# AI procurement needs four separate data checks

Training, retention, access and onward transfers require different controls.

RESEARCH BRIEFING · GOVERNANCE · 11 NOVEMBER 2026

A supplier’s promise not to train on company data does not answer how long that data is stored, who can access it or which connected services receive it. Procurement teams need separate answers to all four questions.

## Key points

• Approve a specific service configuration, rather than a provider name.

• Map data sent to connected tools and additional recipients.

• Test permissions, export and deletion before expanding use.

## A 30-day period is not a universal rule

OpenAI’s documentation says default API abuse-monitoring logs may be retained for up to 30 days, with legal and harm-prevention exceptions. Application state follows endpoint-specific rules; material sent to third-party tools follows those parties’ policies. These are provider statements checked on 11 November 2026. [1]

A buyer should therefore identify the product, endpoint, feature, retention setting and contract. An approved chat configuration does not automatically approve uploaded files, persistent agents or an external connector.

## Ask four questions

Training: may prompts, files, outputs or feedback be used to improve models? Record the applicable default, opt-in and contractual terms.

Retention: what persists in application state, monitoring logs, backups and third-party systems? Obtain deletion periods, exceptions and a test procedure.

Access: which customer roles, support personnel and subprocessors can see the data? Check permissions with accounts representing actual users.

Transfers: where is information processed or made accessible, and to which entities? The EDPB says qualifying transfers outside the EEA must satisfy GDPR Chapter V as well as ordinary GDPR requirements. A hosting-region label cannot answer the full question. [2]

## Separate reading from action

An assistant that retrieves documents has a different authority from one that sends messages or changes records. Approve those capabilities separately. Test whether a user can retrieve restricted material or cause an unauthorised action through generated instructions.

NIST's generative-AI profile identifies risks involving confabulation, information security and human–AI interaction. The framework supports a risk review; it does not certify a vendor configuration. [3]

## Require an operating record

Before expansion, retain the data-flow map, approved settings, access test, deletion test and incident owner. Repeat relevant checks after material model or connector changes.

Report the proportion of approved workflows with verified data flows, the number of unresolved exceptions and time to revoke access. Set thresholds from the actual risk assessment. A completed questionnaire is evidence of a review, not proof that the controls work.

## Sources & research notes

01 / OpenAI / Data controls (accessed 11 November 2026)

https://developers.openai.com/api/docs/guides/your-data
Provider-policy snapshot. Product, endpoint, account settings and contractual exceptions matter.

02 / EDPB / International data transfers (accessed 11 November 2026)

https://www.edpb.europa.eu/sme/be-compliant/international-data-transfers_en
Chapter V conditions apply in addition to ordinary GDPR obligations for qualifying transfers.

03 / NIST / Generative Artificial Intelligence Profile, AI 600-1 (2024)

https://doi.org/10.6028/NIST.AI.600-1
Voluntary risk-management guidance, including confabulation and human–AI interaction; not a certification.

Independent TACITUS analysis. Research checked 11 November 2026. Published findings retain their dates and limits. Proposed curricula, tests, schedules and worked calculations are labelled analytical examples, not client results.