Responsible AI Policy

Vidyaluma — A product of KRAVIA PRIVATE LIMITED

Published policy · Effective from 1 November 2026

Approved: 1 November 2026

T03 · Version 1.0.0

Permitted AI assistance, prohibited uses, human authority, privacy, provider controls and correction.

On this page

1. Product scope and AI purpose
2. Permitted educational assistance
3. Decisions AI must not make
4. AI-assisted grading and published reports
5. Privacy and minimum data
6. Provider approval and change control
7. Retrieval, memory and school isolation
8. Connected tools and consequential actions
9. Accuracy, bias and transparency
10. Child safety and harmful content
11. Reporting, appeals and corrections
12. Audit, retention and evidence

▶ On this page

1. Product scope and AI purpose

1.a — Operator. Vidyaluma is a product of KRAVIA PRIVATE LIMITED. AI is used only within an enabled, reviewed school or service workflow and the permissions of its authorized user. The Policy covers generative models, OCR, extraction, recommendations, AI-assisted reports, retrieval and tool-connected assistants where those features are actually enabled.

1.b — Availability. Listing a use here defines its permissible boundary; it does not assert that every listed feature is live or supplied under every plan. Each production AI function must have a verified purpose, data flow, provider, retention arrangement and user disclosure before activation.

2. Permitted educational assistance

2.a — Drafting and understanding. AI may help authorized users prepare lesson materials, draft feedback, summarize relevant records, extract information, identify academic trends or find permitted school context. The task must need the information used and remain suitable for the intended audience.

2.b — Review. Teachers verify factual accuracy, age suitability, curriculum alignment, source relevance and licensing before distributing generated educational materials. Important outcomes require a qualified reviewer rather than a default approval checkbox. KRAVIA must make the draft/approved status and correction route understandable.

2.c — School choice. Authorized school administrators control optional AI enablement. A non-AI route should remain for essential workflows where supported, and limitations must be disclosed before reliance. Disabling optional AI must not silently change the text of a policy or authorize a different provider.

3. Decisions AI must not make

3.a — Academic and personal decisions. AI must not independently determine final examination results, admissions, discipline, guilt, suspension, expulsion or employment outcomes. It may organize relevant material for a permitted administrative task, but its output is not a binding decision or substitute for the authorized person's judgment.

3.b — Sensitive inference. Vidyaluma AI must not diagnose medical or psychological conditions, infer children's emotions or attention from images, identify students through facial recognition, or generate hidden behavioral/commercial profiles. Educational observations must not be presented as clinical findings or guaranteed predictions.

3.c — Identity and authority. AI must not adjudicate Aadhaar identity, validate school contracting authority, override failed document integrity, manufacture consent or execute an agreement. OVSE, DigiLocker, school-authority review and Leegality retain their separate deterministic and human controls.

4. AI-assisted grading and published reports

4.a — Source-based evaluation. Where grading assistance is enabled, the system must use the authorized answer sheet, marking rubric and relevant approved context. Poor scans, unclear handwriting, missing pages or uncertainty should be surfaced. A confident-looking explanation is not evidence that extraction or grading is correct.

4.b — Teacher approval. Suggested marks and feedback remain drafts until an authorized teacher verifies and approves them. The interface records the source, AI contribution, reviewer and material changes. Bulk approval requires meaningful review controls and must not hide low-confidence or incomplete cases.

4.c — Already published errors. A credible error report leads to appropriate review, preservation of the earlier state, authorized correction and notification of affected recipients where appropriate. The school makes final academic decisions; KRAVIA investigates attributable extraction, calculation, retrieval or publishing defects. Neither party may silently overwrite evidence to avoid accountability.

5. Privacy and minimum data

5.a — Task selection. Only information necessary for the permitted task may be provided to an AI model or retrieval tool. Access must be evaluated at the source and at use. A student's full profile, unrelated messages or another school's records must not be included by default to improve model context.

Skip to content

5.b — Restricted inputs. Health, disability, disciplinary, welfare and government-identity information is excluded from ordinary AI functions by default. Any exceptional permitted task requires a specific lawful purpose, authority, disclosure, access controls and privacy/security review. A general school-administrator role is not sufficient authorization.

5.c — No general-purpose training. KRAVIA and its AI providers must not train or fine-tune general-purpose models on private school records, prompts, conversations or outputs. Pseudonymized or lightly de-identified school content is not an automatic exception. A feedback button does not authorize a provider's training use.

5.d — Service processing versus training. Performing an authorized inference, storing an approved school report or investigating a specific incident is distinct from training. Each still requires necessary scope, access and retention. Provider abuse monitoring or human review must be disclosed and constrained; "no training" does not mean "no provider retention".

6. Provider approval and change control

6.a — Before egress. An AI provider and service tier require contract, privacy, security, processing-location, retention and training review before school content is sent. The verified register must identify actual providers; a model name or API key alone is not evidence of suitable terms.

6.b — Fallback. If a provider fails, an unreviewed backup model must not automatically receive private records. Affected work should pause, preserve progress, explain the failure and offer an authorized alternative where available. Incomplete output must not be published as a successful result.

6.c — Changes. New model versions or providers with material behavioral, data-use or location changes require reassessment and applicable school/user notice. Old approvals are not automatically transferred to a different commercial tier or provider policy.

7. Retrieval, memory and school isolation

7.a — Authorized retrieval. Retrieval-augmented generation must respect current school membership, document permissions, approved purpose and source revocation. Citations point to sources the recipient is authorized to see. A public source and a confidential source must not be treated as interchangeable simply because their text is similar.

7.b — Conversation history. Conversations have defined user and school access rules; they are not a shared cross-school memory. Administrators and support personnel do not automatically receive unlimited access to every personal conversation. Retention and deletion follow purpose, the relevant account and T07.

7.c — Source cleanup. Deleting or revoking a source must affect relevant retrieval indexes, embeddings and caches as well as the original file. A restored backup must reapply restrictions. The system must not reproduce deleted private material through stale retrieval state and then call the original deletion complete.

8. Connected tools and consequential actions

8.a — Permission checks. An assistant may propose only actions within its authorized tool scope. The backend independently checks identity, school, object, operation and policy. Text supplied in a document or model answer cannot elevate privileges or change the system's authority.

8.b — Explicit approval. Sending communications, changing records, approving payments, exporting data or invoking protected external workflows requires the appropriate human authorization and audit record. The approved action must match the actual request, recipients and payload; open-ended approval of an entire conversation is insufficient.

8.c — Nondelegable boundaries. AI cannot edit or publish Legal/Trust documents through a website, execute signatures for a person, supply their OTP, decide school authority or bypass identity/consent integrity controls. Failures must be safe and traceable rather than retried until a protection is accidentally bypassed.

9. Accuracy, bias and transparency

9.a — Limitations. AI can make factual, numerical, contextual and citation errors. Vidyaluma must communicate meaningful uncertainty and limitations, including unsupported languages, poor image quality and unsuitable uses. It must not fabricate a citation or claim evidence was reviewed when it was not.

9.b — Evaluation. Testing should assess representative handwriting, languages, age groups, subject matter and accessibility needs without unnecessary use of real child data. Identified bias or recurring error requires investigation and appropriate restriction, correction or retraining using only lawfully permitted non-school training sources consistent with this Policy.

9.c — Labels. Users should see when a response, assessment or report is AI-assisted, whether it is a draft and who approved a consequential result. Necessary labels and provenance must not disappear from a download in a way that makes unapproved output appear authoritative.

10. Child safety and harmful content

10.a — Child-facing use. Direct student AI access requires school authorization, age-appropriate design and relevant guardian consent where required. The system must not expose a child to adult-oriented features through an unreviewed generic chat interface or use age labels that contradict actual access.

10.b — Safeguards. The service must address exploitation, grooming, bullying, discriminatory abuse, dangerous misleading instructions and other inappropriate material. Safe responses, reporting and appropriate human escalation are required. No safety filter is represented as perfect or a replacement for school safeguarding and emergency channels.

10.c — Sensitive reports. Reported content is reviewed only by appropriate staff for the necessary purpose, with minimized evidence and protection of the reporter. A safety report is not consent to general model training or unrestricted provider access to the student's history.

11. Reporting, appeals and corrections

11.a — Reporting. Applicable AI features must provide an in-app way to report harmful or offensive output without forcing the user to leave the app. Users must also be able to flag inaccurate, biased or fabricated results. A stored report needs appropriate tracking, review ownership and follow-up; an inert button does not satisfy the intended control.

11.b — Response. KRAVIA investigates attributable system defects and may correct an output, restrict a feature, revise filters, withdraw a result or notify affected users. Schools handle academic reassessment. Appeals under L10 must allow a meaningful human review rather than rerunning the same model as the only decision.

11.c — Responsibility. Users should not rely on unverified consequential results, but required human review does not transfer every AI or software failure to the school. Liability, correction and remedies follow cause, control, applicable law and enforceable agreements.

12. Audit, retention and evidence

Skip to content

12.a — Audit events. Necessary audit records identify school/user context, task, relevant provider/model version, source/operation references, approvals, material changes and outcome. Raw prompts, sensitive documents and secrets are not logged by default. Access and retention are limited to their justified purpose.

12.b — Retention. Temporary inputs and processing artifacts are minimized and removed when no longer needed. Approved school outputs follow the relevant school-record schedule. A provider's separate log or backup retention must be verified and disclosed, and deletion must extend to permitted provider copies and retrieval material.

12.c — Accountability. Product, security, privacy and legal owners review material AI changes and substantiated safety issues. Founder approval and a controlled source release are required for this Policy. Production claims require evidence of the actual controls, not only this written commitment or a completed questionnaire.

Related documents
L01 — Privacy PolicyL02 — Terms & ConditionsL04 — Acceptable Use PolicyL05 — DisclaimerL09 — Data Processing AgreementL10 — Grievance, Complaints & Appeals PolicyT02 — Student & School Data Protection NoticeT04 — Data Use & Purpose Limitation StatementT05 — School Control, Access & Sharing StatementT06 — Subprocessor & International Processing NoticeT07 — Data Retention, Export & Deletion Policy

The current published version is shown. Earlier review copies do not create customer acceptance or contractual terms.

Vidyaluma

VidyaLuma

Clear school work, calmer daily decisions, and trusted records in one connected place.

Book a walkthrough

Contact support

Explore
About
Product
Pricing
Guides
Release notes
FAQs

School support
Book a walkthrough
Contact
Verify partner ID
Status

Trust
Security & Service Reliability Statement
Student & School Data Protection Notice
Responsible AI Policy
Data Use & Purpose Limitation Statement

Legal
Privacy Policy
Terms & Conditions
Cookie & Tracking Technologies Policy
Acceptable Use Policy
Disclaimer

Skip to content

School Control, Access & Sharing Statement

Subprocessor & International Processing Notice

Data Retention, Export & Deletion Policy

Vulnerability Disclosure Policy

Accessibility Statement

Billing, Subscription, Cancellation & Refund Policy

Support & Service Commitments Policy

Data Processing Agreement

Grievance, Complaints & Appeals Policy

School-controlled sharing **Role-aware access** **GST included in public pricing**

Built for school teams that need clear, trusted operations. Made in India, with love ❤ Copyright 2026 Vidyaluma. All rights reserved.

Legal Trust Privacy request Delete an account Service status Cookie choices Back to top ↑