ **No tool scores or classifies students by social behaviour or personal traits**
Social-scoring style profiling is separately prohibited.
-
- ☐ **No tool categorises people biometrically by sensitive attributes**
Race, political opinions, religion, trade union membership, sex life, sexual orientation.
-
- ☐ **No facial recognition database has been built from scraped images**
Includes photo libraries assembled from the web or from CCTV.
-
- ☐ **No system exploits the vulnerability of children to distort behaviour**
Relevant when reviewing gamified or persuasive edtech.
-
- ☐ **Written vendor confirmations are on file for every answer above**
Your own assurance is not evidence. Theirs is. [Evidence item E-03](#)
-

IF YOU FIND ONE

Stop using it. Record the date you stopped, who authorised the stop, and what replaced it. A prohibited practice discovered and remediated with a dated record reads very differently to a regulator than one still running.

PART 4 / 11

AI system register

You cannot train people on systems you have not listed. The register is the first artefact any authority will ask for, and it takes about ninety minutes to build if you ask heads of department directly. Start with the rows below and add your own.

SYSTEM	USED BY	PURPOSE	RISK	OWNER	DISCLOSED?
<u>General assistant (c</u>	<u>All teaching s</u>	<u>Planning, resources</u>	Limited ▼	<u>Name</u> _____	Not ye ▼
<u>Writing feedback toc</u>	<u>English, Hum</u>	<u>Formative feedback</u>	Limited ▼	<u>Name</u> _____	Not ye ▼
<u>Similarity / AI-use de</u>	<u>Secondary, D</u>	<u>Academic integrity c</u>	High ▼	<u>Name</u> _____	Not ye ▼
<u>MIS / SIS built-in AI</u>	<u>Admin, SLT</u>	<u>Reporting, attendan</u>	Limited ▼	<u>Name</u> _____	Not ye ▼
<u>Parent-facing chatbc</u>	<u>Admissions, f</u>	<u>Enquiry handling</u>	Limited ▼	<u>Name</u> _____	Yes ▼
<u>Add a system</u>	<u>Department</u>	<u>What it does</u>	Minimal ▼	<u>Name</u> _____	Not ye ▼

TWO COLUMNS PEOPLE FORGET

Shadow use. Add a row for tools staff use without approval. You are the deployer whether or not procurement knew. A monthly amnesty question in briefing surfaces more than a policy ever will.

Disclosure. Article 50 means people should know when they are dealing with an AI system or looking at AI-generated content. The parent chatbot needs a line. AI-generated images in your prospectus need a line. This is enforceable from 2 August 2026.

PART 5 / 11

Readiness audit

Twenty-four questions across six domains. Answer honestly — a low score on a first pass is normal and is itself useful evidence that you assessed yourself. Mark **Evidenced** only where you could hand someone a document today.

A Governance and accountability

0 / 8

A single named person is accountable for AI across the school

Not a committee. A name in a job description.

NO

PARTLY

EVIDENCED

Governors or the board have received a written AI position paper

Minuted, with a date.

NO

PARTLY

EVIDENCED

A route exists for staff to request approval for a new AI tool

And someone actually answers it.

NO

PARTLY

EVIDENCED

An incident log exists for AI-related problems

Wrong output acted on, data exposure, integrity dispute.

NO

PARTLY

EVIDENCED

B Inventory and risk classification

0 / 8

A complete AI system register exists and was updated this term

Part 4 of this instrument.

NO

PARTLY

EVIDENCED

Every system is classified against the Act's risk tiers

Prohibited / high / limited / minimal.

NO

PARTLY

EVIDENCED

Every Article 5 screening question has a written vendor answer

Especially emotion inference.

NO

PARTLY

EVIDENCED

Shadow AI use has been surveyed at least once

Anonymous staff survey counts.

NO

PARTLY

EVIDENCED

C Staff AI literacy — the Article 4 core

0 / 8

All staff who touch AI have completed baseline training

Teaching, admin, leadership, and anyone operating a tool on the school's behalf.

NO

PARTLY

EVIDENCED

Training is differentiated by role and technical background

Article 4 requires proportionality, not one identical session.

NO

PARTLY

EVIDENCED

Attendance and completion are recorded per person

Names, dates, module, duration.

NO

PARTLY

EVIDENCED

New starters get AI literacy in induction

Including supply and contract staff.

NO

PARTLY

EVIDENCED

D Student programme and curriculum

0 / 8

AI literacy appears in schemes of work, not only in assemblies

Named outcomes in at least three subjects.

NO

PARTLY

EVIDENCED

Progression is mapped across year groups

What a Year 4 learns differs from a Year 12.

NO

PARTLY

EVIDENCED

Students know when they may and may not use AI in assessed work

Written, per subject, and understood.

NO

PARTLY

EVIDENCED

Student outcomes reference an external framework

The OECD–EC AILit framework, or an equivalent you can name.

NO

PARTLY

EVIDENCED

E Transparency, data and safeguarding

0 / 8

Anyone interacting with an AI system is told so

Parent chatbots, automated replies, AI-marked feedback.

NO

PARTLY

EVIDENCED

AI-generated content published by the school is labelled

Images, copy, translated material.

NO

PARTLY

EVIDENCED

No child's personal data goes into a tool without a lawful basis and a DPA

Coordinate with your GDPR record of processing.

NO

PARTLY

EVIDENCED

Safeguarding leads have been briefed on AI-specific risks

Synthetic imagery, chatbot dependency, grooming vectors.

NO

PARTLY

EVIDENCED

An evidence file exists in one place

Part 10 lists what goes in it.

NO	PARTLY	EVIDENCED
----	--------	-----------

Policies carry a version, an owner and a review date

Undated policy is weak evidence.

NO	PARTLY	EVIDENCED
----	--------	-----------

A review is diarised at least annually

Sooner if the deferred deadlines move again.

NO	PARTLY	EVIDENCED
----	--------	-----------

Someone tracks changes to the Act and to national guidance

Named person, named sources.

NO	PARTLY	EVIDENCED
----	--------	-----------

0

/ 48

NOT STARTED

Answer the audit above to generate a readiness position you can paste into a board paper or a governors' report.

The staff course

DISCHARGES ARTICLE 4 · SIX MODULES · 4 HOURS 45 MINUTES TOTAL

Built to be delivered as one INSET day, or as six twilights across a term.
Modules 1, 2 and 6 are mandatory for every member of staff who touches an AI system. Modules 3–5 are differentiated: run the teaching track and the

operations track separately. That differentiation is not a nicety — Article 4 asks for measures proportionate to people's technical knowledge and role.

M1 What these systems actually are

All staff · mandatory

45 min

Whole staff

WHY IT EXISTS

You cannot judge an output you do not understand the origin of. Most poor AI decisions in schools trace back to staff believing the system knows things, checks things, or remembers things that it does not.

LEARNING OUTCOMES

- Describe in plain language how a language model produces text, without using the word "thinks"
- Explain why a confident answer and a correct answer are unrelated properties
- Identify three tasks where these tools are strong and three where they are unreliable
- State what happens to data typed into a consumer tool versus a licensed one

CONTENT

- **Prediction, not retrieval.** A live demonstration: same prompt, three runs, three different answers. Ask the room what that tells them about citation.
- **Training data and cutoffs.** Why a model invents a plausible policy reference for your own school.
- **Confidence is a style, not a signal.** Show a fluent, wholly fabricated exam-board reference. This is the moment the room changes.
- **Where the data goes.** Consumer account, school licence, enterprise agreement — three different answers, three different risk positions.

ACTIVITY – HALLUCINATION HUNT (15 MIN)

In subject pairs, ask a model for five specific facts inside your own specification: syllabus codes, mark scheme wording, past paper questions. Verify each against the real document. Pairs report their hit rate to the room. Nobody forgets their own subject's failure rate.

EVIDENCE PRODUCED

Signed attendance sheet · completed hallucination-hunt sheets retained per department · three-line personal reflection collected.

AI ACT ART. 4

AILIT · ENGAGE WITH AI

M2 The law, in the form it reaches a classroom

All staff · mandatory

40 min

Whole staff

M3 Teaching track — using AI without eroding the work

Teaching staff

60 min

Teaching

M4 Operations track — data, procurement and disclosure

Admin, SLT, office, admissions

60 min

Operations

M5 Safeguarding, bias and the harder conversations

All staff · pastoral emphasis

50 min

Whole staff

M6 Your commitments, on the record

All staff · mandatory · closes the loop

30 min

Whole staff

DELIVERY SHAPES THAT WORK

One INSET day. M1, M2 before break. M3 and M4 in parallel tracks after. M5 and M6 to close. Everyone leaves with a signed card.

Six twilights. One module per fortnight across a term. Better retention, worse attendance. Record every session so absentees can complete asynchronously — and log that they did.

New starters. M1, M2 and M6 compressed to 90 minutes in induction. Non-negotiable.

The student programme

MAPPED TO THE OECD–EUROPEAN COMMISSION AILIT FRAMEWORK, PUBLISHED JUNE 2026

Article 4 is about staff. Your students are a separate matter — and the reason the job is worth doing properly. The AILit framework gives you an external reference to cite, which is exactly what inspectors and accreditation bodies want to see instead of a school inventing its own outcomes.

The four domains

The framework organises AI literacy into four domains and nineteen competences, each combining knowledge, skills and attitudes, with a three-level progression. The domains run as a developmental pathway: Engage and Create sit in parallel as hands-on encounters, then Manage adds judgement, and Shape asks students to see human choices behind the technology. It feeds the PISA 2029 assessment, so alignment now is not wasted work.

DOMAIN 1

Engage with AI

Recognising AI when you meet it, understanding roughly how it works, and reading its outputs with appropriate suspicion.

School translation: can a student say what the tool is doing and why it got that wrong?

DOMAIN 2

Create with AI

Using AI purposefully in making, problem-solving and delegating parts of a task — while keeping authorship.

School translation: can a student show which parts are theirs?

DOMAIN 3

Manage AI

Responsible decision-making: privacy, dependency, verification, knowing when not to use it at all.

School translation: can a student justify a decision to switch it off?

DOMAIN 4

Shape AI

Seeing that these systems embody human choices and values, and that those choices could have been made differently.

School translation: can a student name who decided, and propose an alternative?

Progression by year band

BAND	ENGAGE	CREATE	MANAGE	SHAPE
Years 3–4 Ages 7–9	Spot AI in everyday things. Know a computer can be "taught" by examples.	Use a tool with an adult to make something, and say which bit they did.	Never share your name, photo or address with a machine that asks.	People built this. People chose what it does.
Years 5–6 Ages 9–11	Explain that examples shape what it learns. Notice when output is wrong.	Iterate: change the instruction, see the change, keep the better one.	Check a claim against a second source before believing it.	Notice who is missing from the pictures it makes.
Years 7–9 Ages 11–14	Describe prediction versus understanding. Explain training data and its limits.	Decompose a task, delegate the right parts, keep authorship of the argument.	Assess privacy trade-offs. Recognise dependency in their own habits.	Audit a system for bias and propose a design change with a reason.
Years 10–11 IGCSE	Evaluate reliability by task type. Explain why confidence is not accuracy.	Produce work with a declared AI-use statement they can defend orally.	Apply the school's permitted-use rules and justify borderline decisions.	Argue a position on a real governance question with evidence.
Years 12–13 DP / A-level	Interrogate a system's failure modes systematically.	Use AI in extended research with a transparent, auditable method.	Weigh societal costs — labour, energy, concentration of power.	Propose and defend policy. Understand regulation as a design constraint.

WHERE IT LIVES ON THE TIMETABLE

Cross-curricular, with an anchor. Computing or Digital Technology owns the Engage and Shape strands and the progression map. English, Humanities and Science each own one Create or Manage outcome per year. Pastoral owns dependency and privacy. One named person keeps the map — otherwise it dissolves within a year.

Twelve classroom activities

Three per band, each mapped to an AILit domain. Most run in a single lesson and need nothing beyond what a classroom already has. The unplugged ones are deliberate — the strongest AI literacy activities for younger students involve no screens at all.

YEARS 3–4

A1 **Teach the machine**

Unplugged · classification

40 min

Engage

A2 **Machine or person?**

Sorting · recognition

30 min

Engage

A3 **Catch the mistake**

Verification · early critical habit

30 min

Manage

YEARS 5–6

A4 **The biased sorting machine**

Unplugged · training data

50 min

Shape

A5 **Prompt laboratory**

Iteration · authorship

50 min

Create

A6 **Real or generated?**

Provenance detective

45 min

Manage

YEARS 7–9

A7

Rules versus learning

Build both, compare

2 × 50 min

Engage

A8

Occupation audit

Systematic bias investigation

50 min

Shape

A9

Synthetic media defence

Verification workflow

50 min

Manage

YEARS 10–13

A10

Red-team the model

Systematic failure-mode testing

2 × 50 min

Engage

A11

Write the permitted-use contract

Students author the rules

50 min

Create

A12

Regulate it yourself

Policy design under real constraints

2 × 50 min

Shape

Five documents that close the loop. Replace everything in brackets, add a version number, an owner and a review date, and get them approved through your normal route. An undated policy is weak evidence; a dated, owned, reviewed policy is strong.

T1 · Staff AI acceptable use

[SCHOOL NAME] – Staff use of artificial intelligence

Version [1.0] · Owner [NAME, ROLE] · Approved [DATE] · Review [DATE]

1. SCOPE

This applies to all employees, contractors, supply staff and volunteers who use any AI system in the course of work for the school. [SCHOOL NAME] is a deployer of AI systems under Regulation (EU) 2024/1689 and carries the AI literacy duty under Article 4.

2. BEFORE YOU USE A TOOL

- 2.1 Only tools on the approved register may be used with any school or pupil data.
- 2.2 Requests to add a tool go to [NAME] using the approval route at [LOCATION].
- 2.3 You may not enter pupil personal data into any tool that is not on the register.
- 2.4 Never enter safeguarding records, medical information, SEN reports or staff HR matters into any AI tool without written approval from [NAME].

3. HOW YOU USE IT

- 3.1 You remain the author. Anything issued under your name is your responsibility, whatever produced the first draft.
- 3.2 Verify every factual claim, citation, date and reference before use.
- 3.3 Do not use AI to generate a final grade or a reported assessment judgement. AI may support marking; a human makes and owns the decision.
- 3.4 Do not use any tool that infers emotion, mood or attention from faces, voices or body signals. This is prohibited in education settings under Article 5(1)(f).
- 3.5 Apply the same professional judgement to AI output as to anything else you would put in front of a child.

4. TELLING PEOPLE

- 4.1 Where a parent, pupil or colleague could reasonably assume a human wrote or answered something, say if AI was involved.
- 4.2 AI-generated images or text published by the school are labelled as such.

5. WHEN SOMETHING GOES WRONG

- 5.1 Report to [NAME] within [one working day].
- 5.2 Reportable: incorrect output acted on, personal data exposed, a suspected prohibited practice, a disputed integrity decision, a safeguarding concern.
- 5.3 Reports are logged. Reporting in good faith carries no penalty.

6. TRAINING

6.1 All staff complete the AI literacy programme before using AI tools at work.

6.2 New starters complete the induction module within [30 days].

6.3 Completion is recorded and retained.

Signed [.....] Date [.....]

T2 · Student AI use — secondary

[SCHOOL NAME] – Using AI in your work

Years [7–13] · Version [1.0] · Review [DATE]

WHY WE HAVE THIS

We are not banning AI. We are protecting the thing school is for: you getting better at things. Some tasks are hard on purpose. If a tool removes the difficulty, it removes the learning – and you will meet that gap later, in an exam hall or a job, without the tool.

THE THREE LEVELS

Every assessed task tells you which level applies. If it does not say, ask.

LEVEL 1 – NO AI

Nothing generated by AI. This is used where we need to see what you can do unaided. Examples: [in-class assessments, controlled conditions, ...].

LEVEL 2 – AI WITH A DECLARATION

You may use AI for [named purposes: brainstorming, explaining a concept, checking grammar, generating practice questions]. You submit a short declaration saying what you used and for what. No penalty for honest use.

LEVEL 3 – AI AS A TOOL OF THE TASK

The task is about working with AI. Full use, documented method.

YOUR DECLARATION

Attach this to any Level 2 or 3 submission:

Tool(s) used:

What I used it for:

What is entirely my own:

Anything I checked and corrected:

WHAT COUNTS AS CHEATING

Submitting AI-generated work as your own thinking. Using AI at Level 1.

Declaring falsely. That is it – it is not complicated.

IF WE THINK SOMETHING IS WRONG

We will talk to you first. We will ask you about your work – how you got there, why you made a choice, what you would change. We do not treat a detection score as proof. You will always get to explain.

WHAT WE ASK OF YOU

Check what it tells you. It is confident when it is wrong.

Keep your own voice. Your teachers can tell, and it matters more than you think.

Do not put your personal information, or anyone else's, into these tools.

Tell someone if you see AI used to harm another student.

Student signature [.....] Tutor [.....] Date [.....]

T3 · Letter to parents and carers

Dear parents and carers,

I am writing to set out how [SCHOOL NAME] approaches artificial intelligence, and what we are asking of you.

WHERE WE STAND

AI tools are now part of the working world your children will enter. Our position is that pretending otherwise fails them, and that unsupervised use also fails them. So we teach it deliberately: what these systems are, where they fail, when to use them, and when to switch them off.

WHAT WE HAVE DONE

Every member of staff has completed AI literacy training. We keep a register of every AI system in use across the school and who is responsible for each. We have checked all of them against the practices that European law prohibits in schools – including any system that tries to infer a child's emotions, which is banned and which we do not use. Our policies name an owner and a review date.

WHAT YOUR CHILD WILL LEARN

Our programme follows the AI literacy framework published by the OECD and the European Commission in 2026. It runs from Year [3] to Year [13] and covers recognising AI, creating with it, managing its risks, and understanding that people design these systems and could have designed them differently.

WHAT WE ASK OF YOU

- Talk about it. "What did it get wrong?" is a better question than "did you use it?"
- Check the age requirements on tools at home. Most set a minimum of 13 or higher.
- If your child uses AI for homework, ask them to show you what they changed.
- Contact us rather than the internet if you are worried about something.

WHAT TO DO IF SOMETHING HAPPENS

If your child is affected by an AI-generated image or message, contact [NAME, ROLE] on [CONTACT]. We will respond through our safeguarding procedures. Your child will not be in trouble for reporting.

Our full policy is at [LOCATION]. I am glad to discuss it.

[NAME]

[ROLE], [SCHOOL NAME]

T4 · Vendor written confirmation request

Subject: EU AI Act – written confirmations required before procurement

Dear [VENDOR],

[SCHOOL NAME] is an education institution in [COUNTRY] and a deployer of AI systems under Regulation (EU) 2024/1689. Before we can proceed with [PRODUCT], we need written answers to the following. Marketing material is not sufficient for our records; we need a signed response we can retain.

1. Does the system identify or infer emotions, mood, attention or engagement of any person from biometric signals, including facial expression, voice or posture?
If yes, on what basis do you consider this compatible with Article 5(1)(f)?
2. Does the system perform any function listed in Annex III point 3 – determining admission or assignment, evaluating learning outcomes, assessing the appropriate level of education, or monitoring behaviour during tests?
3. What is your classification of this system under the Act, and what documentation supports it?
4. Where is data processed and stored? Please provide your data processing agreement.
5. Is customer data, including pupil data, used to train or improve your models?
If so, can we opt out in writing, and does opting out change the service?
6. What accuracy figures can you provide, and on what population were they measured?
7. What human oversight does the system require in operation, and how is it enforced rather than merely recommended?
8. How are errors reported, what is your remediation commitment, and what is your incident notification timeline?
9. What happens to our data on termination, and within what period?

Please also confirm the name and contact of the person accountable for AI Act compliance at your organisation.

We cannot complete procurement without these. Thank you.

[NAME], [ROLE]
[SCHOOL NAME]

T5 · Board / governors position paper — skeleton

[SCHOOL NAME] – AI governance position

Paper for [BOARD / GOVERNORS] · [DATE] · Author [NAME, ROLE]

1. PURPOSE

To inform the [board] of the school's obligations under Regulation (EU) 2024/1689 and to seek approval for the AI literacy programme and policy set at Annex A.

2. THE OBLIGATION

The AI literacy duty under Article 4 has applied since 2 February 2025. National market surveillance authorities began supervising and enforcing it on 2 August 2026. Transparency obligations under Article 50 also became enforceable on that date. Obligations for high-risk systems – which include several education uses – were deferred to 2 December 2027 by political agreement reached on 7 May 2026. The deferral does not affect Articles 4, 5 or 50.

3. OUR POSITION TODAY

Self-assessment completed [DATE] using [INSTRUMENT]. Score [X] of 48.

Position: [.....].

Strongest domains: []. Weakest domains: [].

4. WHAT HAS BEEN DONE

- AI system register established, [N] systems recorded
- All systems screened against Article 5 prohibitions; [N] issues found, [N] resolved
- Staff programme delivered to [N] of [N] staff on [DATES]
- Policy set T1–T4 drafted and [approved / pending approval]

5. RESIDUAL RISK

[]

6. ASKS

- 6.1 Approve the policy set at Annex A.
- 6.2 Confirm [NAME] as accountable owner in their job description.
- 6.3 Approve [BUDGET] for [].
- 6.4 Diarise annual review for [DATE].

7. RECOMMENDATION

That the [board] approves 6.1 to 6.4.

PART 10 / 11

The evidence file

Article 4 asks you to take measures. Measures you cannot show are measures you did not take. Keep all of this in one place, dated, with a named owner. Twelve items — most schools already have five of them scattered across three drives.

☐ **E-01 · AI system register**

Current, dated, with an owner per system.

☐ **E-02 · Risk classification per system**

Including the reasoning, not just the label.

☐ **E-03 · Vendor written confirmations**

Template T4 responses, signed.

☐ **E-04 · Article 5 screening record**

Every question answered for every system, with dates.

☐ **E-05 · Staff training records**

Per person: name, role, modules, dates, duration, signed commitment card.

☐ **E-06 · Training materials as delivered**

Slides, handouts, scenario cards. Shows proportionality by role.

☐ **E-07 · Induction record for new starters**

Including supply and contract staff.

☐ **E-08 · Approved policy set**

T1–T3 with version, owner, approval date, review date.

☐ **E-09 · Disclosure inventory**

Where AI involvement is declared, and the exact wording used.

☐

E-10 • Incident log

Even if empty. An empty dated log beats no log.

☐

E-11 • Board minute approving the programme

Governance evidence, not paperwork.

☐

E-12 • Completed self-assessment with a date

This instrument, printed. Repeat annually and keep both.

WHAT GOOD LOOKS LIKE TO AN INSPECTOR

Not perfection. A named owner, a current register, evidence that real people were really trained in a way suited to their role, a policy someone can find, and a dated record showing you assessed yourself and knew where your gaps were. A school with a modest score and an honest, dated improvement plan is in a stronger position than a school with nothing written down.

PART 11 / 11

Twelve-week rollout

One term. Assumes a single person with roughly half a day a week, plus one INSET slot. Weeks 1–4 produce the compliance floor. Weeks 5–12 produce the programme that actually changes practice.

WEEK	DO THIS	OUTPUT	WHO
1	Name the accountable owner. Write it into a job description.	Named owner, board informed	Head / Board
2	Build the register. Email every head of department the same three questions.	E-01 draft	AI owner
3	Screen every system against Article 5. Send template T4 to each vendor.	E-03, E-04 started	AI owner + Ops

WEEK	DO THIS	OUTPUT	WHO
4	Classify each system. Stop anything prohibited, with a dated record.	E-02 complete	AI owner
5	Draft T1 and T2. Circulate for comment — comments are engagement, not delay.	Policy drafts	AI owner + SLT
6	Board paper using T5. Seek approval for the policy set and the owner.	E-11	Head
7	Prepare the INSET. Build the scenario cards from your own register.	E-06	AI owner
8	Deliver modules M1–M6. Collect every commitment card before people leave.	E-05	All staff
9	Publish T3 to parents. Add disclosure lines to AI-facing systems.	E-09	Comms + Ops
10	Map the student programme. Assign one AILit outcome per subject per year.	Progression map	Curriculum lead
11	Pilot three activities from Part 8 across three year bands.	Student work samples	Named teachers
12	Re-run this audit. Print it. File everything. Diarise the annual review.	E-12, complete file	AI owner

ROLES, MINIMALLY

ACCOUNTABLE OWNER

One person. Holds the register, the evidence file and the review date. Typically Head of Digital Learning or a Deputy Head. Needs about half a day a week in term one, then a day a term.

DATA PROTECTION LEAD

Owens the GDPR interface. Signs off what data may enter which system. Often external — brief them properly rather than assuming they know the Act.

SAFEGUARDING LEAD

Owens the AI-specific safeguarding annex and the synthetic imagery response protocol. Must be in the room for module M5.

CURRICULUM LEAD

Owns the AILit progression map and makes sure outcomes reach schemes of work rather than assemblies.

HEAD OF SCHOOL

Names the owner, takes the board paper, and makes the programme mandatory rather than encouraged. This is the whole job.

BOARD / GOVERNORS

Approve, minute, and ask for the review a year later. The minute is evidence.

Scope and limits. This instrument is a practical planning and training tool prepared for school leaders. It is not legal advice and does not create a presumption of compliance. Article 4 requires measures proportionate to your context — only your school can judge what is sufficient for its own staff and systems.

Verify before you rely. The timeline in Part 1 reflects the position following the political agreement of 7 May 2026 on the Digital Omnibus, which was still subject to formal adoption. Deadlines, national guidance and enforcement practice change. Confirm the current position with your national market surveillance authority, your data protection officer and, where the stakes justify it, a qualified lawyer.

Sources referenced. Regulation (EU) 2024/1689, Articles 4, 5 and 50 and Annex III point 3 · European Commission Guidelines on prohibited AI practices, February 2025 · OECD and European Commission, *Empowering Learners for the Age of AI: An AI Literacy Framework for Primary and Secondary Education*, June 2026.

AIL-01 · Rev A · Nothing on this page is transmitted anywhere.
Print to keep a record.