

RED6 / PRACTICAL LMS INSIGHT

The Good, the Bad and the AI

A practical governance and procurement guide for deciding which AI features to enable, pilot, constrain or reject in a learning platform.

AI AND LMS GOVERNANCE - WORKING GUIDE

START HERE

This is a working guide, not a longer web page.

The website gives you the principle. This download turns it into a decision system: a read/write risk model, a use-case traffic light, vendor evidence requests, a scored evaluation, a controlled pilot and a one-page operating policy.

Let AI read and assist widely. Let it draft with review. Be slower about letting it change records, and slowest where a decision affects access, compliance or opportunity.

WHAT YOU WILL LEAVE WITH

☐	A fast method for classifying AI features by what they can actually do
☐	An evidence-led vendor questionnaire that goes beyond security assurances
☐	A reusable scorecard for comparing features and suppliers consistently
☐	A pilot plan, stop conditions and short operating policy your team can adopt

Use it once to diagnose the problem. Keep the worksheets for the next review, renewal or vendor conversation.

01 / WORKING SECTION

Classify the feature before discussing the promise

Most arguments about AI become easier when you ask three concrete questions: what can it read, what can it create, and what can it change without approval?

Capability	Typical examples	Primary concern	Default posture
Read	Search, summarisation, question answering, report explanation and content discovery.	Source quality, permissions and inappropriate disclosure.	Enable in bounded sources when answers show their evidence.
Draft	Descriptions, outlines, questions, tags, translations, audiences or learning-path suggestions.	Accuracy, bias, rights and an illusion of finished quality.	Allow with an identified reviewer before publication or action.
Write	Create content, update metadata, change a learner record or configure a rule.	Auditability, unintended change, ownership and recovery.	Require approval, logging, narrow permissions and rollback.
Decide / act	Infer a skill, assign training, affect compliance status, recommend progression or trigger another system.	Fairness, consistency, explainability and material impact on people.	Pilot only with explicit governance; keep high-impact decisions human-led.

DO NOT ACCEPT THE LABEL

An assistant can still be an agent

Classify the permissions and outcomes, not the product name. A friendly chat interface that can assign training or change a record deserves write-level controls.

Feature boundary

Feature and intended user	
Sources it can read	
Records or assets it can create	
Actions it can take	

Human approval point	
How to reverse an error	

02 / WORKING SECTION

Use a traffic light for real learning-platform use cases

Risk comes from the combination of data, permissions and consequence. The same model can be low-regret in one workflow and unacceptable in another.

Signal	Use case	Why	Minimum control
GREEN	Search approved content and return links to the source.	Assists discovery without changing the system of record.	Respect permissions; show sources; capture incorrect-answer feedback.
GREEN	Summarise anonymised survey themes or draft course metadata.	Creates a reviewable convenience rather than a learner decision.	Human review before publication; retain the source.
AMBER	Draft questions, translations, learning paths or audiences.	Quality errors can reach learners or distort intent.	Named subject or platform reviewer; version history; sample testing.
AMBER	Recommend learning from role, history or stated goals.	The recommendation may shape access and opportunity.	Explain inputs; allow refusal or correction; monitor uneven outcomes.
RED	Automatically assign mandatory learning from an inferred skill gap.	A probabilistic inference becomes a compliance process.	Use deterministic rules or require accountable approval and a complete audit trail.
RED	Change completion, competency, access or employment-related records.	Errors have material and potentially difficult-to-reverse consequences.	Do not enable without formal impact review, restricted permissions and tested recovery.

Four questions that can change the colour

Whose data? Public or approved content is different from employee performance, disability, health or customer-confidential information.	Whose permission? A learner, manager and administrator should not expose the same sources or take the same actions.
What consequence? A poor summary wastes time. A wrong assignment or skills inference can affect compliance, workload or opportunity.	Can it be undone? Prefer bounded, reviewable outputs. Require tested rollback before the feature writes to a system of record.

03 / WORKING SECTION

Make the vendor show evidence

Broad questions create broad reassurance. Ask the supplier to trace one realistic request through data, model, permissions, output, logs and deletion.

Twelve questions worth putting in writing

☐	What exact data is sent to the model, including prompts, retrieved content, metadata and user identifiers?
☐	Where is each part processed and stored, and which subprocessors are involved?
☐	Can our prompts, files, records or outputs train any shared or supplier model?
☐	Which LMS permissions are enforced during retrieval, and how is leakage across audiences tested?
☐	Can the user inspect the sources behind an answer or recommendation?
☐	What can the feature create or change without human approval?
☐	Does the audit log record the prompt, source, model/version, output, approver and final action?
☐	Can write access be disabled while read and draft capabilities remain available?
☐	Can every created asset and its source material be exported in a usable standard format?

☐	What happens to AI-created content, logs and embeddings when the feature or contract ends?
☐	Which licence tier and usage charges apply at pilot and full-population scale?
☐	What happens when the model, retrieval source or integration is unavailable or changes version?

PORTABILITY TEST

Ask for the file, not the roadmap

Have the vendor build a representative course or asset with its AI tooling, export it and open it in a separate environment. An assurance that export is supported is not equivalent to a usable production file.

04 / WORKING SECTION

Score the feature consistently

Score the implementation you can buy now, not the roadmap. A low score in human control or permissions can be a veto regardless of the total.

Criterion	1 - Weak	3 - Usable	5 - Strong	Score
Named outcome	Novelty without a baseline.	Task and user are clear.	Measured benefit and current cost are agreed.	—
Source visibility	Output has no evidence.	Some links or citations.	Every material answer traces to approved source content.	—
Permissions	Separate AI access model is unclear.	Main roles are respected.	LMS permissions are inherited, tested and logged by audience.	—
Human control	Acts before review.	Review exists for some actions.	Approval is meaningful, informed and required before material change.	—
Audit trail	Only final outcome is stored.	User and time are logged.	Inputs, model, sources, output, approver and action are reproducible.	—
Portability	Output is locked in.	Partial or manual export.	Usable standard export has been demonstrated with real content.	—

Criterion	1 - Weak	3 - Usable	5 - Strong	Score
Recovery	Errors require supplier intervention.	Manual correction is possible.	Rollback, outage behaviour and ownership are documented and tested.	—
Commercial fit	Price or tier is unclear.	Pilot cost is known.	Full-population, usage and exit costs are modelled in writing.	—

Decision bands

33-40 - Controlled candidate Proceed to a bounded pilot, still using explicit success and stop conditions.	25-32 - Resolve conditions Pilot only after the low-scoring controls have named owners and dates.
17-24 - Demonstration, not deployment Keep it in a test environment while evidence, permissions and recovery are improved.	8-16 - Decline The operating risk and lock-in outweigh the current benefit. Revisit when the product materially changes.

VETO CONDITIONS

Stop regardless of total score

Do not deploy where restricted content crosses permissions, material actions lack accountable approval, outputs cannot be corrected, or the organisation cannot tell what the feature changed.

05 / WORKING SECTION

Run a pilot that can genuinely say no

A pilot is useful only if it uses representative work, measures the current process and has stopping conditions agreed before enthusiasm takes over.

- 01** Choose one bounded task
Name the user, source, output and decision. Avoid a general assistant pilot with no defined job.
- 02** Measure the baseline
Record current time, completion rate, error or rework rate, user effort and any access or quality incidents.
- 03** Build the test set
Include normal cases, restricted content, ambiguous requests, missing data, outdated sources and users with different permissions.
- 04** Set review and escalation
Name who samples outputs, who can pause the pilot and how users report a harmful, incorrect or inaccessible result.
- 05** Run long enough to meet reality
Include ordinary work, busy periods and at least one model, source or integration failure where possible.
- 06** Decide with evidence
Compare against baseline, review exceptions and make enable, constrain, redesign or stop a real option.

Minimum pilot measures

Measure	Baseline	Target	Actual	Owner
Median time to complete the task	_____	_____	_____	_____
Outputs accepted without material correction	_____	_____	_____	_____
Unsupported or incorrect claims	_____	_____	_____	_____

Measure	Baseline	Target	Actual	Owner
Permission or inappropriate-disclosure incidents	—	0	—	—
User confidence and ability to challenge output	—	—	—	—
Cost per completed, accepted task	—	—	—	—

Pre-agreed stop conditions

☐	Restricted content is returned to a user who should not see it.
☐	A material action occurs without the required approval or audit record.
☐	Checking and correction take as long as the original task.
☐	The supplier changes the model, price or data terms beyond the pilot assumption.
☐	Users cannot understand the source, challenge the output or recover from an error.

06 / WORKING SECTION

Adopt a one-page operating policy

A short policy people can use is more valuable than a broad principle nobody can translate into a platform setting.

Policy area	Default Red6-style position
Purpose	AI features must have a named user, task, measurable benefit and accountable service owner.
Approved use	Read and retrieval over approved sources may be enabled when permissions and source links are demonstrated.
Drafting	AI-created drafts require a named reviewer with suitable subject, accessibility and platform knowledge before publication.

Policy area	Default Red6-style position
Material actions	Changes to access, assignments, completion, competency or compliance require explicit approval and a complete audit record.
People decisions	AI output does not make or silently determine employment, progression, performance or mandatory-training decisions.
Data	Users must not enter restricted, personal or client-confidential data unless that exact use has been approved and contracted.
Transparency	Users are told when they are interacting with AI, what sources it can use and how to report or challenge a result.
Monitoring	The owner samples quality, reviews incidents, tracks model or supplier changes and can pause the capability.
Exit	Material outputs, sources and logs remain exportable; deletion and supplier exit are tested before reliance grows.

Decision record for one feature

Feature, owner and intended outcome	
Classification: read / draft / write / act	
Approved data and users	
Required human approval	
Success and stop measures	
Review date and decision	

The goal is not to slow useful AI down. It is to make the safe path obvious and the consequential path accountable.