

## Basic data flow in EmpowerGPT

<b>Step 1</b> <b>User uploads prompt / document</b> EmpowerGPT App & Services(AKS – Germany West Central)	<b>Step 2</b> <b>Prompt + document content processed &amp; routed</b> Private Azure Virtual Network(no public internet exposure)	<b>Step 3</b> <b>(Optional) RAG processing</b> Chunking → Embeddings → Vector DB(all within Germany West Central)	<b>Step 4</b> <b>LLM generates answer</b> Azure AI FoundryEU Data Zone processing only	<b>Step 5</b> <b>Generated answer returned to user</b> EmpowerGPT App(AKS – Germany West Central)
---	--	---	--	---

### Explanation

EmpowerGPT operates entirely within Azure's Germany West Central region and the broader EU Data Zone. When a user submits a prompt or uploads a document, the request is handled by application services running on Azure Kubernetes Service within a private network. If document-based querying is used, the system processes the file through a RAG pipeline: documents are stored in Azure Blob Storage, split into chunks, converted into embeddings using Azure OpenAI, and stored in a PostgreSQL vector database — all within the same region and without public exposure. The enriched prompt is then sent to a selected model hosted via Azure AI Foundry, where inference occurs strictly within the EU Data Zone. Abuse monitoring is disabled, meaning no prompts or responses are logged or retained by Microsoft. The generated response is returned to the application and displayed to the user. All data storage, retention, and deletion — including backups and full teardown on contract termination — are fully controlled within the same regional infrastructure.

The terms and conditions under which we use LLMs have been summarized in the appendix to our terms at [empowergpt.ai/terms](https://empowergpt.ai/terms)