Zenforge

Subprocessor List

Third-Party Data Processors

March 20, 2026

Classification: Public

Contents

1 Subprocessor List

This document lists the third-party subprocessors that Zenforge uses to deliver its service, the categories of data shared with each, applicable retention periods, and the availability of zero data retention (ZDR) agreements.

Last updated: March 2026.

1.1 LLM Providers

These providers process customer signal text to extract structured intelligence. Customer data is sent via API and is **not** used for model training.

Provider	Data Sent	Retention	ZDR	DPA
Anthropic (Claude API)	Signal text, workspace context, extraction prompts	7 days[1]	Yes, via addendum	Available, not yet executed
OpenAI (GPT API)	Signal text, workspace context, extraction prompts	30 days (abuse monitoring)	Yes, on approval	Available, not yet executed
Google Vertex AI (Gemini API)	Signal text, workspace context, extraction prompts	90 days (abuse monitoring, if flagged)[2]	Yes[3]	CDPA (automatic)

None of these providers use API-submitted data for model training under their current terms of service.

1.2 Infrastructure Providers

Provider	Data Stored	Retention	ZDR	DPA
Neon (PostgreSQL)	All application data (primary database)	Point-in-time recovery (currently 6 hours; configurable up to 30 days)	N/A	Under review
Fly.io	Application logs, deployment artifacts	Live: 2 days, API: 15 days,	N/A	Under review

[1] Reduced from 30 days as of 2025-09-15.
[2] Exception from abuse monitoring available on request. In-memory cache (24h) is separate and disablable.
[3] Requires: disable in-memory cache, approved exception request for abuse monitoring, avoid standard Google Search grounding.

Provider	Data Stored	Retention	ZDR	DPA
		Grafana: 30 days		

1.3 Observability Providers

Provider	Data Sent	Retention	ZDR	DPA
Langfuse	Full prompt and completion content for diagnostics. Covered by workspace data erasure.	Per terms	N/A (opt-in)	Under review

1.4 Identity Provider

Provider	Data Stored	Retention	ZDR	DPA
Kinde (OIDC)	User email, name, authentication state	Per terms	N/A	Under review

1.5 Data Residency

- **Primary database (PostgreSQL):** Sydney, Australia (Neon `aws-ap-southeast-2`)
- **Application hosting:** Sydney, Australia (Fly.io `syd` region)
- **LLM processing:** United States (Anthropic, OpenAI, Google API endpoints)
- **Observability (Langfuse):** EU (Langfuse Cloud)
- **Identity (Kinde):** Sydney, Australia

1.6 Data Flow Summary

Customer data flows through the following path:

1. User submits signal content via the Zenforge web application
2. Signal text is stored in the PostgreSQL database (Sydney, AU)
3. During processing, signal text is sent to the configured LLM provider API for entity and fact extraction
4. Full prompt and completion content is exported to Langfuse (EU) for diagnostics
5. Extracted structured data is stored back in PostgreSQL
6. LLM providers retain API inputs per their retention policies (see table above), then auto-delete

1.7 Changes to This List

Zenforge commits to notifying customers of material changes to this subprocessor list (addition of new subprocessors or changes to data categories) with at least 30 days' notice. This commitment will be formalised in customer agreements as part of our SOC 2 readiness programme.