

How we use AI technologies

Catapult uses AI-enabled services in limited, purpose-specific ways to support assessment review, academic integrity workflows, and internal content development. These tools are not intended to replace human judgement, trainer responsibility, assessment validation, or organisational decision-making.

All AI use must align with the following guiding principles:

- **Human Accountability:** Humans remain responsible for all outputs and decisions.
- **Data Protection First:** Confidential and regulated data must be safeguarded at all times.
- **Transparency:** AI use must be disclosed where relevant, internally or externally.
- **Accuracy & Validation:** Outputs must be verified before use.
- **Fairness & Non-Discrimination:** AI must not introduce bias or unfair outcomes.
- **Security by Design:** AI must not weaken company security controls.

Our approach is based on data minimisation. Where AI services are used, Catapult aims to send only the information required for the specific task being performed. Student profile information, platform identifiers, enrolment records, USI details, grades, course progress, account details, and other administrative records are not intentionally sent to AI providers unless explicitly required, approved, and documented.

AI outputs should be treated as decision-support information, not final decisions. Staff remain responsible for reviewing, validating, and applying any AI-generated result in line with assessment, compliance, privacy, and academic integrity requirements.

Where third-party AI providers are used, data is processed outside the Catapult application environment. Vendor handling of submitted content is governed by the relevant provider's terms, privacy policy, data processing arrangements, and configured account settings. Vendor assurances should be reviewed periodically, particularly where student data, assessment data, or customer content may be involved.

Winston plagiarism detection

Our integration with Winston AI is designed around data minimisation.

When a student response is submitted for AI detection, the only content sent to Winston AI is the student's written response text required to perform the AI-detection check. We do not send student profile information or identifiers such as the student's name, email address, date of birth, student number, USI, enrolment details, course progress, grades, account information, or other system records.

The student response is transmitted securely to the third-party AI service for processing and the result is returned to the platform. We treat this response text as potentially sensitive, even though we do not attach identifying student metadata to it.

As Winston AI is an external provider, data will be processed outside of the Catapult environment. Their published privacy information states that submitted text is analysed independently, is not compared against other users' submissions, and is not used to build or train AI models without explicit opt-in consent. Winston AI's privacy policy also states that submitted content is encrypted in transit and at rest, and that content is retained while the associated analysis report exists, with deletion from active storage within 30 days after report/account deletion and possible encrypted backup retention for up to 90 days.

From our side, the key privacy controls are:

1. only the student response text is sent;
2. no student identifiers or platform profile data are included;
3. the integration is used only for the requested AI-detection purpose;
4. results should be treated as a support tool rather than the sole basis for an academic integrity decision.

We recommend that the parent/guardian consent wording makes clear that the student's submitted written response may be sent to a third-party AI-detection provider for analysis, but that the platform does not include student-identifying information when making that request.

OpenAI assessment support

Our OpenAI integration is designed around data minimisation and is used to support assessment review against defined benchmarks.

When the feature is used, the platform sends OpenAI only the information required to evaluate the student's response against the relevant assessment criteria. This includes:

1. the assessment question;
2. the assessment benchmarks or marking criteria;
3. the student's submitted response; and
4. system instructions that tell OpenAI how to assess the response and return the result.

We do not intentionally send student profile information or platform identifiers such as the student's name, email address, date of birth, student number, USI, enrolment details, course progress, grades, account information, or other administrative records.

The student response itself may contain personal information if the student has entered that information into their answer. For that reason, we treat submitted response text as potentially sensitive, even where no student identifier is attached by the platform.

The information is transmitted securely to OpenAI for processing and a result is returned to the platform. The result is used to support assessment review against the supplied benchmarks and should not be treated as the sole basis for an academic or administrative decision without appropriate human oversight.

OpenAI states that, by default, it does not train its models on inputs or outputs from business products, including the API, unless the organisation explicitly opts in to share data, which Catapult has not opted in to. OpenAI's API data controls documentation also states that abuse-monitoring logs may contain prompts, responses, and related metadata, and are retained by default for up to 30 days unless a different approved retention control applies or longer retention is legally or operationally required.