Document officiel d'OpenAI, copie du 18/11/2026. Source : https://developers.openai.com/api/docs/guides/your-data

Data controls in the OpenAI platform

> For the complete documentation index, see llms.txt. Markdown versions of documentation pages are available by appending `.md` to the page URL.

Understand how OpenAI uses your data, and how you can control it.

Your data is your data. As of March 1, 2023, data sent to the OpenAI API is not used to train or improve OpenAI models (unless you explicitly opt in to share data with us).

Types of data stored with the OpenAI API

When using the OpenAI API, data may be stored as:

- **Abuse monitoring logs:** Logs generated from your use of the platform, necessary for OpenAI to enforce our Usage Policies and agreements and mitigate harmful uses of AI.
- **Application state:** Data persisted from some API features in order to fulfill the task or request.

Data retention controls for abuse monitoring

Abuse monitoring logs may contain certain customer content, such as prompts and responses, as well as metadata derived from that customer content, such as classifier outputs. By default, abuse monitoring logs are generated for all API feature usage and retained for up to 30 days, unless longer retention is required by law, or is reasonably necessary to protect our services or any third party from harm.

Eligible customers may have their customer content excluded from these abuse monitoring logs, subject to the limitations below, by getting approved for the Zero Data Retention or Modified Abuse Monitoring controls. Currently, these controls are subject to prior approval by OpenAI and acceptance of additional requirements. Approved customers may select between Modified Abuse Monitoring or Zero Data Retention for their API Organization or project.

Customers who enable Modified Abuse Monitoring or Zero Data Retention are responsible for ensuring their users abide by OpenAI's policies for safe and responsible use of AI and complying with any moderation and reporting requirements under applicable law.

Get in touch with our sales team to learn more about these offerings and inquire about eligibility.

Modified Abuse Monitoring

Modified Abuse Monitoring excludes customer content (other than image and file inputs in rare cases, as described below) from abuse monitoring logs across all API endpoints, while still allowing the customer to take advantage of the full capabilities of the OpenAI platform.

Zero Data Retention

Zero Data Retention excludes customer content from abuse monitoring logs in the same way as Modified Abuse Monitoring.

Additionally, Zero Data Retention changes some endpoint behavior: the `store` parameter for `/v1/responses` and `v1/chat/completions` will always be treated as `false`, even if the request attempts to set the value to `true`.

Besides those specific behavior changes, the endpoints and capabilities listed as No for Zero Data Retention Eligible in the table below may still store application state, even if Zero Data Retention is enabled.

ZDR with Private Safety Processing

Zero Data Retention with Private Safety Processing enables OpenAI to perform automated safety monitoring while preserving Zero Data Retention protections. Endpoint and feature limitations listed on this page still apply.

Customers using ZDR with PSP must configure customer-controlled storage and meet additional technical and operational requirements described in the ZDR with Private Safety Processing guide.

Private Retention with Private Safety Processing (fka Eyes Off)

For customers approved for Zero Data Retention or Modified Abuse Monitoring, we reserve the right to make models ineligible for Zero Data Retention or Modified Abuse Monitoring for specific customers, as notified in advance to the impacted customers in writing. In this instance, customer content will be retained in encrypted abuse monitoring logs in OpenAI-managed infrastructure, but such content will be excluded from human review unless required by applicable law. For additional information, see Appendix A of the Private Safety Processing technical whitepaper.

For customers who have executed an OpenAI Business Associate and Healthcare Addendum, once your org ID is provisioned with Private Retention with Private Safety Processing, BAA-eligible endpoints can be used for processing PHI, even if data is retained. Endpoint and feature limitations listed on this page still apply.

Safety Retention

For customers approved for Zero Data Retention or Modified Abuse Monitoring, we reserve the right to make models ineligible for Zero Data Retention or Modified Abuse Monitoring for specific customers if reasonably necessary to investigate or prevent severe risk activity, as notified in advance to the impacted customers in writing. In this instance, we may retain and human review customer content when using these models that our classifiers detect as potentially violating our Usage Policies or your agreement. Otherwise retention will not be affected. For customers who have executed an OpenAI Business Associate and Healthcare Addendum, once your org ID is provisioned with Safety Retention, BAA-eligible endpoints can be used for processing PHI, even if data is retained.

Configuring data retention controls

Once your organization has been approved for data retention controls, you'll see a **Data Retention** tab within Settings → Organization → Data controls. From that tab, you can configure data retention controls at both the organization and project level.

- **Organization-level controls:** Choose between Zero Data Retention or Modified Abuse Monitoring for your entire organization.
- **Project-level controls:** For each project, select `default` to inherit the organization-level setting, explicitly pick Zero Data Retention or Modified Abuse Monitoring, or select **None** to disable these controls for that project.

Storage requirements and retention controls per endpoint

The table below indicates when application state is stored for each endpoint. Zero Data Retention eligible endpoints do not retain any customer content for application state, subject to the limitations below. Zero Data Retention ineligible endpoints or capabilities may retain application state when used, even if you have Zero Data Retention enabled.

Endpoint	Data used for training	Abuse monitoring retention	Application state retention	Zero Data Retention eligible	Private Retention with PSP and Safety Retention eligible
`/v1/chat/completions`	No	30 days	None, see below for exceptions	Yes, see below for limitations	Yes, see below for limitations
`/v1/responses`	No	30 days	None, see below for exceptions	Yes, see below for limitations	Yes, see below for limitations
`/v1/decisions`	No	30 days	None, see below for exceptions	Yes, see below for limitations	Yes, see below for limitations
`/v1/conversations`	No	Until deleted	Until deleted	No	No
`/v1/conversations/items`	No	Until deleted	Until deleted	No	No

Endpoint	Data used for training	Abuse monitoring retention	Application state retention	Zero Data Retention eligible	Private Retention with PSP and Safety Retention eligible
`/v1/chatkit/threads`	No	Until deleted	Until deleted	No	No
`/v1/agents`	No	30 days	Until deleted	No	No
`/v1/assistants`	No	30 days	Until deleted	No	No
`/v1/threads`	No	30 days	Until deleted	No	No
`/v1/threads/messages`	No	30 days	Until deleted	No	No
`/v1/threads/runs`	No	30 days	Until deleted	No	No
`/v1/threads/runs/steps`	No	30 days	Until deleted	No	No
`/v1/vector_stores`	No	30 days	Until deleted	No	No
`/v1/images/generations`	No	30 days	None	Yes, see below for limitations	No
`/v1/images/edits`	No	30 days	None	Yes, see below for limitations	No
`/v1/embeddings`	No	30 days	None	Yes	No
`/v1/audio/transcriptions`	No	None	None	Yes	No
`/v1/audio/translations`	No	None	None	Yes	No
`/v1/audio/speech`	No	30 days	None	Yes	No
`/v1/audio/voices`	No	30 days	Until deleted	No	No
`/v1/files`	No	30 days	Until deleted*	No	No
`/v1/fine_tuning/jobs`	No	30 days	Until deleted	No	No
`/v1/evals`	No	30 days	Until deleted	No	No
`/v1/batches`	No	30 days	Until deleted	No	No
`/v1/moderations`	No	None	None	Yes	No
`/v1/completions`	No	30 days	None	Yes	No
`/v1/live/sessions`	No	30 days	None, or 30 days if stored	Yes, with limitations below	No
`/v1/realtime`	No	30 days	None	Yes	No

`/v1/chat/completions`

- Audio outputs application state is stored for 1 hour to enable multi-turn conversations.
- When Zero Data Retention is enabled for an organization, the `store` parameter will always be treated as `false`, even if the request attempts to set the value to `true`.
- See image and file inputs.
- Prompt caching may store encrypted key/value tensors in GPU-local storage as application state. This data is stored on the local GPU machines and is not retained after the 24-hour expiration. For `gpt-5.5` and `gpt-5.5-pro`, setting `prompt_cache_retention` to `in_memory` returns an error. For GPT-5.6 models and later model families, `prompt_cache_options.ttl` controls the minimum cache lifetime, not this maximum application-state retention period. To learn more, see the prompt caching guide.

`/v1/responses`

- Except as noted below, the Responses API has a 30 day Application State retention period by default, or when the `store` parameter is set to `true`. In those cases, response data will be stored for at least 30 days.
- When Zero Data Retention is enabled for an organization, the `store` parameter will always be treated as `false`, even if the request attempts to set the value to `true`.
- Background mode with the default completion window stores response data to disk for roughly 10 minutes to run the request and make the response available for retrieval. For projects using Modified Abuse Monitoring,

including enhanced Modified Abuse Monitoring, foreground requests follow standard retention when `store` is omitted or set to `true`. With the default completion window, background responses follow the standard retention period only if you explicitly set `store=true`. If you omit `store` or set it to `false`, the response is deleted after the temporary retention period.

- Projects with Zero Data Retention can use the `2h` and `8h` background completion windows. The API treats `store` as `false` and temporarily stores application state to process the request and make the response available for retrieval.
- Audio outputs application state is stored for 1 hour to enable multi-turn conversations.
- See image and file inputs.
- MCP servers (used with the remote MCP server tool) are third-party services, and data sent to an MCP server is subject to their data retention policies.
- Hosted containers used by Hosted Shell and Code Interpreter may write temporary application state to the container filesystem (backed by ephemeral block storage) while the container is active. Container data is deleted when the container expires or is explicitly deleted.
- Prompt caching may store encrypted key/value tensors in GPU-local storage as application state. This data is stored on the local GPU machines and is not retained after the 24-hour expiration. For `gpt-5.5` and `gpt-5.5-pro`, setting `prompt_cache_retention` to `in_memory` returns an error. For GPT-5.6 models and later model families, `prompt_cache_options.ttl` controls the minimum cache lifetime, not this maximum application-state retention period. To learn more, see the prompt caching guide.
- For server-side compaction, no data is retained when `store="false"`.
- We support Skills in two form factors, both local execution and hosted container-based execution. Hosted skills follow the same container lifecycle as hosted shell: mounted skills and container files remain available while the container is active and are discarded when the container expires or is deleted.
- Data transmitted to third-party services over network connections is subject to their data retention policies.

`/v1/decisions`

By default, abuse monitoring logs are retained for up to 30 days.

The Decisions API supports Zero Data Retention, Private Retention with PSP, and Safety Retention, subject to the eligibility and account configuration requirements described above.

The Decisions API follows the same Private Retention with PSP policy as the Responses API. Flagged content and a sample of requests may be retained for automated safety review. See Private Retention with PSP for details.

The following limitations apply:

- Prompt caching may store encrypted key/value tensors in GPU-local storage as application state. This data is stored on the local GPU machines and is not retained after the 24-hour expiration. To learn more, see the prompt caching guide.
- See image and file inputs for the CSAM retention exception for image inputs.

The Decisions API is eligible for HIPAA use under an executed OpenAI Business Associate and Healthcare Addendum, subject to the applicable account configuration requirements.

`/v1/assistants`, `/v1/threads`, and `/v1/vector_stores`

- Objects related to the Assistants API are deleted from our servers 30 days after you delete them via the API or the dashboard. Objects that are not deleted via the API or dashboard are retained indefinitely.

`/v1/images`

- Image generation is Zero Data Retention compatible when using `gpt-image-2.5-sunburst`, `gpt-image-2.5-sunburst-2026-09-08`, `gpt-image-2.5-flare`, `gpt-image-2.5-flare-2026-09-08`, `gpt-image-2`, `gpt-image-1.5`, and `gpt-image-1-mini`.

`/v1/files`

- Files can be manually deleted via the API or the dashboard, or can be automatically deleted by setting the `expires_after` parameter. See here for more information.

Historical Videos API retention

Before the September 24, 2026 shutdown, the Videos API documentation specified 48 hours for downloading generated videos, followed by 30 days of retention for abuse monitoring. These periods describe the policy documented before shutdown; they do not promise download access after shutdown. See the Videos API shutdown notice.

Image and file inputs

Images and files may be uploaded as inputs to `/v1/responses` (including when using the Computer Use tool), `/v1/chat/completions`, and `/v1/images`. Images may also be uploaded to `/v1/decisions`. Image and file inputs are scanned for CSAM content upon submission. If the classifier detects potential CSAM content, the image will be retained for manual review, even if Zero Data Retention, Modified Abuse Monitoring, or Private Retention with PSP is enabled.

Web Search

Web Search with live internet access is not HIPAA eligible and is not covered by a BAA. Web Search in offline/cache-only mode (`external_web_access: false`) is eligible to be covered by a BAA when used with an API key from a ZDR-enabled project within a ZDR organization. This HIPAA/BAA guidance applies only to the Responses API `web_search` tool. Note: Preview variants (`web_search_preview`) ignore this parameter and behave as if `external_web_access` is `true`. We recommend using `web_search`.

<span id="how-does-data-residency-work" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="sub-processors-and-regional-request-processing" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="additional-requirements-for-non-us-regions" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="how-to-use-data-residency" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="select-a-processing-region-per-request" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="which-models-and-features-are-eligible-for-data-residency" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="support-by-region" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="api-endpoint-tool-and-model-support" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="endpoint-limitations" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="v1decisions-1" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="v1chatcompletions-1" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="v1responses-1" className="scroll-mt-[var(--docs-scroll-offset)]"

<span id="v1livesessions" className="scroll-mt-[var(--docs-scroll-offset)]"

Data residency controls

See the Data residency guide for regional storage and processing options, eligibility, setup instructions, and supported models and endpoints.

Enterprise Key Management (EKM)

Enterprise Key Management (EKM) allows you to encrypt your customer content at OpenAI using keys managed by your own external Key Management System (KMS).

Once configured, EKM applies to any application state created during your use of the platform. See the EKM help center article for more information about how EKM works, and how to integrate with your KMS provider.

EKM limitations

OpenAI supports Bring Your Own Key (BYOK) encryption with external accounts in AWS KMS, Google Cloud (GCP), and Azure Key Vault. If your organization leverages a different key management service, those keys need to be synced to one of the supported cloud KMS providers for use with OpenAI.

EKM does not support the following products. An attempt to use these endpoints in a project with EKM enabled will return an error.

- Assistants (/v1/assistants)
- Vision fine tuning