DIGITAL & CRITICAL TECHNOLOGIES

Artificial Intelligence: a Roadmap for European Strategic Autonomy

The following paper was originally published in French in *Le Grand Continent* and was translated to English with the help of artificial intelligence.

Lenny Benbara, Baptiste Lefort, Eric Hazan

THE COMPONENTS OF AN AI SOVEREIGNTY EFFORT do not run on the same clock. Chips can be bought in eighteen months; power, grid-connected land, large-scale training expertise and decision-making institutions take five to ten years. That asymmetry is missing from a debate now organised around a single question which is whether Europe should sprint for the frontier. Propositions going for the affirmative do so on the strength of two bets. The first is that the race to the most powerful model will remain the decisive contest. The second is that Washington will leave Europe enough time to buy the chips that catching up requires. An alternative is conceivable: a commoditised AI, where it is enough to buy models from an ever wider and more accessible market, as the Chinese Kimi 3 model suggests. Nobody can settle the question today and an independence strategy must therefore rest on slow assets, the ones that will confer real freedom of choice tomorrow. This article sets out a strategy that is indifferent to how the technology evolves and proposes to invest now in what will be useful under every scenario, map the most critical risks in order to address them, build the economic and geopolitical levers that will make those responses credible, prepare the frontier effort without committing to it blind, and equip Europe with the governance instruments needed to decide at the right moment.

1. A debate held hostage by technological uncertainty

Published on 8 July 2026 by the journal *Le Grand Continent*, "Operation Prometheus: France can win the AI race. What would it cost?" is a rare contribution to current European discussions on AI policy. Where most calls for a sovereign AI amount to wishful thinking, it offers a coherent, costed strategy which acknowledges its risks and its limits. It is treated here as the strongest available formulation of the maximalist case, not as an adversary: while we share a common ambition, the disagreement set out below bears rather on sequencing.

The French and European debate on sovereign AI has organised itself around two apparently irreconcilable poles. On one side, a sophisticated fatalism: Europe has neither the capital, nor the talent, nor the industrial base to close the technological gap. Economic rationality therefore dictates buying the best American models and abandoning autonomy ambitions judged ruinous. On the other, an avowed Prometheanism, of which Operation Prometheus is the most fully worked-out expression: $700 billion over three years, 12 gigawatts of compute by 2029 and a "Prometheus Act" suspending the ordinary planning and procurement rules, all to lift France to third place at the frontier.

These opposing positions rest on contradictory answers to the same question, one rarely stated explicitly: how fast will models keep improving?

There is a first possible future: gains become marginal, the distance between the best models and their pursuers narrows, and artificial intelligence gradually becomes a commodity, comparable to electricity, chosen on price. Call this the commoditisation scenario. In that world, the race to the frontier is a money pit without reward, and the fatalists eventually win the argument.

In a second possible future, capabilities keep advancing, the best models are used to train the next ones and to automate research itself. Technological progress becomes exponential and the gap widens rather than closes. Call this the breakaway scenario. In that world, a position at the frontier becomes an attribute of power comparable to nuclear weapons, and Prometheanism is strategically correct.

2. A question with no settled answer

Nobody today knows which scenario will prevail. Those who read the trend as incremental point to the narrowing of gaps on most public benchmarks. The latest wave of low-cost models, including Muse Spark 1.1, released by Meta on 9 July at a fifth of the price of premium models, is the most recent example. They also advance substantive arguments. First, a given capability level replicates at a falling cost: distillation and model compression deliver most of yesterday's performance for a fraction of the original compute. Second, barriers to entry in this commoditised segment are thinning, as training recipes circulate and open architectures proliferate, energy and chips aside. Third, the current model paradigm may be reaching its limits: Yann LeCun has long argued that large language models are a dead end on the road to artificial general intelligence, and that an architectural break will be needed, which would reset the race to zero.

The technological optimists reply that these benchmarks do not measure what matters. They saturate. Having become optimisation targets, they cease to measure anything. This is Goodhart's law, and the best models are now hitting the ceiling of what those tests can assess. The same conclusion was reached at the first Frontier AI Expert Forum convened by the European Commission in April 2026. Benchmarks capture neither the long-horizon tasks entrusted to agents nor capability jumps of the Mythos kind. The optimists add that customer preference for the best models shows clearly in revenue figures, a point corroborated by the European AI Office Expert Forum, which reports that industrial purchasing behaviour already reveals a marked preference for frontier systems over their immediate pursuers. The revenue gap between Anthropic and Mistral has widened rather than narrowed. They further note that the announced plateau still does not appear in the data, that the length of tasks agents can complete end-to-end roughly doubles every seven months according to METR, at an accelerating rate, and that barriers to entry, in compute as in exceptional talent, are among the highest in industrial history.

The situation can be summarised by a paradox that dissolves part of the controversy: efficiency gains make a given capability level cheap year after year (what Mythos cost yesterday will cost next to nothing tomorrow), but they do not make the frontier itself cheap, and its cost keeps rising. Both camps can therefore be right at once: rapid commoditisation of yesterday's capabilities alongside continuous inflation in the cost of tomorrow's. The question is whether tomorrow's capabilities will justify their price and prove necessary. That disagreement runs through the laboratories as much as through the investor base.

3. A method for testing the robustness of European strategy

A methodological consequence follows. Faced with deep uncertainty, a rational strategy does not necessarily seek the optimal decision within a predefined scenario; it combines measures that are robust, reversible and complementary, so that an adverse scenario does not produce a catastrophic collective loss. Decision theory has a name for this criterion: minimax regret. The best strategy is not the one that pays most in a given world, but the one that is catastrophic in none. It consists in buying options rather than making irreversible commitments, and in converting the option into a commitment only once the uncertainty has resolved. It is against this criterion, robustness across both scenarios, that we will test the maximalist plan, the fatalism of the commoditisation camp, and our own proposal.

This reading of Europe's present moment now has institutional backing. The Expert Forum convened by the European AI Office, drawing on more than a hundred specialists in advanced artificial intelligence, reaches two conclusions. The first is that 2026 will be decisive for Europe, which must devise and adopt a coherent strategy at Union and member-state level, taking account of capability, market and infrastructure trajectories and their implications for European sovereignty. The second is that under both scenarios Europe must strengthen its sovereign capacity to "access, choose, control and exploit" frontier models, which presupposes a clear view of the levers it holds, or could build, at each layer of the stack. Neither task has yet been carried out, and the coming months will show whether Brussels and the member states take them up or let them dissolve into another round of communiqués and announcements. The proposal sketched here answers both conclusions: a strategy designed to adapt to both scenarios and to settle the frontier question at the right moment, and the map of dependencies and levers the experts are calling for.

4. The dependence diagnosis is correct, and so is the case against inaction

The geoeconomic diagnosis behind the maximalist case is right. Dependence on frontier models is now a first-order strategic exposure.

The American restrictions imposed in spring 2026 on Anthropic's most advanced models, and their partial relaxation at the end of June, moved access to frontier models into a regime resembling that of semiconductors: revocable without notice, deliberately ambiguous, administratively discretionary. The most plausible reading of the episode is in fact the more troubling one: everything suggests that a security alert, probably linked to the model's hacking capabilities, led the US administration to seek to withdraw access from all users, American ones included. Export controls would have been the only legal instrument available to that end. Europe was collateral damage. Beyond the scenario of targeted coercion, this means that European access to advanced capabilities now hangs on an American assessment of domestic risk, and that no behaviour, no commercial concession and no diplomatic goodwill can secure it.

The diagnosis needs completing with a mechanism that makes the threat structural rather than merely transient and political: compute scarcity. Supplying a frontier model is not like selling software, where the marginal copy costs nothing. Every answer from a large model consumes tokens, the unit of account for AI consumption (the equivalent of the kilowatt-hour for electricity), and behind every token sits real compute. The major laboratories are in a state of chronic shortage and are constantly arbitrating between enterprise customers, consumer subscriptions and their own research needs. In a world of rationed tokens, foreign customers' access is the natural adjustment variable for purely economic reasons, before any political intention enters the picture. The allocation hierarchy observed at the leading laboratories confirms this: available compute serves US government contracts first, then enterprises, then the general public, as the differentiated service interruptions during load peaks show. Restrictions on foreign customers' access did not wait for a deliberate policy from the laboratories. That is precisely what makes the mechanism structural rather than intentional. Anton Leicht put it before this spring's episode: AI will not create abundance, but scarcity.

The asymmetry of levers between the Union and the United States is equally real. Cutting access to a model is instantaneous, whereas blocking exports from ASML, the Dutch company that builds the machines required for advanced chip production, is a slow lever whose effects are felt only once stocks run down. The risk of retaliation on other fronts must also be weighed: the United States can answer an ASML blockade with trade restrictions or by withdrawing security guarantees from the continent and from Ukraine. Europe has a lever, but not a sufficient one.

The warning against a naive reading of the Chinese laboratories is likewise warranted. Their rapid catch-up owes much to distillation, a technique that consists in querying an advanced model at scale in order to train one's own on its answers, in other words in free-riding on compute financed by someone else. That establishes nothing about the frontier being cheap for a new European entrant. One further consequence, rarely drawn, should be added: American laboratories and the US government now have every reason to close the door on distillation by tightening user identity checks and restricting access. The follower's route, catching up with the leaders at low cost by drawing on their models, is closing. That makes an autonomous training capability more valuable still.

5. Three uncertainties in a three-year frontier sprint

Three features of any three-year frontier sprint deserve discussion, and Operation Prometheus makes them unusually visible because it quantifies them.

1. The first concerns timing. The plan sets its priorities on the basis of three years of ramp-up fed by an uninterrupted flow of American chip deliveries, at a moment when its own visibility gives Washington every incentive to halt it. The scheme is thus most vulnerable exactly when it is most exposed. It is hard to argue simultaneously that commercial interdependence cannot protect our access to models and that it will secure the flow of far more strategic components. Two objections to this reading, raised by the plan's authors, deserve to be taken seriously. The first: Washington has so far preferred to sell; the current administration has relaxed part of the inherited restrictions and Nvidia lobbies effectively for open markets. But a favourable discretionary policy remains a discretionary policy, and the risk that it changes is the main risk. The second: a sovereign effort of this kind ought to be shielded while it is being built, and the United States would probably not take it seriously in its early years. That argument might hold for a politically discreet programme, but it does not fit a plan that claims 12 gigawatts, an exemption statute and a multilateral treaty, which is to say maximum publicity.

2. The second feature concerns the economics of technological catch-up. Capital buys catch-up, but slowly, while the cost of the frontier doubles every seven months; over three years the target will therefore have moved by several orders of magnitude. Recent examples abound, and Meta provides the most telling. The company made a colossal financial effort that produced only a repositioning in the mid-market on a price-performance basis: capable and profitable, but far from the absolute frontier. Three years of investment among the largest in the world bought a solid place in the pack, well behind the leaders. The trajectory of xAI and Grok 4.5 confirms the pattern: to reach the level of its competitors' previous generation, xAI needed three years of development, the backing of SpaceX and the acquisition of key software vendors such as Cursor. The pull of a national mission is a powerful lever available only to state structures, as the Manhattan Project or the early years of the French Atomic Energy Commission (CEA) showed, but the construct accumulates organisational burdens that the plan itself identifies as its main limiting factor. Its three-year horizon corresponds to the period in which over-capitalised projects routinely fail, for want of having learned to make trade-offs under resource constraints.

3. The third feature is the moving target. Aiming for 12 GW by 2029 means aiming for where the global leaders stood at the end of 2026. At the current rate of expansion, the frontier will sit several orders of magnitude above that level by the time the plan matures. The European AI Office Expert Forum describes global compute capacity doubling roughly every seven months; the largest planned sites will each exceed 3 gigawatts as early as 2027. Aiming for 12 gigawatts in 2029 means aiming at the leaders' present, and probably not at their future. Even on the plan's own account, a few gigawatts will not cover global economic demand. The project is therefore under-sized technologically while being crushing budgetarily, requiring between 4.5% and 8% of GDP per year. For comparison, the 1974 Messmer Plan, which built the French nuclear fleet, mobilised around 1% of GDP to deploy at scale a technology that was already proven and whose initial transfer was not revocable quarter by quarter. Add a further uncertainty: according to the experts consulted by the Commission, no durable business model has yet been demonstrated at the frontier. Some recent model generations are profitable, but the leaders' spectacular revenues still come with massive capital raises and persistent operating losses. Betting 700 billion on a position whose commercial profitability remains unproven is to stack one wager on another. In the commoditisation scenario, the sprint invests hundreds of billions to conquer a technological frontier whose market value collapses. In the breakaway scenario, it risks hanging on the willingness of its American chip supplier at the most critical moment of its construction.

Mid-July's news gives that uncertainty a concrete face. Moonshot AI has released Kimi K3, a 2,800-billion-parameter model whose weights should be freely downloadable by the end of the month. Early independent evaluations place it at the level of several recent large closed systems, though it remains behind the best, particularly on certain specific tasks **(FIG. 1)**. Its imminent availability in open form and its competitive access price nonetheless raise sharply the probability of a commoditisation scenario, that is, of a quasi-monopoly that can be replicated, is widely accessible and exerts continuous pressure on the price of proprietary models.

This opening can be read as an industrial power strategy: shifting competition from the proprietary model to the ecosystem, accelerating global adoption of Chinese architectures, and compressing the revenues that finance the American laboratories' compute race. Other things being equal, the situation resembles that of other European industrial sectors that faced the same subsidised-price competition, producing a well-identified geoeconomic ratchet: an abundant supply, backed by strategic capital, can render competing production economically unviable well before it achieves a dominant position.

AI model rankings on LiveBench

Scores out of 100 · Twelve models, eight task categories · Ranked by overall score

#	MODEL	OVERALL	REASONING	CODING	AGENTIC CODING	MATHEMATICS	DATA ANALYSIS	LANGUAGE	INSTRUCTION FOLLOWING	COST PER SUCCESSFUL TASK
1	GPT-5.6 Sol Max Effort	82.4	91.7	83.9	65.6	96.2	79.8	87.7	71.8	1
2	Claude Fable 5 Max Effort	80.8	89.7	86.0	46.9	96.0	80.5	90.7	75.8	2
3	GPT-5.5 Thinking xHigh Effort	79.9	89.7	82.1	52.1	95.9	81.6	87.4	70.7	1
4	GPT-5.6 Terra Max Effort	79.8	90.6	78.2	68.0	94.9	79.3	82.9	64.6	0
5	Claude 4.8 Opus Thinking xHigh Effort	78.9	89.7	79.3	56.1	95.3	78.3	81.4	72.4	1
6	Kimi K3 (open)	78.5	90.7	81.4	57.6	84.4	78.7	85.5	71.4	0
7	GPT-5.4 Thinking xHigh Effort	78.0	88.1	77.5	53.8	94.1	79.3	82.6	70.2	0
8	Gemini 3.1 Pro Preview High	77.1	84.0	76.5	45.4	91.0	78.5	85.4	79.1	0
9	Claude 4.7 Opus Thinking xHigh Effort	76.5	87.2	82.1	50.7	92.9	78.3	77.9	66.7	1
10	Grok 4.5	76.3	87.2	68.6	59.8	90.8	73.0	82.8	71.5	0
11	Muse Spark 1.1 xHigh Effort	76.2	87.7	77.2	65.1	87.1	72.5	74.3	69.6	0
12	Claude Sonnet 5 xHigh Effort	74.8	88.7	80.7	51.1	92.9	71.7	75.0	63.9	0

Best score in each category shown in coral.

Source: LiveBench

Kimi K3 thus reveals a third risk for any frontier sprint. After the technological risk of missing a frontier that keeps moving, and the geopolitical risk of losing access to the chips catch-up requires, comes the commercial risk of succeeding too late in a market whose prices have collapsed. The financial closure of such a plan assumes a progressively self-sustaining trajectory in which private demand takes over from the state once the frontier is reached. If near-frontier capabilities remain durably available in open form, the French laboratory could hit its technical target without ever reaching economic break-even. The $700 billion asset would depreciate before it was ever realised.

That consequence applies to our own proposal as well. In a market liable to be durably distorted by pricing and openness strategies backed by geopolitical capital, priority must go to the assets whose value best withstands a collapse in model prices: power and grid-connected land, a multi-model inference fleet underwritten by purchase commitments, orchestration, evaluation capacity and training expertise. Above all, it must be accepted that maintaining a domestic training capability will probably not rest on the market alone. Like a defence industry, it will have to be sustained by multi-year public procurement.

6. The rising strategic weight of the orchestration layer

For several months, value and strategic control have been migrating towards the orchestration layer. The term covers the systems that connect models to organisations' data and processes, handling memory management, permission control, certification of outputs and the routing of tasks according to their cost or complexity. If models are the engines of artificial intelligence, orchestration is the chassis, the transmission and the dashboard.

The signs of this shift are numerous. OpenRouter, a routing platform spanning several hundred models, illustrates how fast the market is repricing this layer: valued at $1.3 billion in May 2026 in a round led by Alphabet's growth fund and Nvidia's venture arm, it is reportedly, according to the *Wall Street Journal*, in acquisition talks with Stripe at around $10 billion barely two months later. Likewise, the laboratories' commercial success no longer rests on selling raw tokens but on deploying agentic products that wrap models in structured workflows. And the emergence of low-cost offerings such as Muse Spark or the open-weight model Inkling is accelerating the adoption of multi-model routers able to select the cheapest solution for each request. By making technological architectures interchangeable, orchestration exerts continuous downward pressure on models. Much of a model's capability is in fact latent, revealed or lost through the quality of the system around it. Two qualifications apply. At the top of the capability hierarchy, the integrated stacks of the leading laboratories retain a clear lead over equivalent assemblies built around open models: interchangeability is fully effective for routine tasks, but not yet for the most demanding autonomous workflows. And routing is first of all a matter of economic optimisation: large French asset managers have deployed internal gateways that route each of their developers' requests to the cheapest model capable of handling it, aligning compute consumed with task complexity.

This shift redefines dependency relationships. It is now at the orchestration layer that user adhesion, institutional memory and switching costs crystallise. An economy that masters this application layer can substitute one model supplier for another at low cost, as one swaps an engine under the same chassis. Conversely, an economy that gives up control of it stays captive even if it has a sovereign model: what holds the customer is no longer the raw algorithm but the surrounding ecosystem. Sovereignty in use is therefore decided first in orchestration.

Two planes should nonetheless be kept distinct. Economically, the orchestration layer is where value added is created. Geoeconomically, models are substitutable for one another through the orchestration layer, which allows a switch from an American supplier to a Chinese, European or open one, provided a fallback exists. That is precisely what the foundation we propose below is meant to guarantee. Two risks specific to this layer remain to be identified and neutralised: a model supplier can degrade competing orchestrators' access to the interfaces they depend on, and orchestration players tied to laboratories by shareholding can tilt the field towards their parent company.

This layer is favourable terrain for Europe, whereas frontier models have so far eluded it. Enterprise integration is a matter of business software, process knowledge and sector expertise, the core competence of firms such as SAP and Dassault Systèmes. Here the entry ticket is measured in human capital and industrial expertise, not in compute. The urgency is no less real: American laboratories are now investing aggressively in this application layer, and the window may close quickly.

A robust strategy on this front also copes with both technological scenarios. If models commoditise, margins and value will take refuge at the orchestration layer. If technological rupture continues (the breakaway), this layer determines an economy's ability to convert raw AI power into real productivity gains. At equal access to large models, the value gap will widen according to the quality of integration, a gap that already separates American firms from their European counterparts.

Orchestration is no panacea. It rounds out the strategic portfolio, but it does not compensate for the absence of access to power models (of the Mythos class), which are indispensable to cyber defence and intelligence. It does, however, correct a persistent bias in the public debate: by fixing attention on the technological frontier, we neglect the segment where European strengths are real and the barriers to entry surmountable. This work does not require outsized funding; it requires activating execution levers that are already available: public procurement that is demanding on sovereignty for critical application layers, multi-supplier obligations to guarantee business continuity, and the structuring of sectoral data commons to protect Europe's advantage on the ground.

7. The robustness test applied to fatalism and to the breakaway

The fatalist approach fails against the challenges identified above. In the commoditisation scenario, it deprives Europe from the outset of the industrial positions it could take in a market that has become accessible. In the breakaway scenario, it leaves Europe defenceless against sudden restrictions on access.

Before setting out our proposal, let us anticipate the first objection it will meet. The robustness requirement we defend must face its most serious criticism: is a fallback strategy built on a player below the technological frontier, such as Mistral, valid only in the scenario where it turns out to be unnecessary? If commoditisation happens, the fallback is easy but superfluous; if the frontier pulls away, a model eighteen months behind becomes obsolete for advanced uses. The objection calls for a two-part answer.

First, the main risk we seek to guard against is not technological lag but loss of access, a latent threat in both scenarios. In the event of a cut-off, the question is no longer competing with the global state of the art but preserving a minimum of autonomy. A sovereign model close to the frontier without reaching it will always be infinitely more useful for the vital functions of the state and for economic continuity than nothing at all. Ukraine's experience is telling: without dominating the frontier of autonomous systems, its army transformed its military capabilities using drones built from accessible technological components and integrated in the field.

Second, we accept that a simple insurance logic is not enough in the breakaway scenario. If frontier capabilities are a first-order attribute of sovereignty, Europe will sooner or later have to possess them or guarantee strategic access to them. That is why our proposal is not confined to protection: it is structured as an option doctrine, in the financial sense. Where insurance compensates a loss after the fact, an option buys today, for a fraction of its true cost, the right to enter the race at a moment of one's choosing. The whole challenge of the trajectory we propose is therefore to determine what must be built now to secure that option, without paying its prohibitive price in advance.

8. Our proposal: an industrial foundation that holds under every technological scenario

The plan we propose takes account of a temporal asymmetry usually ignored in the sovereign AI debate. The components of a frontier effort do not share the same horizon: chips can be acquired within eighteen months once capital and allocation access are secured, whereas power, grid-connected land, large-scale training expertise, evaluation protocols and decision-making bodies require five to ten years to build. The maximalist plan proposes to finance the fast, expensive brick immediately: frontier training compute. The uncertainty around its value seems to us too great. We propose instead to begin with the lower stages of the rocket: building now, at controlled cost, the slow infrastructure that will retain its value whatever the final scenario.

This alternative trajectory rests on a foundation of five priority projects.

1. Power and grid-connected land. Whatever the technological scenario, Europe will need inference on a massive scale, and abundant decarbonised electricity is the only link in the value chain where France holds a global comparative advantage, provided it adapts its energy policy. The competitiveness gap is documented: in 2025, industrial electricity prices in the European Union were on average twice those in the United States and half as high again as those in China. Power, permits and grid connection, rather than capital or chips, were the real bottleneck for compute in Europe. Fixing this means accelerating connections, purging the phantom projects that clog the queues, creating dedicated industrial zones so as to bring permitting times down to a few months, establishing a single point of contact and relaunching nuclear. This component of the maximalist plan should be executed in full, because it holds even if not a single model is ever built on European soil. Its direct budgetary cost is modest, since most of it is regulatory; acquiring public land and providing grid guarantees runs to a few hundred million euros a year.

2. A sovereign inference fleet. That is, a set of data centres dedicated to running models day to day. In the short term, Europe will have to accept reliance on Nvidia here. Refusing to buy American processors on grounds of dependence would confuse stock with flow: the strategic risk lies in future deliveries and revocable software access, not in physical infrastructure already installed. The sovereignty of this fleet does not depend on the origin of its components but on two imperatives: location on European soil and operation exclusively by European players. A data centre located in Europe but operated by a foreign third party can be disconnected remotely and instantly; a centre run by a European firm with imported hardware offers, at worst, several years of respite before technological obsolescence, long enough to reverse the balance of power. The sizing is not utopian: even the most conservative demand projections indicate that by 2028 a large European country's inference needs will far exceed all the capacity currently planned on the continent. The imminent danger is under-provision, not overcapacity. Nvidia and AMD remain unavoidable for now, but hardware sovereignty means, over the medium term, shifting a growing share of inference towards specialised accelerators (ASICs) and, over the longer term, mastering the whole stack: interconnect, memory, packaging, compilers and libraries, not an isolated chip. Developing a European accelerator industry now is the natural extension of the inference fleet.

For this fleet, targeting 2 to 3 GW by 2030 would represent an investment of €66–101 billion over five years, using the unit costs put forward by the Prometheus plan itself (around €33.4 billion per gigawatt owned and €0.8 billion a year in operating costs). The order of magnitude is striking, but it is the rigour of the financial structure that makes it credible.

The economics of such a fleet turn on two critical variables: utilisation rates and the accelerated depreciation of chips, which lose most of their value within four to five years. The first risk can be neutralised through anchor tenants. These strategic customers secure the infrastructure's returns through firm take-or-pay contracts (a commitment to pay for reserved capacity whether or not it is used). Such commitments would be subscribed by public procurement, by the capability custodian of the third project, and by the large European firms whose continuity plans require precisely a domestic fallback capacity. The second risk is managed through the capital structure. The core of the syndicate must bring together long-term investors (Caisse des dépôts, France's state investment bank, the European Investment Bank, infrastructure funds) backed by a state-assumed first-loss tranche to reassure private capital, as the Prometheus plan rightly argues. Calibrated on depreciation risk, that public stake would represent 10% to 20% of the total, or €2–4 billion a year. The remaining capital would be raised from sovereign wealth funds seeking to hedge their own exposure to the Sino-American duopoly: Norway, able to pair the firepower of its €1,800 billion fund with its hydroelectric assets, or the Gulf states, whose interest is already visible in the presence of the Emirati fund MGX in the capital of Campus AI alongside Mistral, Bpifrance, the French public investment bank, and Nvidia.

3. A custodian of training capability. This is a more demanding reformulation of the fallback-provider role. Mistral's strategic function, from the state's point of view, lies not in its model catalogue but in the fact that it is today the only place in Europe where the know-how of training large models, a form of human and collective capital that cannot be bought in eighteen months, exists and renews itself. That know-how becomes more precious still as the distillation route to catching up closes: tomorrow, one will either know how to train or be wholly dependent.

The base of multi-year public procurement (defence, security, core state functions) must be sized and contracted to perform this function: keeping a world-class team in continuous training conditions, within reasonable distance of the frontier, against verifiable capability milestones. Concretely, this means a firm purchase commitment of €1.5–2 billion a year over five to seven years, plus dedicated training compute of €0.5–1 billion a year, or €2–3 billion a year in total. Through this lever, the state buys an option. The formula carries a demanding counterpart: support is conditional and reversible. If the designated provider durably fails to meet the objectives set, the contract can be transferred to another vehicle, a consortium, a new structure or a team reassembled around the same infrastructure. The bet is on a capability, wherever it sits, not on a company. Without a custodian of that capability, the frontier option does not exist, however much capital is available on the day one wants to use it.

4. Sovereign evaluation capacity and dependency mapping. Independently assessing the real performance, the failure modes and the critical uses of the most advanced models costs around €100 million a year at European scale. This evaluation conditions the official classification of uses (those where dependence is acceptable, those that require a fallback, and those that exclude it), the credibility of access negotiations and, above all, the reading of the indicators on which the decision to exercise the option depends. It is the nervous system of the doctrine, what makes it operational and usable.

5. Collective negotiation of access guarantees. It is worth spelling out precisely what would give Europe a real lever, against the received idea that the size of the European market is enough. Where compute is scarce, the loss of European customers affects American laboratories only marginally: they reallocate the freed capacity to domestic demand or to their own research without difficulty. The status of "interested buyer" is therefore structurally harmless.

Europe's real strategic pivot lies elsewhere. It lies first in its infrastructure offer: Europe, and France in particular, can offer American suppliers a resource they badly lack, namely connected sites and abundant decarbonised electricity, in exchange for contractual guarantees of access to the best models. This mechanism has a valuable asymmetric property: once the compute centre is built on European soil, the invested capital becomes hostage to the relationship. A supplier who reneged on its commitments would be left with billions in assets deprived of power, while the US government, if it tried to impose restrictions, would run into the powerful lobbying of its own companies.

The second lever is a coalition of industrial chokepoints. ASML's monopoly does not stand alone; the global semiconductor value chain also depends on optics from the German firm Zeiss, lasers from the German firm Trumpf and memory architectures from the South Korean firms SK Hynix and Samsung. Coordinated between The Hague, Berlin, Seoul and Tokyo, this collective position carries geopolitical weight that commercial demand alone no longer has. To this architecture must be added the classic contractual clauses: notice periods, continuity of service, escrow of model weights with a trusted third party in the event of an access cut-off, and one condition too often set aside, the absolute security of European networks, since no laboratory will entrust its most valuable models to infrastructure judged permeable. None of these guarantees amounts to full sovereignty, and recent events are a reminder that granted access can be revoked. But combined with the inference fleet and the stewardship of training capability, they change the structure of the game: an access cut-off no longer condemns Europe to paralysis; it imposes a prohibitive switching cost on the supplier. Economic deterrence does not remove dependence, it makes it costly to exploit.

Annual public spending, by workstream

Billions of euros per year · Public outlay only · Five workstreams

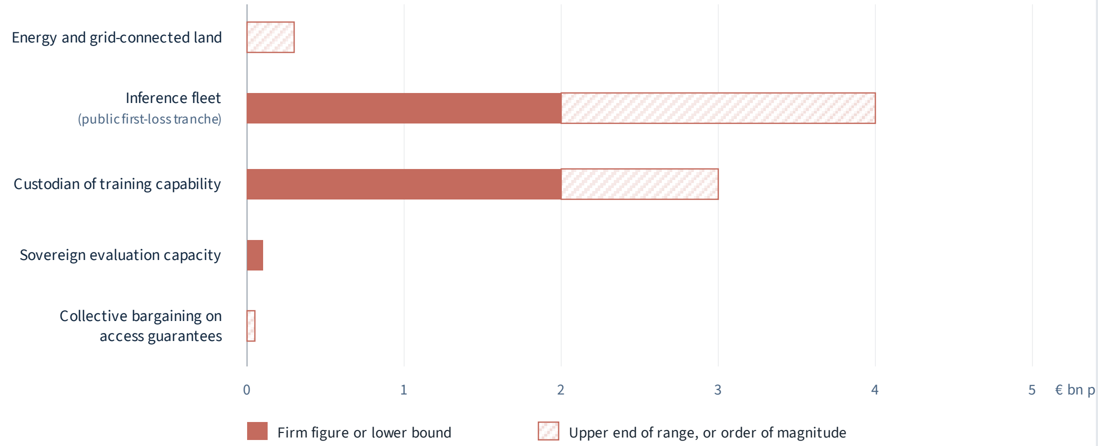

Workstream 2 shows only the public first-loss tranche; the remainder of the inference fleet (€75–115 billion over five years) is financed by private and foreign sovereign capital. Workstream 1 is not precisely costed in the text. Source: GTPO (2026)

In total, the cost of this foundation for France and its partners comes to €5–7 billion a year, or around 0.2% of French GDP **(FIG. 2)**. That figure should be compared with the 1.5% required for the public component of the maximalist plan alone, and with the 4.5% to 8% of its full cost. Over five years, our approach mobilises $90–135 billion, financed overwhelmingly by the private sector and backed by operating revenues. Where the sprint commits largely public money, we propose to mobilise mainly private capital, backed by productive assets and operating income. With this robustness trajectory in place, let us turn to the two maximalist scenarios.

Branch A: if commoditisation sets in, go on the offensive through diffusion

If the next eighteen months confirm the incremental hypothesis, marked by converging performance, commoditised capabilities and a price war of which Meta's repositioning is the leading indicator, the strategic centre of gravity will shift definitively towards diffusion, meaning the effective adoption of AI across the economy. As the Draghi report established, Europe's productivity gap is as much a gap in adoption as in production. The €1,000 billion in annual spending mentioned by Arthur Mensch will be captured by the models actually deployed on the ground, not necessarily by the heaviest frontier architectures. In that context, cost of use, speed of execution, business integration, regulatory compliance, data localisation and command of agentic workflows become the decisive variables, all areas where European players can build a durable comparative advantage. The supplier's nationality retains major strategic importance, moreover, as regards where the value generated in Europe is booked, the exploitation of usage data to improve future models, and resilience in the face of potential restrictions.

In this scenario, several levers must be pulled immediately. First, the public levers that accelerate the diffusion of AI use: the price of electricity, a key driver of usage cost; public procurement acting as an anchor customer; the administration, which remains the continent's largest buyer of intellectual services; the structuring of sectoral data commons in health, industry and mobility; the training effort; and a regulatory overhaul that actively encourages operational deployment.

Second, European preference must be rigorously calibrated so as not to act as a brake on diffusion. A blanket preference would artificially raise the cost of AI for the whole economy at a moment when wide diffusion is essential. Worse, a blind preference would risk equipping Europe's hospitals, networks and administrations with second-rate cyber defence tools while advanced offensive capabilities become commonplace. Preference must therefore be doubly bounded: reserved for the critical uses identified by the foundation above, and conditional on sovereignly verified performance thresholds. No preference mechanism should apply without a performance floor: if the European player proves unable to hold the required level in a given segment, the clause is automatically suspended. This is a key condition for aligning economic preference with the security imperatives it claims to serve.

Building this policy doctrine will require weighing on a balance of forces that is already engaged but still entirely open. The Cloud and AI Development Act, presented on 3 June 2026 as the centrepiece of the Tech Sovereignty Package, proposes a four-tier sovereignty framework for public cloud procurement, but the text is still only a proposal: its real content will be decided in the negotiation between the European Parliament and the Council, where the effective level of stringency remains contested. The recommendation here therefore aims at influencing that negotiation rather than at persuading a Commission that would speak with one voice, given that the internal arbitration between an internal market directorate (DG GROW) more favourable to assertive industrial policy and a competition directorate (DG COMP) more attached to competitive neutrality will continue to shape it.

Branch B: the breakaway scenario and exercising the option

If model performance indicators instead confirm the breakaway scenario, marked by widening gaps on long-horizon and agentic tasks, repeated performance jumps, a drying up of open architectures and a tightening of access controls by American or Chinese decision, the frontier effort becomes fully justified. The foundation will then have been designed precisely to make that ambition possible, at lower cost and lower risk. This branch takes up the substance of the maximalist plan, but triggers it at a moment when the project has become a certain and winning option, in sovereignty, technology and commercial terms alike.

It would differ from an immediate sprint on five points.

First, the starting point is cleared: the energy infrastructure, the connected sites, the inference fleet and the training teams already exist. The additional cost falls exclusively on training compute proper, preserving public funds during the ramp-up phase. Exercising the option does not make the frontier cheap; it guarantees an informed decision, prepared ground and a burden that can be shared.

Second, the trigger rests on observed empirical milestones rather than on a theoretical bet. And the geopolitical context that would force the option to be exercised would also reshape the underlying balance: a tightening of access dictated by Washington would legitimise the coalition, dispel the hesitations of partners paralysed by fear of offending the American ally, and accelerate the emergence of alternative chip architectures.

Third, the coalition would take a tighter, variable-geometry form, structured around the capability custodian rather than around a new body created from scratch. The state would act as guarantor of strategic control, not as scientific prime contractor. The CEA philosophy invoked by the plan is the right one, but its treaty mechanism, guaranteed access rights in exchange for a financial entry ticket, must extend beyond the European Union alone. It should include the United Kingdom, Switzerland and Norway for their expertise and their energy, Japan and South Korea for the semiconductor supply chain, and the Gulf states for their capital. All of these countries share one vital interest: circumventing the Sino-American duopoly.

Fourth, financing would rely on the structure already tested for the inference fleet, with patient capital from large sovereign wealth funds, remunerated through guaranteed usage rights rather than through mere promises of financial return.

Fifth, the strategic objective would no longer rest on the uncertain promise of "joining the leaders in three years", whose fragility recent market dynamics have exposed. The ambition would be to join the lead pack within half a decade, on a realistic learning curve consistent with the industrial history of capital-intensive technological catch-up.

9. Triggers and the governance of the decision

An option doctrine is worth nothing unless a body monitors its critical indicators and is empowered to activate execution. That is precisely the blind spot of the current debate, which accumulates strategic analyses without equipping itself with a decision-making arm. To remedy this, our proposal rests on three concrete provisions, unified by one cross-cutting principle: variable geometry. None of the projects below should be made contingent on unanimity among the twenty-seven member states, whose national interests are too divergent to carry a power policy on this subject. A core of willing states must be the starting point.

First, steering requires a public scenario dashboard, operated by a sovereign evaluation authority working with the European AI Office but independent of its supervision. This instrument would objectively measure the evolution of performance gaps on complex agentic tasks, and not only preference rankings; the frequency and scope of American and Chinese access restrictions; the vitality of the open-model pipeline; the tightening of protections against distillation; the penetration rate of autonomous agents in the real economy; and the industrial maturity of alternatives to Nvidia. This dashboard would be reviewed quarterly against predefined criteria whose breach would trigger the relevant protocols.

Second, governance must rest on a clearly identified and deliberately flexible decision mandate. The states involved in co-financing, European or not, would form a contributors' board, while the European Commission would provide technical and legal support without holding a veto. This is the proven architecture of the continent's great industrial successes, from Airbus to space cooperation, and the opposite of multilateral initiatives dissolved in governance. The decision to launch the frontier effort is political and budgetary. It must rest on costed scenarios kept continuously up to date, so that it is not improvised in the urgency of an access crisis.

Third, the arrangement must include a private component, since the criticality of these technologies affects firms first. European boards must treat model access as a major supply risk. That means implementing dual-supplier inference policies, rigorously testing continuity plans that include switching to models operated under European jurisdiction, and introducing binding contractual notice clauses. This operational toolkit should be built by drawing every lesson from the baptism of fire suffered this spring, with the partial suspension of access to Mythos and Fable.

The price of independence, and the order in which it is paid

The Gaullist formula invoked in support of the sprint, "it is very expensive, but it is the price of independence", cuts both ways. Nuclear deterrence was a bold technological bet, but its strength lay in the ability to convert an initial dependence (Norwegian heavy water exfiltrated in 1940, German ballistic expertise recovered in 1946) into a closed domestic chain within fifteen years, from the Marcoule plutonium plant (1956) to the missiles of the Albion Plateau silos (1971). That closure was possible because France abandoned heavy-water technology, which depended on outside supply, in favour of the gas-graphite reactor line.

The fundamental objection to pursuing a frontier model in the short term concerns not only cost but the sequencing of investment. In our view, the foundation that makes it possible must come first: energy, infrastructure, skills and decision-making institutions. Priority belongs to what takes time. If the maximalist assumptions are borne out, this approach will at least have prepared the ground: the sprint will then find determined partners and a components market potentially freed from today's monopolies.

In an uncertain world, technological and industrial sovereignty does not run through a maximalist bet but through a plan resilient enough to accommodate every technological scenario. That means the capacity to prosper if the technology commoditises, to protect ourselves if access closes, and to enter the frontier race at the right moment. This industrial foundation, proposed as an option value, costs roughly ten times less than the sprint it makes possible, and requires no one's permission in Washington.

Notes

01 Lenny Benbara and Baptiste Lefort, « Pourquoi l'IA c'est Thémis plutôt que Prométhée » [Why AI calls for Themis rather than Prometheus], *Le Grand Continent*, 28 July 2026. legrandcontinent.eu

02 Meta released Muse Spark 1.1 in early July 2026, presented in the press as a low-cost offering aimed at mass usage. See Harshita Mary Varghese and Katie Paul, « Meta debuts Muse Spark 1.1 model with preview open to developers », *Reuters*, 9 July 2026. reuters.com

03 « AI Office publishes frontier AI experts' conclusions on EU competitiveness, sovereignty and security », European Commission, 15 July 2026. digital-strategy.ec.europa.eu

04 On Mythos, its exceptional cybersecurity capabilities and the restriction of access to it (Project Glasswing).

05 METR, « Measuring AI Ability to Complete Long Tasks » (2025): the length of tasks completed autonomously at 50% reliability doubles roughly every seven months.

Disclaimer. This paper was originally published in French in *Le Grand Continent* on 28 July 2026 and translated to English with the help of artificial intelligence.