Two Screenings: Why Europe Screened Huawei but not the Hyperscalers

Stéfane Fermigier · Working paper v0.6 · 14 July 2026

Abstract

Within a single polity and a single decade, the European Union designed two opposite screens for foreign suppliers of critical digital infrastructure. Against Chinese telecom vendors, the 2020 5G Toolbox screened structurally, on ownership and home-state control; against US cloud providers, EU instruments converged on behavioural requirements, and the EUCS certification scheme dropped its ownership and immunity criteria in 2024. We explain the divergence with two channels from a companion mechanism-design model: verification feasibility (the decisive attribute is trivially verifiable in one case and requires a contested apparatus in the other) and capture, which targeted enforcement where the screen's design was publicly legible (5G) and design where it was not (cloud). A design-by-enforcement matrix organizes the record and yields predictions that developments postdating the calibration window (the Cloud Sovereignty Framework, the Cloud and AI Development Act, the 2026 Cybersecurity Act revision) bear out.

Keywords: procurement, screening, regulatory capture, 5G, cloud, digital sovereignty, European Union.
JEL: D82, D72, H57, L96.

1. Introduction: one Union, two screens

Between 2019 and 2026 the European Union confronted, twice, the same problem in the same terms: suppliers of critical digital infrastructure whose home jurisdiction can compel them: to hand over data, to insert access, to cut off service. Twice it designed a screen. The designs are opposites.

Against Chinese telecom equipment vendors, the *5G Toolbox* (NIS Cooperation Group, 29 January 2020) told member states to assess suppliers on explicitly *non-technical* criteria: "the likelihood of the supplier being subject to interference from a non-EU country," evidenced by "a strong link between the supplier and a government," the home country's legislation (especially where "there are no legislative or democratic checks and balances") and "the characteristics of the supplier's corporate ownership." That is a *structural screen*: it tests origin and control, facts that are public, cheap to verify, and that no commercial restructuring can wash away.

Against US cloud providers, the Union's instruments converged on the opposite design. The candidate structural criteria (ownership caps and immunity from extraterritorial law, present in drafts of the EUCS cloud-certification scheme from 2021) were *removed in the March 2024 draft* after sustained industry and member-state pressure, leaving self-declared transparency about applicable law. What remained, at EU level, were *behavioural screens*: where the data sits, who staffs the service, which operational certifications the supplier holds. The market's answer is on the public record: Microsoft's EU Data Boundary satisfies residency requirements in full while its own France director testified to the French Senate (10 June 2025) that he could not guarantee French data would never be disclosed under a US CLOUD Act order; AWS's "European Sovereign Cloud" (GA January 2026) offers EU incorporation, EU staff, and EU leadership under a US ultimate parent.

The pairing is known to be strange. The political-science literature has documented each episode (the Toolbox's adoption, EUCS's reversal), and a recent research programme calls the gap between sovereignty rhetoric and policy substance a general puzzle (Falkner et al. 2024); the closest study of cloud certification notes that hyperscalers themselves invoke "Huawei-type scenarios" as the risk

their European customers fear, and explicitly defers the 5G comparison to future research (Blancato and Carr 2026). No account, descriptive or formal, explains why the same Union screened one dependency structurally and declined to screen the other. That is this paper's job.

The lens. We import, without re-deriving, two results from a companion mechanism-design model of sovereignty procurement (Fermigier 2026). First, a *feasibility* result: a requirement screens a hidden property only if it targets an attribute on which the genuine and the non-genuine suppliers' costs differ *and* whose nominal/effective gap the buyer can verify; attributes failing either condition are non-screens outright, at any price and any audit intensity. Second, a *capture* result, from the companion's political-economy extension: the firms a screen would exclude spend influence to degrade it, and the degradation can target the screen's *design* (which attribute is verified, at what accuracy) or its *enforcement* (whether the verification actually runs). These two results generate the paper's organizing objects: for each episode, (i) the decisive attribute and its verification accuracy α (ii) the type-cost gap the screen exploits; (iii) the coalition configuration on each side; (iv) the *locus* of capture (design or enforcement) and why capture chose it.

The explanation, in brief. On the feasibility channel, the two cases sit at opposite ends of the verification spectrum. For a Chinese state-linked vendor, the decisive attribute, state control, is *presumed and public*: no subsidiary structure hides it, no audit is needed to establish it; $\alpha \approx 1$ at near-zero cost, and the type-cost gap is total (a state-linked vendor cannot un-link itself). For a US hyperscaler, the decisive attribute, effective control despite EU-incorporated subsidiaries, must be verified against sophisticated legal engineering; and the attributes that *are* cheap to verify (residency, operations) are precisely those on which hyperscalers outperform everyone, i.e., non-screens by the companion's washing theorem. Structural screening of Chinese vendors was administratively trivial; structural screening of US cloud required exactly the verification apparatus that EUCS's structural criteria would have created; that apparatus is what was removed.

On the capture channel, the coalitions differ in composition, stakes, and available strategies. The pro-Huawei coalition existed (European telecom operators facing replacement costs, joined in Germany by parts of the security bureaucracy), but it was a coalition of *customers* with compensable, transient losses, facing a security-framed, highly salient design decision. It *lost at the design stage* (the Toolbox adopted origin criteria intact) and *won at the enforcement stage*: by mid-2023, 24 of 27 member states had legal powers but only 10 had imposed restrictions; Sweden banned outright (October 2020), France ran sunset licences under its 2019 law, Germany reached what a close observer calls a "largely symbolic" arrangement in July 2024 (Rühlig 2025). The pro-hyperscaler coalition is of a different order: the suppliers themselves, their European partner and reseller ecosystems, a dependent domestic industry, and a bloc of member states, all facing a low-salience, technically-framed *certification* decision. It *won at the design stage outright*: the EUCS criteria were deleted before any audit ever ran.

Why capture chose different stages is the paper's analytical contribution beyond assembling the record. A structural screen with $\alpha \approx 1$ is *publicly defensible*: its design is a salient, legible political object (banning "high-risk vendors" reads as security policy), so capturing the design is expensive: the capture must happen where discretion lives, in national implementation, quietly. A screen whose decisive attribute is hard to verify offers the opposite cost structure: its design is decided in technical standardization venues of low salience, where deleting an ownership criterion reads as a drafting choice; capture at the design stage is cheap, and enforcement never needs capturing because there is nothing left to enforce. The general statement: *the capturable margin is the one where discretion is high and salience is low; verifiability shifts both.* We state this as a proposition-level claim organizing the cases; its formal development into a theorem belongs to the companion's political-economy extension.

Predictions, and the post-2024 record. The lens yields three refutable predictions. (P1) *Venue migration:* structural criteria, defeated in a capturable certification venue, should reappear in venues where the buyer's discretion is inherent and salience higher: procurement. This is what happened: the Commission's Cloud Sovereignty Framework (October 2025) ranks sovereignty, separately from security certification, as a procurement evaluation tool, already applied in a €180M tender requiring its second-highest level; the proposed Cloud and AI Development Act (COM(2026) 502, 3 June 2026) writes EU ownership and control into law from its third assurance level. (P2) *Enforcement hardening:* where the 5G screen was captured (national implementation), the counter-move should bind enforcement in law; the January 2026 Cybersecurity Act revision proposes exactly that, making high-risk-supplier derisking of mobile networks mandatory and extending a binding risk framework across the NIS2 sectors. (P3) *Within-regime variance:* enforcement-stage capture should track domestic exposure (operator replacement costs and China trade) rather than alliance politics alone; the cross-country record supports this (Walter and Trampusch 2024). We take the confirmation of P1 and P2 by events *postdating* the EUCS reversal as the strongest evidence the lens is capturing something real.

What the explanation is not. The leading alternative is alliance structure: Europe screens the rival and spares the patron. Cross-country evidence indeed makes security dependence on the United States the best single predictor of Huawei bans (Christie, Jakobsen and Jakobsen 2024), and we concede the point for the *level* difference between the two cases. But alliance structure is constant within the EU and across the decade, so it cannot explain the within-EU implementation variance (operator exposure and China trade do), the *time path* of the cloud screen (structural criteria present in drafts for three years, deleted in 2024, re-emerging in procurement instruments in 2025–26), or why the deletion happened at the design stage of a certification scheme rather than anywhere else. Explaining levels is not explaining variation; our channels are built for the variation.

Roadmap. Section 2 states the analytical frame. Sections 3 and 4 reconstruct the two episodes from primary documents. Section 5 develops the two-channel explanation and the design × enforcement matrix. Section 6 states the predictions and confronts them with the post-2024 record. Section 7 addresses alternatives: alliance structure, technical differences, the missing-European-alternative argument. Section 8 concludes.

2. The frame (imported)

The companion paper models a buyer screening suppliers for a hidden property when compliance can be produced *nominally* (a shell, a local region, a certification binder) without being produced *effectively*. Each candidate requirement targets an attribute i characterized by two numbers: the *type-cost gap* g_i (how much more it costs a non-genuine supplier than a genuine one to satisfy the attribute *effectively*) and the *verifiability* α_i of the nominal/effective gap. Three results carry over here, stated without proofs.

- *(R1) Non-screens.* A requirement screens only if $g_i > 0$ *and* $\alpha_i > 0$. If the types can satisfy the attribute at equal cost (residency, operations: the things global suppliers do best), an audit of it detects nothing: the mimic complies genuinely there and fakes only where the types differ. The failure is structural: no price, penalty, or audit intensity repairs it (a mechanism-invariance result).
- *(R2) The worst-type bar.* When the decisive attribute *is* verified, the accuracy required is set by the highest-stake would-be mimic: by the largest pretender's ecosystem rents rather than by the average bidder.
- *(R3) Anti-screening burdens.* A requirement that costs the *genuine* type more than the mimic (fixed compliance overhead that a large incumbent's certification machinery absorbs and a small genuine supplier cannot) widens the mimic's advantage: bulky referentials sort by size instead of by the screened property.

On the political side we import a *distinction*, with no theorem behind it: influence aimed at a screen can target its *design* (which attribute is verified, at what accuracy, in which venue) or its *enforcement* (whether the verification actually runs against particular suppliers). The distinction descends from the capture literature (design-stage capture of the rule vs collusion at the audit), and the companion's political-economy extension leaves the *choice of locus* as an open modeling question. This paper supplies its empirics. Each episode below is read through four objects: the decisive attribute; its (g, α) the coalition that the screen would exclude, with its stakes and structure; and the locus where that coalition spent its influence.

3. Case I. The 5G Toolbox: structural design, captured enforcement

Design. The EU Toolbox on 5G Cybersecurity (NIS Cooperation Group, CG Publication 01/2020, 29 January 2020) instructs member states to assess suppliers on non-technical criteria quoted here from the document: "the likelihood of the supplier being subject to interference from a non-EU country," including "a strong link between the supplier and a government" of a third country; the third country's legislation, "in particular where there are no legislative or democratic checks and balances"; "the characteristics of the supplier's corporate ownership"; and third-country pressure on the "place of manufacturing." In the frame's terms this requires the origin/control attribute: for a Chinese state-linked vendor the type-cost gap is total (a state-linked vendor cannot commercially un-link itself) and the verification accuracy is near one at near-zero cost (the linkage is presumed and public; no legal engineering hides it). The design decision was salient and security-framed: a legible political object.

A coalition against the screen existed and should be named precisely, because a naive telling omits it: European telecom *operators*, facing replacement costs and delays, opposed restrictions, and in Germany parts of the security bureaucracy long resisted a hard line; cross-country work shows the bans that happened were carried by cross-party parliamentary security coalitions *over* initial operator and governmental reluctance (Walter and Trampusch 2024). At the design stage this coalition lost: the Toolbox adopted the origin criteria intact.

Enforcement. It won the next stage. By June 2023, three and a half years in, the Commission's own implementation report counted 24 of 27 member states with legal powers adopted or in preparation but *only 10 having imposed restrictions*, with 3 more underway. The national record diverges exactly as domestic exposure predicts: Sweden banned outright (October 2020); France ran its 2019 authorization law with limited-duration sunset licences; Italy used golden-power hurdles; Germany (the largest market, the deepest operator exposure) arrived only in July 2024 at an arrangement a close observer characterizes as "largely symbolic"; Hungary stayed open (Rühlig 2025: "unitary toolbox, divergent policies"). The screen's design was intact; its bite was discretionary, national, and quiet: the capturable margin.

The counter-move. The response to enforcement capture was to move enforcement *into* design: the Commission's January 2026 Cybersecurity Act revision (IP/26/105) proposes making the "mandatory derisking" of European mobile networks from "high-risk third-country suppliers" binding (explicitly "building on the work already carried out under the 5G security toolbox") and extends a binding risk-and-restriction framework across the eighteen NIS2 critical sectors, with fines up to 7% of worldwide turnover. Binding rules are discretion removed; the proposal is, in the frame's language, the security coalition buying back the enforcement margin at the design stage. It is, predictably, contested.

4. Case II. Cloud sovereignty: captured design

The EUCS trajectory. The candidate structural screen existed on paper for three years: successive drafts of the EU cloud-certification scheme (2021–2023) contained, at the highest assurance level,

immunity-from-extraterritorial-law requirements backed by EU-control criteria. The documented line-up pitted a sovereignty-sceptic member-state bloc against a French-Italian-Spanish (and intermittently German) bloc favouring the criteria, with industry position papers on the record. The fight was over exactly these clauses, and on 22 March 2024 the revised draft *removed them*, retaining self-declared transparency about applicable law. The scheme remains unadopted; the deletion happened in a technical standardization venue of low public salience, before any audit had ever run. In the frame's terms: the decisive attribute (effective control behind EU-incorporated structures) has a *large* type-cost gap but a *contested, apparatus-requiring* verifiability (ownership chains, key custody, continuity clauses), and the deletion removed the apparatus at the design stage. Nothing was left to capture at enforcement.

The national exception proves the mechanism. France's SecNumCloud, formally a *security* qualification whose 3.2 revision embeds a page of structural criteria (ownership caps of 24%/39%, ANSSI's "protection vis-à-vis du droit extra-européen," key custody), demonstrates that the structural page, where it survives, sorts: no hyperscaler-parented offering can pass it, while EU-controlled operators qualify. Its limits are equally instructive: the page travels inside a 278-requirement security referential whose burden (18–24 months of instruction, audit costs on the order of €200k per three-year cycle, a 65% success rate; Sénat, rapport n° 830, 2025) prices small genuine suppliers out of demonstrating sovereignty at all (the frame's R3), and qualification of a Thales-owned operator running licensed US technology buys legal access-immunity; it buys neither supply continuity nor code integrity.

The market under behavioural screens. With EU-level requirements reduced to residency and operations (attributes with $g \approx 0$, non-screens by R1), the market's best performers on those attributes duly satisfied them: Microsoft's EU Data Boundary was completed in February 2025 in full residency compliance, and on 10 June 2025 the company's France public-affairs director told a Senate commission of inquiry that he could not guarantee French data would never be disclosed to US authorities; AWS's European Sovereign Cloud went live in January 2026 with EU incorporation, EU staff and EU-citizen leadership under a US ultimate parent. Nominal compliance thrived; the exposure the requirements were named for did not move.

Venue migration. What did move, after the design-stage defeat, was the *venue*. The Commission's Cloud Sovereignty Framework (v1.2.1, October 2025) ranks sovereignty explicitly as a *procurement* evaluation dimension, separate from security certification, with SEAL levels from "No Sovereignty" to "Full Digital Sovereignty… subject only to EU law"; it was applied within months in a €180M EU-institutions tender requiring the second level. The proposed Cloud and AI Development Act (COM(2026) 502, 3 June 2026) writes EU ownership and control into legislation from its third assurance level, the top level barring third-country control including over the software. The structural criteria outlived March 2024 by changing address: they moved from a capturable certification venue to venues where the buyer's discretion is inherent and the political object is legible.

5. The explanation: two channels and a matrix

The feasibility channel. The two decisive attributes sit at opposite ends of the verification spectrum. State linkage of a Chinese vendor is presumed, public, and immune to restructuring: $\alpha \approx 1$, verification cost near zero, type-gap total. A screen on it is administratively trivial and its design is legible to any parliament. Effective control behind a US hyperscaler's EU subsidiaries must be *established* against sophisticated legal engineering (ownership chains, golden shares, key custody, continuity under compulsion), so the screen needs an apparatus, the apparatus needs a mandate, and every element is contestable in drafting. Meanwhile the attributes that are trivially verifiable on the cloud side (residency, operations, certifications) are exactly those with no type-cost gap: non-screens

(R1) that select *for* the most operationally capable non-sovereign suppliers. Feasibility alone thus predicts the observed asymmetry of *default* designs; it cannot, however, explain why criteria drafted and defended for three years were deleted. For that, capture.

The capture channel. The coalitions differ in kind rather than degree. Against the 5G screen stood *customers*: operators with replacement costs that were large, one-off, transient, and compensable (and partially compensated by delay). Against the cloud screen stood the *suppliers themselves* plus their ecosystem: hyperscaler partner networks, a dependent European industry, and a bloc of member states. The stakes are permanent, measured in ecosystem rents (R2′s worst-type logic says these are exactly the stakes that set the bar), and the coalition is organized on both sides of the table.

The matrix, and why capture chose different stages. The organizing regularity:

	Design stage	Enforcement stage
5G / Chinese vendors	Capture **fails**: salient, security-framed, and the criterion's $\alpha \approx 1$ makes both the rule and any weakening of it *legible*	Capture **partially succeeds**: national implementation is discretionary, technical, quiet (until CSA 2.0 attempts to buy the margin back)
Cloud / US providers	Capture **succeeds**: low-salience certification venue; deleting a hard-to-verify criterion reads as a drafting choice	Moot: nothing structural left to enforce

The mechanism we propose, at proposition level: *influence flows to the margin where discretion is highest and salience lowest, and verifiability moves both.* A high-α criterion is legible, so weakening it at the design stage is visible and expensive; the discretion pocket is enforcement. A low-α criterion lives or dies in technical drafting where salience is minimal; design capture is cheap and sufficient. The corollary explains the counter-moves as well as the captures: hardening 5G enforcement into binding law (CSA 2.0) *removes discretion*, relocating the fight to a legible design stage the security coalition can win; migrating cloud criteria into procurement (CSF, CADA) *raises salience and changes the pivotal players*: a buyer scoring a tender is not a standardization committee. Formal development of locus-endogenous capture belongs to the companion's political-economy extension; this paper's contribution is to show the record demanding exactly that model.

6. Predictions and the post-2024 record

The lens was, in effect, calibrated on data through early 2024 (both design outcomes). Three predictions follow, and events since then bear on each.

- *(P1) Venue migration.* Structural criteria defeated in a capturable venue should reappear where discretion is inherent and salience higher: procurement and legislation. *Confirmed by events postdating the calibration window:* CSF (October 2025) and its €180M SEAL-2 tender (April 2026); CADA (June 2026) with ownership-and-control tiers in law. *Live falsifier:* if CADA's structural tiers are deleted in trilogue (the same fate as EUCS's criteria, in a higher-salience venue), the salience half of the mechanism fails and we will say so.
- *(P2) Enforcement hardening.* Where capture won at enforcement, the counter-move should bind enforcement into design. *Confirmed:* the January 2026 Cybersecurity Act revision (mandatory mobile-network derisking; binding framework across 18 sectors). *Live falsifier:* CSA 2.0 dying quietly in Council would show the enforcement margin cannot be bought back: capture holding both stages.

- *(P3) Within-regime variance.* Enforcement-stage outcomes should track domestic exposure (operator replacement costs, China trade); alliance position is constant within the EU and cannot explain the variance. *Confirmed in the cross-country record* (Walter and Trampusch 2024; the implementation map of Section 3).

7. Alternative explanations

Alliance structure. The strongest rival: Europe screens the rival and spares the patron. Cross-nationally, security dependence on the United States is indeed the best single predictor of Huawei postures (Christie, Jakobsen and Jakobsen 2024), and we concede the *level* difference between the cases to it without resistance. But alliance position is constant within the EU and across the decade; it does not explain the within-EU enforcement variance (P3), the three-year presence of structural cloud criteria in drafts before their deletion, or their re-emergence in 2025–26 instruments while the alliance stood unchanged. A constant explains no variation, and the variation is where the information is.

Technical asymmetry. Network equipment executes code inside critical infrastructure; cloud "merely" holds data; perhaps the risk difference explains the design difference. We read this argument as *inside* the frame rather than against it: whatever raises the perceived stakes of the 5G case raises the design stage's salience, which is one of the matrix's two coordinates. And the technical argument alone predicts nothing about venue migration or enforcement capture, which is where the record's action is. (It also understates the cloud exposure: the update stream of licensed closed software is code execution too, the companion's code-integrity channel.)

No European alternative. Rone (2024) attributes EUCS's failure partly to Europe's limited cloud capabilities: with no domestic hyperscaler, an exclusionary screen threatened buyers with having nothing to buy. Through the companion's welfare lens this is the outside option shrinking the region where a mandate is worth its cost, which *raises the return to capture* (a screen that hurts buyers is cheaper to argue down). We treat it as reinforcing the capture channel rather than competing with it, with one observable implication: as EU-controlled capacity grows (the qualified SecNumCloud set, sovereign-cloud entrants), the capture channel's cheapest argument weakens, consistent with the criteria's 2025–26 re-emergence.

The bundling misplacement. Part of the EUCS outcome is architectural: sovereignty criteria were embedded in a *security certification* scheme (a venue built for a different objective, with a different community and low salience), so the design-stage fight took place on capture-friendly terrain. The unbundled successor (CSF: sovereignty ranked separately from security certification) is both the remedy and further confirmation that venue is a choice variable in these fights rather than scenery.

8. Conclusion

One Union, one decade, two opposite screens; the mystery dissolves once two questions are asked of each episode: *how verifiable is the decisive attribute*, and *where could the losing coalition spend its influence*. Structural screening of Chinese vendors was cheap to verify and expensive to fight in daylight, so it passed at design and was eroded in national enforcement; structural screening of US cloud required a verification apparatus that could be quietly amputated in a standardization committee, so it never reached enforcement at all. The design × enforcement matrix that organizes this record is, we think, exportable (investment screening, research security, export controls face the same locus-of-capture logic), and it hands the formal literature a concrete demand: a model in which the *stage* of capture is the lobbyist's choice variable, disciplined by the verifiability and salience of the rule itself. The companion programme takes that as its next assignment. The account is falsifiable on its stated terms (CADA's structural tiers and the CSA 2.0 proposal are its live tests), and if the record moves against it, the matrix, unlike the requirements it studies, is cheap to verify.

The design-and-enforcement matrix of Section 5 summarizes the record in four cells: capture failed at 5G design and partially succeeded at its national enforcement; it succeeded at cloud design, leaving nothing structural to enforce.

References

(All items verified online 2026-07-03 against Crossref/RePEc/publisher/official records; see the project's verification log.)

- Blancato, F. G. & Carr, M. (2026). The trust deficit: EU bargaining for access and control over cloud infrastructures. *Journal of European Public Policy* 33(3):936–967. https://doi.org/10.1080/13501763.2024.2441418
- Christie, Ø. S., Jakobsen, J. & Jakobsen, T. G. (2024). The US way or Huawei? An analysis of the positioning of secondary states in the US–China rivalry. *Journal of Chinese Political Science* 29(1):77–108. https://doi.org/10.1007/s11366-023-09858-y
- Clayton, C., Maggiori, M. & Schreger, J. (2026). A framework for geoeconomics. *Econometrica* 94(1):105–136. https://doi.org/10.3982/ECTA23206
- Dinc, I. S. & Erel, I. (2013). Economic nationalism in mergers and acquisitions. *Journal of Finance* 68(6):2471–2514. https://doi.org/10.1111/jofi.12086
- Falkner, G., Heidebrecht, S., Obendiek, A. & Seidl, T. (2024). Digital sovereignty — rhetoric and reality. *Journal of European Public Policy* 31(8):2099–2120. https://doi.org/10.1080/13501763.2024.2358984
- Fermigier, S. (2026). Screening for sovereignty: optimal procurement when compliance is cheap to fake and costly to verify. Working paper (companion).
- Friis, K. & Lysne, O. (2021). Huawei, 5G and security: technological limitations and political responses. *Development and Change* 52(5):1174–1195. https://doi.org/10.1111/dech.12680
- Lysne, O., Elmokashfi, A., Schia, N. N., Gjesvik, L. & Friis, K. (2019). Critical communication infrastructures and Huawei. TPRC47 conference paper, SSRN 3426222.
- Paine, F., Townsend, R. R. & Xu, T. (2025). Should governments restrict foreign investments in startups? NBER Working Paper 33803.
- Rone, J. (2024). "The sovereign cloud" in Europe: diverging nation state preferences and disputed institutional competences in the context of limited technological capabilities. *Journal of European Public Policy* 31(8):2343–2369. https://doi.org/10.1080/13501763.2024.2348618
- Rühlig, T. (2025). The "Huawei saga" in Europe revisited: German lessons for the rollout of 6G. *Notes du Cerfa* 187, Ifri, June 2025.
- Walter, M. F. & Trampusch, C. (2024). Economic statecraft by design and by default: the political economy of the 5G-Huawei bans in the United States, United Kingdom and Germany. *Competition & Change* 28(5):539–560. https://doi.org/10.1177/10245294241262087

Official documents and primary sources

- NIS Cooperation Group (2020). *Cybersecurity of 5G networks — EU toolbox of risk mitigating measures.* CG Publication 01/2020, 29 January 2020.
- European Commission (2023). *Communication — implementation of the 5G cybersecurity toolbox.* C(2023) 4049 final, 15 June 2023.
- European Commission (2025). *Cloud Sovereignty Framework*, v1.2.1, 20 October 2025; and news item "Commission advances cloud sovereignty through strategic procurement," 17 April 2026 (Cloud III DPS: €180M over six years, SEAL-2 minimum).
- European Commission (2026). *Proposal for a Regulation on harmonised rules for cloud and artificial intelligence development (Cloud and AI Development Act).* COM(2026) 502, 3 June 2026, procedure 2026/0138(COD).

- European Commission (2026). Press release IP/26/105, 20 January 2026 (revision of the Cybersecurity Act and cybersecurity package).
- Reuters — Foo Yun Chee (2024). “EU drops sovereignty requirements in cybersecurity certification scheme, document shows,” 3 April 2024 (EUCS draft dated 22 March 2024).
- ANSSI (2022). *Prestataires de services d’informatique en nuage (SecNumCloud) — référentiel d’exigences*, version 3.2, 8 March 2022.
- Sénat (2025). *L’urgence d’agir pour éviter la sortie de route : piloter la commande publique au service de la souveraineté économique.* Rapport n° 830 (2024–2025), commission d’enquête sur la commande publique, S. Uzenat (prés.), D. Wattebled (rapp.), déposé le 8 juillet 2025.
- Sénat (2025). Audition de Microsoft France (A. Carniaux, P. Lagarde), commission d’enquête sur la commande publique, 10 juin 2025.