

EdTech in Schools: A Framework for Use

How this framework works

Classification in this framework operates across three independent layers. A platform may be prohibited on any one of them regardless of its standing on the others.

Architecture	How does the tool behave? Does it optimise for engagement, simulate relationship, or produce output on the child's behalf? These architectural features are grounds for prohibition regardless of the evidence on learning outcomes.
Evidence	What does independent research show about learning outcomes? Vendor-reported data, engagement metrics and short-term gains do not count as evidence. Where no independent comparative trial exists, the evidence base is treated as absent.
Policy	What does the school choose to permit? Some classifications reflect a policy judgment rather than an empirical finding. The prohibition on 1:1 devices and on digitised homework are policy positions grounded in safeguarding rationale, not learning outcomes evidence. They apply regardless of what individual platform assessments show.

Two rules

- 1) **No 1:1 devices.** Technology used in schools should be shared, purposeful and time-limited. A device assigned to a single child for the school day becomes a default presence rather than a learning tool.
- 2) **No digitisation of homework.** A device used at home cannot be adequately monitored or filtered for safeguarding purposes.

Two tests

Before any platform reaches the pedagogical question, two prior compliance tests apply. Where a platform fails either, it should not be deployed regardless of perceived educational value.

1. GDPR compliance	The school, as data controller, must verify that pupil data is processed lawfully, fairly and transparently. This means: the platform has a clear lawful basis for processing children's data; data is not harvested, monetised or shared with third parties for non-educational purposes; data retention and deletion policies are clear and enforceable; and the vendor has signed a Data Processing Agreement compliant with UK GDPR. Where a vendor cannot demonstrate compliance clearly and in writing, the platform should not be deployed. The risk exposure sits with the school, not the vendor. Note: published research raises serious concerns about data flows in widely-used platforms including Google Classroom even under education accounts. Schools should seek explicit written assurance from any platform with a commercial parent company before deployment.
2. Security and access controls	The school must verify that appropriate access controls are in place to ensure children cannot be exposed to inappropriate content, profiling, advertising or contact with unknown adults. This includes content moderation policies and their enforcement, advertising and recommendation algorithm controls, and third-party integrations that may introduce uncontrolled content. Where due diligence has not been completed, the platform should not be deployed.

Tool or agent?

Technology in schools falls into two categories. The first is a tool: something a child uses to produce an output that requires their own thinking, effort and judgment. The second is an agent: something that responds to the child, adapts to the child, or produces output on the child's behalf. Tools can have a place in education. Agents should not.

Prohibited: no exceptions

The following are prohibited regardless of supervision level, claimed use case, or age group.

Fundamental architecture misaligned with child development

All pupil-facing conversational and generative AI systems. This includes AI tutors, chatbots, and any platform where the child interacts with a language model. The prohibition is architectural: no supervision level, use case or age group changes what these tools are. It is supported by convergent evidence across two domains: in learning outcomes, dependency effects are documented (gains that collapse once AI is removed), suggesting the tool offloads the learning rather than enables it; in safeguarding, the concern is emotional disclosure redirected from adults with duty of care to systems outside any safeguarding structure.

Core features problematic as commercially deployed

All platforms where engagement-optimised design is the primary retention mechanism. This means platforms built around streaks, social comparison leaderboards, or reward loops where the loop is not incidental but structural. A classroom quiz tool with a timer and a score is bounded and teacher-controlled. A platform that sends notifications to maintain streaks, ranks children against each other, or uses coins and avatar rewards to drive daily return is engagement-optimised regardless of the subject matter it covers. Some research shows gamification can produce short-term motivational gains. The concern is not with game elements used purposefully by a teacher, but with platforms where those mechanics are structural, non-optional and designed primarily to maximise return visits rather than deepen learning. No such product is currently available without those mechanics intact.

Provisional status and updates

Conditional status is provisional and should be reviewed whenever a vendor introduces major new features or when significant new evidence becomes available. If a platform adds AI-tutoring or engagement-optimised design, it should be reassessed against the prohibited-architecture conditions. This framework should be revisited periodically to ensure it remains up to date with evolving technology and research. The onus to check sits with the school at the point of procurement and at any major product update, not on a continuous monitoring basis.

Applying this framework

Apply the following steps in order. Stop at the first failure. A tool that fails any step should not be used, regardless of claimed educational benefit.

Step 1	GDPR compliance. Has the platform a clear lawful basis for processing children's data, a signed Data Processing Agreement, and no data sharing for non-educational purposes?
Step 2	Security and access controls. Are content moderation, advertising controls, and third-party integrations verified?
Step 3	Architectural test. Does the child direct the tool, or does the tool respond to the child? If the tool is conversational, engagement-optimised, or produces output on the child's behalf: Do not proceed.
Step 4	Bounded. The tool has a defined, limited function. It does not optimise for the child's continued engagement.
Step 5	Purposeful. It is used for a specific task with a clear learning outcome that could not be achieved as well without it.
Step 6	Time-limited. It is used for a defined period, not as an ambient or default medium of instruction.
Step 7	Supervised. A trained adult has selected the tool, is present during its use, and is interpreting the output. The child is not left alone with the platform.
Step 8	Opportunity cost. What valuable non-digital activity is being replaced? Is that trade-off justified by strong independent comparative evidence?

Notes

The two rules are policy positions grounded in safeguarding rationale, not learning outcomes evidence. The relevant question is not what filtering software can achieve at home but where responsibility sits.

Before using any borderline tool, ask two questions. Is the child doing the thinking, or is the tool doing it for them? Is the tool responding to the child, or is the child directing the tool? If the answer to either question is unclear, do not use it.

YouTube is a useful example of why this distinction matters. A teacher selecting a specific video and playing it to the class on a shared screen passes the instrument test: the child watches, the tool does not respond to the child. That specific use is conditional. But the surrounding architecture is prohibited: YouTube's recommendation algorithm is engagement-optimised and profiling-based, and cannot be meaningfully disabled on a standard school account. Any individual pupil access to YouTube on a device fails the architectural test entirely. YouTube is therefore conditional in theory and prohibited in practice for most classroom settings, because the conditions required to make it safe cannot be reliably met.

Oversight does not make a prohibited tool acceptable. A conversational AI system supervised by a teacher is still a conversational AI system. Supervision matters for permitted tools. It cannot rehabilitate prohibited ones.

Efficiency is not a justification for technology use. If a tool saves teacher time but reduces the quality of a child's learning experience, the tool should not be used. The question is never whether technology makes a task faster or easier. The question is whether it makes learning deeper.

SEND and assistive technology. Compensatory tools for pupils with identified needs (text-to-speech, speech-to-text, AAC devices) operate under different principles. The distinguishing principle is this: a compensatory tool addresses a barrier to access that prevents a child from engaging with learning at all; it does not substitute for a cognitive process the child is capable of performing. A speech-to-text tool for a child with dysgraphia enables written output the child cannot otherwise produce. An AI writing assistant for a child without that barrier removes the cognitive labour of composition that constitutes the learning. These are not equivalent uses of responsive technology. The former is assessed separately using professional judgment and appropriate specialist input, and is not subject to the blanket prohibitions above.