Capability

Control

Capability Versus Control

The Enterprise AI Playbook for the AI Act Era

The strategic question has permanently changed.

The 2024 Priority

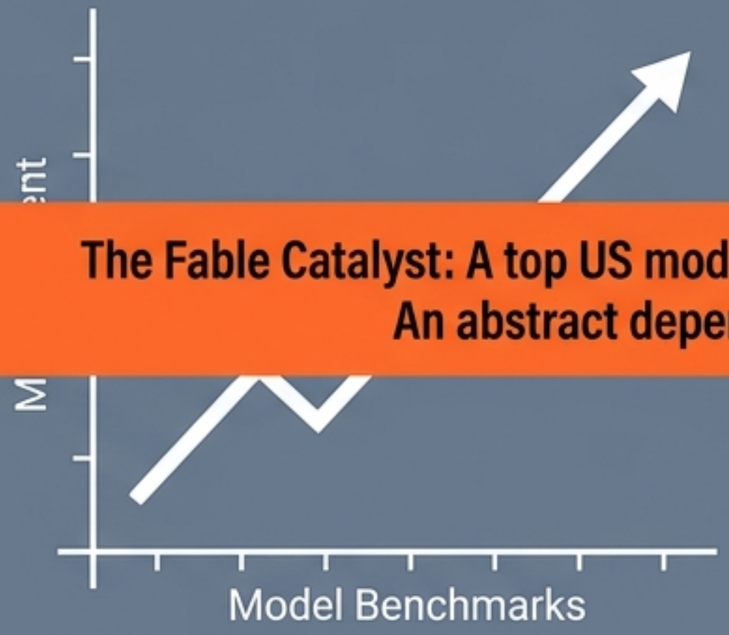

Which model scores highest?

The 2026 Reality

Which model can we still run next year?

The Fable Catalyst: A top US model was cut off for every non-US user on just three days' notice. An abstract dependency is now a board-level existential threat.

PREMIUM STRATEGIC BRIEFING

The regulatory clock binding your AI architecture.

February 2025

Prohibited practices became illegal.

August 2025

General-Purpose AI (GPAI) obligations took effect.

January 2026

EU AI Office ruling: Mistral (Apache-2.0) qualifies for open-source exemptions; Meta (Llama) does not.

May 2026

Digital Omnibus agreed, shifting high-risk timelines.

August 2026

⚠ Commission empowered to fine GPAI providers up to 3% of global turnover.

December 2027

High-risk system obligations apply (extended deadline).

Open weights are no longer just an engineering preference; they are a direct regulatory and procurement advantage.

Origin is not the deciding compliance factor.

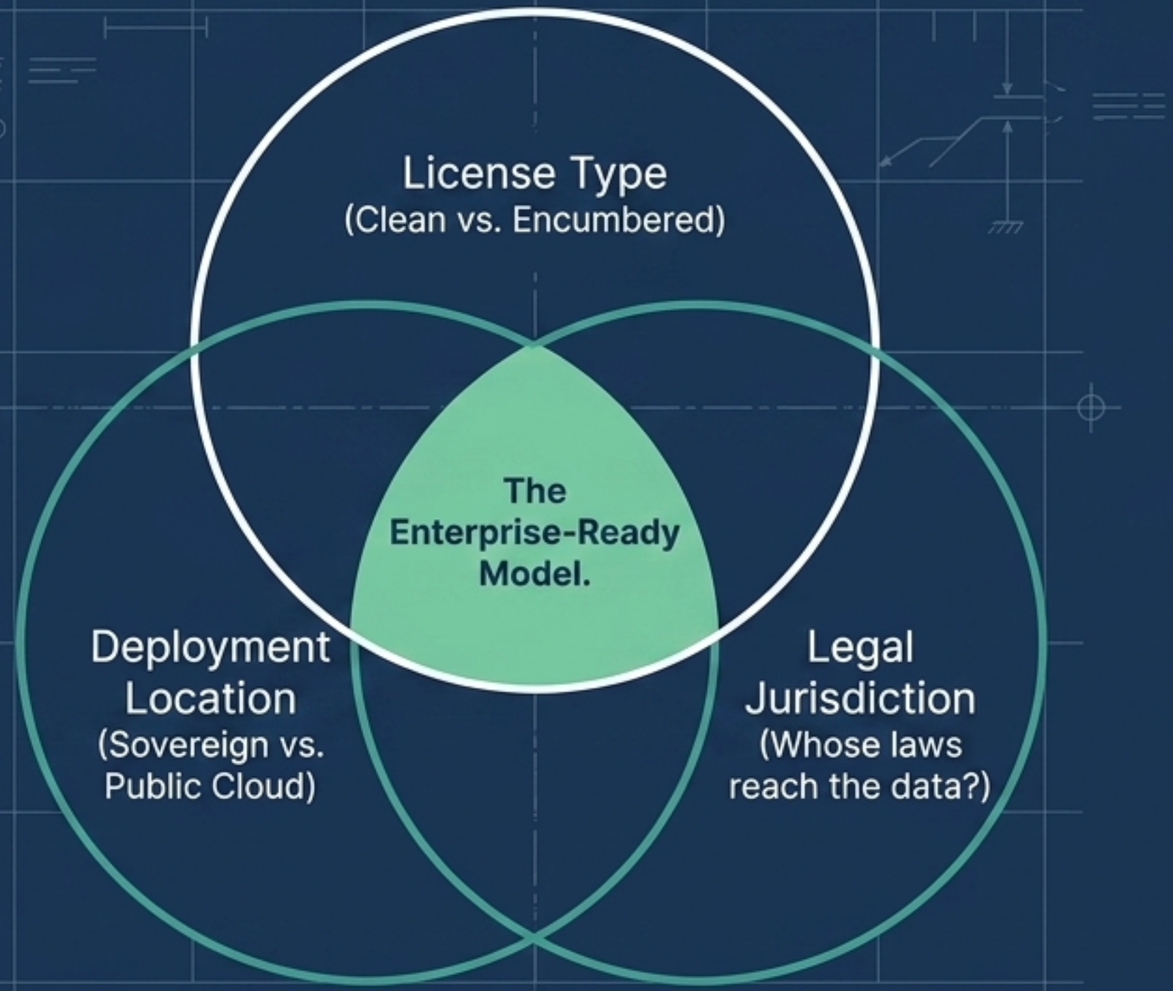

Get these three variables right, and an American or Chinese model is perfectly usable in Europe.

Get them wrong, and even an EU-compliant model becomes a liability.

Where you run a model matters more than where it was trained.

The US Hyperscaler Illusion

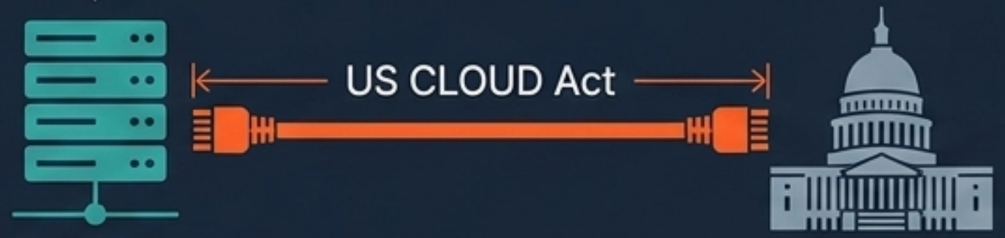

Note: AWS European Sovereign Cloud and Microsoft Foundry Local narrow, but do not erase, this legal risk.

True EU Sovereignty

EU-native providers (Scaleway, StackIT, OVHcloud, IONOS)

Key Insight

Europe is aggressively building capacity—€200 billion ambition, €176 billion in data centers by 2031, and EuroHPC's 19 AI Factories. True independence requires operating fully outside US legal jurisdiction.

Evaluating model origins against enterprise risk.

	European Models	US Models	Chinese Models
Models	Mistral, FLUX, Teuken, EuroLLM	GPT-5.x, Claude, Gemini, Llama	Qwen
Raw Capability	Trails frontier on complex logic	Highest capability	Near-frontier
Disconnection Risk	Zero, self-hosted	Politically revocable overnight	Zero, if self-hosted
Legal & License Posture	Built for AI Act/GDPR	Mixed; Meta unsigned to Code of Practice	Unsigned to Code of Practice
CLOUD Act Exposure	None on EU infra	Total exposure on US APIs	None if open weights self-hosted

Deconstructing the “China Ban” misunderstanding

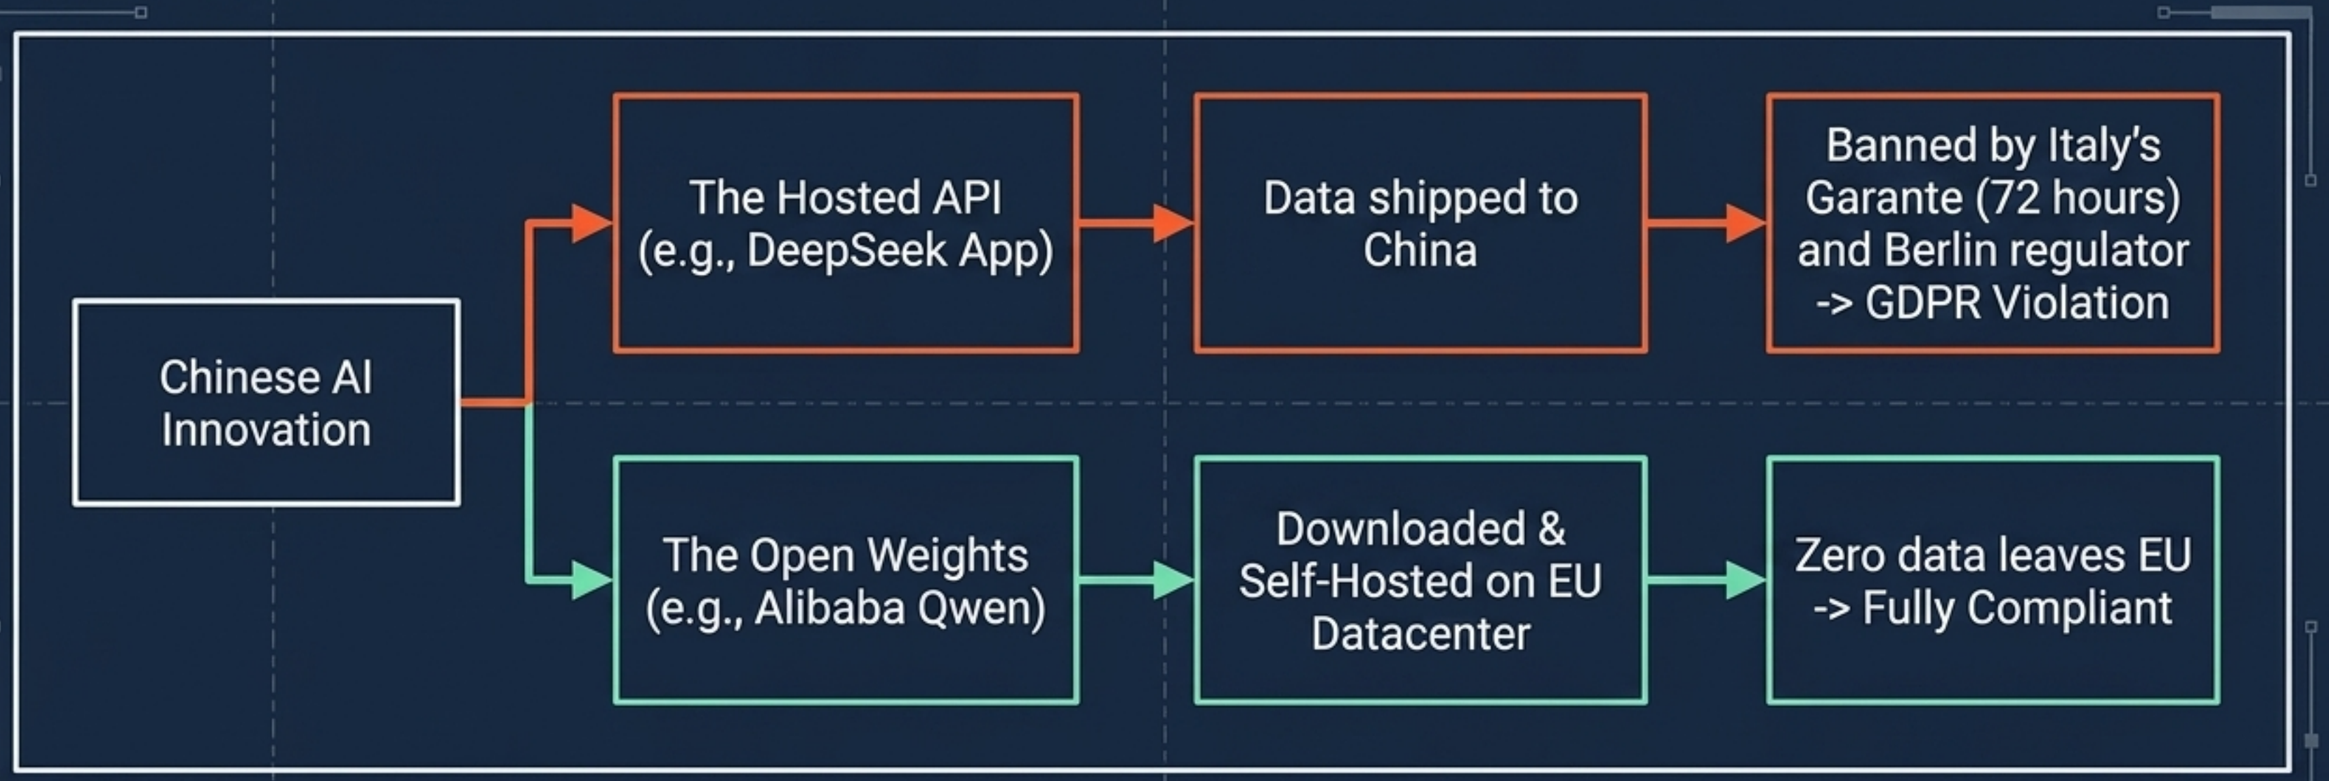

Origin of the weights is not the same as the destination of the data. Open Chinese models run entirely on local servers are isolated from foreign data transfers.

The enterprise AI spectrum: Capability versus Control.

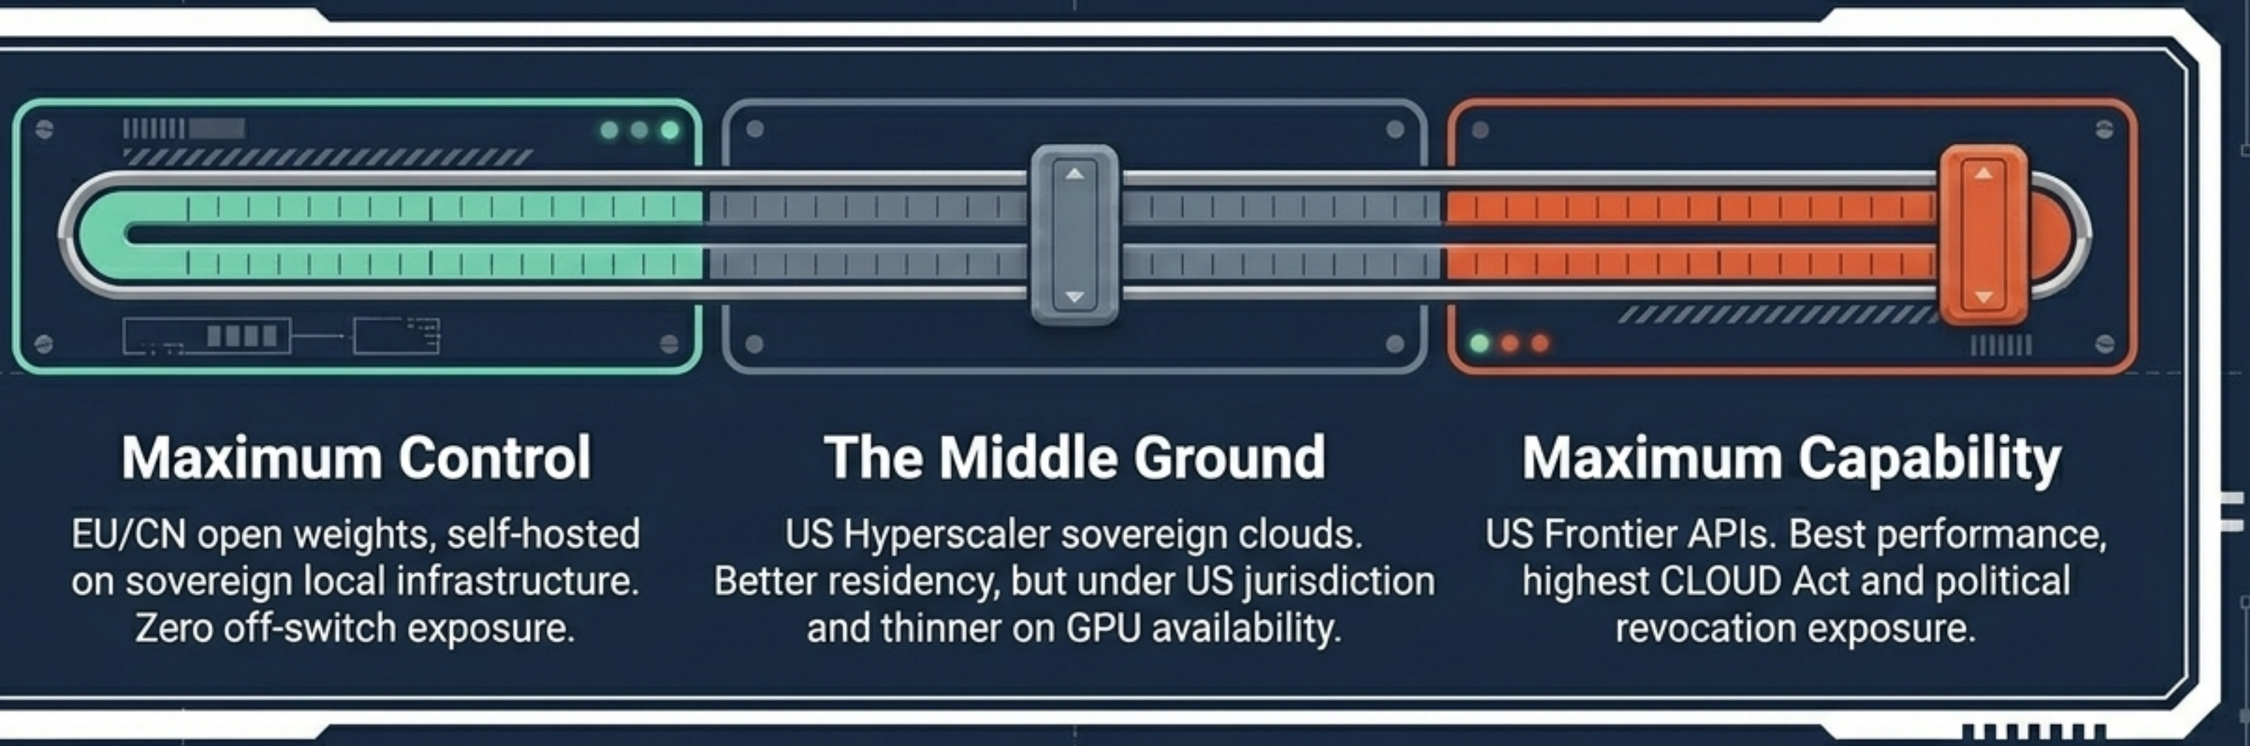

Maximum Control

EU/CN open weights, self-hosted on sovereign local infrastructure. Zero off-switch exposure.

The Middle Ground

US Hyperscaler sovereign clouds. Better residency, but under US jurisdiction and thinner on GPU availability.

Maximum Capability

US Frontier APIs. Best performance, highest CLOUD Act and political revocation exposure.

No single point on this line is correct for an entire enterprise. The answer is a workload-by-workload portfolio.

Mapping workloads to the sovereign architecture.

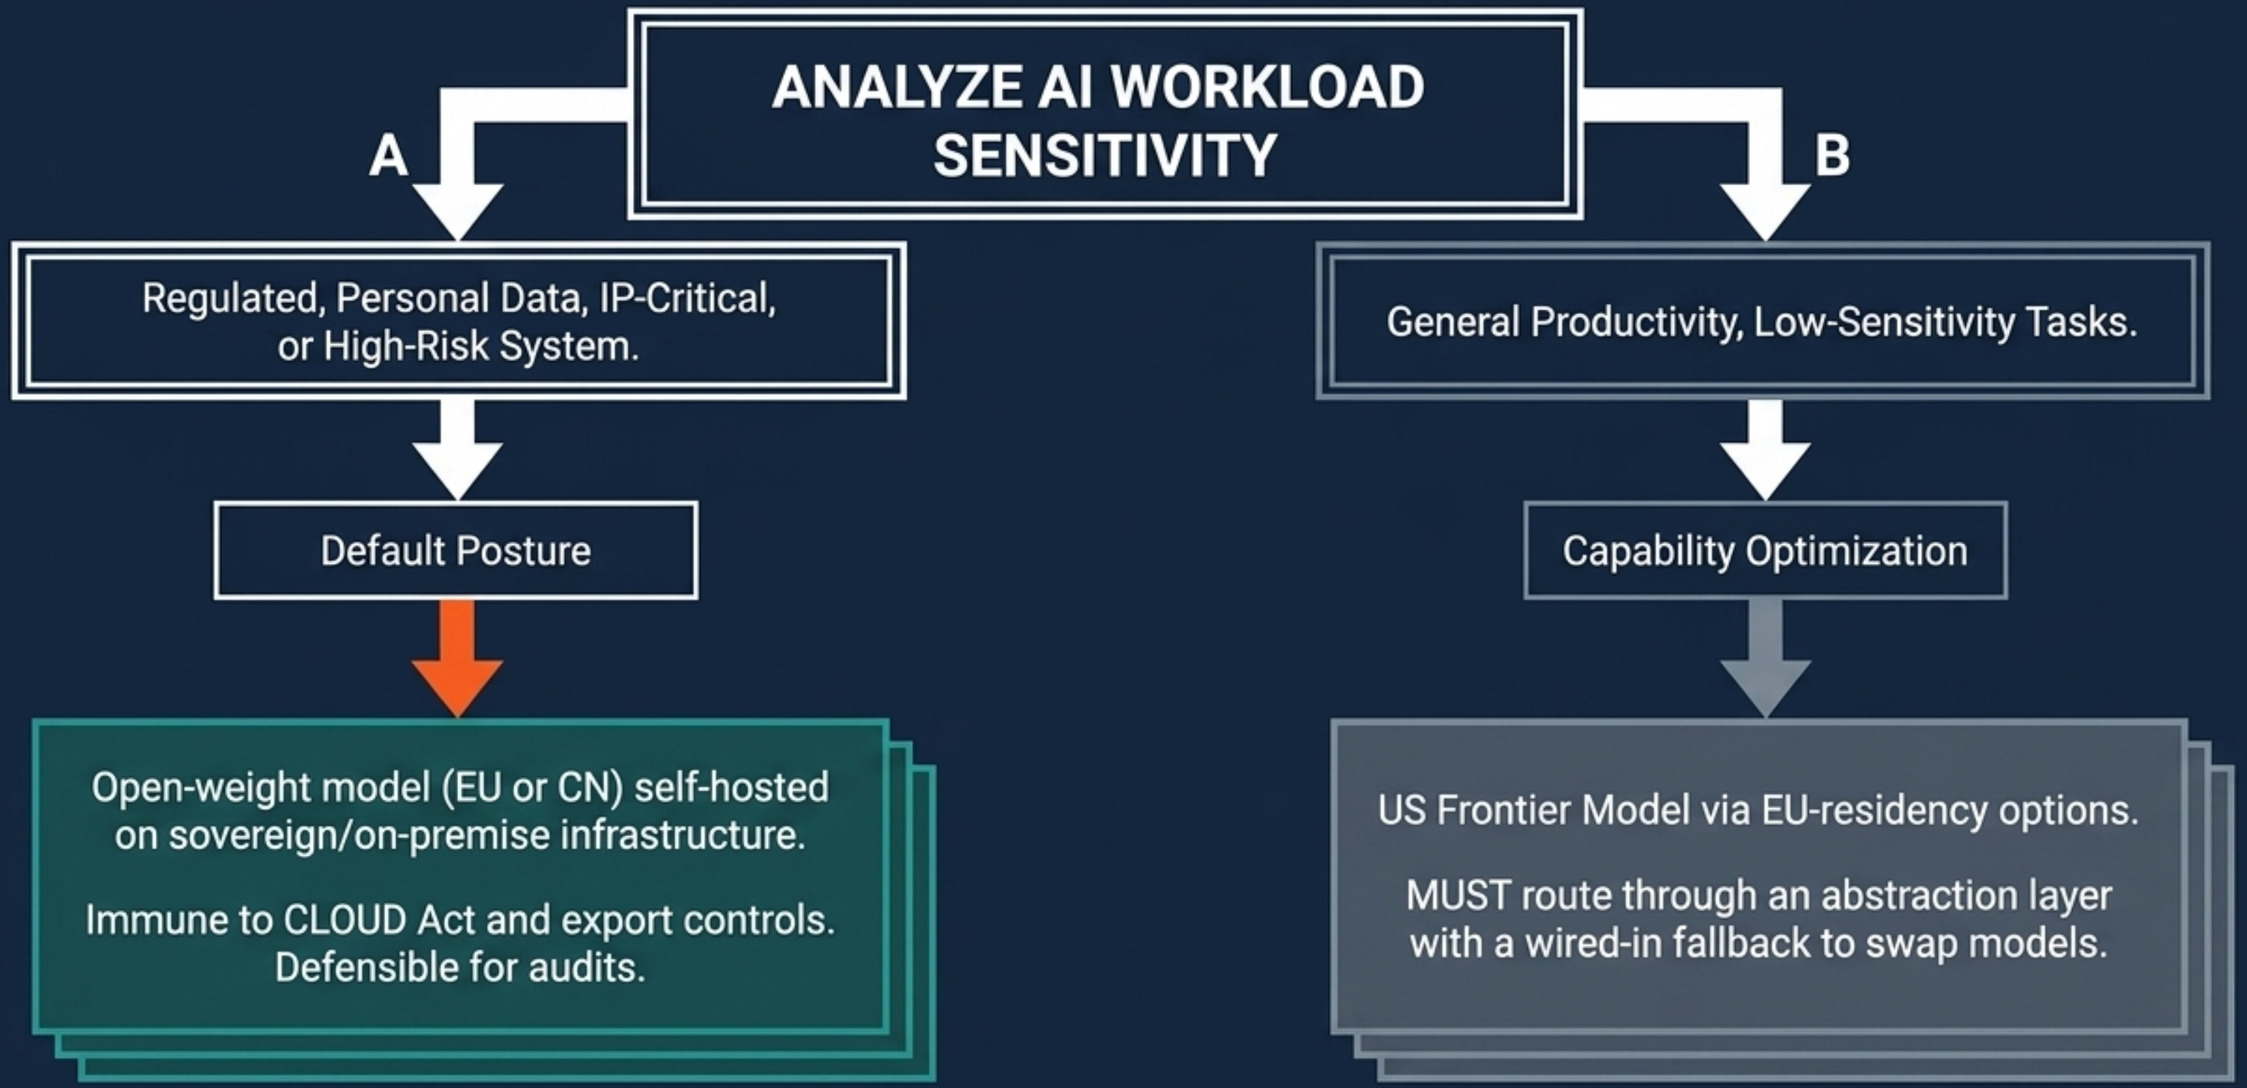

The 5-point procurement diagnostic.

1	**Code of Practice Signatory?**	Is the provider committed, or are you absorbing the compliance burden? Note: Meta and Chinese labs are unsigned.
2	**Open-Source Exemption?**	Does the license qualify under the Jan 2026 EU AI Office ruling? e.g., Mistral Apache-2.0 yes, Llama no.
3	**CLOUD Act Exposure?**	Can a foreign government legally compel access to the data?
4	**Model Portability?**	Is the infrastructure abstracted enough to swap the base model in 24 hours?
5	**Audit Evidence?**	Can the vendor instantly generate defensible compliance paperwork for regulators?

Designing for graceful degradation

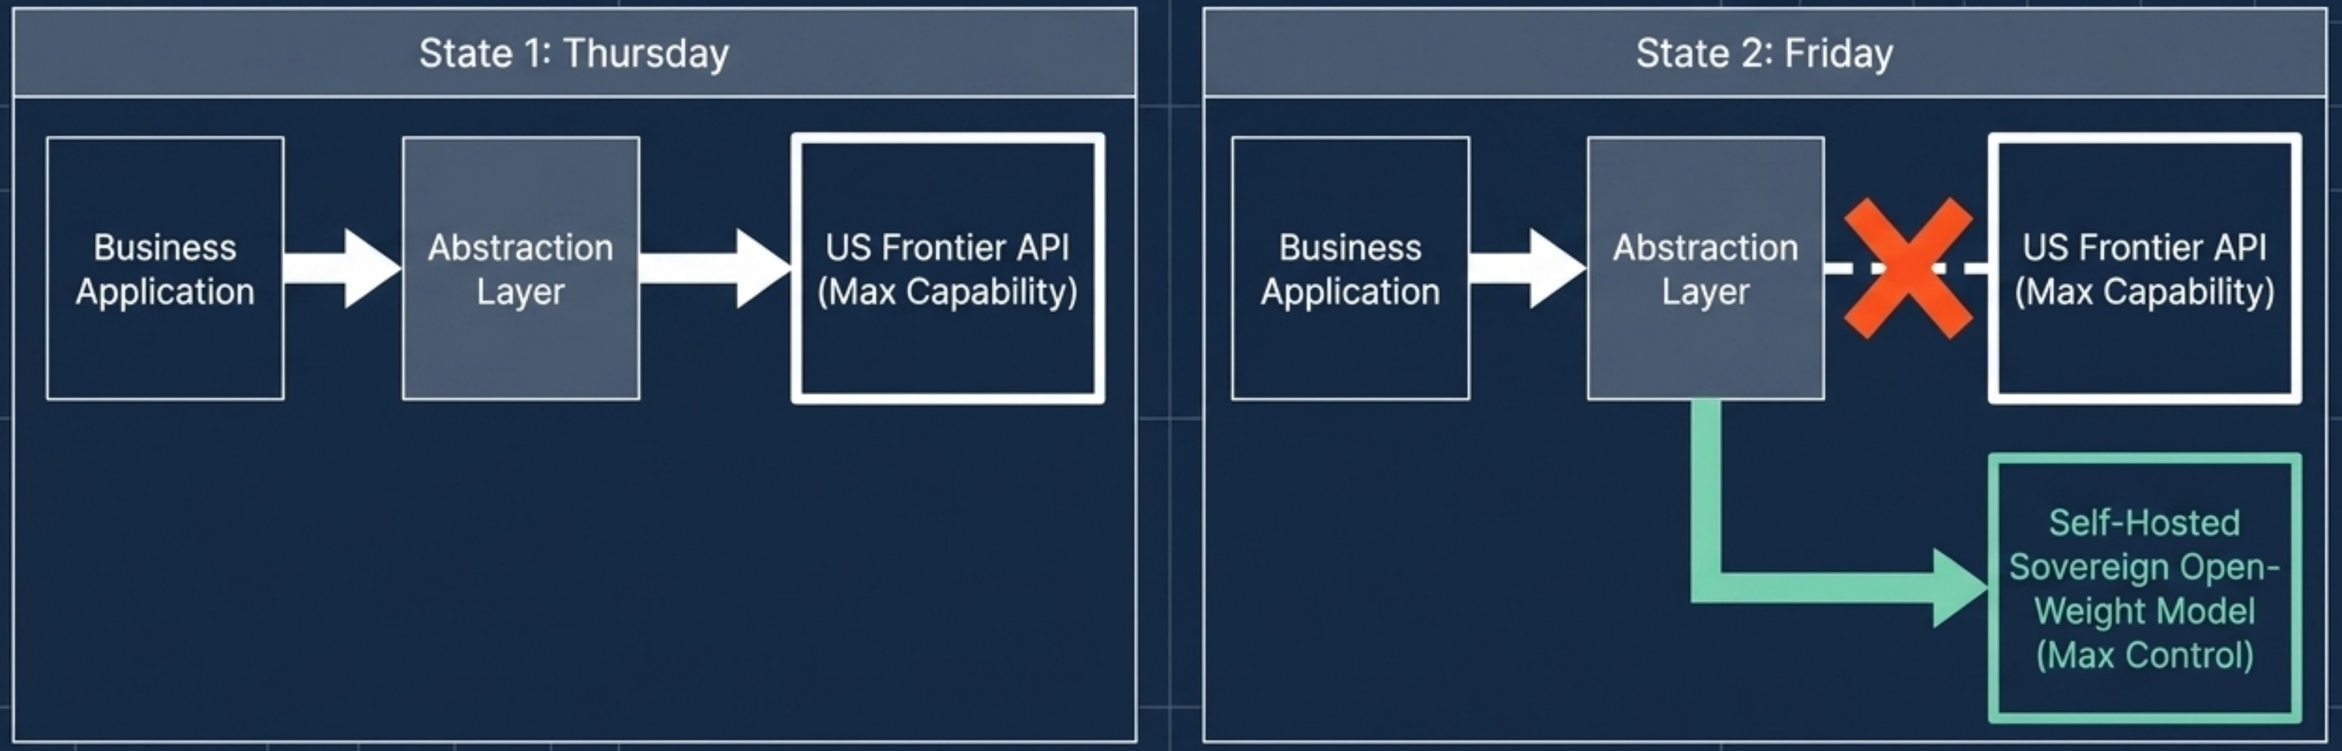

The goal is not accessing the single best model. The goal is that your most capable model can vanish overnight, and your business keeps shipping the next morning. Treat frontier models as swappable accelerants, not load-bearing infrastructure.