

CUSTOMER BRIEF · DATA SECURITY AND DATA SOVEREIGNTY

Data security and data sovereignty

This document explains how Zap AI handles your data, where that data is stored and processed, who can access it, and the controls available to you. It is written for prospects and customers evaluating Zap AI, and for their security and procurement teams.

How Zap AI works

01 · HOW ZAP AI WORKS

Zap AI adds analysis, recommendations, and reporting help to the Zap Data Hub platform. It does not receive a continuous feed of your data. Instead, when a user prompts the AI, the AI requests the results of a specific report or query, and Zap sends those results to the AI provider for that prompt. Zap currently uses OpenAI as its sole AI provider, under a signed Data Processing Addendum.

Two points matter for how you assess this. First, the AI only receives data after an administrator has enabled data access in the platform. Second, the data sent is a narrow slice per prompt, but it can contain real, sensitive data, because the purpose of the feature is to answer questions about your data.

02 · WHERE YOUR DATA IS STORED AND PROCESSED

There are three locations to understand, and we are direct about each one:

AT REST

Your data stays in your elected region at rest.

The Zap Data Hub warehouse and reporting model that hold your data remain hosted in the region you select, for example the United Kingdom or South Africa.

PER-PROMPT PROCESSING

Query results are processed in the United States or Europe.

When a user prompts the AI, the results of the relevant query or report are sent to our AI provider for processing and a response is returned. Processing takes place in the United States or Europe. Section 04 sets out which regions can also hold that data at rest locally. These results can include personal data and commercially sensitive figures.

RETENTION

AI prompts and responses are retained in Australia.

Zap keeps a copy of AI prompts and responses, including the query results sent to the AI, in its central data warehouse in Australia. This is used to operate, support, monitor, and improve the AI features, and to measure usage for billing.

We state this plainly because it matters for data sovereignty assessments. When you use Zap AI with data access enabled, query results containing your data are processed in the United States or Europe, and retained by Zap in Australia. Your core warehouse data, at rest, stays in your elected region. Zero data retention with our AI provider is available on request, and is covered in section 05.

03 · WHO CAN ACCESS YOUR DATA, AND FROM WHERE

Authorised Zap personnel whose role includes querying customer usage may access the retained AI prompts and responses, including the query results. This access is role-based and spans Zap's operating locations, which currently include Australia, the United Kingdom, the United States, and South Africa, and may extend to other locations as Zap grows. We disclose this broadly on purpose, so there is no gap between what we say and what our people can do.

Where access from one country reaches data held in another, that cross-border access is covered by Standard Contractual Clauses where required, and all access is subject to confidentiality obligations and Zap's ISO 27001 certified information management system.

04 · WHERE DATA IS STORED, AND WHERE IT IS PROCESSED

These are two different things, and our AI provider sells and controls them separately. It is worth separating them before reading the table, because they do not always land in the same country.

Storage is a safe in your country. Anything the provider keeps at rest stays there. **Processing** is the compute that reads the data and writes the answer. Every prompt makes that trip out and back, and it takes a few seconds. The safe stays where it is. The trip does not.

REGION	STORAGE AT REST IN-REGION	PROCESSING IN-REGION
United States	Yes	Yes
Europe (EEA and Switzerland)	Yes	Yes
United Kingdom	Yes	No, runs in Europe or the US
Australia	Yes	No, runs in the US
Canada, Japan, India, Singapore, South Korea, UAE	Yes	No, runs in the US
South Africa	No option	No option, runs in the US

The United States and Europe are the only two regions where both storage and processing can be in-region. Everywhere else, storage can be local and processing runs in the United States or Europe. Regional storage is enabled on request rather than by default.

The wording in your agreement reflects this: your data is **processed in the United States or other supported regions, including in your region where available and enabled**. We route in-region where we can. It is not a commitment that processing never leaves your country, and we would rather say so here than have you discover it later.

05 · ZERO DATA RETENTION


Zap holds an enterprise agreement with OpenAI that includes zero data retention. It is **available on request and is not switched on by default**, and it runs a restricted feature set, so it is worth discussing against what you actually need the AI to do.

Where zero data retention is activated, the provider keeps nothing: no prompts, no outputs, no copies. The only moment your data is outside your elected region is the few seconds of processing itself.

One thing this does not change: Zap's own operational record of prompts and responses is still retained in Australia, as described in section 02. That is Zap's system rather than the provider's, and zero data retention does not remove it.

How your data flows

The diagram below is our canonical reference for how data moves when you use Zap AI.



Your controls

06 · YOUR CONTROLS

You decide whether the AI gets access to your data at all:

- ➔ **Documentation-only mode.** If an administrator does not enable data access, the AI answers from product documentation only. In this mode no customer data is sent to the AI provider, and nothing leaves your environment.
- ➔ **Disable AI entirely.** The AI module is optional. You can run the core Zap platform without it, and disable it at any time if your data handling requirements change.

07 · COMMON QUESTIONS

QUESTION	ANSWER
What data does the AI actually receive?	The results of specific queries or reports, requested by the AI in response to a user prompt. Not a bulk feed of your warehouse. The results can contain real, sensitive data.
Is data encrypted in transit and at rest?	Yes. Data sent to and from the AI provider is encrypted in transit using TLS. Data at rest in the Zap warehouse and in the AI provider environment is encrypted.
Is my data made public?	No. Data sent through the AI provider is not public, and is not shared with other Zap customers or other organisations.
Is my data used to train public AI models?	No. Under OpenAI's API terms and the Data Processing Addendum Zap has signed, data submitted through the API is not used to train OpenAI's models. This is a contractual commitment.
Can I use AI without sending it my data?	Yes. In documentation-only mode the AI answers from product documentation and receives no customer data. An administrator controls whether data access is enabled.
What contracts and safeguards are in place?	A signed Data Processing Addendum with OpenAI that includes the EU Standard Contractual Clauses and, where applicable, the UK International Data Transfer Addendum.
Can I see my AI usage?	Yes. Token usage per user is shown in the platform, so you can track and manage how your team uses AI.

For specific security or data protection questions, your Zap account manager can arrange a more detailed response from our team.