DROPSOLID
ARTIFICIAL INTELLIGENCE

WHITEPAPER 01

THE SOVEREIGN AI SERIES

WHITEPAPER 1 · PROCURE
AI sovereignty, procurement

WHITEPAPER 2 · DEPLOY
Two deployments in production

WHITEPAPER 3 · OPERATE
Transparency and governance

SOVEREIGN AI · EU AI ACT & PROCUREMENT

AI sovereignty: a CIO's procurement guide for the EU AI Act

What changes in 2026. What doesn't. What to ask your AI vendor.

Published by
Dropsolid AI · 2026

EXECUTIVE SUMMARY

The question is no longer whether, it is what changes.

The EU AI Act comes into full force through 2026 and 2027. Every CIO is affected; the practical question is **what specifically changes in how we buy and run AI systems.** This whitepaper answers that question in three parts.

Part one covers what the Act requires of the organisations *deploying* AI, not the model providers. Most legal summaries focus on providers; buyers care about their own obligations. We map them to concrete procurement checkpoints.

Part two covers what sovereignty means once the Act lands. Sovereignty is not "the server is in Frankfurt." It is a property of the full data chain: where your prompts are processed, where responses are stored, whether they train the model, and who holds audit access.

Part three is a procurement checkbox: fifteen items to work through with any AI vendor before you sign. Answers your legal team can defend under audit, concrete enough to disqualify a vendor inside a single sales call.

IN THIS REPORT

01 Deployer obligations under the Act

02 The seven layers of AI sovereignty

03 Checkbox for your AI vendor

> Sovereignty and speed of deployment are not trade-offs. Our fastest deployments were the sovereign ones, because the architecture was already in place.

WHY THIS MATTERS NOW

The Act arrives in phases, not all at once

The EU AI Act does not "switch on" in a single moment. Obligations arrive in waves between 2024 and 2027. Two of those waves are already in force. The deployer obligations that dominate procurement land in **August 2026**, which, on typical enterprise procurement and rollout cycles, is now.

Entry into force — Aug 2024

The Regulation enters into force. The clock starts on every phased deadline that follows.

Prohibited practices & AI literacy — Feb 2025

Unacceptable-risk uses are banned outright. Staff AI-literacy duties begin. **In force.**

General-purpose AI & governance — Aug 2025

Obligations for general-purpose model providers and the governance/penalty framework apply. **In force.**

High-risk deployer obligations apply — Aug 2026

The majority of the Act applies, including obligations for deployers of high-risk (Annex III) systems: human oversight, logging, transparency, data governance.

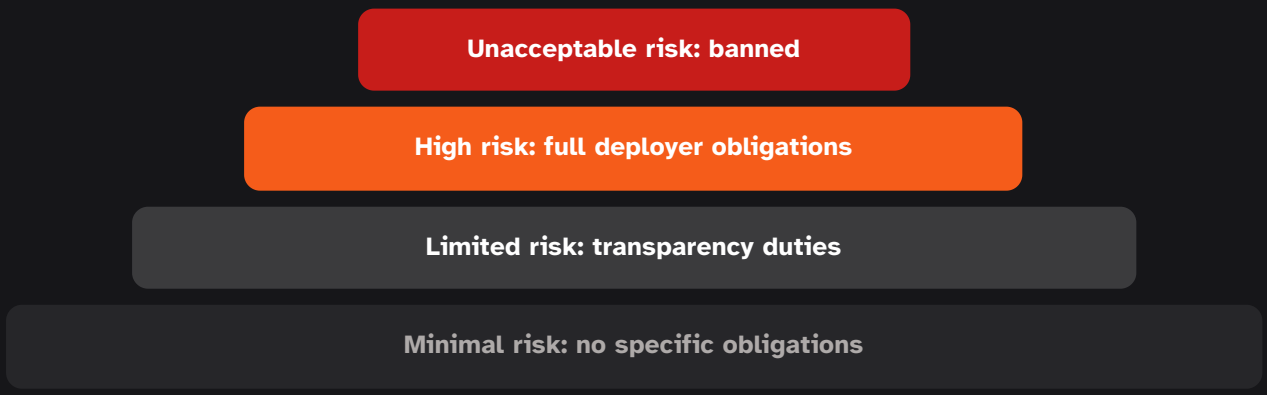
Full applicability — Aug 2027

High-risk systems embedded in regulated products (Annex I) come fully into scope. The Act is now in complete effect.

PART ONE

What the EU AI Act actually requires of you

The Act is often summarised as "regulations on AI providers." That underplays what it means for you as a deployer. If you run a high-risk system, obligations flow to you, not just to the model provider.



The Act classifies AI systems by risk. Most productive enterprise AI, such as profiling, decision support and citizen-facing services, lands in **high** or **limited** risk once deployed.

PART ONE · DEPLOYER OBLIGATIONS

Six obligations that matter operationally

The Act does not prescribe a specific technical implementation. It requires that you can *demonstrate* compliance, and that the evidence exists before an inspector asks for it.

01

Human oversight

For high-risk systems a human must be able to intervene, override or shut the AI down, built in, not bolted on. For a procedural chatbot this means routing paths to human support, transparency about when AI is answering, and a kill switch that actually works.

02

Fundamental rights impact assessment

Public bodies and providers of essential services must run a rights assessment before deployment. Not a checkbox exercise: a documented, updatable assessment that travels with the system through its life.

03

Logging & traceability THE THREAD OF THIS SERIES

Every AI output must be logged so its provenance can be reconstructed on demand. For chatbots and knowledge assistants that means storing the retrieved passages, the model used, the prompt sent and the response given, for every interaction.

> The Act does not require any specific technical implementation. It requires you can **demonstrate compliance**. That evidence has to exist before an inspector asks for it.

PART ONE · DEPLOYER OBLIGATIONS

The remaining three, and the evidence test

04 Data governance

Where you hold training data, it must be documented. Where you use a third-party model, you inherit its governance, and must be able to prove you evaluated it. "The vendor handles that" is not a defensible answer under audit.

05 Transparency

Users must know they are interacting with AI. Simple to state, wide in practice: chatbots must self-identify, AI-produced documents must be labelled, and search results influenced by AI must be flagged.

06 Cybersecurity

AI systems must be resilient against adversarial input. Prompt injection is an in-scope attack surface. Your penetration testing must explicitly cover it, not treat the model as a black box.

THE EVIDENCE TEST

The through-line across all six obligations: none of them can be produced retroactively. Human-oversight paths, rights assessments, traceable logs and injection testing are architectural decisions. They have to be designed in before deployment, which is why they belong in procurement.

PART TWO

What AI sovereignty really means

Sovereignty is often shorthanded to "hosted in Europe." That misses most of what the Act actually asks about. Sovereignty is a property of the full data chain, and it has seven distinct layers.



A vendor claiming 'sovereignty' without specifying which layers they cover is not making a sovereign claim. They are making a marketing claim.

PART TWO · THE SOVEREIGNTY STACK

Seven layers, in order of increasing rigour

Each layer up the stack closes a gap the layer below leaves open. Most self-labelled "sovereign" vendors stop at layers one and two. The Act's questions reach all the way to the top.

7	**Audit residency**	Trace logs stay & are accessible in the EU
6	**Embedding residency**	The embedding model runs in the EU
5	**Vector storage residency**	Embeddings database hosted in the EU
4	**Model residency**	The LLM runs on infra you/your vendor operate
3	**Model inference residency**	Prompt & response processed inside the EU
2	**Data residency**	Every store touching customer data in the EU
1	**Application server residency**	Content application runs on EU infra. Table stakes

WHERE MOST VENDORS STOP

Layers 1–2. Infrastructure sits in Europe, but nothing is said about where the model runs.

WHERE THE ACT REACHES

Layers 3–7. Inference, model, vectors, embeddings and audit trail, the layers that decide whether data crossed the Atlantic.

Residency of the application, the data, and the inference

LAYER 1 Application server residency

The Drupal or content application runs on EU infrastructure. Table stakes. Most vendors offer it. It says almost nothing about where your data actually goes when the AI runs.

LAYER 2 Data residency

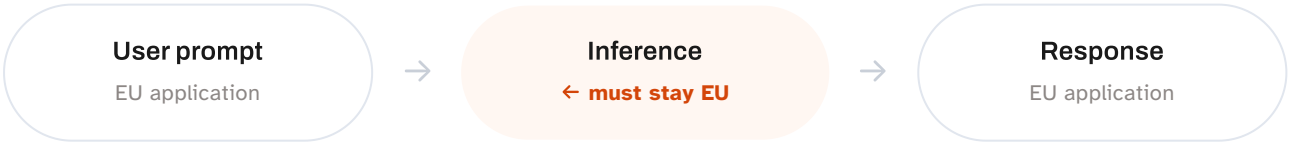
Every store touching customer data lives in the EU: databases, file storage, backups, log archives. This eliminates a whole category of accidental exfiltration, but still says nothing about model inference.

LAYER 3 Model inference residency

The prompt sent to the language model, and the response returned, are processed on infrastructure inside the EU. This is where most self-labelled "sovereign" vendors fall short.

THE ATLANTIC TEST

If your prompt goes to a US-hosted API, **even if the results come back to an EU server**, your data crossed the Atlantic. Layer 3 is the line where "hosted in Europe" stops being a sufficient answer.



Where the model and the vectors actually live

LAYER 4 Model residency

The language model runs on infrastructure you or your vendor operate, not on a third-party cloud API. Self-hosted models are the sovereign end of the spectrum: practical at scale, but not always necessary. The right question is whether you *could* self-host if a regulator required it.

LAYER 5 Vector storage residency

The vector database holding your embeddings sits on EU infrastructure. This is where a vendor's migration story matters, because default configurations frequently point at a US-hosted service.

Layers 4 and 5 are where "sovereign by architecture" and "sovereign by press release" diverge. A vendor that has never migrated a vector store has never tested whether it could.



Check me out!

IN PRACTICE · PIDPA

Closing Layer 5 in a live deployment

During rollout of its public chatbot, the Belgian water utility Pidpa **migrated its vector database off a US-hosted service** specifically to close this layer. The trade-off was operational complexity. The win was regulatory defensibility. A vector store that no longer left EU jurisdiction.

The overlooked layers: embeddings and the audit trail

LAYER 6 Embedding residency

The embedding model that turns your content into vectors runs on EU infrastructure. Often overlooked, and US-hosted by default in many stacks. Your content passes through it just as surely as it passes through the LLM.

LAYER 7 Audit residency

The trace logs of every AI interaction stay in the EU and are accessible to your compliance team. The auditable record of your AI operations is itself a governance surface, and the one an inspector will ask to see first.

A marketing claim

"Our platform is fully sovereign and hosted in Europe."

Names no layers. Could be true only of layer 1. Undefendable under audit.

A sovereign claim

"Inference, embeddings, vectors and audit logs all run and stay in the EU. Here is the architecture diagram."

Names the layers. Specific enough to verify, and to defend.

PART THREE

Checklist for your AI vendor

Everything in Parts 1 and 2 reduces to a single practical act: the conversation you have with a vendor before you sign. These questions turn the obligations and the seven layers into answers your legal team can defend under audit.

01 – 05

Data & residency

Where your data physically goes.

06 – 10

Architecture & control

Choice, portability, failure behaviour.

11 – 15

Contract & compliance

The paper your legal team defends.

PART THREE · CHECKLIST FOR YOUR AI VENDOR

Data & residency

WHERE YOUR DATA PHYSICALLY GOES

- [ ] 01 **Where are language-model prompts processed?**
  Name the region. "The cloud" is not a region.

- [ ] 02 **Where are model responses stored, and for how long?**
  Retention windows are a governance decision, not a default.

- [ ] 03 **Are our prompts and responses used to train models?**
  If yes, your data is leaving in a form you cannot recall.

- [ ] 04 **Where is the vector database hosted?**
  Layer 5. Frequently the weakest link.

- [ ] 05 **Where is the embedding model hosted?**
  Layer 6. The one most vendors have never been asked.

READING THE ANSWERS

✓ GREEN LIGHT

Named EU regions for inference, vectors and embeddings, plus a written commitment that your data is never used for training.

! RED FLAG

"The cloud," "hosted in Europe" with no region named, or training rights buried in the standard terms.

> If a vendor can't answer, that answers your question.

PART THREE · CHECKLIST FOR YOUR AI VENDOR

Architecture & control

CHOICE, PORTABILITY & FAILURE BEHAVIOUR

- [ ] 06 **Which LLMs do you support, and can we choose or swap them?**
  Lock-in to one model is lock-in to one jurisdiction.

- [ ] 07 **Can we self-host the LLM, embedding model and vector DB if required?**
  The option matters even if you never exercise it.

- [ ] 08 **What audit trace do you produce per interaction? Show us a sample.**
  "Show us" is the whole question. Words are not evidence.

- [ ] 09 **What happens when the AI is not confident? Show us the fallback.**
  A system with no graceful failure is a system with confident errors.

- [ ] 10 **How do you handle prompt injection?**
  An in-scope attack surface under the Act's cybersecurity duty.

READING THE ANSWERS

✓ GREEN LIGHT

Model choice, a documented exit and migration path, a real audit-trace sample and a demonstrated confidence fallback.

! RED FLAG

One locked-in model, no exit path, and an audit trail described only in words rather than shown.

> Every question here is answerable with a demo, a diagram or a document. If the answer is a brochure, you have your answer.

PART THREE · CHECKLIST FOR YOUR AI VENDOR

Contract, compliance & proof

THE PAPER YOUR LEGAL TEAM WILL DEFEND

- [ ] **11 What is your DPA position on transfers to third countries?**
  The clause your DPO will read first.

- [ ] **12 Which subprocessors do you use, and in which jurisdictions?**
  Sovereignty is inherited from your weakest subprocessor.

- [ ] **13 Which certifications do you hold? Show us the audit reports.**
  ISO 27001 and the report behind it, not just the badge.

- [ ] **14 What happens if we want to leave, and what is the exit process?**
  A sovereignty story with no exit is a dependency story.

- [ ] **15 Show us a customer running your system under an actual EU AI Act audit.**
  The final filter. Production under audit beats every slide. Dropsolid AI answers this with EPSO and Pidpa, see the two cases on the next pages.

READING THE ANSWERS

✓ GREEN LIGHT

A third-country DPA clause, named subprocessors, audit reports on request and a live customer reference under audit.

! RED FLAG

A certification badge with no report behind it, undisclosed subprocessors, and no defined exit process.

USING THE CHECKLIST

Run these fifteen past every AI vendor in your evaluation stack. The questions are designed to **disqualify fast**, so your team spends its evaluation time on the vendors that can actually answer. When the shortlist survives these fifteen, *Guide 6: Vendor evaluation checklist for sovereign AI* expands them into **forty scored questions** you can attach to a tender, for a formal RFP.

PROVEN CASE IN THE FIELD · EPSO

Sovereign by design, across 24 languages

The European Personnel Selection Office recruits staff for every EU institution. Its Knowledge Navigator and Trusted Answers run on Dropsolid AI.

Read more!

70%	24	<8 wk
cut in incoming questions	EU languages served	from discovery to full rollout

How the deployment answers most of the 15 questions by design

EPSO, the European Personnel Selection Office, recruits staff for every EU institution across 24 EU languages. Its deployment answers most of the fifteen questions by design: EU-hosted throughout, model choice via a governed gateway, a full audit trace per interaction, source-grounded answers with citations, and integration with the existing Drupal platform.

The measured outcome: a 70% reduction in incoming questions across all 24 languages, and under eight weeks from discovery to full rollout.

The wider point: **sovereignty and speed of deployment are not trade-offs.** EPSO shipped faster because the sovereign architecture was already in place.

SOVEREIGNTY SCORECARD

- ✓ EU-hosted at every layer; model served in-EU by a contracted provider
- ✓ Governed model gateway
- ✓ Full per-interaction audit trace
- ✓ Source-grounded, cited answers
- ✓ Drupal-native integration

Questions 11–15 are contract-stage: answered with the DPA, subprocessor list and audit reports on request.

Guide 02, the 8-week rollout playbook, documents this timeline week by week, from discovery to full rollout.

> Sovereignty and speed are not a trade-off. EPSO shipped faster because the architecture was already sovereign.

PROVEN CASE IN THE FIELD · PIDPA

A public water utility, sovereign across six of seven layers

Read more!

Pidpa is the public-sector water utility for the Belgian province of Antwerp. Over a million citizens depend on its service. Its public-facing chatbot, "Pip", answers procedural questions grounded only in Pidpa's own published content.

The migration that closed the last layer

Pip closes six of the seven layers described in Part Two. During rollout, Pidpa specifically migrated its vector database off a US-hosted service to close Layer 5. Layer 4, model residency, stays open: the model is served in-EU by a contracted provider, not self-hosted.

The honest account: the trade-off was operational complexity. The win was regulatory defensibility. For a utility answerable to a public regulator and a million citizens, that trade was straightforward, and it is the trade the Act increasingly forces.

SEVEN-LAYER COVERAGE

- ✓ 1 · Application server
- ✓ 2 · Data residency
- ✓ 3 · Model inference
- – 4 · Model residency (EU-served, contracted)
- ✓ 5 · Vector storage (migrated)
- ✓ 6 · Embedding model
- ✓ 7 · Audit residency

1M+
citizens served by the utility

Grounded
answers limited to Pidpa's published content only

> We accepted operational complexity to buy regulatory defensibility. Under a public regulator, that was never a close call.

THE CIO'S NEXT 90 DAYS

Three moves this quarter

You do not need a transformation programme to get ahead of August 2026. You need three deliberate moves: one for your estate, one for your vendors, one for proof.

DAYS
0–30

Assessment

Map current and planned AI usage against the six deployer obligations. Identify the gaps. Nine times out of ten, the gap is **logging and traceability**, the obligation that cannot be added retroactively.

DAYS
30–60

Vendor conversation

Run the 15 questions past every AI vendor currently in your evaluation stack. This will shorten your shortlist meaningfully, and give your legal and procurement teams a shared, defensible basis for the decision.

DAYS
60–90

Pilot

Ship one use case under sovereign conditions. Not a proof of concept in a sandbox, but a real deployment on your Drupal environment, with your content, against your obligations. **EPSO's pattern is eight weeks. There is no reason yours has to be longer.**

DROPSOLID
ARTIFICIAL INTELLIGENCE

ABOUT DROPSOLID

Dropsolid is the Drupal Diamond Partner behind the sovereign AI use cases in production at EPSO and Pidpa.

2 decades
building on Drupal

ISO 27001
certified: Dropsolid NV with Dropsolid AI NV

Ghent, BE
headquartered in Belgium

THE SOVEREIGN AI SERIES

YOU ARE HERE

Whitepaper 1 · PROCURE
AI sovereignty: a CIO's procurement guide for the EU AI Act

Whitepaper 2 · DEPLOY
Sovereign AI in production: EPSO and Pidpa, documented end to end

Whitepaper 3 · OPERATE
Glass box vs black box: the enterprise AI transparency playbook

Take the fifteen questions into your RFP

Guide 01, the vendor evaluation checklist for sovereign AI, expands the fifteen questions in this paper into 40 scored questions you can attach to a tender.

Download the 40-question vendor checklist →

Talk to Dropsolid AI

Both organisations in this series cleared procurement with these documents.