EuroStack

Shifting to European Digital Supply: A "How To" Approach

SIX KEY QUESTIONS: SCOPE, AI, THE BUSINESS CASE, CONTRACTS, RENT OR BUILD AND THE CONSORTIUM

Stefane Fermigier for EuroStack · 9 September 2026

Stefane Fermigier, Abilian, sf@abilian.com

EuroStack

Contents

Shifting to European Digital Supply: A "How To" Approach

This brief sets out an approach to six potential questions that companies and the investors behind them most often ask themselves when they weigh moving their digital supply to European providers:

1. **What do we put in scope?**
2. **How far does AI belong in it?**
3. **How do we construct the business case?**
4. **How do we structure the contract?**
5. **Do we rent or build?**
6. **How do we put together a consortium of buyers?**

They were first put to us by a fund weighing a move to European suppliers for a group of the companies it holds, and they now arrive with some regularity as the demand side searches for European alternatives to its current setup. The answers apply to any large European buyer, and to any fund, corporate investor or industrial group behind it. The brief also takes up EuroStack's recommendation that buyers form buy groups and procurement consortia, which create critical mass and improve bargaining power.

This is intended as one substantive contribution for discussion, nobody's official position. It written for a technically competent rather than a generalist audience. We are vendor-neutral and recommend no supplier.

Part 1 summarizes the answers for the decision layer. Part 2 works each answer through: operating teams and the companies in an investor's portfolio will find what to do there. It closes with a seventh answer, on what an investor should require of a company at the moment it funds it. That is the cheapest instrument an investor has, because it acts before any lock-in exists. Part 3 carries the checklists, the proposed studies, what the brief rests on and how far its key figures can be quoted.

One idea underlies all answers. The decision that determines your exposure to supplier power and foreign influences is not where the servers sit. It is who controls the software your estate runs on, from the operating systems up through the control plane to the interfaces your applications are written against, and what leaving would cost you. A data centre in Europe running a US provider's proprietary software at every one of those levels leaves your switching cost, and therefore the price you can be charged, exactly where it was. Every answer below follows from that.

Part 1. Executive summary

Why this is on the agenda. Large European companies, and the investors behind them, are asking what it would take to move their cloud and software supply to European providers, and whether it would pay. The commercial case stands on its own, which makes it defensible before a board that will not sign a political forecast. On a study for Cigref, the association of large French IT buyers, prices for the cloud and software European buyers cannot easily leave have risen 8.7% a year for three years. Behind the escalation is a mechanism: a supplier can raise its price up to what it would cost you to walk away, so you concede that premium every year whether or not you ever move. Your exposure, throughout this brief, is that premium plus the damage a supplier could do by cutting you off. Two things decide it: who controls the software your business runs on, and what leaving it would cost. Moving the servers into a European data centre while keeping the same provider's proprietary software on them changes neither.

What we recommend, and for whom. Answers 1 to 5 are for the buyer, its CIO and its CFO. Answer 6 is for an investor asked to anchor a buyers' consortium, and answer 7 for any investor with a portfolio. Start by conceding what is true: the technology of the large US cloud providers, the hyperscalers, is more often better integrated and better marketed than better or cheaper. The integration gap is real where the value is highest, in managed platform services and AI. The answer to the board's question, whether to leave the hyperscalers, is layer by layer. Take control of the control plane wherever it runs, move the commodity layers where European providers are cheaper, stop adding to the proprietary tier, and let the exit terms discipline what stays. Sort your systems by where a move removes the most exposure for the least effort, and take the decisions in the order the questions were asked, with a seventh for investors:

1. Move the control plane first, the software that runs everything else, because a large exposure has a small fix there; take the commodity layers as a quick win in parallel, and build nothing new on the services you cannot leave.
2. Put AI in scope at the layer where projects actually fail, and budget it as a running cost.
3. Build the business case on the bill you will receive over the years your exit cost keeps you in place.
4. Screen suppliers on who controls them, and promise them volume without tying yourself to their products.
5. Rent infrastructure from European providers by default, and build a data centre of your own only where you can keep it busy.
6. If you anchor a consortium, bring your holdings in as its first contracted members and lock the governance before the first euro goes in.
7. Write the covenants into the next shareholders' agreement, and put a yearly exit-cost disclosure on the board agenda of every company you hold.

1. What do we put in scope? Start with the control plane. The control plane is the software that runs your systems: who can log in and to what, where the encryption keys are kept, the monitoring, and the tooling that deploys everything else. It is the one place where a large exposure has a small fix. The exposure is large because all of it is configured in your provider's own formats and little tooling moves it. Leave in a hurry and you rebuild it under pressure with the business down, since when the login system stops, everything stops. Almost nobody leaves. Your supplier prices every renewal knowing that. The fix is small because a planned replacement rewrites no business logic. You rebuild, once and under your own control, the record of who your users are and what each may do, the custody of the keys, and the deployment and monitoring tooling. Then you re-point your applications at it. That re-pointing is configuration work, so the cost grows with the number of systems that authenticate against your provider and is independent of the size of your codebase. We cannot size it for you; it is the first thing to cost on your own systems. Go first here because every other system authenticates against it, so each later move is simpler once this one is done.

Take the cheap layer as a quick win, in parallel. Computing capacity, storage, databases and the standard plumbing around them are where European supply already exists. On published prices, European providers are 2.5 to 6 times cheaper for computing capacity, about 2.4 times cheaper for object storage and at parity on managed databases, and far cheaper for moving data out; your negotiated rates narrow the gap, so run the comparison on them. Run it as a normal sourcing exercise while the control-plane work goes on, budget the parallel running and migration effort like any other move, then take the saving; it funds the programme without being it.

What stays, and in what order. The rest of your systems fall into three kinds:

- A provider's most proprietary services, sold as serverless and managed AI, are woven into your application code; moving them means rewriting the applications, for which no European provider sells an integrated substitute. Do not plan to replace them. Stop adding to them and keep what you have confined to parts of the system you could rewrite later.
- The office and collaboration suite is the one large bill with proven European substitutes and a known sequence: back-end services first, then the client applications, and the workstation operating system last, the layer with the least mature substitutes and the longest tail of exceptions. Whether it is your largest software line is your own figure; in many industries it is.
- Line-of-business systems, ERP included, stay out of any joint purchasing with other companies and get the exit clauses of answer 4.

2. How far does AI belong in it? Fully, at the layer where projects actually fail. Your teams are held back by something other than chips, data centres or model quality: between a third and a half of the teams building with these models have no working system in daily use, whichever models they buy. The price of a given level of capability has fallen roughly fifty-fold in three years. Frontier models keep improving at about the same price

while open-weight models, many of them Chinese, sell last year's capability at a fraction of the cost; with more than one origin required and a tested ability to switch, that competition works for you. What is missing sits between a model and a working application: the software that connects a model to your systems, gives it memory and tools, enforces who may see what, and searches your own documents. Put that layer in scope and budget for it. Two budget lines follow:

- Budget inference as a running cost, metered per token or per hour of processor time, since it behaves nothing like a per-seat licence.
- Insist that whatever your AI systems accumulate, the indexes over your documents, the conversation history and the permission settings, can be exported. An assistant built into your incumbent's suite learns your business from the content stored there and is worth nothing the day that content moves.

3. How do we build the business case? On the bill you cannot walk away from. Separate two kinds of move. Where a European provider already sells the same service cheaper, the case is a price comparison with the migration effort added, which usually closes on its own. Where the move means rebuilding something, such as the control plane or a collaboration suite, it is an investment: the cost comes first and the benefit later. A migration buys a credible exit threat first, which disciplines the price of everything you keep buying; cost reduction, if it arrives, arrives late and as a consequence. The migration evidence in this brief is public-sector, because corporate migrations are rarely published. A three-year comparison is biased against any such investment, because the spending falls inside the window and most of the return outside it. Whether it pays within three years depends on what you are charged per seat today, so run it over ten years as well and show both columns. The usual comparison leaves out three things:

- Escalation. Apply your own rate, computed from your last three years of renewal quotes, to the whole horizon.
- The margin you concede. A supplier can charge above the competitive price only by up to what it would cost you to leave, so lowering your exit cost lowers what you pay on the contracts you keep. We cannot yet put a euro figure on that bound, so it is a direction and not yet a column.
- Insurance. A working alternative limits the loss when a supplier fails or cuts you off, and the case for it rests on no political forecast.

The calculation to run. Instead of forecasting savings, compute the price at which the move breaks even: what the incumbent would have to charge you, per seat, per hour of computing or per million tokens, for the move to pay over three, five and ten years. Put that figure next to your current invoice, which you already hold. The inputs you assume are your escalation rate and your own migration effort; Annex B's worked illustration is a workforce suite, so run your own for computing capacity. One more rule: a service that arrives at no extra charge inside a suite has no budget line and is never compared with anything, so charge your business units an internal price for it and put a periodic review on the calendar.

4. How do we structure the contracts? Screen suppliers on control, and commit volume, never architecture. Data residency, local staff, local incorporation, a local headquarters in the newer requirement texts, and generic certifications are often necessary for reasons of their own. A foreign-controlled supplier meets every one of them at the same cost as anyone else. Those texts stop short of asking who ultimately controls the company, the one question on which suppliers differ. In June 2025 Microsoft's legal director in France told a Senate inquiry, under oath, that he could not guarantee that data entrusted to Microsoft would never be handed to the US authorities without French agreement. Establish five things and keep everything else off the scoresheet:

- who ultimately owns the supplier;
- which law governs the contract;
- who holds the encryption keys;
- what the parent's home law lets a foreign authority compel;
- what the contract says happens to your service if it does.

Every extra clause weakens the ones that matter, since a supplier caught cheating loses the same contract whether it cheated on one clause or on five. A heavy compliance file also screens out the small European suppliers you were trying to reach. Award on price at the commodity layer, where you can leave. For a supplier you will rely on and have verified, a decent margin gives it something to lose if caught cheating, so awarding on lowest price there removes its stake in behaving, and you end up paying for the enforcement yourself.

Commit volume, never architecture. A European supplier needs a demand floor before it will invest. The usual way to give one is a committed-spend agreement, a promise to spend so much over so many years on the supplier's own catalogue, which is exactly how lock-in arrives. Commit instead to capacity, in seats, cores, terabytes or tokens, specified against interfaces that more than one supplier serves, such as standard SQL for databases, the S3 interface for storage, OpenID Connect for logins and OpenTelemetry for monitoring. The supplier gets a floor it can plan on and you keep the right to place the load elsewhere, so being wrong about your 2029 load costs you money instead of your freedom of action. The contract structure that makes such a floor bankable for a small supplier without re-creating the lock-in is the second study we propose in Annex D. Four terms carry the weight, with the full set in Annex C:

- exit assistance, priced and scoped in the contract;
- data export, free and tested;
- each bidder naming who could replace it;
- a tender written for the service you need and weighted on the supplier's willingness to build to it.

5. Do we rent or build? Rent from European providers, and build a data centre of your own only where you can keep it busy. This answer is about physical infrastructure, servers and data centres; it says nothing against building the applications that are your

business in-house. Buy cloud services from European providers, on exit terms your lawyers could actually enforce. Where a system needs hardware under your own custody, rent space in a European operator's data centre, or a private cloud it runs for you, before you build anything. Whatever you rent, keep under your own control the four things that decide who can switch: the login system, the encryption keys, the data, and the interfaces your applications are written against. And keep a tested reserve, a separately hosted fallback you could switch to, which is the cheapest form of ownership.

When building pays. A data centre of your own running the same proprietary software leaves the supplier's grip where it was. It also runs below the utilisation a provider reaches by pooling many customers, so you carry the idle hours. Build anyway in three cases:

- No supplier you could use can be verified to the standard a workload needs, so custody of the hardware is the only real protection.
- The workload's physical profile favours ownership: sustained high utilisation, data whose custody cannot be delegated, hard latency limits, or energy and heat you can supply below market.
- A consortium can contract enough pooled demand to keep the facility busy.

The build decision turns on a number no build-or-rent model we have found computes: the committed load, in megawatts, at which owning beats renting once pooled demand and the exit cost of the chosen design are counted. Compute it on the anchor investor's own load and energy data (answer 6) before committing capital.

6. How do we put together a consortium? Be the anchor yourself, and lock the governance before the first euro goes in. A buyers' consortium of this kind has seven roles: an anchor buyer, paying members, members who contribute engineering instead of fees, an operator, integrators and software suppliers, an infrastructure and energy provider, and financial investors. An investor with industrial and energy holdings can credibly fill three of them at once: the financial investor directly, the infrastructure and energy provider through its industrial holdings, and the anchor buyer through the holdings it brings in as the first contracted members. The anchor is the first large customer, whose committed demand makes joining rational for everyone else; clubs form when someone credible commits first and dissolve when it withdraws too early. A facility is financed only against contracted demand. A fund does not consume cloud and will not sign a capacity contract itself; its holdings consume it, so the contracts come from them and from the other members. The fund's part is to bring its holdings in, through its board seats in the mature ones and the influence a shareholder has over a young company still building its systems, and to put in the equity. In return it gets first-mover terms, seats in the governance, and a lower exit cost across its own holdings.

Five rules for the anchor. Competition law and state-aid rules for a club anchored by a public investor are outside this brief; take advice on them before anything is signed. Within those limits:

- Put equity into the club's own vehicle, against contracted demand, and use demand everywhere else. In the lab's model, buying every US-controlled supplier into European ownership would remove about as much exposure as procurement and portability rules do at no capital cost. So do not buy cloud and software suppliers: have your holdings commit demand and write the contracts, because equity there tends to buy an acquisition target.
- If the consortium takes a stake in a supplier anyway, require a governance lock: foundation ownership, founder super-voting shares, a golden share or state control. This rule binds an investor whose reason for the stake is a supplier it can keep relying on; for a purely financial investor a sale is an exit, and the rule does not apply. Without one, a supplier fed by the club's demand is a well-fed acquisition target. With one, it cannot be taken over against its will, while you can still sell your stake. Across 125 prominent European technology companies, those with a lock stayed under European control in 89% of resolved cases, against 55% without, an association and not an estimated effect.
- Give the related-party conflict its own rules. Holding three roles at once means the anchor buys from a facility it finances and powers, so what its holdings pay must be set at arm's length and approved by members without a stake.
- Apply the supplier screen of answer 4 to the consortium's own procurement, including members' own bids.
- Contract your holdings' commitments for longer than the formation phase, since an anchor that leaves early takes the club with it. Where the club finances a facility, those commitments are long by design; what stays open is the architecture, so leaving at the end of the term is cheap.

7. What should we require of the companies we fund? A covenant at the next investment, and a board item in the companies you already hold. The cheapest instrument an investor holds works through the companies it funds, at the moment of funding, because a company that has not yet built its systems has nothing to migrate. For a new investment, the brief comes down to seven covenants in the shareholders' agreement, which cost little at that stage. Cloud credits, the free capacity hyperscalers give start-ups, are how the dependence usually arrives; book them as a loan repaid in switching cost. The seven covenants:

1. The login system and the encryption keys under the company's own control.
2. Open interfaces named in the plan.
3. A stop-adding rule for the proprietary tier, with exceptions justified on the record.
4. A yearly disclosure to the board of what leaving each material supplier would cost.
5. Free and tested data export in every material contract.
6. Portable formats for what its AI systems accumulate.
7. An internal price on anything that arrives bundled.

The yearly disclosure does most of the work; it is also the one item a minority investor or a board member can ask for without writing any terms. For a company twenty years old, the same covenants are a migration programme. The return shows up in the earnings of the companies you hold, since a company that could walk away from its suppliers concedes them less margin, and later in their ability to buy together. Section 7 has the checklist.

Part 2. The answers in full

2.1 Scoping the estate, and structuring it across companies

Which of our infrastructure, platform and software services could and should move to sovereign supply? How do we structure that across several companies, and what do simplified technical requirements and a volume structure look like?

The answer in one line. Rank each service on two separate quantities: what leaving it would cost you today, which measures the exposure, and how deeply it is wired into your application code, which sets the price of removing it unless another supplier can serve the same interface unchanged. Attack where the first is high and the second is low. Starting where leaving is already cheap removes little, since a service you can walk away from is not one you are being overcharged for. The estate, here and throughout, is every system the company runs, from the infrastructure up to the applications.

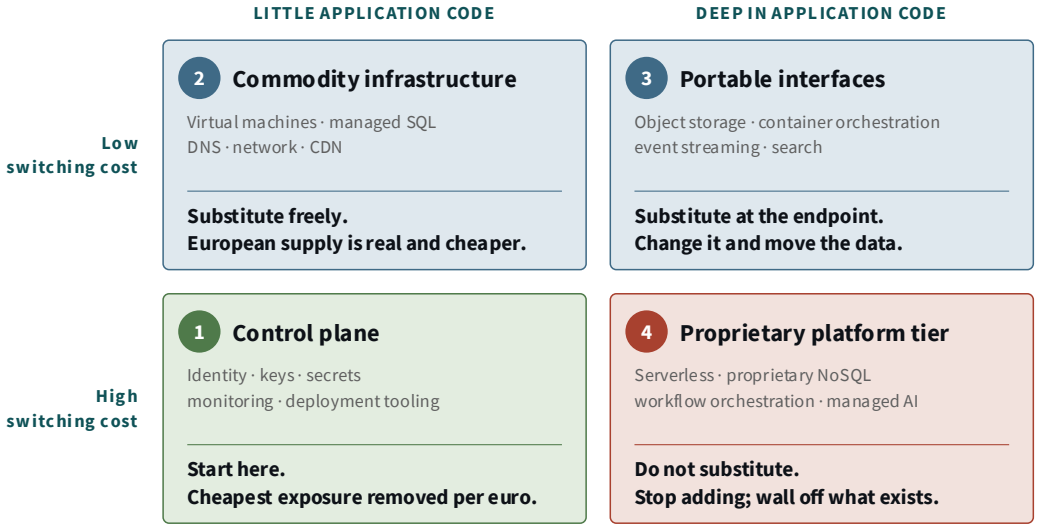

Figure 1. Where the exposure actually sits. Thirty-five managed services from Amazon Web Services' catalogue were scored on these two axes under a published rule; the examples shown are drawn from that scoring, with the action to take read off the cell. The numerals are the order of attack. Identity comes first inside the control-plane cell, because everything else authenticates against it. Portable means more than one supplier can serve the same client unchanged, which is not the same as an open standard; how that property arises, and what each way leaves exposed, is worked through in the text.

Why that corner first. Identity, key management, secrets, monitoring and deployment tooling are the *control plane*: the layer that administers the estate, as distinct from the interfaces the applications are written against. The ordering rests on a distinction between two costs that usually coincide. The first is the switching cost, what leaving would cost you today. It is high here because the access-control model, key custody and the operational

configuration are expressed in the provider's own terms, little tooling moves them, and an identity outage is a total outage. Almost nobody leaves, which lets the supplier price every renewal against that. That cost is the exposure, since it is the margin the supplier can extract and the loss you would take in a cutoff. The second is the price of removing the exposure, set by how deeply the service is wired into application code. The control plane is barely wired into it at all. Making it yours means standing up a directory, a role model, a key store, and a set of pipelines and dashboards under your own control, once, without touching a line of business logic. The work is disruptive, but bounded, and independent of the size of the codebase.

The intuition that a high switching cost makes moving hard is correct where the switching cost comes from code, which is the proprietary application tier: it carries the same switching cost and no bound on the price, because leaving means rewriting the application. At the commodity layer both the price and the exposure are low. The control plane is the corner where the exposure removed per euro of effort is largest. Identity comes first within it for an operational reason as well: every other service authenticates against it, so its outage is a total outage and everything downstream waits on its migration.

Where European supply is already real. European providers cover compute, block and object storage, relational databases, container orchestration and message streaming. They win on price (roughly 2.5 to 6 times cheaper than hyperscalers on like-for-like compute, and 10 to 86 times cheaper on outbound bandwidth, with several charging nothing). The interfaces are open standards, re-implemented de facto interfaces, or open-source software every provider runs, so moving means changing an endpoint and moving data.

One distinction the figure compresses. The portable-interface cell holds three different kinds of thing. Some of its interfaces are formal open standards with an independent specification body behind them: SQL, POSIX, DNS, OIDC and SAML, IMAP and SMTP and CalDAV. Others are one project's interface that several parties have since re-implemented: the S3 object API, the Kafka wire protocol, the OCI container specifications. The third kind is an open-source product with one implementation that many suppliers run; the container orchestrator, Kubernetes, is the largest example. Every production deployment of its API is the upstream code, repackaged by a distributor or hosted by a provider. Its code and its API are governed by a US-registered foundation, so a layer standardised on it has moved a dependency into that foundation's governance instead of removing it. The second and third kinds deliver what matters commercially, more than one supplier able to serve the same client unchanged, and neither delivers specification independence. They fail differently. A re-implemented interface belongs to whoever governs the originating project. If that project moves the interface, the other implementations are the buyer's recourse. A single-implementation product belongs to whoever governs the code. If its licence or its direction turns, the recourse is a fork, which is credible only where the ecosystem around the project is dense enough to carry one. Several of the bodies behind the formal standards are US-domiciled too, which is a far smaller exposure than a single vendor's API, though still an exposure. The ordering follows:

where a standard exists, explicitly or implicitly, in a protocol or a data format, an interface with several implementations is always preferable to a single product, however open the product is. A single-implementation open-source product is acceptable where no standard covers the need, as with most line-of-business applications, because the alternative there is a proprietary product with the same property and no fork option.

Stop adding to the proprietary tier. In the proprietary serverless and managed-AI tier, migration means rewriting business logic, with no European equivalent as an integrated offer. The correct objective is non-accumulation: stop putting new systems on it, and put clean internal boundaries around the existing estate so the dependency stays confined to replaceable modules instead of spreading through the codebase.

Productivity and collaboration. This is separately scoped and better understood, because several European public administrations have already done it at scale. France's La Suite numérique serves more than 500,000 civil servants monthly across fifteen ministries; Germany's openDesk, the federally governed suite, is in regional rollout (Fermigier 2026k). Schleswig-Holstein has published both a strategy and a feasibility study for around 25,000 workstations. Corporate migrations exist as well; they are rarely published, which is why the public cases carry the evidence here. Reference architectures and sizing heuristics exist for estates of roughly 250, 2,500 and 25,000 seats. The layer also decomposes into three steps that can be taken independently and each deliver value alone: the back-end services, the client applications, and the client operating system. Taking two or three of them at once multiplies the disruption and the risk.

The difference in kind between the third step and the first two decides the sequence. Where the back end has European and open substitutes at production maturity, the client operating system and the app-store distribution attached to it are the most concentrated position in the whole measured stack. In our programme's classification of Europe's building blocks, end-user computing comes out as captured, the programme's term for a block with effectively no European supply. Where a market share is published for that block, non-EU providers hold effectively all of it. This concerns the client estate; on servers Linux is the substrate almost everywhere and the problem does not arise. So the workstation goes last for a reason stronger than its long tail of exceptions: at that layer there is not yet a substitute to sequence toward, so a programme that opens there spends its credibility on the one step it cannot finish. Four things remain hard and should be sequenced around: AI assistance grounded on your own content and permissions, secure badge-release printing, enterprise telephony, and calendar availability across the boundary while both systems run in parallel.

Structuring it across companies. Sort the estate into three bins before anything else. Pool what is standardisable and shared: identity federation, object storage, container platform, managed databases, observability, backup and continuity, key custody, inference serving. Pool the requirements but keep separate contracts where members want the same guarantees at different volumes. Leave line-of-business applications and anything with its own regulatory perimeter out entirely. The consistent finding across the cooperative-

purchasing record is that buying groups form around standardisable requirements and stall on heterogeneous ones, and that governance disputes break them once they scale.

Simplified technical requirements. Specify interfaces, never products, and take a formal standard over a project's interface wherever one exists.

Formal standards, where they cover the need: put relational data behind standard SQL wire protocols; identity behind OIDC and SAML, with the directory of record and key custody inside your own boundary; mail, calendar and contacts behind IMAP, SMTP, CalDAV and CardDAV; printing behind IPP; and naming and routing behind DNS.

De facto interfaces, where no standard covers it: take the S3 API for object storage, the Kafka wire protocol for streaming, the OCI specifications for container images and runtimes, and OpenTelemetry for monitoring. For each of these, require in the tender that at least two independent implementations exist and name them, because that is the property doing the work; the interface's own governance stays an open exposure and should be recorded as one.

Open-source products with one implementation, where no standard covers the need: require at least two independent suppliers running the upstream project, no dependence on one distributor's extensions, and a licence under which a third party could continue the code. The project's governance stays an open exposure, recorded as one, with fork capability as the hedge against it.

One clause outweighs all of them: data export in documented formats, on demand, at no incremental charge, with a procedure that has actually been tested.

> **Where we are not yet firm.** There is no volume model for a multi-company consortium, and no service-by-service map of which offering has a production-ready European substitute today. Both are tractable and remain to be done.

2.2 How far AI belongs in scope

> *How far should AI be in scope, beyond traditional cloud services?*

The answer in one line. Fully, and aimed one layer below where the public debate is aiming, because the constraint on European AI deployment is not compute or model quality.

The constraint, on the evidence. In Mozilla's global developer survey published in July 2026, 63% of teams using closed models reached live production and 51% of teams using open models did. Between a third and a half stall short of production either way, including teams buying the most capable models with the full vendor bundle attached. The Commission's own June 2026 assessment of European AI capacity reports the same failure for Europe in qualitative terms: absent return on investment, failed technology transfer, experiments that do not become reliable production systems. It then directs its investment recommendations to compute, chips, data centres, energy and skills.

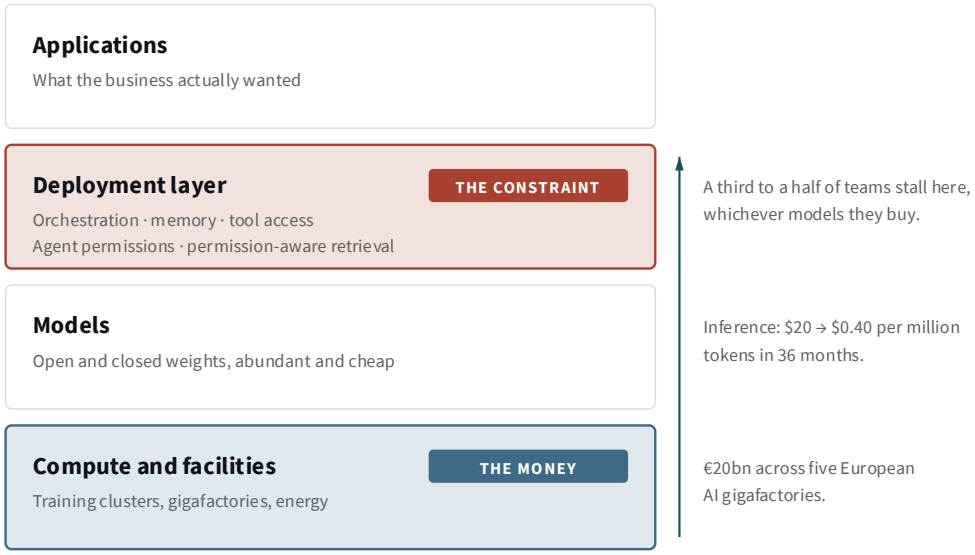

Figure 2. The AI money is aimed one layer too low. European investment is directed at the bottom band. A developer survey and the Commission's own qualitative assessment both put the binding constraint two bands above it, where no programme line exists. The bands are a reading of those two documents and not a measurement.

The funded inputs are now cheap. The price of a given level of model capability, yesterday's frontier, fell from roughly $20 to $0.40 per million tokens over thirty-six months, while the frontier itself keeps improving at about the same price; competition from open-weight models, many of them Chinese, drives the price of a given capability down. The best open-weight models sit about three percent behind the best closed models on preference ratings, at a fraction of the cost per task. A European buyer in 2026 does not fail to deploy because capable models are scarce.

The missing layer. It sits between a trained model and a working application: orchestration, memory, tool access, sandboxing, agent permissions, and retrieval over the organisation's own documents, mail and chat that enforces each user's permissions at query time. That layer has no line in the European investment programme and no cell in the value chain the Commission's study reproduces. It is also where commercial value is moving. Once models commoditise, pricing power returns to whoever controls the integration layer.

What this means for a consortium.

Budget inference and the deployment layer explicitly. Inference is a running cost that scales with use, priced by the token or the GPU-hour. It behaves nothing like a per-seat subscription; a plan that carries model licences and no inference line will be wrong by the amount that matters. The deployment layer above it is engineering, the shared item: every member needs the same orchestration, retrieval and permission machinery, which makes it the strongest candidate in the whole brief for pooled funding.

If AI capacity is offered to you as an investment, know what it will be used for. This is an investor's question. In our reading, the €20bn Europe has earmarked across five AI gigafactories is not sized to a training strategy, and if the facilities fill at all they will fill with inference and serving. That is the basis for pricing one: a facility whose demand is inference demand should be underwritten against inference contracts, and the demand-floor logic of section 5 applies to it unchanged.

Book AI assistants as switching costs. An assistant grounded on your incumbent suite's content and permission graph grows more valuable the more content sits there, and is worth nothing the day the content moves. That is a switching cost by construction and belongs on the ledger as one. Adoption is then a decision with a term attached.

Keep the assistant separable from the content platform. An assistant that reads your documents through a documented interface can be replaced. One welded into the platform cannot; it welds the platform in place too.

Second-source while it is cheap. Introductory API pricing is expected to expire around 2027 to 2028. Open-weight supply is also increasingly Chinese and a revisable policy choice by foreign labs. Require more than one viable origin and a demonstrated ability to switch between them; portability terms and a working second source cost far less to establish before that than after.

> **Where we are not yet firm.** We have the strategic argument and not the scoping. What a corporate consortium concretely needs at the deployment layer, and what it costs, remains to be scoped.

2.3 Making the business case

> *How do we make the case in practice, when hyperscaler technology is often better integrated and, at least in the short term, cheaper? What else belongs in the case, and what does a best-practice business case look like?*

The answer in one line. Build it against the bill you will actually receive over the period in which you cannot leave, and separate the parts of the estate where the saving is immediate from the parts where the move is an investment.

The estate splits into two cases. At the commodity layer the comparison is a price comparison and the European offer wins it today, so the business case there is ordinary procurement and does not need this section. Everything below is about the other case: a migration, where the benefit is structural and the cost arrives first (i.e. an investment).

Start by conceding what is true. Hyperscaler technology is more often better integrated and better marketed than better or cheaper, and the integration is real. The gap is largest exactly where the value is: managed platform services and AI. European components exist at close to competitive prices; but no supplier assembles them into a single integrated, ecosystem-backed offer. A business case that denies this is easy to dismantle in its first review. The precise nature of the gap matters, because an assembly problem admits

instruments that an absence does not: pooled demand, interface terms, and integration work the buyers commission themselves.

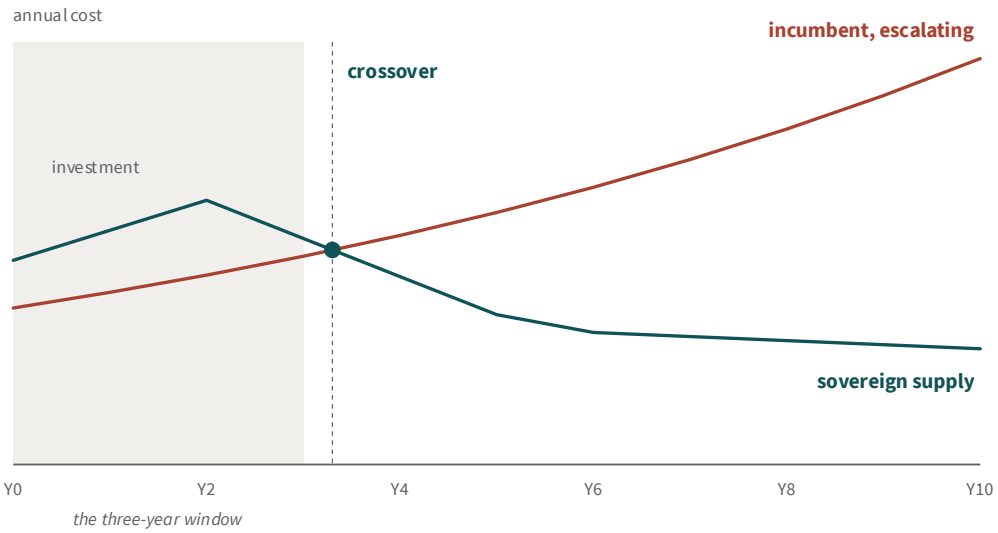

Figure 3. Why the three-year comparison can pick the wrong answer. The curves show shapes and carry no costing. The investment phase falls inside a three-year evaluation window, so a test on that horizon can favour the incumbent even when the ten-year answer is the opposite.

Then correct the comparison. It needs three additions.

Escalation exposure. Prices on Europe's locked cloud and software base rose 8.7% a year over the last three years, on an Asteres study commissioned by Cigref, the French association of large corporate IT buyers. One published public-sector series, Schleswig-Holstein's open-source strategy, records spending on a single supplier rising more than fourfold over a decade. The UK competition authority found the two largest cloud providers earning returns sustained above their cost of capital for years, while fewer than 1% of customers switch provider in any given year. Use your own escalation rate, computed from three years of your own renewal quotations; it is a better input than any published figure and you already have it.

The rent ceiling. A supplier can charge you above competitive cost only as far as leaving would cost you. Switching cost therefore belongs on the finance side, as the direct determinant of the margin you will concede over the life of the relationship. Two documented anchors give a sense of scale: one firm's compute exit cost about $700,000 of hardware against a $3.2m annual cloud bill, with the residual cost dominated by engineering work on interfaces; another reported $74.6m of infrastructure savings over two years after moving about 90% of its data off a hyperscaler. Both figures are US dollars, near enough to euros at recent rates for the order of magnitude to carry.

Insurance value, and this is the column that does not require you to forecast politics. The case need not rest on the probability of any particular legal or political event, which is fortunate, because that probability is not estimable. On the lab's disruption model, a hyperscaler cutoff is the worst case, and half the serious cases are not American. Holding alternatives pays in both states of the world: less margin conceded while nothing happens, a smaller loss when something does. That makes the column defensible in front of a board that will not sign up to a geopolitical forecast.

The practical device. Instead of forecasting savings you cannot defend, invert the model and derive the incumbent cost at which the move breaks even over a given horizon. Then compare against a figure you already possess.

The unit is whatever the workload is billed in. Getting the unit wrong sends the whole exercise astray. A workforce estate (mail, files, collaboration, workstations) is per seat per year. A customer-facing platform has no meaningful seat count and should be run per core-hour or per virtual machine, per terabyte-month of stored data, per million requests, or per million tokens for an inference workload. Some estates need two or three of these run separately, because a seat-denominated comparison on a workload whose cost scales with traffic will simply not correspond to the invoice. The worked illustration in Annex B is the workforce case, a back-end collaboration stack per seat per year. That illustration yields two readings that count for more than its numbers. Migration cost is close to independent of estate size, so the fixed cost of doing it at all dominates and smaller members face the hardest arithmetic. The three-year comparison is biased toward the incumbent, because the investment phase falls inside that window and most of the return falls outside it. Whether the move pays within three years depends on where the incumbent's price sits against the three-year break-even. Present three-year and ten-year columns side by side and explain the difference.

The decision that never happens. A service that arrives inside a bundle at no incremental price has no budget line, so no renewal event fires for it, so it is never compared with anything. A challenger's price never enters a decision, because the decision is never convened. The remedy is nearly free and internal: charge business units the imputed price of bundled components, and put a periodic evaluation duty on a clock. The same instrument exists in the regulatory record. The Commission's 2025 Teams commitments require the suites to be sold without the component at a lower price, with the corresponding offer shown wherever the bundled one is advertised. That is itemisation and disclosure, imposed from outside on a seller who had no reason to volunteer them. Without a decision event, an alternative-supplier policy is inert however good the alternative is.

What you are actually buying. You are buying three things. A credible exit threat, which disciplines the price of everything you continue to buy, including from the incumbent. Knowledge of your own dependencies, otherwise acquired during an incident. The ability to fix things, or to pay someone to. For the migrations, cost reduction, if it arrives, arrives late and as a consequence of the first.

One structural rule for the scoring. Express hard requirements as gates applied before any weighted scoring, never as heavy weights inside it. A weighted sum will faithfully produce a defensible-looking wrong answer whenever a requirement has been encoded as a preference.

2.4 Contract design and the procurement process

> *What is best practice in contract design and procurement? How do we avoid lock-in while giving suppliers enough commitment, and how do we handle deep uncertainty about future loads and usage, given AI?*

The answer in one line. Screen on control, keep the requirement list short, and denominate your commitment in standards instead of in a supplier's catalogue.

Screen on the facts that differ. Verify ultimate ownership and the control chain, including any parent that a locally registered subsidiary conceals. Verify the law governing the contract, the reach of foreign legal instruments, key custody and administration, and continuity of supply under foreign compulsion. These are public-record or attestable facts, expensive to falsify and detectable when falsified.

Prerequisites, kept in their place. Data residency, local staffing, local incorporation and generic operational certifications are often necessary, sometimes legally required, doing real work on latency, data-protection compliance and operational responsiveness. They cannot separate suppliers, which is a distinct job. A foreign-controlled supplier satisfies every one of them in substance and at the same cost as a domestic one, so they discriminate on nothing that bears on control. Newer requirement texts add a local headquarters to the list, which closes the shell-subsidiary case and leaves ultimate control exactly where it was. Require them where they are needed for their own sake; do not weight them in the award, and do not let them stand in for a control test. Two live examples. Since February 2025 Microsoft claims to store and process its European customers' data, support data included, inside the EU, so on its own account the residency requirement is met in full (Microsoft 2025). In June of the same year, a Senate commission of inquiry questioned the legal director of its French subsidiary under oath. Would data entrusted to Microsoft ever be handed to the US authorities, after a US government injunction, without the explicit agreement of the French authorities? He answered that he could not guarantee it (Sénat 2025). Amazon's European Sovereign Cloud is locally incorporated and staffed under a foreign ultimate parent. Both are exactly the configuration a location-and-practice test cannot exclude.

Keep the list short. Enforcement cost rises faster than the number of clauses, because a supplier caught cheating on any clause forfeits the same thing whether it cheated on one or on five. Every additional requirement therefore dilutes the credibility of all the others unless the audit budget rises more than proportionately. A second effect runs the same way. A long referential imposes fixed compliance overhead that a large incumbent's compliance function absorbs easily and a small genuine supplier cannot, so a heavy scheme sorts suppliers by size and works against the objective it was written for.

Do not award on lowest price where you rely on the supplier. A supplier is deterred from cheating by what it stands to lose if caught, so a margin above the participation floor buys enforcement you would otherwise have to audit for. The same premium is also a larger prize for a supplier whose sovereignty claim is cosmetic. The rule that follows: pay a premium to discipline a supplier you have verified, and rely on exit rights against suppliers you have not. At the commodity layer, where you can leave, price is the right criterion.

Commitment without lock-in. A subscale European supplier needs a demand floor to justify investment, and giving one normally means accepting lock-in. It does not have to. Lock-in comes from the vendor-specific catalogue a commitment is denominated in, which is why hyperscaler committed-spend agreements shape workloads: the commitment is drawn down against proprietary services. Denominate the consortium's commitment in capacity measured against a named interface instead: seats, cores, terabytes or tokens, each specified against the interfaces of section 1. You can then give a supplier a volume floor it can plan on while retaining the ability to place it elsewhere. The take-or-pay structure that makes such a floor bankable for a subscale supplier without re-creating the lock-in is Annex D item 2.

The clauses that carry the weight are in Annex C. Four of them matter most here. Exit assistance comes as a priced deliverable with a defined scope, invoked at your discretion, because without it the exit path is theoretical. Data export costs nothing incremental, and arrives with a tested procedure and a scheduled retest. Named second sources are elicited by asking each supplier which organisations could take over if you ended the relationship in three years and what you would hand them; a good answer names competitors. A supplier that will not answer has told you something. The tender goes out for the service while the product follows, with award weight on the supplier's willingness to take your technical requirements into its roadmap.

For regulated members, already law. DORA Article 28 requires financial entities to hold documented and tested exit strategies for their ICT third-party arrangements, irrespective of who owns the provider. That is the clause set above and the tested reserve of section 5, arrived at by a regulator from risk supervision and by this brief from the economics. Any consortium member inside the DORA perimeter already owes these artefacts, which makes them the cheapest place to start and the natural template for members outside it.

Load and usage uncertainty. Four instruments pay, in this order.

1. Portability turns a wrong volume forecast into a pricing error. If you can move in practice, being wrong about 2029 costs money and not optionality. This is the main answer, and the reason the interface requirements matter more than the volume negotiation.
2. Short commitment terms on capacity, long terms on interfaces. The exception is a facility the club itself finances, where the offtake commitments are long by design and the interface terms carry the exit.

3. Split the commitment by predictability: baseload on committed capacity at the best price, the uncertain tail on-demand at a second source. Standard energy-procurement practice, and it transfers directly.
4. For AI, require portable formats for what accumulates. The asset that locks you in is the accumulated retrieval indices, evaluation sets, conversational memory and permission configuration. Require portability at award, while introductory pricing still makes second-sourcing cheap.

If the consortium scores suppliers, use a two-stage rule: a weakest-link threshold gating eligibility, because a supplier is no more sovereign than its worst dimension and an average hides the single fatal one, then a weighted score to rank the survivors. Hold unknowns distinct from zeros, publish coverage alongside the score, and freeze the rule before scoring anybody. This is the structure the European Commission's own cloud sovereignty framework uses.

A score cannot certify a stack. Scores attach to nodes, and a stack is a path through many of them. Every supplier in a chain can score well while the chain still has exactly one way of operating, because the alternatives at each layer turn out to be controlled by the same party once you follow them down. Score suppliers to choose between them, and assess the stack separately, on whether its layers have alternatives that are independent of each other.

2.5 European provider or own data centre?

> *How do the two models compare: our data on a European cloud provider's servers, or a data centre of our own, built with sovereign cloud providers through a consortium of companies? When is each the more efficient, and what is critical mass for a facility of our own?*

The answer in one line. Buy from European cloud providers by default, under enforceable exit terms; co-locate with a European operator where a workload's physical profile requires it; own the layers that decide control plus a tested reserve; and build physical capacity of your own where its utilisation justifies it, which for most buyers means at consortium scale.

That answer is firm on its first three clauses and conditional on the fourth, because the quantity a build decision turns on has not, to our knowledge, been computed: the committed load in megawatts at which owning beats renting. The rest of this section sets out what does decide the question, and where the missing number would bite.

Two facts reframe the question. First, owning the metal leaves most of the exposure in place. A facility running a foreign provider's proprietary software, from the operating systems and the control plane up to the interfaces your applications are written against, leaves your switching cost, and therefore the price you can be charged, entirely intact. The metal is yours and everything that operates it is not. Second, ownership transfer buys less than the demand-side instruments. On the lab's model of the European technology stack,

moving every US-controlled layer into European ownership reduces exposure by about as much as procurement, portability and open substrates do without transferring a single share. The residual that neither removes is Asian and commercial, in fabrication, memory and raw materials, which is also the reason not to frame the exposure as a purely American and purely legal one. The ordering holds under the model's reweighting; the levels behind it are stylised and are not quoted here.

The decision rule. Ask what specifically you are trying to be independent of, and by when.

Objective	Where it points	Why
Price escalation and supplier lock-in	Any hosting model works. Decide on cost and capability, and put the exit terms in the contract.	Portable applications and documented data formats restore the exit threat at the application layer whatever runs beneath.
Jurisdictional exposure for specific data	European cloud first; then co-location or managed private cloud from a European operator; an own facility only at consortium scale.	The infrastructure operator retains the legal relationship, the operational access and the commercial leverage, and each rung up the ladder keeps more of them with you.
Continuity under disruption	A separate, independently hosted, tested reserve capability, which can sit somewhere different from the main estate.	Loss of service announces itself, and a working alternative insures it.

The middle row almost always applies to a subset of the estate, which permits the split: classify the data first, then host each class where its own requirements point.

The default, and the ladder above it. Buy cloud services from European providers. Where a workload's physical profile requires custody of the metal, take co-location or managed private cloud from a European data-centre operator, which keeps the hardware under your control without the build. Build it at consortium scale, or in the two narrower cases below. The reason is the utilisation criterion below. A company that builds its own capacity runs it below the utilisation a provider reaches by pooling many customers. Where no European offer exists, at the managed-platform and AI layers, co-location does not close the gap, because running your own platform software on rented metal is the self-build this rule avoids. The answer there is the non-accumulation rule of section 1 and pooled demand to make an offer emerge.

Build in one of three cases. In the first, the exposure is dominated by silent access risk and the supplier pool cannot be verified to the required standard, so portability is a consolation. The real choice is then metal under your own custody, co-located or owned, or not placing the workload at all. In the second, the physical profile favours ownership: sustained high utilisation, data whose custody cannot be delegated, binding latency

constraints, or energy and heat integration a member can supply below market. In the third, the consortium can contract a demand floor that beats the utilisation break-even.

Outside the custody case, the real criterion is utilisation. Sovereignty does not enter it. Owning beats renting when you can keep the asset busy. The consortium's economic contribution to a build case is that pooling several companies' imperfectly correlated demand raises achievable utilisation and reduces the variance the facility must be sized against. A provider pools more demand than any consortium, which is the system-level case for buying ahead of building.

Critical mass. For a buyers' club this is exact and computable (Fermigier 2026h): participation is self-reinforcing, so a threshold exists below which membership collapses to zero; an anchor above a computable size removes the collapse case. The record is consistent. The European public-sector clubs still operating hold large fractions of their eligible membership: roughly 80 of Denmark's 98 municipalities in one case, more than 140 Walloon municipalities in another. Those that dissolved did so when public support lapsed before self-sustaining mass had formed. The lesson for an anchor: commit for longer than the formation phase and say so publicly.

For a facility there is no answer yet. Critical mass there is a committed-load figure in megawatts against a utilisation floor. Build-or-rent models for data centres exist and are a developed literature. We have not found, and have not built, one that carries the two terms this decision turns on: the demand-pooling effect across consortium members, and the switching cost the architecture concedes. Absent those, any megawatt threshold quoted today is a rule of thumb and should be labelled as one. The gap has a narrow reach. It does not affect the buy-from-providers default, which turns on exit terms and on where control sits. Nor does it affect the decision rule above, which sorts workloads before any sizing question arises. It affects one thing: how much committed load a build needs before it is the cheaper option. That is item 1 of Annex D. It is well defined; an anchor's own load and energy data would turn it into a real number instead of a range.

The shape we would recommend. Own the control plane and the data: identity, keys, the directory of record, and the interface contract your systems are written against. These are cheap relative to metal and are where the exposure lives. Buy capacity where it is cheapest, on the interface terms of section 1, from more than one supplier. Where a workload needs the metal under your custody, co-locate with a European operator before building. Own a tested reserve, the cheapest form of ownership and the one delivering most of the reversibility a full build would. Consider owning metal where the conditions above are met and where the anchor's own energy and development capabilities put the input cost below market, which is an unusual advantage where it exists.

2.6 Consortium roles, value creation and investment partners

Assuming a data centre, who are the consortium's players, what are their roles and value-creation paths, and who could the investment partners be? What

changes when the anchor investor is also a project developer or an energy provider?

The answer in one line. The structure has seven roles, of which an investor that also develops infrastructure and supplies energy can credibly hold three; the role that decides whether the consortium forms at all is the anchor. The roles are drawn for the build the question assumes; under the default of section 5 the operator is a European provider running its own facility, and the roles are otherwise unchanged.

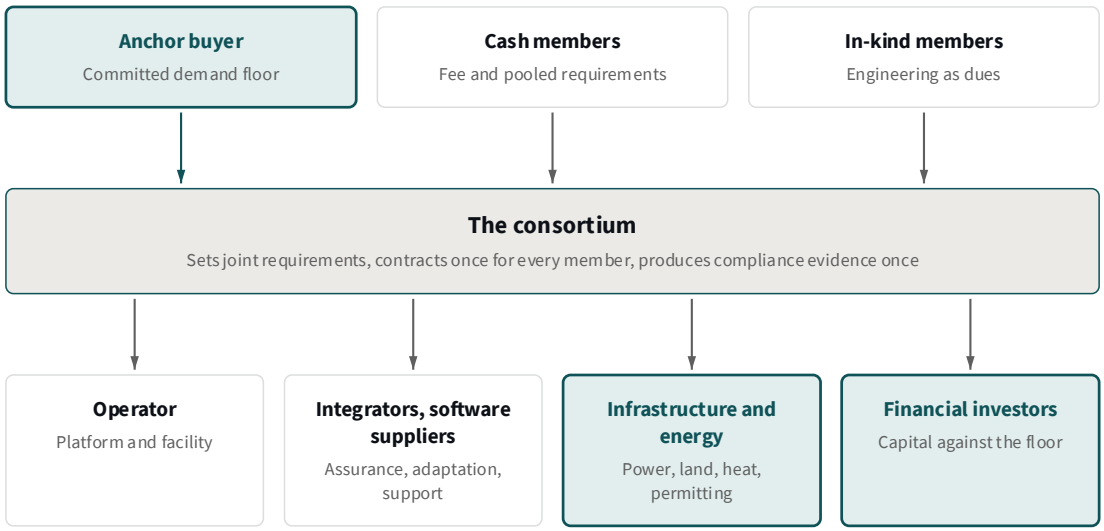

Figure 4. Seven roles, three of which one investor can hold at once. The shaded roles are the ones an investor with industrial and energy holdings can credibly take. An investor whose holdings are the anchor buyers underwrites the demand floor through them, which makes a facility bankable.

Role	Contributes	Captures
Anchor buyer	The committed demand floor that makes participation rational for everyone else	First-mover terms, governance seats, exposure reduction on its own estate
Cash members	Fee, pooled requirements, volume commitment	Cost below solo provision, a supplier they can replace, compliance evidence produced once for the whole membership
In-kind members	Engineering capacity delivered in place of fees	Lower net cost, influence over the roadmap, visible technical standing
Operator	Platform and facility operation	Utilisation-driven margin, with the club's committed volume as the bankable asset

Role	Contributes	Captures
Integrators and software suppliers	Assurance, adaptation and support around the shared platform	Recurring service revenue; the scarce asset is regional integration expertise
Infrastructure and energy	Power, land, heat offtake, development, permitting	The return on the scarcest input in any AI-era facility build, and the one least exposed to foreign supply
Financial investors	Capital against the committed-demand floor	Infrastructure-like returns, conditional on that demand being contractual and never merely intentional

Why that position is unusual. Holding the anchor, the energy and development, and the investor roles at once resolves the two problems that normally kill these vehicles at formation. It also creates a related-party conflict of its own, an anchor buying from a facility it finances and powers, which needs governance rules before competition law is even reached. The anchor problem is that few firms will commit before others have. The record shows this is decisive: clubs form when someone credible commits ahead of the others and dissolve when that commitment is withdrawn too early. The bankability problem is that an infrastructure investor needs contracted demand, which a consortium of merely intending members cannot supply. An investor whose holdings become the anchor buyers can underwrite the demand floor through them. The fund signs no capacity contract itself; its holdings do. The fund brings them in through its board seats in mature holdings and through the influence a shareholder has over a young firm still building its systems. The condition attached is that their commitments must be contractual and must outlast the formation phase, because a commitment of that kind recruits the members waiting to see whether anyone else moves.

Which holdings fill which role. The three-role argument above distinguishes such a consortium from a generic one. It becomes concrete only when named for a particular investor. Which industrial holding supplies power, land and heat offtake, which supplies development and permitting capability, and which would be an early cash or in-kind member of the club instead of a supplier to it. That mapping converts the argument from a structural claim into the investor's own balance sheet. It is a short conversation and it should come before any of the studies in Annex D.

Two proven operating models for the club itself run at scale in European public administration, and either transfers to a corporate consortium. In the first, the club procures from a vendor ecosystem. One Danish municipal association counts roughly 80 of the country's 98 municipalities and buys development and maintenance from more than sixty vendors under a single standardised framework agreement. That agreement fixes terms once instead of once per member-vendor pair, and collapses a coordination cost that otherwise scales with the members multiplied by the vendors. In the second, the club owns its supplier. A Walloon equivalent is co-owned by more than 140 municipalities and

maintains their common components against their own requirements, so delivery cost is internalised and the fee becomes an owner's budget contribution. Under the default of section 5 the operator in either model is a European provider running its own facility, the consortium's contribution is the contracted demand floor, and a consortium-scale build is the exception that Annex D item 1 exists to test. The choice between them, the make-or-buy decision for the service layer, deserves an answer.

Three ways value is destroyed.

1. **Investing without a governance lock.** Across 125 prominent European technology companies, those with a governance lock (foundation or steward ownership, founder super-voting shares, a golden share, or state control) remained under European control in 89% of resolved cases against 55% without; none of the locked companies was acquired against its will. For a purely financial investor a sale is an exit, and this argument does not bind it. For an investor whose reason to build the consortium is a supplier it can keep relying on, a sale to a foreign acquirer is the loss of that reason at whatever price, and the lock protects the objective the capital was committed for. Make the lock a condition of investment; it costs little relative to the capital.
2. **The consortium's own procurement**, including from members' own bids. The screen in Annex C applies internally.
3. **The anchor withdrawing too early.** Contract the commitment for longer than the formation phase. Where the club finances a facility, the offtake commitments are long by design and the interface terms of section 1 carry the exit.

Investment partners. Beyond the anchor investor, four classes come naturally. Other national investment and pension vehicles carry the same exposure and the same tolerance for long payback. Utilities and infrastructure funds know the facility as a familiar asset once demand is contracted. The corporate members' own commitments convert into the equity story. European public co-funding is available too. If it is taken, report its private component separately, because a leverage ratio under a fixed co-funding rate measures the funding rule and nothing else (Fermigier 2026g).

The fund's own capital. Two rules follow for a fund's own allocation, and neither depends on the consortium forming. First, value a holding's exit cost as a margin item and therefore a valuation item. A supplier can charge a holding above competitive cost only as far as leaving would cost that holding, so exit cost reduction across a portfolio shows up in earnings and belongs in the valuation. Second, put equity into the club's own vehicle, against contracted demand, and buy as a customer everywhere else: at layers where alternatives already exist, procurement rules and contract terms reach further per euro, while capital tends to buy an acquisition target. The governance lock is the first of the three failure modes above.

> **On fund structures.** We hold the economics any vehicle has to satisfy and not yet a vehicle design. Available now: the layer-sorted case for where capital is the

right instrument, the governance-lock condition and the anchor-mass calculation (Fermigier 2026h). Designing the structure itself remains to be done, and it is the work an investor's own experience would shape most. Competition law and the state-aid treatment of a buyers' club anchored by a public investor are outside this brief and need their own advice before anything is signed.

2.7 Advising the companies you fund

Not one of the six questions. It is the channel with the shortest path from an investor's decision to a changed outcome, and it works before the exposure exists.

The answer in one line. For a company whose systems are already built, this is a migration programme. For one still building, it is seven covenants in the shareholders' agreement, and a board pack, which cost little at that stage and are worth a great deal by the time anyone thinks to ask.

Why timing dominates. Switching cost accumulates without ever being chosen. An engineer under delivery pressure picks the fastest managed service. Cloud credits, the free capacity hyperscalers give start-ups, make that choice free at the time; book them as a loan repaid in switching cost. The cost of that choice is paid by a different person several years later. Every instrument in this brief exists to reintroduce that missing feedback after the fact. Each of them is more expensive than having imposed the constraint at the start. A company at Series A has almost no estate to defend, so the stop-adding rule of section 1 costs it little. A company at twenty years has an estate whose exit cost is the subject of Annex B. The fund is one of the few actors positioned at the cheap end of that curve, because it is in the room when the architecture is still a paragraph in a plan.

Require seven things. They follow in the order an investment agreement would list them.

1. **Identity and key custody, in-house.** The directory of record and the keys sit in a tenant the company controls, never in the identity provider's or the cloud vendor's. This is the control plane's foundation, cheap on day one and expensive to unwind later, because everything else authenticates against it.
2. **Interfaces named in the plan.** Relational data behind standard SQL wire protocols and identity behind OIDC and SAML, which are formal standards; object storage behind the S3 API and monitoring behind OpenTelemetry, which are project interfaces with several implementations each; and, where a layer offers only a single-implementation open-source product, the upstream project without one distributor's extensions. The company would build on these anyway in most cases, and naming them converts a default into a commitment.
3. **A stop-adding rule for the tier**, with a written board-level justification where it is used, and an internal boundary around it so the dependency stays confined to replaceable modules. Not a prohibition: proprietary serverless is often the right call for a small team. The rule is that the choice is made on the record and walled off.

4. **Annual exit-cost disclosure.** What would it cost this company to leave each material supplier. One page, self-assessed, on the triage of Annex A. This is the instrument that does most of the work, because it is the missing feedback.
5. **Data export, free and tested**, in every supplier contract above a threshold. The clause is cheap to obtain at signature and costly at renewal.
6. **Portable formats for AI assets** at the point the company first buys AI: retrieval indices, evaluation sets, conversational memory, tool and permission configuration. That is the asset that locks, and the window for asking cheaply is now.
7. **Itemised internal charge-back**, so that a component arriving at no incremental price still generates a decision somewhere. Without an event, the company has no occasion to compare it with anything.

What this returns to the fund. Two things come back, neither of them an IT saving. A holding with a low exit cost concedes less margin to its suppliers over the life of every contract it signs, which is an earnings effect. A holding whose requirements are written against standards is a holding that can join a buyers' club later. The consistent finding in the cooperative-purchasing record is that pooling forms around standardisable requirements and stalls on bespoke ones, so the checklist above is also how a future consortium's membership is cultivated years before it is convened.

> **Where we are not yet firm.** The checklist is derived from results this brief cites and has not been tested against a portfolio. What it costs a company at each stage of maturity, and which of the seven items actually bind, is a question a fund with a portfolio can answer and we cannot.

Part 3. Annexes and sources

Annex A. The service triage

Figure 1 is the instrument. Sort every service on two axes, how deeply it is wired into application code and what leaving it would cost, and read the action off the cell. The examples there are from Amazon Web Services' catalogue, thirty-five services scored under a published rule; the same triage applies to the others. The numerals in the figure are the order of attack and are unrelated to the quadrant labels the source study uses.

How the second axis decomposes decides which instruments reach it. In the source's calibration, switching cost is roughly thirty percent the mass of data that has to be moved and seventy percent logical lock-in, meaning a proprietary interface specification, an absence of migration tooling, or both. Policy and contract terms that act on data transport reach about a fifth of the measured switching cost across the scored catalogue, and less than a third of it on every deeply locked service. This is the arithmetic behind the idea that opens the brief: where the data sits, and what moving it costs, is the smaller part of the exposure. Removing egress charges, as the Data Act does from January 2027, is correct and reaches a minority of the problem. The rest of the instruments have to attack the interface term.

Three working notes follow. The portable-interface cell is named for the property that makes exit cheap: more than one supplier can serve the same client unchanged. That property arises in two ways, from independent implementations of an interface and from an open-source product that many suppliers run. Both are weaker than an open standard. An interface governed by a single project can change, leaving the buyer's recourse to the other implementations rather than a specification process it can participate in; a single-implementation product can change its licence or its direction, leaving the recourse to a fork. The tender should say which of the three it is buying.

Identity comes first inside the control-plane cell, because everything else authenticates against it, so everything downstream waits on its migration and its outage is a total outage. Exit cost compounds across an estate faster than it adds, because managed services bind to one another: requiring open standards at the service boundaries lets components be replaced one at a time, so the estate need not move as a block.

Annex B. The break-even device

Rather than forecasting savings, derive the incumbent cost at which the move breaks even over a given horizon, then compare against a figure you already have. The denominator is the workload's own billing unit (section 3); the illustration here is a workforce estate, so it is per seat per year. An illustrative back-end collaboration stack gives:

Estate	3 years	5 years	10 years
250 seats	401	269	170

Estate	3 years	5 years	10 years
2,500 seats	293	229	182
25,000 seats	207	160	126

The table reads in EUR per seat per year. Above the figure, the move is cheaper over that horizon; below it, the incumbent is.

The inputs. The prices are February 2024 list prices for Nextcloud, Collabora, OpenXchange, Element and Univention (Standtke and Tiede 2024). The day rate is a blended EUR 1,200. The one-off effort is 144, 364 and 953 person-days for the three estate sizes, sized on the coarse progression the Swiss study used. The productivity loss is one user-day per user at EUR 300. The day rate, the effort and the user-day cost are illustrative and are the first things to replace.

How to run it on your own numbers. Replace the day rate, the effort estimates and the supplier quotations with your own, and compute your escalation rate from three years of your own renewal quotations. Then read the two structural facts before the specific numbers. Migration cost is close to independent of estate size, so the fixed cost of doing it at all dominates and small members face the hardest arithmetic; this is a direct argument for pooling. The three-year column is the one most likely to favour the incumbent, because the investment sits inside the window, so present both horizons and explain the difference.

Costs a model must not hide, because they are what turn a business case around:

- infrastructure and redundancy
- backup, continuity and off-site data transfer
- security capability and the people to run it
- parallel running of both stacks for the duration
- data migration tooling and effort
- training
- recruitment and retention
- contributions back to the software you depend on
- the incumbent's exit terms

That last item comes first in sequence, because enterprise agreements are frequently structured so that reducing seats mid-term saves nothing, which means savings begin at the renewal date whatever the technical plan says.

Caveats on the illustration. These figures come from the lab's own back-end services guide (Fermigier 2026k). Its cost model uses February 2024 published list prices from the Bern University of Applied Sciences study for the Swiss Federal Chancellery (Standtke and Tiede 2024) and illustrative effort assumptions, with the columns above undiscounted. The 250-seat row omits three components whose lowest published tier starts at 300 seats, so its ten-year figure sits below the 2,500-seat one; below about 300 seats much of this

market publishes no price at all. Discounting raises them, because a migration pays at the start and the incumbent pays across the horizon: at a four percent rate the ten-year figure rises by about a tenth, at eight percent by about a fifth. They demonstrate a method and do not state a cost.

Annex C. Procurement checklist

Requirements that screen, to be verified:

- ultimate ownership and control chain, including concealed parents
- governing law of the contract
- reach of foreign legal instruments to compelled disclosure or alteration
- key custody and administration
- continuity of supply under foreign compulsion

Prerequisites only. Require where needed on their own merits, give them no award weight, and never let them stand in for a control test: data residency, local staffing, local incorporation, generic operational certifications.

Contract clauses:

- Exit assistance as a priced deliverable, with defined scope and duration, invoked at your discretion.
- Data export in documented formats, on demand, at no incremental charge, with a tested procedure and a scheduled retest.
- Named second sources, recorded in the contract context. Elicit them by asking who could take over in three years and what you would hand them.
- Second-sourcing capability required above a workload threshold, demonstrated by rehearsal on a non-production perimeter.
- Support obligations as response and resolution targets by severity, with escalation into the upstream software project defined.
- Commitments denominated in capacity against a named portable interface, never in a supplier's proprietary catalogue.
- Portable formats for accumulated AI assets: retrieval indices, evaluation sets, memory, tool and permission configuration.
- Contributions to the upstream software you depend on, so that fixes you pay for do not become a private variant you maintain forever.
- No metering traps: verify the pricing model cannot become per-seat by another name.
- Source code escrow only where the build pipeline and documentation come with it.
- For members inside the DORA perimeter, the exit-strategy documentation and testing that Article 28 already requires, reused as the consortium's own template.

Award rules:

- hard requirements as gates before scoring, never as weights inside it
- declare weights before scoring
- no lowest-price award where the supplier is one you will rely on
- tender the service and let the product follow
- require exit-cost disclosure from every bidder at award

Supplier scoring rule, if you use one: weakest-link threshold to gate eligibility, weighted score to rank the survivors, unknowns held distinct from zeros, coverage published alongside the score, rule frozen before anyone is scored.

Annex D. Two studies the decisions wait on

Two of the decisions in this brief wait on work that has not been done. Each study would produce a reproducible model, with every number regenerated from published inputs and the assumptions visible.

1. Own or rent. *Settles the rent-or-build question of section 5, and produces the megawatt figure*. The study builds a capacity cost model comparing an owned facility (capital cost, operating cost, refresh cycle, utilisation curve) against co-location with a European operator, against rented European cloud capacity, and against hyperscaler on-demand and committed spend. It carries two terms we have not found in a published model: the pooling effect, where several members' imperfectly correlated demand raises achievable utilisation, and the switching-cost term, so the comparison is cost per unit net of the margin the architecture concedes. It outputs a break-even utilisation and a committed-load floor in megawatts, by workload class, with AI inference separated because its load shape differs. An anchor's own load and energy data would turn this from an illustration into a decision input.

2. Commitment without capture. *Settles the hardest part of the contracting question of section 4*. We can currently price exit rights and price premia. We cannot yet specify the take-or-pay structure that gives a subscale supplier a bankable demand floor while capping the buyer's realised switching cost. The output is a contract template with the covenants that decide it, and a real-options treatment of AI load uncertainty.

Annex E. What this rests on, and what it does not

The basis. The brief rests on the Abilian Econ Lab's working papers on the economics of European digital dependence, listed in the references. They are working papers, not peer-reviewed, so the brief's readers will be assessing the answers on their own terms. The external figures quoted here carry citations and confidence flags, and no proof. The four figures in this brief are drawn from the numbers stated in the text.

Where the case is an argument. The comparison in section 5 between transferring ownership and the demand-side rules rests on an ordering from a graded model, which

holds under the model's reweighting; the levels behind it are stylised and are not quoted. The split between exit rights and ownership screening in section 4 is sound as reasoning; its formal boundary conditions are not yet fully stated. The claim that switching cost caps extractable margin is correct in direction and does not yet convert to a euro figure; the study that would convert it uses the Data Act's removal of switching charges in January 2027 as a natural experiment.

Figures to quote with their provenance. The 8.7% annual price growth on Europe's locked base comes from an Asteres study commissioned by Cigref, an association of large corporate IT buyers, so it is an interested party's measurement and should be named as one. The fourfold-in-a-decade series is public-sector expenditure rather than unit price, so it overstates a pure price effect and should be replaced by the reader's own renewal history. The governance-lock retention figures (89% against 55%) are an association across resolved cases and not an estimated effect.

What we do not claim. The brief claims no causal identification anywhere. Where the models explain a documented record without testing against a counterfactual, they say so. It does not claim that any regulatory instrument delivers protection: enforcement is a separate question from enactment; certification carries costs that fall unevenly across supplier sizes. Nor does it claim that jurisdiction is the dominant channel of exposure. The channels we can measure with confidence are economic: extracted margin, switching cost, and the concentration of capability in one supplier.

References

An entry carrying a note says where the brief uses it. The rest are the background the argument stands on, listed for a reader who wants to go past this document; they are not quoted here.

The programme's own work

Public working papers are linked; the rest are internal and are marked as such.

- Fermigier, S. (2026a). *Cheaper and Losing: A Layered Model of Cloud Dependence and the Limits of Subsidy*, v0.6. Abilian Econ Lab. https://econ.lab.abilian.com/papers/pdf/cheaper-and-losing-wp.pdf. The price comparisons, the European share series, and why subsidy is a weak lever at either layer.
- Fermigier, S. (2026b). *Screening for Sovereignty: Optimal Procurement when Compliance is Cheap to Fake and Costly to Verify*, v0.15. https://econ.lab.abilian.com/papers/pdf/screening-for-sovereignty-wp.pdf. Which requirements can screen, why long referentials work against the objective, and the two faces of the price.
- Fermigier, S. (2026c). *A Two-Dimensional Taxonomy of Cloud Service Integrations*, v1.1. https://econ.lab.abilian.com/papers/pdf/cloud-services-matrix-wp.pdf. The triage of Figure 1, thirty-five services scored, and the finding that data-transport policy reaches a fifth of the measured exit cost.
- Fermigier, S. (2026d). *Exit or Own? Portability Mandates and Ownership Screens in Sovereign Procurement*, v0.6. https://econ.lab.abilian.com/papers/pdf/exit-or-own-wp.pdf. Why exit rights and ownership screening cover different sides of the exposure.
- Fermigier, S. (2026e). *The Unpriced Bundle: Decision Suppression in Software Procurement*, v0.1. https://econ.lab.abilian.com/papers/pdf/unpriced-bundle-wp.pdf. Why a bundled component is never compared with anything, and why discounts are inert against that.
- Fermigier, S. (2026f). *Path Sovereignty: Dependence, Control, and Options in the European Digital Stack*. https://econ.lab.abilian.com/papers/pdf/power-network-wp.pdf. The measured dependency graph behind the section 5 comparison.
- Fermigier, S. (2026g). *Leverage Without Additionality: What Co-Funding Ratios Identify in Public Technology Programmes*, v0.2. https://econ.lab.abilian.com/papers/pdf/leverage-without-additionality-wp.pdf. Why a reported leverage ratio restates the funding schedule.
- Fermigier, S. (2026h). *Buyers' Clubs for the Open-Source Flow: Collective Provision in the FLOSS Value Network*, v0.7. https://econ.lab.abilian.com/papers/pdf/collective-provision-wp.pdf. Critical mass, the anchor buyer, and the two membership currencies of Figure 4.

- Fermigier, S. (2026i). *Sovereignty as a Flow: Capture and the Durability of European Technology*. https://econ.lab.abilian.com/papers/pdf/capture-wp.pdf. The governance-lock retention figures.
- Fermigier, S. (2026k). *Open Source Back-End Services: A Migration Guide for Directory, Storage, Mail, and Real-Time Communication*. https://econ.lab.abilian.com/papers/pdf/backend-services-wp.pdf. The hosting decision, the reference architectures, the procurement clauses, and the break-even model of Annex B.
- Fermigier, S. (2026l). *The Open Source Workstation: A Migration Guide for Applications, the Desktop, and the Managed Linux Estate*. https://econ.lab.abilian.com/papers/pdf/workstation-wp.pdf. The third of the three independent migration steps, and the sequencing argument for taking it last.
- Fermigier, S. (2026m). *Cuts and Options: A Theory of Strategic Autonomy in Production Networks*. https://econ.lab.abilian.com/papers/pdf/theory-wp.pdf. Why per-layer certification cannot certify a stack.
- Fermigier, S. (2026n). *The EU's AI money is aimed one layer too low*. https://econ.lab.abilian.com/papers/pdf/eu-ai-model.pdf. The argument and sources behind Figure 2.
- Fermigier, S. (2026p). *Scoring the Seller: An Economic Foundation for a Supplier-Side Sovereignty Index*, v0.6. Internal working paper. The source of the aggregation rule cited in section 4, whose properties the paper now states as a proposition with its proof and its matching checks.

Market structure, pricing and competition

- Asterès (2025). *Technological dependence on American software and cloud services: an assessment of the economic consequences in Europe*. Study commissioned by Cigref, May 2025.
- Asterès (2026). *From technological dependency to economic capture: the cost of cloud and software services inflation to Europe*. Study commissioned by Cigref, May 2026. Source of the 8.7 percent annual growth on Europe's locked base.
- Competition and Markets Authority (2025). *Cloud services market investigation: final decision report*, 31 July 2025. Source of the switching rate below 1 percent a year and of the returns sustained above cost of capital.
- European Parliament (2025). *European Software and Cyber Dependencies*. Study PE 778.576, Policy Department for Transformation, Innovation and Health, December 2025.
- European Commission (2025). *Commitments in the Microsoft Teams proceedings*, September 2025.
- European Commission (2026). *Study on the EU's critical digital capacities deployment beyond 2027: final report and Annex 4 (Artificial Intelligence) and Annex 6 (Cloud-Edge-IoT).* Prepared by Technopolis Group for DG CNECT. https://data.europa.eu/doi/10.2759/1036940

- getdeploying (2026). Cloud provider price-comparison tables (compute, storage, GPU, egress). https://getdeploying.com
- Hansson, D. H. (2024). *Our cloud-exit savings will now top ten million over five years*, 17 October 2024. https://world.hey.com/dhh/our-cloud-exit-savings-will-now-top-ten-million-over-five-years-c7d9b5bd
- Dropbox (2018). *Form S-1 registration statement*, US Securities and Exchange Commission. Source of the reported infrastructure savings after moving off a hyperscaler.

Public-sector migration evidence

- Kern, S., Voss, N., Kempe, C., Blank-Babazadeh, M., Hinz, A. and Gabriel, M. (2022). *Linux-Arbeitsplatz für die öffentliche Verwaltung*. Kiel: Der Ministerpräsident des Landes Schleswig-Holstein, with Dataport AöR.
- Lecrivain, D., Hemmel, P. and Facia, G.-C. (2022). *Poste de travail Linux: état de l'art et conduite du changement*, v1.2. Paris: Direction générale des Finances publiques.
- Staatskanzlei Schleswig-Holstein (2024). *Open Innovation and Open-Source Strategy of Land Schleswig-Holstein*, v1.00, 20 November 2024. Source of the fourfold-in-a-decade expenditure series, and of the four-phase cost shape behind Figure 3.
- Standtke, R. and Tiede, M. (2024). *Studie zu Open-Source-Alternativen von Microsoft Services und Produkten in der Schweizerischen Bundesverwaltung: Backend-Services*, v1.8, February 2024. Bern: Berner Fachhochschule, Institut Public Sector Transformation.
- Stürmer, M. (2024). *Technologische Perspektive der digitalen Souveränität*. Report to the Swiss Federal Department of Foreign Affairs. arXiv:2406.03266.

Collective procurement and open-source institutions

- ACM, AFM, AP, DNB and RDI (2026). *Towards digital autonomy: increased control over digital dependencies*. Joint position paper of five Dutch supervisory authorities.
- Frey, F. (2023). How we lead successful open source collaborations in the Danish public sector. *IEEE Software* 40(4), 19–24. The Danish municipal association of Figure 4.
- Viseur, R. and Jullien, N. (2023). Communesplone, an original open-source model of resource pooling in the public sector. *IEEE Software* 40, 46–54. The Walloon co-owned supplier of Figure 4.
- Linåker, J. and Muto, S. (2026). Advancing digital government through open source indicators. *Journal of Systems and Software* 242, 113031.
- Bitkom (2026). *Cloud Report 2026: Status Quo & Trends in Wirtschaft und Politik*. Bitkom Research.

Artificial intelligence

- Mozilla (2026). *The State of Open Source AI*, v1.0, July 2026. https://stateofopensource.ai. Source of the production-rate figures.
- European Commission (2025). *InvestAI: mobilising €200 billion for artificial intelligence*. Press release IP/25/467, 11 February 2025.

Economics

- Klemperer, P. (1995). Competition when consumers have switching costs. *Review of Economic Studies* 62(4), 515–539.
- Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: competition with switching costs and network effects. In *Handbook of Industrial Organization*, vol. 3. Elsevier.
- Fischbacher, U., Gächter, S. and Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters* 71(3), 397–404. The conditional-cooperation profile behind the anchor argument.

Sovereignty instruments compared in Annex F

- aDRI (2026). *Indice de Résilience Numérique (IRN), référentiel v1*. Digital Resilience Initiative. CC BY-NC-ND 4.0. The buyer-side grid of Annex F; studied and cited, never copied.
- Caffarra, C., et al. (2025). *The EuroStack Initiative: Buy European, Sell European, Fund European* (the white paper). https://eurostack.eu/the-white-paper/. The doctrine behind the framework.
- EuroStack (2025). *A Proposed Framework for a "Buy European" Regulation of Strategic Digital Procurement*. https://eurostack.eu/a-proposed-framework-for-buy-european-regulation/. The five-dimension provider screen of Annex F.

Legislation and official frameworks

- Regulation (EU) 2023/2854 (Data Act), Chapter VI on switching between data-processing services. Switching charges are removed from 12 January 2027.
- Regulation (EU) 2022/2554 (DORA), Article 28 on ICT third-party risk and exit strategies.
- European Commission, *Cloud Sovereignty Framework*, v1.2.1, 20 October 2025. The two-stage scoring rule cited in section 4.
- Microsoft (2025). *Microsoft completes landmark EU Data Boundary, offering enhanced data residency and transparency*, 26 February 2025. https://blogs.microsoft.com/on-the-issues/2025/02/26/microsoft-completes-landmark-eu-data-boundary-offering-enhanced-data-residency-and-transparency/. The residency programme cited in section 4.
- Sénat (2025). Audition de Microsoft France, 10 juin 2025, in *Commission d'enquête sur les coûts et les modalités effectifs de la commande publique*, rapport n° 830

(2024-2025), tome II. https://www.senat.fr/rap/r24-830-2/r24-830-2.html. The testimony cited in section 4.

About the Abilian Econ Lab

The Econ Lab is the economics research laboratory of Abilian, a French open-source software company. It runs a single research programme, the economics of European digital sovereignty, which asks through which mechanisms dependence on foreign technology providers turns into economic and political exposure, and which interventions actually reduce it. It answers with measured dependency graphs, formal theory, mechanism analysis, service-level architecture measurement, and the economics of open-source adoption and supply, in a field where most public debate runs on anecdote.

Author. Stefane Fermigier holds a PhD in mathematics and is adjunct professor of open-source software at Université Paris Cité. He founded Nuxeo, founded and runs Abilian, and co-chairs the CNLL, the French open-source business association.

Comments are welcome: sf@abilian.com · econ.lab.abilian.com