grammarly

Grammarly Business Trust and Security FAQs

Table of Contents

Introduction

Grammarly delivers secure, trustworthy AI communication assistance for organizations that want to innovate and scale quickly. We uphold industry-leading security standards, keep customer data private and secure, and are committed to responsible AI usage and development. With Grammarly, businesses can speed up workflows and achieve better results with the peace of mind that their data is protected.

This document explains how Grammarly serves its customers, what data it collects, and the measures it takes to safeguard data and ensure customer privacy. The material in this document pertains to Grammarly Pro and Enterprise subscriptions purchased through Grammarly's Sales team and may not apply to those using Grammarly Free or those using Pro subscriptions purchased through Grammarly's website. Please contact us if you are interested in purchasing a volume license with our full suite of enterprise security features.

About Grammarly

Grammarly is the world's leading AI writing assistance company, trusted by over 30 million people and 70,000 professional teams. From instantly creating a first draft to perfecting every message, Grammarly helps people at 96 percent of the Fortune 500 get their point across—and get results—without compromising security or privacy. Our global team combines advanced machine learning with human expertise, breaking new ground in natural language processing and AI to offer unmatched writing assistance to individuals and enterprises. Grammarly is a leading AI SaaS brand in the workplace, offering ubiquity, contextual awareness, and deep investments in enterprise-grade security and privacy safeguards that set its service apart. Grammarly is one of *TIME*'s 100 Most Influential Companies, one of *Fast Company*'s Most Innovative Companies in AI, ranked seventh on the *Forbes* Cloud 100, and one of *Inc.*'s Best Workplaces.

Grammarly does not sell customers' data or rely on an ad-based revenue model. Grammarly makes money by selling subscriptions. Our users are our customers; they are not our product. Grammarly cares deeply about our customers' privacy and safety.

Grammarly has over fifteen years of experience in AI, with multiple dedicated teams focused on championing user privacy, maintaining strict adherence to information security, and applying research and expertise to minimize bias and apply customer feedback.

Privacy, safety, and security are embedded in our product development process. We prioritize upholding our users' trust with every new feature we develop, from concept to delivery. Fairness and safety are always paramount in our decisions and designs. Grammarly builds products and models with checks and balances to prioritize privacy, safety, and security.

Privacy

What data is sent to and processed by Grammarly?

When people use Grammarly, Grammarly processes the following types of data for the purpose of providing writing suggestions and, when enabled, generative AI–powered responses.

User text: User text is the raw text written by a user and checked by Grammarly's writing assistance service. Grammarly checks only text written in editable text fields. Grammarly products do not process any data from fields marked "sensitive," such as those requesting passwords, URLs, or credit card information.

Grammarly account information: This is information provided by a user during the registration process, such as their username, email address, contact, and language preferences.

Log data: Log data is technical information about service use, such as browser type and version, device information, and statistical data about product interactions, such as rates of accepted or rejected writing suggestions. Grammarly uses this information for quality and regression diagnosis.

User customization data: This data is optional user and administrator inputs that customize the suggestions provided by Grammarly, which operates at different levels. These inputs include additions to the personal dictionary at the individual user level and style guide entries at the team admin level.

No institutional customer text is utilized in training our algorithms or machine learning models. User text sent to Grammarly through the Grammarly browser extension, Grammarly for Windows and Mac desktop app, or Grammarly Keyboard on mobile devices is processed solely for the purpose of providing writing suggestions. User text is not stored, it is erased, once writing suggestions have been provided.

Any information used to power Grammarly's generative AI features, such as prompt type, prompt text, and the context in which it's used, is de-identified and captured for the sole purpose of providing customers with Grammarly's generative AI experience. We use a very small number of thoroughly vetted partners for processing. Grammarly does not allow third parties to use customer data to train their models or improve their products, and no data is retained by Microsoft Azure, the cloud service provider we use for secure subprocessing and large language model (LLM) support. Individual users and administrators can control the apps and websites where Grammarly operates in order to limit data processing.

How is data collected?

Data is collected through Grammarly's software clients, which include browser extensions, desktop applications (Grammarly for Windows and Mac), a Microsoft Office add-in, a web-based editor, mobile keyboards, and an iPad application. Grammarly products do not collect any data from fields marked as sensitive, such as those requesting passwords or credit card information. Users can determine whether Grammarly is running in a particular text field on a site or an application by looking for the Grammarly logo, a minimized Grammarly status dot, or an icon showing the number of suggestions needing attention. User text is used only to provide writing suggestions and is not stored after that purpose has been served.

What additional protections are in place for sensitive data?

Grammarly Enterprise plans include additional data loss prevention (DLP) application controls. These controls ensure client apps sanitize personally identifiable information (PII) that comes from end users before it is ever sent to Grammarly processing servers. This guarantees that certain types of data do not leave the customer's environment. Currently, Grammarly sanitizes PII tokens that contain email addresses, URLs, phone numbers, credit card numbers, US social security numbers, UK national insurance numbers, Canada social insurance numbers, and international bank account numbers (IBAN).

How is this data sent, processed, and stored?

All data is transferred to the United States for processing using Amazon Web Services, one of the world's leading data center providers.

- Data in transit is protected by up-to-date encryption protocols (including SSL/TLS 1.2).
- Data at rest is encrypted using the industry-standard AES-256 algorithm.
- Passwords are hashed using the bcrypt algorithm.
- Grammarly utilizes AWS Key Management Services (KMS) for database encryption and key management. Access to the cryptographic keys is restricted to authorized personnel.
- Each Grammarly customer's data is segregated logically from other users' data.
- Any writing that an individual or organization reviews with Grammarly will never appear in another customer's writing suggestions.

What data does Grammarly store?

Grammarly stores documents created in the Grammarly Editor until they are deleted by the user or upon request after contract termination or expiration. Enterprises can choose to block their organization's access to the Grammarly Editor.

All other user text processed by Grammarly (i.e., anything not saved in the Grammarly Editor) is not stored after it is processed to provide writing suggestions. Grammarly does not retain user text for any reason and does not train or improve algorithms with customer data.

Does Grammarly sell the data I share?

No. Grammarly has not sold, does not sell, and will not sell customer data. We make money from selling subscriptions and do not allow third parties to advertise to our customers. For more information, please see the Grammarly Business Security Systems and Data Flows diagram.

Is Grammarly a keylogger?

No. A keylogger records every keystroke and sends data to a third party for the benefit of that party and does so without the user's knowledge. Grammarly does not do this.

Grammarly is blocked from running in read-only and sensitive fields, such as payment forms, passwords, addresses, and URLs. Grammarly does not record every keystroke.

We make it clear when Grammarly is active by turning the Grammarly button green. Users can always control where Grammarly can and cannot operate. When you use Grammarly for Windows and Mac, you can easily turn Grammarly off at any time and then turn it back on within a document or site using the easy-to-find Grammarly icon, which appears as a *G*. When you use the Grammarly extension, you can click the Grammarly icon located on the top bar of your web browser.

How does Grammarly keep data shared with an LLM stored securely and ensure it's not being shared with the broader public?

As with all Grammarly features, Grammarly applies stringent principles and safeguards to the flow of data for its generative AI features to maintain privacy and security. User data is de-identified by Grammarly before being passed through Microsoft Azure. Certain forms of personally identifiable information are also scrubbed from user data before it's passed to Microsoft Azure OpenAI Service.

To deliver generative AI, Grammarly uses Microsoft Azure OpenAI Service as its LLM. Microsoft was rigorously vetted and upholds the same high standards and credentials as Grammarly. Microsoft is contractually prohibited from using any user text from Grammarly to train its models and is not allowed to retain any data Grammarly sends.

Can someone at Grammarly access my data?

No one at Grammarly can read your text at will. Grammarly tightly controls access to user data within the company. Only those who have an approved need to access certain data are granted access via specific, audited permissions.

Access to data requires review and approval by responsible managers and a member of the Security team. For more information, please see the Grammarly Business Security Systems and Data Flows diagram.

Security

What security certifications does Grammarly have?

Grammarly has completed and annually maintains a SOC 2 (Type 2) attestation. This examination, conducted by Ernst & Young, validates that Grammarly meets the strict SOC 2 standards for security, availability, confidentiality, and privacy of customer data.

Grammarly is certified by the Department of Commerce for the Data Privacy Framework (DPF), for EU-US, Swiss-US, and UK extension, providing a transatlantic personal data transfer mechanism.

Grammarly has obtained and annually maintains ISO 27001, 27017, 27018, and 27701 certifications, demonstrating our commitment to information security, cybersecurity, and privacy protection.

Grammarly complies with HIPAA Security, Privacy, and Breach notification rules, as attested to independently by Ernst & Young annually.

Grammarly complies with the Payment Card Industry's Data Security Standard, which validates that payments are handled with industry-standard security. Read Grammarly's attestation of PCI compliance, which is renewed annually.

In addition, Grammarly is a Trusted Cloud Provider of the Cloud Security Alliance (CSA), a nonprofit dedicated to promoting secure cloud computing. Grammarly has completed CSA's Consensus Assessments Initiative Questionnaire (CAIQ), which details Grammarly's security practices. Read Grammarly's CAIQ.

What other certifications does Grammarly have?

Grammarly also maintains state-level certifications including TX-RAMP and compliance with the California Consumer Privacy Act. Please get in touch with your Grammarly representative if you're interested in learning more.

Is Grammarly GDPR compliant?

Yes. Grammarly complies with the EU General Data Protection Regulation (GDPR) for the collection, use, and retention of personal information. For more detail, see Grammarly's Privacy Policy. Committed to the GDPR principles, Grammarly never collects personal data without a lawful basis, limits the amount of collected and processed data, and deletes the data when it is no longer needed for the services outlined in Grammarly's Privacy Policy. Users can request a personal data report through this link.

Can users delete their account and data?

Enterprise customers and individual users can end their Grammarly subscription at any time. At the end of an Enterprise subscription, any documents stored in the Grammarly Editor are deleted from end-user accounts. These user accounts then become Grammarly Free accounts. In accordance with the Privacy Policy, when an individual user deletes their account, the user's personal data is deleted from the internal and external services. In certain limited circumstances, Grammarly may retain data to comply with legal obligations and for fraud detection and prevention.

How can I ensure Grammarly is not training on my organization's user text and that these claims have been authenticated by a third party?

All customer text in business accounts purchased through Grammarly's Sales team and accounts at educational institutions will not be used for product improvements or model training purposes, automatically.

This can be seen in an admin's account hub under **Administration** > **Data Settings**. By default, Grammarly will not train on any customer text for any account in the EU and UK.

Pro accounts purchased through Grammarly's website may switch off **Product Improvement & Training** under **Data Settings** if they wish to opt out.

What does Grammarly do to protect against unauthorized developer and user access?

Grammarly supports multi-factor authentication (MFA) and requires FIDO2 for all employees. Engineers do not have persistent access to production.

End users who are members of a Grammarly Enterprise or Pro account can log in to Grammarly through their identity provider if single sign-on (SSO) has been set up for their account. Any role-based access controls (admin/account manager/user) set up by the administrators of a company's Grammarly Enterprise account will also apply.

What does Grammarly do internally to control access to data?

Grammarly manages internal systems with SSO and mandatory MFA. Only company-managed devices can connect to the Grammarly corporate network, via required biometric or physical key authentication methods that meet FIDO2 specifications. Grammarly adheres to the principle of least privilege, giving employees only the access necessary to perform their work. Grammarly engineers do not have persistent access to production systems. Any access granted is time-bound, requires business justification and approval, is logged via an audit trail, and is automatically removed after a predefined number of hours.

What does Grammarly do to train employees to follow security best practices?

All new employees are required to complete online data security and privacy training. Grammarly employees are provided continuous education throughout the year through phishing training, a Security Champions program, and weekly privacy and security knowledge sharing. A dedicated in-house offensive security team carries out end-to-end red team operations on an ongoing basis to reinforce security awareness and practices.

Grammarly's engineers follow our established secure development guidelines, and all changes undergo design and peer review before going to production.

What third-party audits and tests does Grammarly undergo?

In addition to the previously mentioned security attestations and certifications, Grammarly contracts globally recognized third-party penetration testing firms and undergoes security reviews of our entire cloud and corporate infrastructure and applications.

- Grammarly uses BitSight's vendor assessment system as part of a broader program to evaluate our security maturity as well as the security of our suppliers.
- Grammarly's technical security posture is assessed on the HackerOne platform on an ongoing basis.

What is Grammarly's incident management procedure?

In the case of a security incident, Grammarly's documented incident management procedure establishes channels for identifying and communicating the incident to Grammarly's Security team. The Security team defines the type of event and its severity and then responds to it according to approved service-level agreements (SLAs) based on industry best practices. Grammarly's Legal team is consulted on all incidents to assess the necessity and manner of reporting and remediation.

Security events that impact privacy are subject to additional analysis and response by Grammarly's Legal team.

Grammarly's incident management procedure can be found on Grammarly's security portal. You can obtain access to the portal through your Grammarly representative.

What is Grammarly's breach notification policy?

Grammarly will notify customers promptly upon becoming aware of a breach—where feasible, within forty-eight hours.

What subprocessors does Grammarly use and for what purposes?

Grammarly relies on a limited number of subprocessors for specific services and functions.

As part of our vendor approval process, we conduct multistep security and privacy assessments, a detailed review of each vendor's compliance posture, and an in-depth legal review of each vendor's data practices. Grammarly repeats this due diligence annually. For more information, please read "Does Grammarly use subprocessors?"

What LLM provider does Grammarly use for its generative AI features?

For its generative AI features, Grammarly uses the Microsoft Azure service as a subprocessor and LLM provider. We use a dedicated Azure instance, AI, and machine learning technologies, including GPT 3.5-Turbo.

Responsible AI

Does Grammarly have processes to ensure its AI product and service offerings are being designed ethically?

Grammarly has a dedicated Responsible AI team focused solely on ensuring all features are built with fairness and safety. This team is actively involved in building new features from concept to delivery.

Every new Grammarly feature undergoes a rigorous risk-assessment process, including a hands-on review by Grammarly's linguists to identify potential risks, such as issues with bias and fairness. Following the assessment, teams are required to make updates to address discovered risks. Users can also provide feedback directly in Grammarly's product to indicate when something is incorrect or offensive.

What processes does Grammarly have in place to ensure algorithmic decisions supported by its offerings do not create discriminatory or unjust impacts across different demographic lines (race, sex, etc.)? How is Grammarly developing monitoring and accounting mechanisms to avoid unintentional discrimination?

Grammarly's Responsible AI team is involved in all aspects of product development to prevent bias and ensure fairness throughout the product development process. This team also provides self-serve tools for other teams to conduct these reviews independently. Grammarly has also created a content filtering and moderation solution that is both algorithmic and machine learning based.

All features go through a rigorous risk assessment before launch, and any discovered issues must be resolved. The Responsible AI team continues to monitor for potential issues after deployment. In these reviews, Grammarly specifically measures for bias based on demographic information contained within generated text.

All inputs sent to Grammarly's generative AI provider, Azure OpenAI, are used transiently only. There is no storage or any training on any piece of data that flows from Grammarly.

Does Grammarly have any existing best practices for the detection, identification, and mitigation of unfair bias in AI models?

Using a combination of technologies, Grammarly filters generative AI and natural language suggestions with the aim of preventing issues such as biased content or hate speech. Grammarly has a proprietary content filtering and moderation solution that is both algorithmic and machine learning based, designed to limit instances of higher-risk language that could cause harm. We also aim to prevent issues such as hate speech and harmful misinformation through AI hallucinations and design generative AI prompts and post-process output to ensure safe results. Grammarly's integrations and models help generate more effective text and reduce risks in generative text output.

Grammarly is committed to building models using quality datasets that undergo bias and fairness evaluations. We design and develop products with our team of analytical linguists, who apply research and expertise to minimize bias and apply user feedback. Grammarly has internal responsible AI standards that product teams follow. We maintain a risk assessment process to evaluate for and resolve ethical issues in all products and perform regular quality evaluations. Additionally, we have a sensitivity support process in place to address user-reported issues.

Where can I obtain further documentation about Grammarly's trust posture?

More information can be found at Grammarly's Trust Center.

You can access Grammarly's information security documentation through the online Grammarly security portal. Please contact your Grammarly representative to obtain access.

Grammarly Enterprise Trust and Security FAQs

grammarly