

GradeBot Privacy Policy

Last Updated: November 2025

1. Scope of This Policy

GradeBot is an educational assessment and document-evaluation platform used by instructors, institutions, and authorized personnel ("Users"). GradeBot processes educational documents strictly for evaluation, scoring, and instructor workflows. Students do not access the service directly, and GradeBot does not create student user accounts.

This Privacy Policy describes how we collect, use, store, and share information when you use GradeBot.

2. Definitions

"Grading Session" means a period of activity during which an instructor uploads documents, configures grading or review options, and runs automated evaluation or feedback processes for a given cohort, assignment, or applicant set. A grading session ends when the instructor explicitly exits the session in the interface or when 24 hours have passed since the session was created, whichever occurs first.

"Student Data" means any information contained in documents uploaded by Users that relates to an identified or identifiable student or applicant, including names, IDs, emails, essays, assignments, application materials, and associated metadata.

3. Information We Collect

3.1 Instructor and Account Information

We collect the following information about Users who access GradeBot:

- Name and email address
- Microsoft Entra ID subject identifier and related authentication claims
- Organization or institution identifiers (where configured)

- User role (for example, instructor, admin, read-only)
- Assignment, cohort, and rubric configurations
- Unit purchase and usage tracking, including license and quota usage

3.2 Student and Applicant Information (Short-Term Only)

Instructors may upload documents that contain Student Data, such as:

- Student or applicant names, IDs, and email addresses
- Essays, assignments, application materials, and related attachments
- Other educational content included in uploaded documents

GradeBot does not collect Student Data directly from students. All Student Data is provided by instructors or authorized institutional personnel.

3.3 Usage, System, and Audit Information

We collect system and usage information, including:

- API and model usage metrics (for example, tokens used, operation types, and estimated cost)
- Audit logs for sensitive operations (for example, exports, bulk actions, configuration changes, login events)
- Session identifiers and workflow context for grading sessions
- Error reports and performance logs containing technical details about system operation

Audit logs and usage records are designed not to include full student documents or extracted student content.

4. How We Use Information

We use information for:

- Authentication and secure access control
- Configuring and running grading sessions
- Processing, evaluating, and scoring uploaded documents against rubrics
- Generating feedback, summaries, and screening recommendations

- Maintaining auditability, security, and operational monitoring
- Providing support, troubleshooting, and service improvements

Student Data is used only to perform the grading, evaluation, and feedback workflows requested by the institution or instructor. We do not use Student Data for marketing, advertising, or training generalized machine learning models.

5. Third-Party Service Providers

5.1 AI Processing

GradeBot uses third-party AI providers to perform automated text analysis and grading, including:

- OpenAI Enterprise APIs
- Google Gemini Enterprise APIs

When an instructor runs a grading session, content may be processed by integrated AI service providers (such as OpenAI or Google Gemini) solely to generate the requested output. These providers, under their enterprise terms, do not train on customer data or retain it beyond what is necessary to provide the service. These providers process the data solely to generate responses for the requested operation. Under their enterprise terms, they do not use customer content to train their foundation models or retain it beyond the processing necessary to provide the service.

5.2 Document Processing and OCR

GradeBot may use:

- Azure Document Intelligence for text extraction from PDFs and images.

Uploaded documents or their images may be transmitted to Azure Document Intelligence solely for OCR and content extraction during a grading session.

5.3 Authentication

GradeBot uses:

- Microsoft Entra ID for identity and access management.

Authentication tokens and identity claims are processed to verify user identity and enforce access controls.

5.4 General Restrictions

We do not sell Student Data, share Student Data with advertisers, or connect Student Data to social media platforms.

6. Data Storage and Retention

6.1 Short-Term Student Data Storage

Student Data is stored in encrypted Azure Blob Storage and related ephemeral caches only for the limited duration necessary to complete a grading session.

- Uploaded documents are stored under structured paths (for example, by session and submission) in encrypted Azure Blob Storage.
- Extracted text, OCR results, and intermediate processing artifacts are stored separately in short-lived storage.

Retention limits:

- Uploaded student or applicant documents are automatically deleted within **24 hours** of grading session creation or sooner when the session ends.
- Extracted text, OCR results, and intermediate processing data are automatically deleted within a maximum of **24 hours** or when the grading session ends, whichever occurs first.

Student documents and extracted student content are **not** included in long-term backups, analytics systems, or log archives.

6.2 Long-Term Instructor and System Data

The following data may be retained long-term:

- Instructor account profiles and roles
- Assignment, rubric, and cohort configurations
- System usage metrics and billing-related information
- Security and audit logs that track actions such as logins, exports, and configuration changes

These long-term datasets do **not** contain student documents or extracted student content. They may, however, reference identifiers such as session IDs, assignment names, and pseudonymous user identifiers.

7. Browser Storage and Client-Side Data

The GradeBot web application uses browser storage as follows:

- **sessionStorage** is used to temporarily store:
 - Authentication tokens
 - Active grading session identifiers and metadata
 - Assignment and cohort identifiers and names
 - Temporary evaluation results and navigation state

Data in sessionStorage is cleared when the user logs out or closes the browser. GradeBot does not use localStorage for persistent tracking and does not set its own tracking cookies. Microsoft authentication libraries and Microsoft Entra ID may set their own cookies solely for authentication and security.

We do not use cookies or similar technologies for advertising, cross-site tracking, or behavioral profiling.

8. Data Security

We implement a range of technical and organizational measures to protect information, including:

- TLS/HTTPS for all network communications
- Encrypted storage on Azure services, including Azure Blob Storage and PostgreSQL
- Scoped SAS tokens for secure file access
- Azure AD / Entra-based authentication and JWT validation
- Role-based access control for different user roles (such as instructor, admin, read-only)
- Structured audit logging of sensitive operations
- Containerized deployment and environment-based configuration management

While these safeguards are designed to be robust, no system can guarantee absolute security.

9. Access, Control, and Data Requests

9.1 Instructor and Institutional Control

Instructors and institutions may:

- Update or correct their account details
- Configure assignments, rubrics, and cohorts
- Run grading sessions and export results (for example, CSV, Excel, or PDF)
- Request deletion of their GradeBot accounts and associated instructor-level configurations

When an account is deleted, associated assignments, grading sessions, and related instructor-owned resources may also be deleted. Some audit log entries or non-student system records may be retained for security and compliance purposes, but they do not contain student documents or extracted student content.

9.2 Student Access and Requests

Students do not access GradeBot directly. Any student requests to access, correct, or delete educational records must be submitted through their educational institution. GradeBot will assist the institution, as appropriate, in responding to such requests under applicable law (for example, FERPA).

9.3 Institutional Requests

Institutions may request:

- Confirmation of data retention and deletion practices
- Exports of instructor account data, configuration data, or audit logs that relate to their tenancy
- Confirmation that Student Data has been processed and deleted according to this Policy

10. International Data Handling

GradeBot is deployed on Microsoft Azure infrastructure in regions configured at deployment time. Student Data is processed and stored only in those selected regions. Institutions are responsible for choosing deployment regions that meet their regulatory obligations.

11. Children's Privacy

GradeBot is intended for use by instructors, staff, and authorized institutional personnel. It is not designed for direct use by students or minors. Any Student Data processed by GradeBot must be uploaded and managed by authorized institutional personnel.

12. Changes to This Privacy Policy

We may update this Privacy Policy from time to time to reflect changes in the platform, legal requirements, or data protection practices. When we make material changes, we will notify institutional administrators or post an updated notice in the application indicating the effective date.

13. Contact Information

For privacy, security, or FERPA-related questions about GradeBot, or to exercise data-related rights through your institution, contact:

admin@gradebot.ai