Abilian
ECON LAB

Ask Me Something Hard: Why Three Fifths of Europe's Sovereignty Score Cannot Tell Suppliers Apart

Eighteen instruments measured criterion by criterion, and the eleven questions that separate a genuine supplier from a controlled one

Stefane Fermigier (Abilian) · sf@abilian.com · 18 September 2026 · DRAFT POSITION PAPER

Executive summary

A sovereignty score exists to tell a public buyer something it cannot see: whether a supplier offering a European service is genuinely beyond the reach of a foreign state, or is a subsidiary presenting the appearance of one. Europe now has at least eighteen instruments that ask a supplier questions and turn the answers into a level. The **Cloud Sovereignty Framework** and its SEAL scale is the one that matters most, because it has been in use since the Commission's 2025 tender and because the **Cloud and AI Development Act**, now before the Council and the Parliament, builds a parallel ladder beside it.

This note measures what those instruments ask, criterion by criterion, and applies one test to each question: **can a foreign-controlled supplier answer it as cheaply as a European one?** Where the answer is yes, the question tells a buyer nothing about the thing it is trying to find out, however precisely the answer can be checked. That test is a result in the economics of procurement. It has a name. A requirement screens only where the two types' costs diverge and where verification reaches that divergence.

Applied to the Commission's own calculator, the test produces a number.

What each question of the Commission's sovereignty assessment can do	Share of the weighted score
Separate a genuine supplier from a controlled one	**40.3%**
Inform a reader, and separate nothing	33.2%
Answer no sovereignty concern the buyer has, on the narrower of two defensible scopes	24.4%
Describe the assessment apparatus	2.1%

Three fifths of the score a public buyer computes cannot distinguish the supplier it is worried about from the supplier it wants. The pattern runs across the whole field, so it is more than a drafting slip in one instrument. Across the thirteen supplier-side instruments examined, every one spends between half and seven tenths of its questions on criteria with no separating power. The questions they ask most are the ones that separate least. Place of incorporation is asked by all

thirteen. Data location is asked by twelve. Personnel location by eleven. Governing law by ten. A foreign-controlled provider satisfies every one of them at the same cost a European provider does, which is why it does.

Four further findings bear on the text now before the Council and the Parliament.

- **The instruments disagree with each other, by construction.** Restrict any two of them to the criteria both ask, so that the facts and the coverage are identical and only the grouping differs. They then order a pair of suppliers in opposite directions in a median **33.7%** of cases. Identical groupings never disagree. Unrelated suppliers disagree half the time. A certificate that does not name the grouping it was computed under is not a portable statement.
- **The Commission's calculator declines to score what tells the most.** Nine of its forty-eight questions are neutralised: every available answer scores the top level, so the question cannot lower a supplier's band. Three of those nine are change-of-control risk, financial independence from non-EU capital, plus the origin of the intellectual property. Those are exactly the structural ownership facts a controlled supplier cannot cheaply fake.
- **The questions that separate can mostly be answered without asking the supplier.** Of the twenty-two criteria that pass the test, eleven come from company registers, repositories and probes against the running service. A declaration form needs eleven questions, against the thirty to sixty the instruments ask today.
- **Machine probes separate the types where the declared criteria fail to.** Twelve suppliers were read on four criteria. Two of the four separate cleanly, with **no foreign-controlled supplier cleared by either**. The other two are reported with their failure. The clearest single reading is a German company selling a sovereign cloud to the German administration whose own mail runs on Microsoft.

The recommendations in section 10 are three. Require the structural ownership chain the calculator currently neutralises. Require the assurance certificate to name the criterion grouping it was computed under. Separate what is scored from what is merely published, so that the criteria which cannot separate suppliers stop consuming the weight that should sit on the criteria that can.

About this note

It extends *Measure Twice, Buy Once* (Fermigier 2026a), which compared the Cloud and AI Development Act's Union assurance levels with the Commission's SEAL scale and found that neither contains the other. That note asked whether the two ladders correspond. This one asks a prior question: whether either of them, or any of the sixteen other instruments now in the field, measures something a supplier cannot cheaply fake.

SEAL is the Sovereignty Effectiveness Assurance Level of the Commission's Cloud Sovereignty Framework, a scale from SEAL-0 to SEAL-4 computed across eight objectives, published with an implementation guidance document and a spreadsheet calculator. **UAL**, the Union assurance levels, are the four-level scale of Annex II to the proposed Cloud and AI Development Act, COM(2026) 502 final of 3 June 2026. Every number in this note is computed by a script from the published texts. Annex A lists the scripts and their inputs.

Findings

1. Three fifths of the SEAL weighted score sits on criteria that cannot separate a genuine supplier from a foreign-controlled one. The split between “cannot separate” and “answers no sovereignty concern” depends on a scope choice; the three fifths does not.
2. Of ninety-four distinct criteria asked across the field, twenty-two pass the separating test. Every instrument examined asks between half and seven tenths of its questions outside that set.
3. The five criteria the field asks most are all in the failing set.
4. Two instruments restricted to the criteria both ask reorder suppliers in a median third of cases, on grouping alone.
5. The Commission’s calculator neutralises nine of its own questions, three of which are the ownership facts with the most separating power.
6. Eleven of the twenty-two separating criteria require no declaration from the supplier.
7. On twelve suppliers, two of four machine-readable criteria separate the types with no false pass. A third is silent for most of the sample and a fourth reveals nothing about type.

1. The job a sovereignty score is doing

A buyer procuring a cloud service cannot observe whether the supplier in front of it is genuinely outside the reach of a foreign state. It can observe where the servers are, who signs the contract, which certificates hang on the wall. The thing it wants to know, effective control, is hidden. A supplier that lacks it can present the appearance of holding it at modest cost: an incorporated European subsidiary, a European region, a European-law contract.

This is the structure of a market for lemons. When the hidden quality is cheap to imitate and costly to verify, buyers cannot distinguish, the pool prices at the imitated level, and genuine suppliers have no way to separate themselves. Sovereignty procurement has a name for the imitation: sovereignty-washing. The economics are standard.

A screen breaks the trap: a requirement that the two kinds of supplier cannot satisfy at the same cost. The formal statement is precise and it cuts both ways. A requirement screens only where the types’ compliance costs diverge **and** where verification reaches that divergence. An attribute the two types hold at the same cost says nothing about which type is which, at any level of audit accuracy, however cleanly it can be checked (Fermigier 2026b, Theorems 2 and 3).

Two consequences follow. Both are uncomfortable for the instruments now deployed.

Checkable is not the same as informative. Data residency is the standard illustration. Where a service stores data is a fact an auditor can establish with certainty. It is close to uninformative about whether a foreign state can compel access to it, because a foreign-controlled provider offers European regions at the same cost a European provider does. Perfect verifiability of an attribute with no cost gap yields nothing.

Padding is not free. Audit effort across requirements is superadditive: the cost of verifying a set of criteria exceeds the sum of verifying each alone. An instrument padded with criteria that cannot separate does not merely waste the answers. It consumes the audit budget and the reader’s attention that the separating criteria need, diluting the credibility of the score as a whole. A lean instrument weighted toward the gap-bearing attributes dominates a comprehensive one.

2. Applying the test

Each criterion asked anywhere in the field was put to three questions.

1. **Does it answer a concern the buyer actually has?** Six concerns cover the field: compelled foreign access, continuity if the supplier is cut off, the cost of leaving, ownership and rent, data handling and key custody, plus whether an open option stays open.
2. **Does it have a cost gap?** Would a foreign-controlled supplier have to spend materially more than a European one to hold it genuinely?
3. **Can verification reach the gap?** Is there a channel, a register, a repository, a probe, a document, a demonstration, that detects the difference between the claim and the fact?

A criterion that answers no buyer concern is out. A criterion with no cost gap can be published and cannot set a level, because it reveals nothing at any audit accuracy. A criterion with a gap that only the supplier's own word reaches fails the accuracy threshold and is published rather than scored. What remains sets a level.

The three judgements are recorded per criterion with their reasons, and a decision table computes the verdict from them. The arrangement puts the judgement where it can be contested one criterion at a time. A contested judgement then changes the outcome mechanically.

3. Measuring the field

Twenty-three instruments were identified and eighteen publish their criteria at a level that can be mapped. They were reduced to ninety-four distinct criteria, stated in terms no framework owns, with every source row recorded: four hundred and fifty of them, each cited by at least one criterion and each checked to exist in its source.

Instrument	Owner	Scores	Rows mapped
Cloud Sovereignty Framework calculator (SEAL)	European Commission	a supplier	48 questions
Union assurance levels, Annex II	European Commission	a supplier	11 criteria × 4 levels
C3A, Criteria enabling Cloud Computing Autonomy	BSI, Germany	a supplier	30 criterion groups
Sovereign and Resilient Cloud Services	CISPE	a service	70 controls
SecNumCloud 3.2, chapter 19	ANSSI, France	a supplier	30 requirements
Compliance Document 24.06	Gaia-X	a service offering	25 criteria
Toetsingsinstrument Soevereiniteit Clouddiensten	DICTU, Netherlands	a supplier	15 criteria
Cloud Sovereignty Framework Self Assessment	SUSE	a supplier	32 questions
Sovereignty Score	Posetiv	a supplier	10 objectives
Referentiel IRN v1.1	Digital Resilience Initiative, France	a buyer	30 criteria
Digital Sovereignty Readiness Assessment	Red Hat	a buyer	21 questions
Kriterien zur Bewertung Digitaler Souveränität	ZenDiS, Germany	a buyer	20 criteria

Instrument	Owner	Scores	Rows mapped
openCode badges	ZenDiS with BSI	a repository	5 badges

Five more contribute to the register and not to the criterion mapping, because they publish dimensions without criteria: STACKIT's ES³, the EDSO standard, the Bechtle index, the LOTEC pillar model. The EUCS candidate scheme has no stable public text.

The convergence deserves stating before the disagreements. Ten criteria are asked by ten or more of the thirteen supplier-side instruments. Every instrument asks about incorporation, ownership, data location, personnel, the exit path. Nobody is inventing a private vocabulary. The difficulty is not that the field disagrees about what to ask.

4. Three fifths of the score cannot screen

Apply the test of section 2 to each of the forty-eight questions in the Commission's calculator. Spread each objective's published weight evenly across its questions, which is what the calculator itself does. The result:

What the question can do	Share of the weighted score
Separate the types	40.3%
Inform, and separate nothing	33.2%
Answer no sovereignty concern	24.4%
Describe the assessment apparatus	2.1%

The weights are those the calculator applies. The framework document itself prints a different vector (strategic 15, supply chain 20, security 10 percent, where the calculator has 20, 10 and 15); under it the separating share is 41.3% and the rest 58.7%, so the three fifths do not depend on which of the Commission's two texts is right.

The 24.4% is mostly the security and environmental objectives. Neither is a small or improper subject. Both have their own instruments and their own law. The German federal office drew the same line when it built C3A, which omits the security objective on the stated ground that the C5 catalogue already covers it.

That line is arguable, so it is a parameter here and not a ruling. A buyer procures one service and bears the joint outcome. Several criteria that look like security are control criteria in disguise: who holds the encryption keys, who can read the logs, who can patch without the vendor. This note scores all three, under the concerns they actually answer.

Running the test with security and environmental admitted settles what the argument costs.

	Security and environmental out	Security and environmental in
Criteria that separate the types	22	the same 22
Share of the weighted score that separates	40.3%	40.3%
The rest, as out of scope	24.4%	none
The rest, as informative disclosure	33.2%	57.6%

Admitting them moves every one into disclosure and none into the separating set. They fail the separation test on their own merits whatever the scope, since both kinds of supplier hold an ISO 27001 certificate, publish a carbon figure and run a vulnerability process at the same cost. **The scope argument changes how the remaining three fifths is described and leaves the finding untouched.**

The 33.2% is the more serious number, because those criteria look like sovereignty criteria and are not. They are the answers a foreign-controlled provider gives most readily, since giving them costs it nothing.

5. The field converges on what does not separate

The five criteria most asked across the field all fail the test.

Criterion	Instruments asking it, of 13	Why it does not separate
Place of incorporation	13	The European subsidiary is the standard form of the imitation, so incorporation is what the imitation buys
Data location, including metadata and backups	12	Both kinds of supplier offer European regions at the same cost
Location of operations and support personnel	11	Both can hire in the Union at the same cost
Governing law and competent court	10	Both offer European law and a European court; neither prevents compulsion on the parent
Subprocessor list	10	Publishing the list is costless; the dependency itself is the criterion that separates

The thirteen are the supplier-side instruments of the basis: the nine of section 3′s table that score a supplier or a service, plus four it does not show, the JOTED specification, the EuroStack directory's screening gate, the Cloud III submission form and this lab's own register. Without the lab's register, incorporation is asked by twelve of twelve.

Incorporation deserves a sentence of its own, because dropping it would be wrong. It is informative as part of the ownership chain: the registered entity together with its ultimate parent, the domicile of its directors, plus any veto or appointment right held from outside. Scored inside that chain it separates. Scored beside it, as an independent criterion with its own weight, it rewards precisely the structure the chain exists to detect.

6. Two instruments, the same questions, opposite answers

Every deployed instrument groups its criteria into objectives, scores each objective, then takes the worst objective as the level. That shape has a property its users may not have priced. **The level depends on the grouping.** Splitting an objective in two can only lower a supplier's level. Merging two objectives can only raise it. A supplier facing a published rule can therefore lobby for coarser objectives at no cost in evidence.

The consequence is measurable. Take any two instruments, restrict both to the criteria both ask, so that coverage is identical, the facts are identical, the only difference being how the criteria are grouped. Then draw pairs of suppliers and count how often the two instruments rank them in opposite directions.

	Supplier pairs ordered in opposite directions
Two identical groupings	0%, necessarily
The 45 real pairs, median	33.7%
The worst real pair, SEAL against Annex II	44.4%
Bands of unrelated suppliers	50.7%

Every one of the forty-five pairs disagrees sometimes. The published groupings are more alike than random regroupings of the same shape, which is the convergence of section 3 showing up again. They still reorder a third of supplier pairs.

Two qualifications. The rates are computed on suppliers drawn independently across criteria; when a supplier's criteria move together, so that a good supplier is good across the board, the median rate falls to 11.7% and then to zero as suppliers become uniform. Disagreement is therefore a property of suppliers that are strong on some criteria and weak on others, which is the case a sovereignty score exists to catch. The exercise also measures ranking disagreement and never error, since neither instrument is right.

7. What the calculator declines to score

Nine of the calculator's forty-eight questions are neutralised. Every answer offered scores the top level, so the question cannot lower a supplier's SEAL band whatever the supplier answers. Six of the nine are questions this note's test also discards, which is a point of agreement between two independent readings.

The other three are not.

Neutralised question	What it asks
SOV-1.2, change of control risk	The likelihood of takeover or transfer to a non-sovereign owner
SOV-1.4, financial independence	Reliance on capital from outside the Union
SOV-2.5, origin of intellectual property	Where the intellectual property was created and registered

These are structural ownership facts. They are recorded in company registers and in patent registers. They are expensive to falsify and cheap to check. A foreign-controlled supplier cannot hold them genuinely without ceasing to be foreign-controlled. On the test of section 2 they are among the most informative criteria the calculator contains, yet the calculator has arranged that they cannot affect a score.

The same reading, applied to two other instruments, agrees with the test. The Cloud and AI Development Act's Annex II places each criterion at the assurance level where it first bites. The criteria the drafters require from level 1 have measurably less separating power than those held back to level 2 and above. The German C3A catalogue splits its criteria into basic and additional. The additional ones, which the BSI judged too demanding to require of everyone, have more. Drafters know which of their requirements are hard. Two of these three instruments then decline to score the hard ones.

8. Twelve suppliers, one probe

The test of section 2 rests on judgements about what a controlled supplier can cheaply do. Those judgements are contestable, so here is one of them confronted with evidence.

Twelve suppliers were probed on 18 September 2026 on four criteria the test rates as separating, their type fixed in advance from ownership recorded with its basis: six European-controlled, four foreign parents

selling a sovereign offering, two foreign with no such offering. A criterion is judged by the errors it makes. A false pass, a foreign-controlled supplier the criterion clears, is the error a buyer acts on. A false failure costs a genuine supplier a place on a list and costs the buyer nothing.

Criterion, as probed	False passes	False failures	Unreadable	Reading
Who operates the supplier's own mail and name service	0	2	0	separates
Whether the service publishes source under a licence	0	3	0	separates
Which certification authorities the domain permits	0	3	7	mostly silent
Third-party origins the supplier's home page loads	1	5	2	does not separate

Two of the four separate. Neither clears a single foreign-controlled supplier. Each costs some European suppliers a pass, which is the error a buyer can afford. The sharpest single reading is a German company selling a sovereign cloud to the German administration, whose own mail is operated by Microsoft.

The other two are reported because they were run. Seven of the twelve publish no certification-authority record at all, so that criterion has nothing to read for them, and silence is neither a pass nor a failure. The front-end criterion fails outright: every European supplier whose page could be read loads a third-party origin outside the Union. Its single pass belongs to a foreign-controlled supplier whose page the detector found nothing on. That is a result about the probe as much as about the suppliers.

Three limits apply, and a larger sample repairs none of them. The probe reads the supplier's own corporate footprint and not the footprint of the service being sold, so it is a statement about revealed preference. The type labels come from ownership, so this tests whether one criterion predicts a type another criterion establishes. And twelve suppliers is twelve suppliers.

9. What to ask instead

Twenty-two criteria pass the test. Eleven of them require nothing from the supplier: they come from company registers, from repositories, from probes against the running service. Asking a supplier to declare what a register already settles substitutes the weakest evidence for the strongest.

That leaves eleven questions for a declaration form.

1. Is the service, its provider, or any entity in its group subject to a law of a third country that can compel disclosure of, or alteration to, customer data or the service?
2. Is any component subject to an export-control regime that could restrict its use, transfer or continued supply?
3. Where a parent or affiliate sits in a third country, what separation prevents it from reaching the service, and who audits that separation?
4. If the originating vendor of a component stopped supplying it, could the provider maintain and patch the service itself, for how long, and using whose staff?
5. Is there documentation, source and know-how sufficient for a competent third party to run the service without the provider?
6. Can the customer hold the encryption keys so that the provider cannot read the data?
7. What does it cost a customer to leave, and has a real customer completed a migration off the service?

8. Can any component be disabled, degraded or altered remotely by a party other than the provider?
9. For each dependency outside the Union, is there an identified alternative and a tested switchover?
10. Can any supplier of a critical asset withdraw it for a reason other than non-payment or breach?
11. If the provider were cut off, could the customer keep operating, run by whom, on what, and for how long?

Every one of them asks for a document or a demonstration, and every one is expensive for a supplier whose true answer is the wrong one. That is the property being bought.

10. What the co-legislators should do

Recommendation 1. Require the ownership chain, and stop neutralising it. Annex II should state a criterion, at Union assurance level 2 and above, on the ultimate parent and controlling ownership of the provider, the domicile of its directors, and any veto or board-appointment right exercisable from a third country. Each element is a public-register fact, so the compliance cost is a search rather than an audit. The Commission's own calculator should restore scoring to its change-of-control, financial-independence and intellectual-property questions, which requires no change to the questions themselves, only to the levels their answers earn.

Recommendation 2. Make the certificate name its grouping. A Union assurance level, and any SEAL band quoted alongside it, should be published with the criterion grouping under which it was computed. Two instruments over identical criteria reorder a third of supplier pairs on grouping alone (a measure of disagreement between the instruments, which says nothing about which of them is right), so a level quoted without its grouping is not a portable statement and cannot support mutual recognition under Article 17. The cost of the remedy is one line on a certificate.

Recommendation 3. Separate what is scored from what is published. Criteria that cannot distinguish a controlled supplier from a European one should be published as disclosure and excluded from the level. Data location, personnel location, governing law and subprocessor lists belong in the published record, where a buyer can read them and a court can rely on them. They do not belong in the arithmetic, where they currently absorb weight that the separating criteria need and give a foreign-controlled provider an inexpensive route to a respectable score. The evidentiary value of those criteria is unaffected by removing them from the sum.

Two asks already on the table remain sound and are strengthened by this reading. The software-supply-chain criterion at Union assurance level 2 should keep its bill of materials requirement, since supply-chain transparency is the evidence channel through which the dependency criterion is verified. The derogation route for third-country control at level 3 should also be narrowed. Third-country control is the single criterion with the largest separating power in the field, so a derogation on it is a derogation on the one thing the ladder reliably measures.

Annex A: how the numbers were produced

Every figure in this note is computed from the published texts by scripts in the repository, each of which prints its own checks.

- `rsi/verify/build_basis.py` parses the Commission's sovereignty assessment calculator, reads the criterion identifiers of the other instruments, and checks the criterion basis both ways. Every identifier cited exists in its source; every source row is cited by a criterion. It reports 94 criteria over 450 source rows.

- `rsi/verify/derive_register.py` applies the test of section 2. The three judgements per criterion are recorded in a data file with their reasons; the verdict is computed from a decision table. It produces the 22 separating criteria, the weight decomposition of section 4, and the per-instrument coverage.
- `rsi/verify/composition.py` states the grouping properties for any instrument of this shape and checks them against the ten deployed groupings whose published identifiers encode them. It produces the disagreement rates of section 6, the permutation control, and the sensitivity of the rate to how uniform suppliers are.
- `rsi/verify/revealed.py` compares the separating test against the judgements the instruments reveal about their own criteria, producing section 7.
- `rsi/verify/probe.py` queries the mail exchange, name service, certification-authority record, home page and repository of the twelve suppliers of section 8, records each result with its date, and reports each criterion as a screen judged by its errors. Replayable offline.

The judgement in this method is concentrated in one place: whether each criterion shows a cost gap between the two kinds of supplier. Those judgements are recorded individually with their reasons, and each can be contested on its own, in which case the register recomputes. The weight decomposition of section 4 would move if enough of them were overturned. The qualitative finding would survive a considerable number of reversals, since it rests on the criteria the field asks most and those are the clearest cases.

References

Bundesamt für Sicherheit in der Informationstechnik (2026). *C3A: Criteria enabling Cloud Computing Autonomy*. April 2026. https://www.bsi.bund.de/ (accessed 18 September 2026).

CISPE (2026). *Principles for Sovereign and Resilient Cloud Services*, with Appendix 1, Sovereignty Controls, and Appendix 2, Resilience Controls. April 2026. https://framework.sovereignty.cispe.cloud/ (accessed 18 September 2026).

DICTU (2026). *Toetsingsinstrument Soevereiniteit Clouddiensten*, version 1.0.1. January 2026. https://www.dictu.nl/ (accessed 18 September 2026).

European Commission (2025). *Cloud Sovereignty Framework*, version 1.2.1, October 2025. https://commission.europa.eu/document/download/09579818-64a6-4dd5-9577-446ab6219113_en?filename=Cloud-Sovereignty-Framework.pdf

European Commission (2026a). *Cloud Sovereignty Framework: implementation guidance*, with the annexed sovereignty assessment calculator.

European Commission (2026b). *Proposal for a Regulation on cloud and AI development*, COM(2026) 502 final, 3 June 2026, with Annexes I to III.

Fermigier, S. (2026a). *Measure Twice, Buy Once: The Cloud and AI Development Act's Union assurance levels against the Commission's own SEAL scale, criterion by criterion, and eight amendments for the trilogue.* Draft position paper, 18 September 2026. Abilian Econ Lab. https://econ.lab.abilian.com/papers/two-ladders/

Fermigier, S. (2026b). *Screening for Sovereignty: Optimal Procurement when Compliance is Cheap to Fake and Costly to Verify.* Draft working paper v0.20, 2026-10-21. Abilian Econ Lab. https://econ.lab.abilian.com/papers/screening-for-sovereignty/

Gaia-X (2024). *Gaia-X Compliance Document*, release 24.06. https://gaia-x.eu/ (accessed 18 September 2026).

Agence nationale de la sécurité des systèmes d'information (2022). *Prestataires de services d'informatique en nuage (SecNumCloud), référentiel d'exigences*, version 3.2, 8 March 2022. https://cyber.gouv.fr/ (accessed 18 September 2026).

Posetiv (2026). *Sovereignty Score: methodology and rubric*. https://sovereignty.posetiv.co.uk/ (accessed 18 September 2026).

Red Hat (2026). *Digital Sovereignty Readiness Assessment*. https://github.com/redhat-cop/viewfinder-upstream (accessed 18 September 2026).

SUSE (2026). *Cloud Sovereignty Framework Self Assessment*. https://www.suse.com/cloud-sovereignty-framework-assessment/ (accessed 18 September 2026).

Zentrum für Digitale Souveränität der Öffentlichen Verwaltung (2026). *Kriterien zur Bewertung Digitaler Souveränität*, consultation draft. https://souveraenitaetscheck.de/ (accessed 18 September 2026).