Generative AI in a Truly Sovereign Digital Europe

A field guide to real sovereignty for LLMs

Christian Gapp · lynq.ing · SFSCon 2026

Who's talking

- **Christian Gapp**: developer, co-founder of **lynq.ing**
- We build an LLM-based product in South Tyrol and Austria, for EU customers

What follows is **what we found on the way**, about **LLMs** (large language models, text). The same questions apply to image, audio and video generation.

⚠️ Not legal advice

These are **findings from building a product**, not legal opinions.

The market moves **fast**: anything here may be out of date by tomorrow.

That's why **no providers are named**: check your own vendors yourself.

Typical questions in European AI projects

"Can we use the **(best) models legally** (GDPR)?"

Most projects · GDPR-compliant today. But courts can **cancel the legal basis**, and a foreign government can **cut Europe off**.

"Is the **whole stack** European?"

Sovereignty-minded projects · Hardly any can say **yes** today.

The typical vendor answer to both:

"Yes, we have **European sovereign** solutions."

What "European sovereign" should mean

Sovereignty is about **trust and risk control**: how much do you trust others, and how much risk will you take?

- **EU control:** EU-owned & EU-operated, so non-EU laws have no direct reach
- **No hidden third parties:** only known EU companies handle your data
- **It keeps working:** even if a foreign government cuts Europe off
- **Your data stays yours:** prompts not stored, not passed on, not used for training
- **You can leave:** open-weight models and a standard API, so switching is quick

Where your prompt goes, layer by layer

Prompt



Model	closed: its maker can stop or change it open-weight: you choose who runs it
Inference service	may **store** your prompts or **pass them on**
Cloud / operator	a non-EU law can **force** it to hand over your data
Data centre	who owns it, and who runs it?
Chips / hardware	supply and export controls

Every layer counts: who controls it, and whose laws can reach it?

Why it matters: law follows *control*, not location



Laws like this: for example the US CLOUD Act, China's National Intelligence Law

What's at stake: prompts carry **business secrets, personal data, customer documents**

Not about one country: any parent company under such a law can be *ordered* to hand over data it controls, even from its EU subsidiaries

The sovereignty spectrum

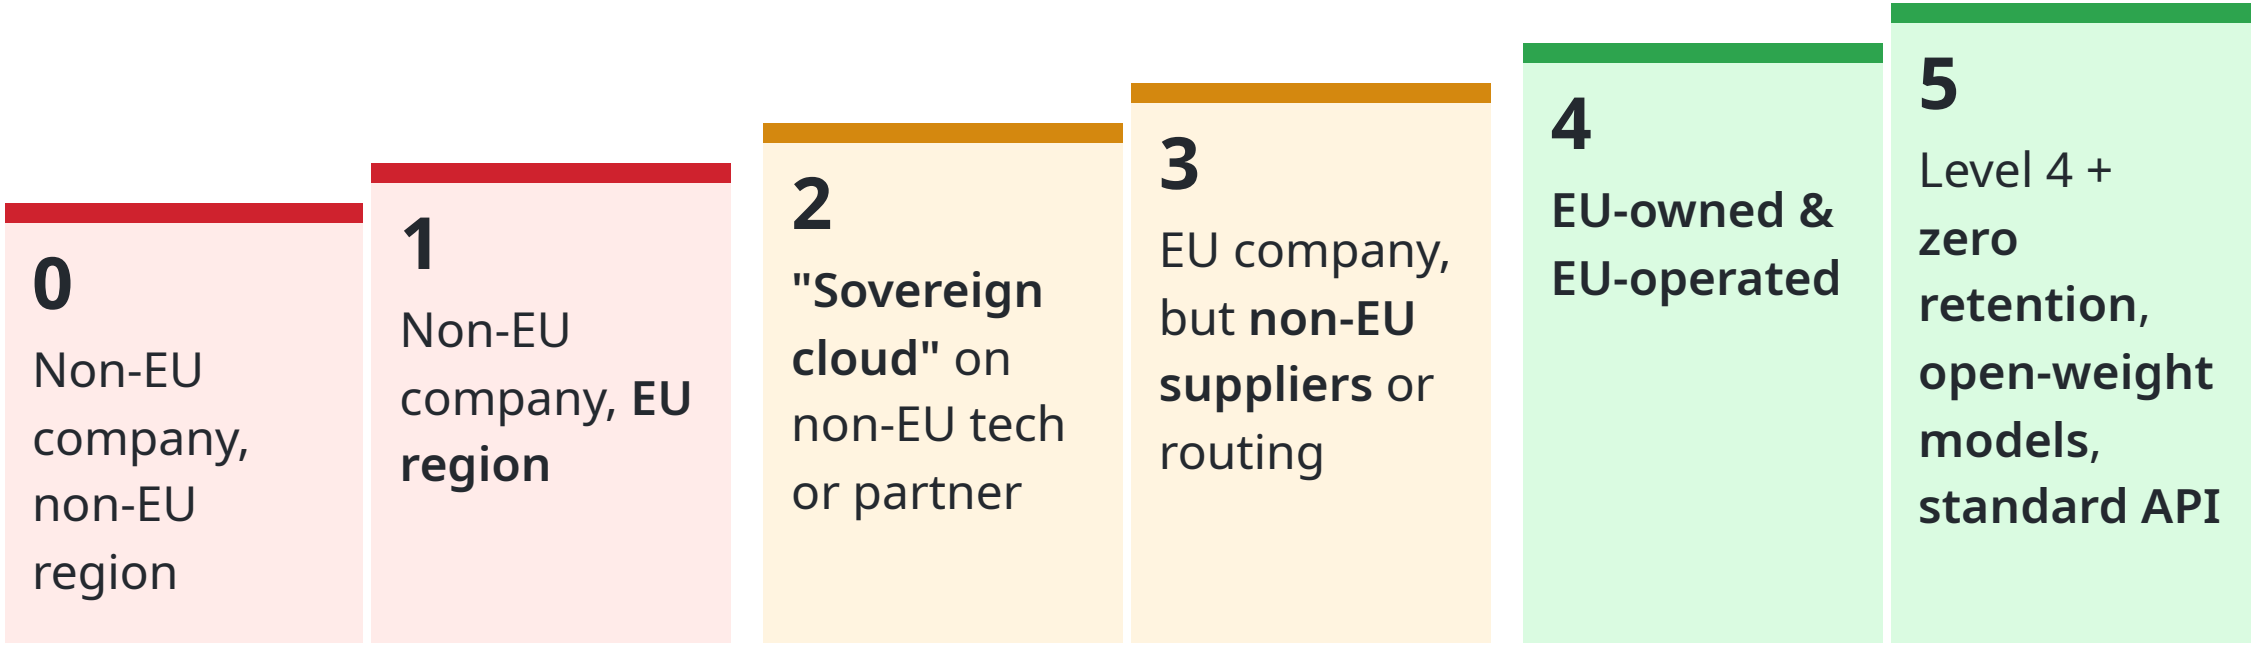

Our practical scale for LLM services. For procurement, see the EU Cloud Sovereignty Framework (SEAL-0 to SEAL-4).

Ask your AI vendor

EU CONTROL

- ☐ Who controls you? Who signs?
- ☐ Who runs it? Who can see prompts, from where?

NO HIDDEN THIRD PARTIES

- ☐ All sub-processors, with countries?
- ☐ Any fallback outside the EU?

IT KEEPS WORKING

- ☐ Still running if Europe is cut off?

YOUR DATA STAYS YOURS

- ☐ What is kept (logs too), how long? Used for training?
- ☐ Proof: audits, certificates, or just promises?

YOU CAN LEAVE

- ☐ Open weights, open licence? Standard API?

🚩 **Signs of sovereignty washing:** "sovereign" with a non-EU parent · "hosted in the EU" only · non-EU fallback · proprietary API

Run it yourself? Possible, but it has costs.

◄ less control · less work — more control · more work ►

EU provider, open-weight models

★ often the best balance

The provider runs it: check who controls it

Rented cloud GPUs

you run it, on a cloud's hardware

GPU-hours, even when unused; security, updates

Self-hosted

you run it, on your own hardware

Hardware, a spare server, admins, a secure room; the largest models won't fit

> Self-hosting fits some workloads. It is **not the default answer**.

What we found while building lynq.ing

Model router / aggregator most let you choose one provider	✗ the router is **one more company** in between
European model lab great models	✗ the API runs partly on non-EU-owned clouds
Big non-EU cloud, EU region top quality, data in the EU	✗ the company is under non-EU law
EU "sovereign cloud" offer	✗ high monthly minimum; some models on non-EU clouds
Young EU inference platforms	✗ partners run the GPUs or even the platform, sometimes non-EU ones
What we chose	✓ **EU-owned & EU-operated, zero retention**

Provider types, names removed, from our 2026 evaluation. Things change, so check for yourself.

The honest trade-offs

- **Models**: no top frontier models; open-weight is *good enough* for us, not for all
- **Cost**: competitive per token, but only a short trial, no free tier or credits
- **Resilience**: our fallback *was* the sovereignty leak, so we removed it.
  If our provider is down, we are down.
- **Every tool counts**: sovereignty doesn't stop at the LLM, *and we're not done yet*

What saved us: an **open, standard API**, so switching providers is mostly a config change, plus testing.

Sovereignty is about trust and risk control

1 **Law follows control, not location.**
An EU data centre doesn't keep non-EU law out.

2 **"Sovereign" is a scale, not yes or no.**
Find where an offer sits. Watch for sovereignty washing.

3 **Openness lets you check, and leave.**
Open-weight models, standard API, FOSS, EU operators.

→ **Before you sign: ask your AI vendor the questions on slide 9.**

Thank you!

Before you sign, ask your AI vendor:

- [ ] Who controls you? Who signs?
- [ ] Who runs it? Who can see prompts, from where?
- [ ] All sub-processors, with countries?
- [ ] Any fallback outside the EU?
- [ ] Still running if Europe is cut off?
- [ ] What is kept (logs too), how long? Used for training?
- [ ] Proof: audits, certificates, or just promises?
- [ ] Open weights, open licence? Standard API?

Not legal advice. Findings from building lynq.ing, as of November 2026.

Slides: lynq.ing/talks/sfscon-2026

Christian Gapp · christian@gappc.net

lynq.ing · info@lynq.ing