API data usage policies

Updated
June 14, 2023

Starting on March 1, 2023, we are making two changes to our data usage and retention policies:

1. OpenAI will not use data submitted by customers via our API to train or improve our models, unless you explicitly decide to share your data with us for this purpose. You can opt-in to share data.
2. Any data sent through the API will be retained for abuse and misuse monitoring purposes for a maximum of 30 days, after which it will be deleted (unless otherwise required by law).

The OpenAI API processes user prompts and completions, as well as training data submitted to fine-tune models via the Files endpoint. We refer to this data as API data.

By default, OpenAI will not use API data to train OpenAI models or improve OpenAI's service offering. Data submitted by the user for fine-tuning will only be used to fine-tune the customer's model. However, OpenAI will allow users to opt-in to share their data to improve model performance. Sharing your data will ensure that future iterations of the model improve for your use cases. Data submitted to the API prior to March 1, 2023 (the effective date of this change) may have been used for improvements if the customer had not previously opted out of sharing data.

OpenAI retains API data for 30 days for abuse and misuse monitoring purposes. A limited number of authorized OpenAI employees, as well as specialized third-party contractors that are subject to confidentiality and security obligations, can access this data solely to investigate and

verify suspected abuse. OpenAI may still have content classifiers flag when data is suspected to contain platform abuse. Data submitted by the user through the Files endpoint, for instance to fine-tune a model, is retained until the user deletes the file.

Note that this data policy does not apply to OpenAI's Non-API consumer services like ChatGPT or DALL·E Labs. You can learn more about these policies in our data usage for consumer services FAQ.

Frequently asked questions

What technical protections and security certifications does OpenAI have in place surrounding the public API?

The OpenAI API is SOC 2 Type 2 compliant and has been audited by an independent third-party auditor against the 2017 Trust Services Criteria for Security.

Where is API data stored?

Content is stored on OpenAI systems and our sub-processors' systems. We may also send select portions of de-identified content to third-party contractors (subject to confidentiality and security obligations) safety purposes. Our 30-day data retention policy also applies to our sub-processors and contractors. You can view our list of sub-processors and their locations for details.

When I call the API, is the data encrypted in transit?

The OpenAI API is only available over Transport Layer Security (TLS), and therefore customer-to-OpenAI requests and responses are encrypted.

Do you offer EU data residency?

Not presently. We continue to explore the feasibility and impact of enabling localized data storage.

Can I use the API for HIPAA workloads?

We are able to sign Business Associate Agreements in support of customers' compliance with the Health Insurance Portability and Accountability Act (HIPAA). To qualify for this, you must have an Enterprise Agreement with OpenAI and a qualifying use case. Please reach out to our sales team if you are interested.

I've received a data deletion request, how do I ask OpenAI to delete data?

We only retain data sent through the API for up to 30 days for abuse and monitoring purposes. If you would like to delete your account before then, please follow these steps.

Does OpenAI have a Data Processing Addendum (DPA)?

Yes. Please complete our DPA form to execute our Data Processing Addendum.

Can we self-host?

We do not offer on-premise hosting. You may purchase dedicated capacity by reaching out to our sales team.