

European Foundation Models and the Value of Sovereignty and Regulation Ready Deployment

A Decision Framework for Procurement Access Language Coverage Compliance Cost and Capital

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

European foundation-model companies are often described as sovereignty assets. The label has economic value only when it creates customer access, reduces a material dependency, improves evidence for regulated deployment or secures a scarce operating right. Incorporation in Europe, European hosting and support for European languages are different attributes. Each should be tested separately against the procurement, legal, technical and financial requirements of the customer.

This paper develops a decision framework for valuing European foundation-model businesses as the European Union moves from rule formation to enforcement and expands sovereign compute capacity. It links six diligence ledgers: customer procurement, language and domain performance, data and intellectual-property rights, regulatory documentation, compute and deployment control, and cash requirements. It also distinguishes obligations that attach to general-purpose model providers from requirements that arise at the downstream system and use-case level.

Four analytical cases are used: API-led European challenger, regulated-sector specialist, sovereign private-deployment platform and multilingual public-service infrastructure. A wholly hypothetical illustration produces annual revenue from EUR 145 million to EUR 310 million, contribution from EUR 32 million to EUR 158 million and cash before financing from negative EUR 112 million to positive EUR 18 million. A separate probability-weighted illustration produces enterprise value of EUR 1,448 million. These figures are assumptions used to demonstrate the framework. They are not market observations, forecasts or valuation conclusions.

The analysis finds that sovereignty should be valued through verified procurement access, retained contribution and controlled deployment rather than through a narrative premium. Regulation readiness can shorten diligence, preserve eligibility and reduce remediation risk, while creating recurring documentation, evaluation, cybersecurity and governance costs. Multilingual breadth can widen the addressable market, although commercial value depends on performance in the customer language, domain and workflow. Investment and transaction structures should therefore release value and capital against evidence: accepted use cases, signed procurement routes, model and data rights, documented compliance, operational portability, contracted compute and collected cash.

JEL Classification: G12, G24, G32, G34, L11, L86, O31, O33

Keywords: European foundation models, sovereign AI, AI Act, regulation ready deployment, multilingual artificial intelligence, public procurement, AI valuation, EuroHPC

Introduction

Europe's foundation-model market sits at the intersection of technology, procurement, regulation, language and industrial policy. The European Union's AI Act creates obligations for providers of general-

purpose AI models, including technical documentation, downstream information, a copyright policy and a public summary of training content. Additional duties apply to models with systemic risk. The Commission states that the provider obligations have applied since 2 August 2025 and that enforcement powers apply from 2 August 2026 [7-12].

Europe is also expanding compute capacity. The AI Continent Action Plan links AI Factories, proposed AI Gigafactories, data access, skills and adoption. EuroHPC describes access routes for eligible startups, industry and research users, including large-scale allocations for foundation models [13-18]. These programmes may improve access to compute and technical support, but an allocation does not automatically become owned capacity, transferable value or recurring commercial supply.

The valuation question is therefore narrow and evidence led. An investor should determine which attributes of a European model company create paid access, reduce delivery risk or preserve strategic options, and what they cost to maintain. Sovereignty, compliance and multilingual coverage belong in the cash-flow model only through observable customer behaviour, enforceable rights and funded operating requirements.

1 Define sovereignty as a set of control rights

Sovereignty is not a single company attribute. The diligence perimeter should separate legal domicile and ownership, model and software rights, training-data provenance, compute location, operational control, support access, encryption and key control, subcontractors, governing law, incident response and exit portability. A company can satisfy one element and depend on a non-European provider for several others.

The customer requirement must be stated before the premium is measured. A public authority may require data location, auditability, continuity and procurement eligibility. A regulated bank may emphasise operational resilience, vendor concentration, model governance and evidence for supervisory review. An industrial buyer may prioritise local deployment, intellectual-property protection and predictable latency. The relevant sovereignty stack differs by use case.

The investment model should assign each control right to a contract, asset, process or verified capability. A marketing claim without evidence supports no base-case premium. A right that can be reproduced by several providers may improve eligibility without creating scarcity. A right that is difficult to replace, accepted by customers and retained through the forecast can support revenue quality or lower risk.

The control analysis should also distinguish authority from location. A European data centre can host infrastructure whose capacity, software administration or cryptographic control remains with a non-European supplier. Conversely, a distributed technical component can be governed through enforceable customer rights, tested continuity and replacement routes. The investment committee should document the practical consequence of each dependency, the time and cost to replace it, and the customer segment for which the dependency is disqualifying.

2 Map European procurement access

Procurement access should be analysed as a funnel. The stages are eligibility, framework admission, security and legal diligence, technical validation, funded pilot, production award, accepted service, invoice and collected cash. Each stage has a conversion rate, cycle time, cost and evidence standard. Broad interest should not be treated as contracted demand.

The Apply AI Strategy describes a European approach that seeks adoption across strategic sectors and supports a buy-European orientation in parts of the public sector [19,20]. Standard procurement clauses and public-buyer guidance also show that buyers may seek transparency, data, audit, risk and exit provisions [21-23]. Their existence can create an access route for a regulation-ready supplier while increasing the cost of proof and contractual performance.

The company should maintain a bid ledger by country, sector and authority. It should record the legal route, required certifications, language, deployment boundary, data class, evaluation protocol, competitive set, price, implementation effort and outcome. Valuation should use observed conversion and collections, with option value for earlier-stage programmes shown separately.

Bid economics need their own forecast accuracy test. Management should compare the original probability, expected award date, implementation budget and collection schedule with actual outcomes for completed opportunities. Persistent delay or cost overrun should change the pipeline conversion and working-capital assumptions. Framework admission without call-off activity should remain a market-access indicator. It should not be presented as a customer contract or recurring revenue.

3 Build the compliance obligation map

The AI Act allocation of responsibility depends on role and use. A foundation-model developer can be a general-purpose model provider. A party that modifies or rebrands a model may become a provider under defined conditions. A downstream business can also carry system-level duties based on the intended use and risk classification. The diligence file should map each product and deployment to the relevant entity, role and obligation [7-12].

For a general-purpose model provider, the evidence room should include technical documentation, downstream information, copyright policy, training-content summary, model evaluation, cybersecurity and incident processes where applicable. The General-Purpose AI Code of Practice provides a voluntary route for demonstrating aspects of compliance, while the Commission guidelines describe its interpretation of scope [8-11]. Legal advice remains necessary for company-specific conclusions.

Compliance cost should be separated into build, recurring operation and event-driven remediation. Documentation, evaluation, rights management, monitoring, security, incident response, external assurance and regulatory engagement can consume engineering and leadership time. The model should state which costs are incremental, which replace existing quality controls and which scale with models, countries or deployments.

The obligation map should include evidence dependencies between provider and deployer. A downstream customer may need model limitations, evaluation, data and integration information to meet its own governance duties. A supplier that cannot provide usable information can lose the opportunity even where the model performs well. Contract drafting, product documentation and engineering records should therefore operate as one delivery system rather than as separate legal and technical workstreams.

4 Treat regulation readiness as an operating capability

Regulation readiness has value when evidence can be produced accurately, consistently and on time. A policy library alone does not demonstrate readiness. The company needs versioned model records, accountable owners, evaluation results, change approvals, downstream documentation, rights evidence, incident logs and a process for correcting deficiencies.

The operating design should connect product release to legal and risk gates. Every model version should have a defined intended scope, evaluation plan, data and compute record, security review, documentation pack and release decision. Material changes should trigger a controlled assessment of provider status and customer commitments. The same identifiers should appear in technical, contractual and financial records.

Valuation should focus on cycle time and loss avoidance. Evidence can shorten customer diligence, preserve a tender position and reduce emergency remediation. It can also delay releases or constrain product choices. The investor should compare actual release and procurement timelines, measure rework and test whether the system supports growth without a linear increase in specialist headcount.

Readiness should be tested under a live diligence request. The company should reproduce the documentation for a selected model and deployment within an agreed period and reconcile it to source records. Exceptions should be logged with ownership and remediation dates. This exercise measures operating capability, reveals manual dependencies and provides a more useful cost basis than a static inventory of policies or certificates.

5 Measure language coverage at workflow level

A list of supported languages provides limited valuation evidence. Language performance varies by task, dialect, domain, register, document type, safety requirement and retrieval corpus. A model can perform well on general translation and fail on legal reasoning, customer-service policy or technical instructions. The company should therefore evaluate the complete workflow in the language used by the customer.

The evaluation protocol should fix the task population, model version, prompt and tools, retrieval source, human-review policy and acceptance threshold. It should include linguistic quality, factuality, task completion, safety, latency and cost. Native-language reviewers should assess difficult and high-consequence cases. Public benchmarks can orient diligence; customer acceptance and production outcomes determine commercial value.

The financial model should separate language breadth from language depth. Breadth can increase eligible opportunities. Depth can improve conversion, retention and workflow ownership. Each language requires data, evaluation, product support, documentation and sales capability. The company should show which revenue cohorts finance those costs.

Commercial coverage also requires local product design. Dates, currencies, personal names, administrative terms, retrieval sources and escalation workflows can vary by country and institution. The company should quantify the work needed to maintain each market after launch. Language revenue should be evaluated after localisation, evaluation, support and sales cost, because nominal model coverage can create a low-margin obligation when customer-specific work remains extensive.

6 Reconstruct European compute economics

Compute economics should reconcile training, fine-tuning, evaluation and inference. The investor should inspect provider invoices, accelerator reservations, utilisation, energy and colocation terms, cloud credits, data movement, support and termination rights. A headline accelerator count does not establish usable capacity or economic control.

EuroHPC access can provide valuable resources and expert support to eligible users [13-18]. The diligence question is how the allocation affects cash, timing, intellectual property, publication commitments, location, availability and continuity. Free or subsidised access can reduce a project cost while remaining non-transferable and time limited. Commercial workloads may require a different route.

Unit economics should be measured per accepted task. Token, accelerator-hour and storage prices are inputs. The full measure includes retries, retrieval, tools, guardrails, human review, idle capacity and service-level performance. The company should demonstrate how optimisation, smaller models, routing and private deployment affect contribution and customer price.

Capacity should be reconciled to the service level promised to customers. The schedule should show reserved and available accelerators, maintenance, allocation priority, network and storage constraints, geographic limits and recovery capacity. A model that assumes every nominal accelerator hour is productive will overstate contribution. Forecast utilisation should be supported by workload telemetry and should include a downside case for delayed procurement and customer ramp.

7 Separate public support from recurring economics

European grants, compute allocations, co-investment and procurement can reduce early capital needs or accelerate capability. Their economic treatment depends on conditions, matching requirements, repayment or clawback, permitted use, intellectual-property terms, reporting and duration. The valuation should not capitalise headline programme amounts that the company has not secured and cannot deploy.

The funding ledger should distinguish awarded, contracted, claimed, received and spent amounts. It should identify restrictions and the operating milestone supported. A compute allocation belongs in the capacity schedule. A grant belongs in the cash-flow model under its recognition and collection terms. A procurement award belongs in revenue only when the contractual and delivery evidence supports it.

Public support can create strategic signalling and ecosystem access. These effects should remain in an option case until they produce a measurable operating result. The board should test whether the company has a viable cost and demand structure when temporary support ends.

Management should also disclose the counterfactual financing need. If a grant, allocation or anchor procurement is delayed, the company needs a plan for scope, timing and liquidity. This test prevents temporary support from masking an unfunded research or delivery programme. It also clarifies which public support changes enterprise economics and which item simply advances expenditure that the company would otherwise finance privately.

8 Diligence data rights and copyright execution

Training and fine-tuning depend on data access and permitted use. The AI Act requires a copyright policy and a public training-content summary for general-purpose model providers [7-12]. The company should be able to connect source categories, licences, rights reservations, collection methods, filtering and exclusions to a model version.

The evidence room should contain contracts, licences, provenance records, crawler policies, rights-reservation handling, deletion and correction processes, dataset versions and accountable approvals. Purchased or partner data requires confirmation that the relevant model-development and commercial uses are permitted. Public availability alone does not establish a right to train.

The investor should model remediation pathways. A dataset may be replaceable at manageable cost, or it may be central to a language or domain capability. The analysis should estimate retraining, evaluation, delay, customer and liability effects. Representations, indemnities, escrow and insurance can allocate part of the risk; they do not repair missing operational evidence.

Rights diligence should extend into the release pipeline. Data added after the principal training run, synthetic-data sources, feedback logs and customer fine-tunes can create separate restrictions and confidentiality duties. The company should be able to remove or isolate a contested source without losing control of the entire model-development record. That capability affects remediation cost, customer confidence and the reliability of warranties in a transaction.

9 Test private and sovereign deployment

Private deployment can support customers that require greater control over data, operations or continuity. The commercial product may include customer cloud, sovereign cloud, on-premise infrastructure, managed private capacity or a hybrid architecture. Each route changes margin, implementation effort, upgrade responsibility and support risk.

The diligence team should test portability rather than accept an architecture diagram. A representative workload should be deployed in the proposed environment with measured quality, latency, throughput,

security, monitoring, recovery and cost. Model weights, licences, container images, dependencies, encryption keys and support rights should be verified.

Private deployment can protect retention and increase contract value, while producing slower implementations and heavier working capital. The financial model should separate licence, implementation, managed service, capacity and support revenue. It should also reserve for parallel versions, customer-specific changes and long-term maintenance.

Deployment margin should be tested by customer configuration. A standard container on customer cloud can have very different economics from a dedicated on-premise environment with local integration and multi-year support. The company should maintain a configuration catalogue, minimum price, implementation capacity and upgrade policy. Bespoke variants should carry an explicit maintenance reserve and an accountable decision on whether the strategic relationship justifies the divergence.

10 Reconcile cybersecurity and operational resilience

A foundation model can expose model weights, training assets, customer data, interfaces and deployment infrastructure. The AI Act includes cybersecurity expectations for general-purpose models with systemic risk, and NIST provides risk-management and generative-AI guidance that can inform control design [24-27]. Customer requirements may add sector-specific resilience and incident obligations.

The company should maintain threat models, access controls, secure development practices, model and data inventories, dependency management, testing, monitoring, incident response and recovery evidence. Private deployments require a clear division of responsibility between company, cloud provider, integrator and customer. Contractual allocation should match operational capability.

Cybersecurity belongs in valuation through cost, sales access, interruption risk and contingent exposure. A certification can support diligence, although scope and operating evidence matter. The investor should inspect incidents and near misses, remediation time, insurance and customer notifications rather than rely on policy documents.

Resilience evidence should include exercises, not only design documents. Recovery tests, access reviews, dependency failures and incident simulations should record time to detection, containment and restoration. The company should show how a customer environment is protected when responsibility is shared. Repeated tests provide evidence for underwriting interruption risk and for deciding whether security capability improves procurement access or merely satisfies a minimum condition.

11 Value talent and research continuity

European model companies compete for research, engineering, product, security, legal and commercial talent. Key-person dependence can be acute when a small group controls training infrastructure, data pipelines, evaluation and release decisions. The company should map critical knowledge, employment and invention rights, retention, succession and access control.

Research productivity should be reconstructed from experiment records. The board needs planned and actual compute, data, elapsed time, capability, evaluation, release outcome, customer adoption and cash effect for each material programme. Publications and benchmarks are useful technical signals. They do not establish a repeatable commercial operating system.

Valuation should distinguish replacement cost, continuity risk and evidence of organisational learning. A company that can reproduce a model, correct failures and deliver successor releases may deserve more confidence than a company dependent on one founder or an undocumented pipeline. Retention packages should align with transferable capability and customer outcomes.

Compensation and retention should be modelled as recurring economics. A valuation that treats scarce technical talent as an acquired asset while omitting the cost of keeping and replacing that team will overstate free cash flow. The board should identify roles that can be distributed through process and tooling, roles requiring named specialists and the time needed to train successors. Employment and invention records should match the work actually performed.

12 Assess distribution and ecosystem leverage

European foundation models can reach customers through direct APIs, cloud marketplaces, integrators, sovereign-cloud operators, sector platforms and public frameworks. Distribution can lower acquisition cost and accelerate procurement. It can also place pricing, customer identity and service control with a partner.

The diligence file should identify the party controlling the contract, billing, data, routing, support, renewal and termination. Marketplace availability is not the same as recurring demand. The company should reconcile partner reports to customer use, invoices and cash and should measure contribution after revenue share, credits and support.

Ecosystem value should be tested through portability. The company should demonstrate that it can move or add a route without losing critical customers, capacity or product function. Exclusive or concentrated relationships require downside cases for repricing, reduced placement or termination.

Partner concentration should be measured across lead generation, contracted revenue, capacity, billing and support. A company can appear diversified by customer while depending on one route for most new business or settlement. The downside case should model reduced marketplace placement, changed commercial terms and a migration period. Direct customer relationships, reproducible deployments and alternative channels can reduce this risk when they have been tested.

13 Build a sovereignty evidence scorecard

A scorecard should convert narrative claims into decision evidence. Recommended dimensions are ownership and control, model and software rights, data provenance, compute and hosting, operational authority, security, documentation, language performance, procurement acceptance, portability and continuity. Each dimension should state the claim, evidence, exception, owner and review date.

Scores should not be added into a single synthetic sovereignty percentage. A critical failure can outweigh several strengths. Missing model rights or an unworkable exit route can impair the proposition even when hosting and legal domicile are strong. The investment committee should define mandatory gates for the target customer segments.

The scorecard should connect to the financial model. Procurement acceptance affects eligible pipeline and conversion. Compute control affects capacity and margin. Language evidence affects addressable cohorts. Compliance maturity affects cycle time and recurring cost. This preserves traceability between diligence and value.

The scorecard should be refreshed after every material change in ownership, model, data, deployment, customer use or law. Historical evidence can become stale quickly when a new model version uses different data or infrastructure. A dated exception register prevents an investment committee from relying on a control that applied only to a predecessor release. The review cadence should match release frequency and customer consequence.

14 Construct four European deployment cases

The API-led challenger case assumes broad developer and enterprise access through managed interfaces. Growth depends on model quality, developer adoption, reliable capacity and competitive pricing.

Sovereignty supports selected customers but does not dominate the proposition. Revenue can scale quickly while price competition and compute dependence remain material.

The regulated-sector specialist case assumes a narrower model and product set for finance, healthcare, government, industry or critical infrastructure. The company earns value through domain evidence, controls, private deployment and accountable support. Sales cycles are longer and documentation cost is higher, while retention and contract value can improve after acceptance.

The sovereign private-platform case assumes customers buy deployment control, portability and continuity. Implementation and support are material revenue lines. Margin depends on standardised deployment and clear responsibility boundaries. The multilingual public-service case assumes deep language and public-sector workflow coverage, framework access and measurable citizen or employee outcomes. These cases are analytical states rather than forecasts.

Scenario probabilities should not be fixed for the investment horizon. They should move after procurement awards, language evaluations, regulatory findings, capacity agreements, product releases and funding events. The board can use trigger-based transitions between cases. This approach makes the forecast responsive to evidence and prevents an optimistic central case from surviving after its key conditions have failed.

15 Build the integrated financial model

Revenue should begin with customer and procurement cohorts. The schedule should show eligible opportunity, validated pilot, signed order, accepted production, invoice and collection. Pricing should separate API usage, licence, implementation, capacity, support and outcome-linked elements. Credits, partner shares and service adjustments should reconcile to effective revenue.

Direct cost should follow the delivery architecture. It includes model access or owned compute, hosting, data, tools, implementation, evaluation, security, support and human review. Central expenditure should separate research, compliance, sales, administration and major training programmes. Public funding and compute support should appear under verified terms.

The board should receive revenue, contribution, compliance and language cost, research capital, cash before financing and minimum liquidity for every case. Correlated sensitivities should combine procurement conversion, price, model performance, compute cost, delivery delay and funding availability. Every management response should have an owner and evidence deadline.

Working capital can be decisive in public and regulated deployments. The company may incur compute, implementation, evaluation and security cost before acceptance and collection. The model should show milestone billing, acceptance dependencies, receivables, supplier terms and required liquidity by programme. A profitable contract can still create a financing gap when delivery precedes cash by several quarters.

16 Apply valuation methods to evidence states

Discounted cash flow can value retained customer cohorts and the capital required to serve them. The forecast should model procurement timing, renewal, delivery cost, compliance, compute and research directly. A higher discount rate cannot replace missing cash-flow mechanics.

Comparable-company analysis requires alignment of business model and capital intensity. Model developers, application companies, cloud providers and services businesses monetise different layers. Revenue multiples should be reconciled to contribution, growth, cash consumption, customer concentration, rights and replacement capital. Transaction evidence should be adjusted for date and strategic terms.

A sum-of-the-parts approach can separate contracted cohorts, reusable technology, deployment and services economics and strategic options. Option value should be assigned only where the company holds a defined right, has a funded route to evidence and can exercise before the opportunity expires. Probability weights should be explicit and reviewed after each evidence event.

The valuation date and currency assumptions should remain explicit. European companies can earn and spend across several currencies and can receive non-cash capacity or support. The model should translate cash flows consistently, identify financing and lease-like commitments and avoid adding public-programme headlines to operating value. Reconciliation between methods should explain which cash flows, assets and options each method captures.

17 Illustrative sovereignty and compliance economics

The hypothetical cases in this paper demonstrate how different business designs can affect value. The API challenger produces EUR 310 million of revenue, EUR 158 million of contribution and EUR 18 million of cash before financing. The regulated specialist produces EUR 260 million of revenue, EUR 126 million of contribution and negative EUR 38 million of cash before financing.

The sovereign private platform produces EUR 215 million of revenue, EUR 86 million of contribution and negative EUR 74 million of cash before financing. The multilingual public-service case produces EUR 145 million of revenue, EUR 32 million of contribution and negative EUR 112 million of cash before financing. Higher implementation, language, compliance and working-capital requirements drive the difference.

These figures are management assumptions for an illustration. They are not observations, forecasts or conclusions. A real valuation requires company contracts, customer acceptance, model and data rights, compute agreements, programme awards, invoices, collections and forecast history.

The illustration also shows why revenue scale alone is an incomplete decision measure. The multilingual public-service case has meaningful revenue and the highest cash need because capability, procurement and delivery costs arrive before collections. The API case has stronger near-term cash but greater exposure to price and model competition. A real committee should test the cases against the company's contracts, capacity and funding timetable.

18 Value procurement and language evidence separately

The first evidence state contains eligible procurement routes and accepted pilots. It has limited value because conversion and recurring economics remain uncertain. The second state contains repeatable regulated deployments, signed contracts and controlled documentation. The third adds scaled multilingual cohorts with measured contribution and renewal. The fourth adds durable platform distribution and strategic access.

The hypothetical probability-weighted example assigns enterprise values of EUR 380 million, EUR 980 million, EUR 2,000 million and EUR 3,400 million across these states. Applied probabilities of 25, 35, 25 and 15 per cent produce EUR 1,448 million. The example demonstrates method and is not a valuation conclusion.

The committee should prevent double counting. Procurement access, compliance readiness, language coverage and sovereignty can influence the same customer cash flow. A premium should arise once through the forecast or a clearly defined strategic option, with the supporting evidence recorded.

Probability weighting should be accompanied by a transition budget. The amount required to move from an accepted pilot to repeatable deployment may include product work, security, documentation, language evaluation, implementation staff and working capital. If that amount exceeds available funding, the higher

evidence state is not reachable under the stated plan. This constraint should reduce option value or prompt a financing condition.

19 Structure investment around evidence gates

Capital can be staged against model, compliance, customer and compute evidence. An initial tranche can fund a defined deployment and documentation programme. Later tranches can depend on accepted customer use, signed procurement access, demonstrated unit economics, secured capacity and a controlled release. The instrument must preserve sufficient runway to reach the gate.

Transaction consideration can include holdbacks, escrow, warranties, indemnities and contingent value linked to verified outcomes. The buyer should define who controls product, pricing, compute and compliance after closing so that the milestone remains achievable and measurable. Legal, accounting and tax advice is required for the chosen structure.

The funding plan should include downside reserves. A delay in procurement or capacity can extend cash burn. A regulatory or rights deficiency can require remediation. A capital structure that assumes every programme converts on schedule can destroy value even when the technology remains strong.

Milestones should be objective and resistant to management discretion. Accepted service, independently reproducible evaluation, signed minimum commitments and collected cash are stronger than product launch or pipeline announcements. The transaction documents should define data sources, measurement periods, dispute processes and treatment of customer delays. The board should retain enough capital to protect customers if a later tranche is not released.

20 Govern value after investment

The board should govern a compact set of linked ledgers. Customer evidence covers procurement stage, acceptance, invoice, renewal and cash. Product evidence covers model version, language and domain evaluation, incidents and release. Rights evidence covers data, software and intellectual property. Infrastructure evidence covers capacity, utilisation, cost and continuity. Compliance evidence covers documentation, obligations and remediation.

Monthly reporting should compare forecast and actual outcomes by cohort and programme. It should identify changes in EU guidance, customer requirements, model scope and deployment architecture. Model or system changes should trigger targeted evaluation and documentation rather than a complete uncontrolled rebuild.

Capital release should follow verified milestones. The board should record decisions, owners and dates and should retain the evidence underlying valuation updates. This creates an audit trail between strategy, technical delivery, regulatory readiness and collected cash.

The governance pack should remain decision oriented. Large inventories of metrics can obscure the few variables driving value. For each quarter, the board should identify the procurement cohorts, capability releases, control exceptions, capacity position, cash runway and financing decisions that change the valuation case. Supporting ledgers provide auditability. The principal report should explain the economic consequence and required action.

21 Decision tests for investors and boards

The first test is customer specificity. Which buyer requires the claimed sovereignty, language or compliance attribute, and what contract evidence shows willingness to pay? The second test is control. Which rights and operations remain with the company across model, data, compute, deployment and support?

The third test is proof. Can the company reproduce its technical, linguistic, regulatory and financial claims from versioned records? The fourth test is economics. What contribution remains after compute, language support, evaluation, compliance, implementation and channel cost? The fifth test is funding. Which programme, reserve and financing are required to reach the next evidence state?

The final test is portability. Can the customer and company continue operating if a cloud, compute, model, data or distribution dependency changes? A valuation case is robust when these questions resolve into contracts, measurements, controlled processes and cash rather than narrative labels.

A further test is strategic independence. The company should identify decisions that it can take without consent from a supplier, investor, programme authority or distribution partner. Consent rights may be commercially reasonable, but they affect control and speed. The investment case should record each material constraint, its duration and remedy, and whether it conflicts with the sovereignty proposition sold to customers.

Conclusion

European foundation-model valuation should translate sovereignty and regulation readiness into verifiable operating economics. The relevant evidence is customer procurement, accepted multilingual and domain performance, model and data rights, documented obligations, controlled deployment, compute access, contribution and collected cash.

The AI Act and related guidance create real documentation and governance requirements. European compute and adoption programmes can improve the opportunity set. Their value varies by company, customer and contractual route. The investor should separate current economics from conditional options and should cost the recurring capability needed to preserve access.

A disciplined process uses evidence gates. It values contracted cohorts, funds defined capability, protects against rights and delivery deficiencies and updates the case after observable events. This approach supports investment and strategic decisions without relying on an undefined sovereignty premium.

European sovereignty and deployment evidence register

The evidence register should include legal entities, ownership and control rights, model versions, data provenance, language and domain evaluations, customer procurement stages, contracts, invoices, collections, compute allocations and commitments, deployment boundaries, regulatory documentation, incidents, financing and forecast vintages. Every investment-committee claim should link to dated evidence and an accountable owner.

Frequently asked questions

What does sovereign AI mean in a valuation

It should be defined as specific ownership control rights and operating capabilities across model data compute deployment security support and continuity. The relevant value depends on the customer's requirement and verified willingness to pay.

Does EU domicile create a valuation premium

Domicile can support eligibility or customer preference. A premium requires evidence that it changes procurement conversion retention contribution or risk. Domicile should be assessed separately from ownership compute data and operational control.

How does the AI Act affect a foundation model provider

The applicable obligations depend on the company's role model and market activity. General-purpose model providers face documentation downstream information copyright and training-content-summary requirements, with additional duties for systemic-risk models. Company-specific legal advice is required.

How should multilingual capability be valued

Value should follow accepted customer workflows by language and domain. A supported-language list is insufficient. The diligence file should connect controlled evaluation to contracts contribution renewal and support cost.

Can EuroHPC access be treated as owned compute

The valuation should follow the awarded terms, duration permitted use continuity and transferability. An allocation can reduce project cost without creating owned or permanently available capacity.

What evidence supports regulation ready deployment

Strong evidence includes a role and obligation map versioned model records data and rights provenance evaluations release approvals downstream documentation security incident processes and customer acceptance.

What is the principal downside case

The most informative downside usually combines slower procurement conversion weaker language acceptance higher compliance and implementation cost constrained compute and delayed financing. Company contracts and operating evidence determine the relevant assumptions.

Which Matchpoint practice is relevant to this work

This research connects to Matchpoint Partners' strategy and execution valuation capital raising transaction diligence and M&A work for technology companies investors and strategic partners.

Appendix B. Procurement access and customer economics file

The file should trace each opportunity through eligibility, framework admission, diligence, pilot, production, invoice and cash. It should record language, domain, deployment, evaluation, price, direct cost, working capital and accountable owner.

Appendix C. Language compliance and deployment dossier

The dossier should contain model scope, language and domain evaluations, data and rights evidence, compute and hosting, regulatory documentation, security, release approvals, customer acceptance, cost, schedule and realised commercial outcomes. Forecast and actual evidence should remain separate.


Appendix D. Sovereign deployment scenario control sheet

Each case should state procurement conversion, price, volume, language coverage, deployment route, compliance cost, compute cost, utilisation, customer retention, working capital, financing and probability. Management should document correlations and operating responses.

Appendix E. Transaction evidence room


The evidence room should include corporate records, intellectual-property assignments, model and data documentation, customer and partner contracts, cloud agreements, financial reconciliations, security evidence, regulatory analysis, insurance, people arrangements, financing and the signed valuation model.

Figure 1 European foundation model value architecture




Proposed framework; every conclusion requires company-specific legal technical commercial and financial evidence.

Figure 2 Hypothetical annual revenue and contribution by European deployment case




Management assumptions in EUR millions; figures are not market observations forecasts or valuation conclusions.

Figure 3 Hypothetical capability investment and cash before financing



Management assumptions in EUR millions; negative values indicate funding need before any minimum liquidity reserve.


Figure 4 Hypothetical probability weighted value by evidence state



Management assumptions in EUR millions; the chart is not a valuation conclusion.

Figure 5 First hundred days regulation ready deployment sequence

Convert sovereignty claims into procurement evidence controlled delivery and cash



Proposed sequence; timing should follow transaction customer regulatory and financing constraints.

Table 1 European foundation model sovereignty perimeter

Component	Required evidence	Valuation question	Principal risk
Ownership and control	entities shareholders governance and reserved matters	who controls the asset and strategic decisions	control rights differ from domicile
Models and software	repositories hashes licences assignments and release records	what capability is owned and transferable	restricted or short lived rights
Data and evaluation	provenance licences language tests and customer acceptance	can capability be reproduced and defended	missing rights or unsupported performance
Compute and deployment	agreements allocations utilisation hosting keys and exit tests	can service continue under required controls	non-transferable capacity or concentration
Compliance system	role map documentation evaluations incidents and approvals	can the company evidence obligations at scale	recurring cost delay or remediation
Customer access	frameworks contracts invoices collections and renewals	which claims create paid demand	pipeline mistaken for cash economics

Proposed diligence structure; evidence requirements depend on the company customer and transaction.

Table 2 European deployment case design

Case	Customer access	Core differentiation	Cost pressure	Capital posture
API challenger	broad developer and enterprise routes	model quality speed and price	compute and competition	continuing research funding
Regulated specialist	sector buyers with formal diligence	evidence assurance and domain workflow	documentation and sales cycle	staged against accepted deployments
Sovereign private platform	customers requiring operational control	portability local operation and continuity	implementation and version support	capacity and working capital reserve

Case	Customer access	Core differentiation	Cost pressure	Capital posture
Multilingual public service	public frameworks and local services	language depth procurement fit and outcomes	evaluation support and collections	programme and milestone funding

Analytical states; these are not forecasts.

Table 3 Procurement and customer economics ledger

Field	Required evidence	Control	Decision use
Eligibility and route	framework rules tender documents entity and certifications	dated requirement map	accessible opportunity
Validation and acceptance	protocol outputs review and sign off	versioned model workload and language	conversion and utility
Contract and billing	signed order price term invoice and service level	entity and term reconciliation	revenue and renewal
Direct delivery cost	compute data implementation evaluation support and channel	invoice and allocation reconciliation	contribution and cash
Collection	receivable settlement and bank receipt	aged balance and exception owner	liquidity and customer quality

Proposed minimum evidence fields for opportunity-level analysis.

Table 4 Hypothetical annual operating cases

Case	Revenue	Contribution	Research language and compliance	Cash before financing
API challenger	310	158	118	18
Regulated specialist	260	126	132	-38
Sovereign private platform	215	86	126	-74
Multilingual public service	145	32	105	-112

Management assumptions in EUR millions; figures are not market observations forecasts or valuation conclusions.

Table 5 Hypothetical regulated specialist contribution bridge

Item	EUR millions	Evidence required
Revenue	260	customer contract acceptance invoice and cash ledger
Compute and model delivery	-72	workload telemetry provider invoice and capacity allocation
Language evaluation and compliance	-24	people supplier and programme records
Implementation channel and support	-38	project time partner settlement and service evidence
Contribution before central cost	126	reconciled customer cohort schedule

Management assumptions in EUR millions; excludes central research sales administration financing and tax.

Table 6 Hypothetical valuation evidence states

Evidence state	Enterprise value	Probability	Weighted value
Eligible routes and accepted pilots	380	25%	95
Repeatable regulated deployments	980	35%	343
Scaled multilingual cohorts	2000	25%	500
Platform distribution and strategic access	3400	15%	510
Total		100%	1448

Management assumptions in EUR millions; this is not a valuation conclusion.

Table 7 Investment and transaction evidence gates

Gate	Required evidence	Capital or consideration response	Post decision measure
Rights and control	assignments licences data rights and reproduced build	closing condition or protected tranche	exceptions and controlled release
Procurement access	eligibility accepted pilot signed order and collection	base value after verification	conversion contribution and cash
Language and domain quality	controlled evaluation and customer acceptance	milestone funding	accepted workflow and support cost
Compliance and security	role map documentation tests incidents and remediation	reserve holdback or staged release	cycle time exceptions and closure
Compute and portability	capacity terms utilisation cost and exit test	funding reserve or price adjustment	accepted task cost and continuity

Proposed framework; actual instruments require current legal tax accounting and financial advice.

References

- [1] IFRS Foundation. IFRS 13 Fair Value Measurement. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [2] IFRS Foundation. Educational material on measuring unquoted equity instruments within IFRS 13. <https://www.ifrs.org/content/dam/ifrs/supporting-implementation/ifrs-13/education-ifrs-13-eng.pdf>
- [3] International Valuation Standards Council. International Valuation Standards. <https://www.ivsc.org/standards/>
- [4] IFRS Foundation. IFRS 15 Revenue from Contracts with Customers. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/>
- [5] IFRS Foundation. IAS 38 Intangible Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
- [6] IFRS Foundation. IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
- [7] European Union. Regulation EU 2024 1689 Artificial Intelligence Act. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [8] European Commission. General purpose AI obligations under the AI Act. <https://digital-strategy.ec.europa.eu/en/factpages/general-purpose-ai-obligations-under-ai-act>

- [9] European Commission. Guidelines on obligations for general purpose AI providers. <https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers>
- [10] European Commission. General Purpose AI Code of Practice. <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>
- [11] European Commission. Guidelines for providers of general purpose AI models. <https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers>
- [12] European Commission. Template for providers to summarise training content. <https://digital-strategy.ec.europa.eu/en/faqs/template-general-purpose-ai-model-providers-summarise-their-training-content>
- [13] European Commission. AI Continent Action Plan. <https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>
- [14] European Commission. AI Continent Action Plan fact page. <https://digital-strategy.ec.europa.eu/en/factpages/ai-continent-action-plan>
- [15] EuroHPC Joint Undertaking. AI Factories Access Modes. https://www.eurohpc-ju.europa.eu/ai-factories/ai-factories-access-modes_en
- [16] EuroHPC Joint Undertaking. Large Scale Access to AI Factories. https://www.eurohpc-ju.europa.eu/large-scale-access-ai-factories_en
- [17] EuroHPC Joint Undertaking. AI Factories Access Calls Terms of Reference. https://www.eurohpc-ju.europa.eu/document/download/bd7aa666-bdf3-4436-b5ec-0b34a781e817_en
- [18] EuroHPC Joint Undertaking. AI Gigafactories Call. 30 July 2026. https://www.eurohpc-ju.europa.eu/eurohpc-joint-undertaking-launches-ai-gigafactories-call-2026-07-30_en
- [19] European Commission. Apply AI Strategy. <https://digital-strategy.ec.europa.eu/en/policies/apply-ai>
- [20] European Commission. European approach to artificial intelligence. <https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>
- [21] European Commission. Legal and procurement guidance for AI in the public sector. <https://interoperable-europe.ec.europa.eu/collection/ai-public-sector/legal-and-procurement-guidance>
- [22] European Commission. Procurement of AI Community. <https://public-buyers-community.ec.europa.eu/communities/procurement-ai>
- [23] European Commission. AI procurement clauses template for non high risk systems. https://public-buyers-community.ec.europa.eu/system/files/2023-10/AI_Procurement_Clauses_Template_NON_HIGH_RISK_EN.pdf
- [24] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework. <https://www.nist.gov/itl/ai-risk-management-framework>
- [25] National Institute of Standards and Technology. Generative AI Profile NIST AI 600-1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [26] National Institute of Standards and Technology. Secure Software Development Practices for Generative AI. <https://www.nist.gov/publications/secure-software-development-practices-generative-ai-and-dual-use-foundation-models-ssdf>
- [27] European Union Agency for Cybersecurity. Artificial Intelligence cybersecurity challenges. <https://www.enisa.europa.eu/publications/artificial-intelligence-cybersecurity-challenges>
- [28] European Data Protection Board. Opinion 28 2024 on aspects of data protection related to AI models. https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en
- [29] European Commission. Data Act. <https://digital-strategy.ec.europa.eu/en/policies/data-act>
- [30] European Commission. Data Governance Act. <https://digital-strategy.ec.europa.eu/en/policies/data-governance-act>
- [31] European Commission. European data spaces. <https://digital-strategy.ec.europa.eu/en/policies/data-spaces>

- [32] European Commission. Digital Operational Resilience Act. https://finance.ec.europa.eu/digital-finance/digital-operational-resilience-act_en
- [33] European Banking Authority. Report on the use of machine learning for internal ratings based models. https://www.eba.europa.eu/sites/default/files/document_library/Publications/Reports/2021/1023883/Report%20on%20machine%20learning%20for%20IRB%20models.pdf
- [34] European Central Bank. Guide on outsourcing cloud services to cloud service providers. <https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.cloudguide.en.pdf>
- [35] European Commission. Guidelines on transparency obligations for providers and deployers of AI systems. <https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems>
- [36] European Commission. European AI Office. <https://digital-strategy.ec.europa.eu/en/policies/ai-office>
- [37] European Commission. AI regulatory sandboxes. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-sandboxes>
- [38] European Commission. Testing and Experimentation Facilities. <https://digital-strategy.ec.europa.eu/en/policies/testing-and-experimentation-facilities>
- [39] OECD. OECD AI Principles. <https://oecd.ai/en/ai-principles>
- [40] OECD. OECD AI Incidents Monitor. <https://oecd.ai/en/incidents>
- [41] Mistral AI. Models and developer documentation. https://docs.mistral.ai/getting-started/models/models_overview/
- [42] Mistral AI. Deployment documentation. <https://docs.mistral.ai/deployment/self-deployment/overview/>
- [43] Mistral AI. La Plateforme pricing. <https://mistral.ai/products/la-plateforme>
- [44] Aleph Alpha. PhariaAI model and platform documentation. <https://docs.aleph-alpha.com/>
- [45] LightOn. Company filings and investor information. <https://investors.lighton.ai/>
- [46] Hugging Face. Open LLM Leaderboard. https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard
- [47] MLCommons. MLPerf Inference benchmarks. <https://mlcommons.org/benchmarks/inference-datacenter/>
- [48] Stanford Institute for Human Centered Artificial Intelligence. AI Index Report 2026. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
- [49] World Intellectual Property Organization. Artificial Intelligence and intellectual property. <https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/index>
- [50] FinOps Foundation. FinOps Framework 2025. <https://www.finops.org/wp-content/uploads/2025/05/English-FinOps-Framework-2025.pdf>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.