An Open Internet Stack: What to Build, in What Order, and Why

Stefane Fermigier

Draft working paper v0.6, 2026-07-17. Measurement and policy analysis on the programme's published instruments; every computed number is reproduced by the verification suite from the repository's data, and every market and macro figure carries a fetched public source with URL, access date, and confidence flag (`data/market-anchors-v1.csv`); the taxonomy carries one revision by its author (Block 15, Section 4.4); the per-euro intervention ranking is anchored at instrument-family level from public cost records, and the block-level ranking awaits per-block cost data.

Abstract. Building-block lists for European digital infrastructure name what the stack needs; they cannot say which block to fund first, or with which instrument. This paper converts one such list, the fourteen Technological Building Blocks (TBBs) of Fermigier (2025), into a criterion that discriminates block by block, composing measured structure with sourced market economics: each block is scored on a 52-category dependency graph with typed jurisdictional-reach edges, an exact cut calculus, service-level exit-friction scores, and a four-class rule (adequate, thin, captured, absent) frozen before any number was computed. Each block also carries its economic rationale and, where public data exists, a sourced European market size and EU-provider share (45 anchor rows, each with source, URL, access date, and confidence flag; eight of thirteen scored blocks have a published market size, ten of thirteen a share or proxy, and the remaining cells say "no published figure found").

Four results. First, the classification discriminates where the money is: six blocks classify captured (identity, cybersecurity supply, operations tooling, development tooling, market platforms, end-user computing and edge), and where a share is published for a captured block, non-EU providers hold from 60 percent of the market (cybersecurity), through about 61 percent of euro-area card payments (the international card schemes), to effectively all of it (operating systems and app-store distribution). The seven thin blocks include Europe's largest measured positions (telecom services EUR 407bn, cloud EUR 61bn at 15 percent EU-provider share, chip demand USD 54bn against roughly 10 percent EU production share). Exactly two captured verdicts rest entirely on jurisdictional reach into forkable open-source substrate: stewardship gaps, priced by public cost anchors three orders of magnitude below build-or-buy.

Second, contracting the dependency graph onto the blocks gives the stack an economic geometry, computed and cross-checked in the suite: seven blocks are net suppliers, with dependencies concentrated on four of them, the silicon chain, the development toolchain, operations tooling, and market platforms (weighted in-degrees 15, 14, 12, 11 against out-degrees of 1 to 3). The application and AI blocks are pure consumers (in-degree 0), cloud is the hinge (in 11, out 18), and the platform block's supplier position is market access (10 of its 11 inbound incidences are distribution edges), locating the rent at the door to users. The demand layer's measured edges concentrate on cloud and applications, touch identity exactly once (one enterprise-demand edge), and never land on AI foundations, operations, or development tooling, so procurement sees the top of a cone whose exposure arrives by composition.

Third, the downstream stake is macroeconomic and sourced: 80 percent of European business cloud and software spending flows to US vendors (EUR 265bn per year, Asterès 2025), and open-source software carries a Commission-estimated EUR 65 to 95bn of economy-wide EU GDP impact (Blind et al. 2021), drawn into the stack through the unpriced substrate of the two stewardship-gap blocks.

Fourth, the verdicts are stable under a 500-draw perturbation of every ordinal coding the rule reads, and a pre-stated reach-edge perturbation locates the residual weight: four of the six captured verdicts ride on one or two single-coder reach edges and track edge survival.

Beneath all of this sits a uniform floor: every block has option number $N = 1$ and coercion number $\varkappa = 1$, and five blocs (the United States, China, Japan, Korea, Taiwan) are each solo cutters for every block, because the unique minimal trace crosses the semiconductor floor. No block can be adequate in isolation, per-block checklists cannot certify the stack, and the pre-stated above-floor variant shows the five-bloc structure is the floor's contribution, with the United States the sole residual solo cutter everywhere except the end-user block, which China also cuts. The classes map onto instrument families (build-or-buy, substrate stewardship, demand aggregation and maintenance, floor capacity), each mapping conditional on compute-floor policy; public cost records give the per-euro criterion coarse family-level anchors in pairwise disjoint ranges (from about EUR 10^5 per stewardship engagement to over 10^10 for floor capacity), and the block-level ranking awaits per-block cost data.

JEL: F52, L86, O25, D85. **Keywords:** digital sovereignty, open internet stack, technological building blocks, dependency networks, market structure, option sets, open-source software, public procurement.

Executive summary

This paper scores the Technological Building Blocks of European digital infrastructure (Fermigier 2025) on a measured 52-category dependency graph under a rule frozen before computation, and puts sourced market economics behind each block. Every computed figure below is reproduced by one verification script; every market figure carries a fetched public source with URL, access date, and confidence flag.

Findings. (1) The criterion discriminates, and the discrimination lines up with market structure. Six blocks classify captured (identity; cybersecurity supply; operations tooling; development tooling; market platforms; end-user computing and edge) and seven thin; none is adequate. Where a share is published for a captured block, non-EU providers hold 60 to 70 percent (cybersecurity, ECSO), about 61 percent of euro-area card payments with 13 countries fully reliant (ECB), and effectively the entire operating-system and app-store layer (StatCounter; DMA designations). The thin blocks hold Europe's largest measured market positions: telecom services EUR 407bn (2024), cloud EUR 61bn (2024) at a 15 percent EU-provider share against 70 percent for the three US hyperscalers, chip demand USD 54bn (2025) against roughly 10 percent EU production share. Captured is a measurement class of the frozen rule; it describes measured exposure and alleges no intentional act of control. (2) The block-level contraction of the dependency graph (new at v0.5, cross-checked against the category graph) shows where dependencies concentrate: seven blocks are net suppliers, four of them carrying most of the inbound weight (the silicon chain, development toolchain, operations tooling, and market platforms); applications and AI are pure consumers; cloud is the hinge that supplies the application layer while consuming every foundation block. The platform block supplies nothing in production terms: its supplier position is distribution, the priced door between applications and users. The demand layer's measured edges never touch AI foundations, operations, or development tooling, and touch identity exactly once (one enterprise-demand edge): what procurement never sees directly, it consumes by composition. (3) Two of the six captured verdicts (operations and development tooling) rest entirely on jurisdictional reach into open-source substrate that is already forkable: stewardship problems, misread as build problems until forkability is priced, and the cheapest anchored gap class on public cost records. (4) The downstream stake is sourced: 80 percent of European business cloud and software spending goes to US vendors (EUR 265bn per year), price rises on the locked base are draining a further EUR 93bn per year by the 2026 estimate (Asterès for Cigref, both flagged as commissioned studies). Open-source software carries an estimated

EUR 65 to 95bn of economy-wide EU GDP impact (Commission study), drawn into the stack through the unpriced commons substrate of blocks 7 and 8. (5) All classifications hold under a 500-draw perturbation of every ordinal coding; the pre-stated reach-edge perturbation (v0.4) shows four of six captured verdicts riding on one or two unreplicated reach edges. (6) The floor, in one paragraph: every block has $N = 1$ and $\varkappa = 1$, and five blocs are each solo cutters for every block through the semiconductor floor, so no per-block checklist can certify the stack; above the floor every block still hangs on a single US-severable path. This is the standing caveat on everything above, and block-level instruments and floor policy are complements.

Instruments, mapped to measured gaps. No instrument on this list delivers sovereignty by purchase or decree; the matrix states which gap each instrument class addresses and claims no more. Build-or-buy addresses the captured blocks with no forkable substrate (4, 5, 10, 15); the identity layer compounds worst, with the highest measured exit friction (mean 0.77). Substrate stewardship (European hosting of forges and distribution, governance participation, maintained fork capability) addresses the two commons-captured blocks (7, 8). Demand aggregation and maintenance funding address the seven thin blocks, whose unreached European categories name the assets procurement can compound. Capacity investment addresses the floor (2a), and conditions everything else.

Costs, coarsely. Public records price the four families in pairwise disjoint ranges: substrate stewardship at EUR 10^5 to 10^6 per component engagement (Sovereign Tech Agency published amounts), demand-side instruments at 10^7 to 10^8 (a completed state migration budget; the Commission's sovereign-cloud tender ceiling), the build-or-buy anchor at 10^9 (IPCEI-CIS public envelope), floor capacity at 10^10 (Chips Act State-aid decisions). These are envelopes and ceilings for instruments aimed at gaps. Whether a gap actually moves is a separate, unpriced event, and the ordering holds at family level only. The block-level per-euro ranking is specified and waits for per-block cost data (Section 9).

---

1. Introduction: lists do not rank

Europe now has several careful enumerations of what a sovereign digital stack requires. The Technological Building Blocks position paper names fourteen blocks, eleven in three layer groups plus three cross-cutting concerns; the EuroStack Initiative Foundation publishes a layered stack model; the CNCF landscape catalogues the cloud-native supply at product granularity; the Sovereign Tech Agency maintains a funded portfolio of critical open-source components. Each answers the question "what belongs to the stack." None answers the questions a budget forces: which block first, with which instrument, and why this block before that one.

The gap is structural, no matter how good the list, and it is economic before it is legal. A list scores nodes: a block is present or missing, funded or unfunded. The quantities that decide funding order live elsewhere: in the market structure of each block (who supplies it, at what concentration, with what switching costs), in the dependency position of each block (what rests on it, what it rests on), and in the paths that connect them, because a block whose every realization runs through a foreign-controlled operating system, identity layer, or fab is exposed however well the block itself is staffed and funded. The companion theory paper proves the path half of this gap is irreducible for per-layer tests: a stack can satisfy a supplier-plurality checklist at every layer while a single foreign actor can cut every end-to-end realization (Fermigier 2026c, the certification theorem). A credible recommendation must therefore compose market evidence with a measured graph, an option calculus, and pre-stated thresholds.

This paper supplies that conversion for the TBB taxonomy, in both directions. Downward, it crosswalks the fourteen blocks onto the 52-category compute core of the programme's measured dependency graph (Fermigier 2026a), computes per block the option number, the coercion number, jurisdictional reach exposure, forkability, and service-level exit friction where the block is service-shaped (Fermigier 2026b), classifies each block under a rule frozen before computation, and contracts the category graph onto the blocks to obtain the block-to-block dependency structure (Section 7). Upward, it grounds each block economically: why the block is a distinct economic object (Section 2), what its European market is worth and who holds it, from public sources only (Section 2.2), and what value rests on it downstream (Section 8). The taxonomy is the author's own; where the measured graph disagrees with the taxonomy, the disagreement is recorded as an open revision question, and one such revision has happened: on the crosswalk's numbers, the author added a fifteenth block for end-user computing and edge (Section 4.4). The published position paper keeps its fourteen and is cited as published.

The headline is the discriminating structure. Six blocks classify captured and seven thin, and the captured class lines up with the blocks where published shares put non-EU providers at 60 to 100 percent of the measured market. Two captured verdicts dissolve into stewardship problems once forkability is priced, and the contraction shows which blocks supply the stack, which consume it, and which tax the passage between applications and their users. Beneath the discrimination sits a uniform floor: every block has $N = 1$ and $\varkappa = 1$ with the same five blocs each a solo cutter, which is the certification theorem's empirical bite and the reason block-level certification cannot close the case (Section 6.5 states it in one paragraph and prices what remains above the floor). Section 9 maps the classes to instrument families; the mapping states which measured gap each instrument class addresses and claims no more than that.

2. The blocks as economic objects

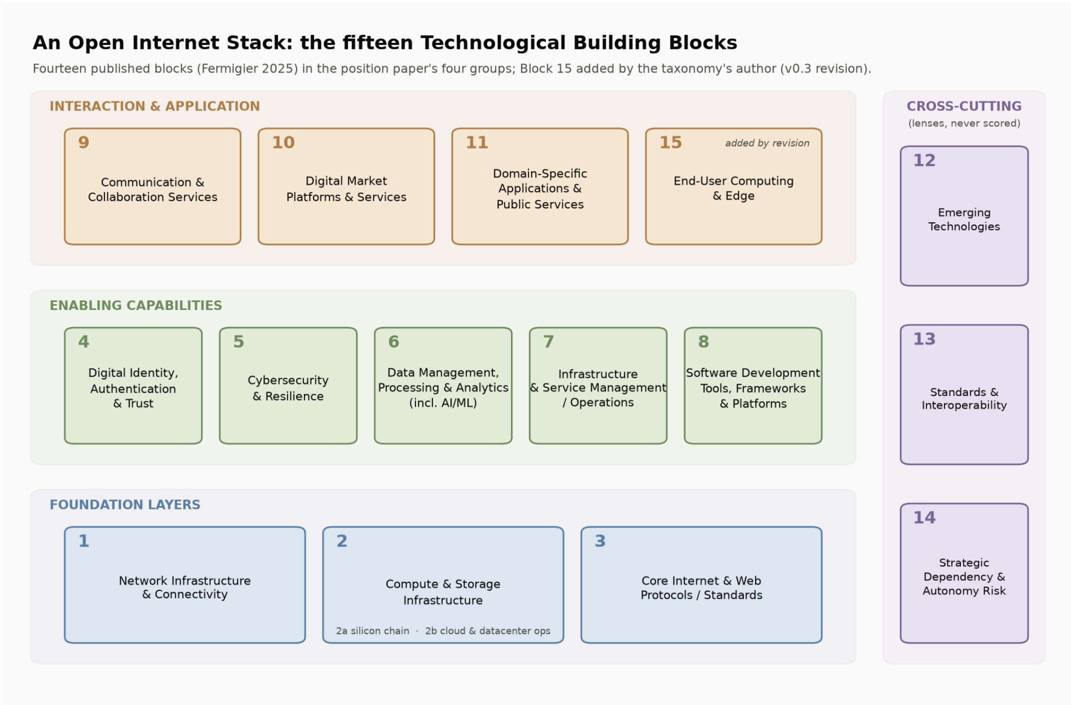

A block earns its place in the taxonomy on economic grounds, on one or more of three: it is a distinct market with its own concentration and switching costs (a rent locus); it is infrastructure in the economic sense, an input that every layer above requires and that repricing propagates through; or it is a locus of market failure that supply and demand alone will not correct (public goods, commons under-provision, natural-monopoly physics). The position paper's four groups sort the blocks by function; this section states the economics per block and then attaches sourced magnitudes. Every figure below carries a row in `data/market-anchors-v1.csv` with source, URL, access date (all 2026-07-17), and confidence flag; figures from commissioned or interested-party sources are flagged as such in their rows.

2.1 Why each block is a block

Foundation Layers (1 to 3). Network Infrastructure & Connectivity (block 1) is the classic regulated infrastructure: high fixed costs, local natural monopoly in access networks, and returns regulated in exchange for universal service. It is also the group's largest European position, and the one layer Europe structurally supplies itself at the service level: European telecoms services turned over EUR 407bn in 2024 inside a EUR 1.09tn digital communications ecosystem, with operator retail revenues of EUR 247bn and capex of EUR 64.6bn (Connect Europe 2026). The dependency question lives below the service layer, in equipment and trust anchors: the top three RAN vendors in Europe are Ericsson, Nokia, and Huawei, and Chinese-vendor 5G RAN reached 41 percent of European mobile subscribers on 2022 installed-base counts (rows b1-ran-global, b1-cn-ran-reach). Compute & Storage Infrastructure (block 2) contains two economically distinct objects, which is why the analysis scores it as two sub-blocks (Section 4.2). The silicon chain (2a) has the most extreme scale economies in manufacturing: a single leading-edge fabrication plant is an asset above EUR 10^10, global supply is an oligopoly at every layer, and the EU holds roughly 10 percent of global production against a 20-percent-by-2030 target, with nearly 80 percent of the suppliers it depends on headquartered outside the EU (European Commission 2026c). European chip demand was USD 54.1bn in 2025, about 7 percent of the world market (ESIA 2025), and the one EU-controlled chokepoint, lithography, alone turned over EUR 32.7bn (ASML FY2025, row b2a-asml). Cloud and datacenter operations (2b) is an increasing-returns business (scale, data gravity, service breadth): the European market reached EUR 61bn in 2024, growing around 24 percent a year, with the three US hyperscalers at 70 percent and European providers at 15 percent, a share that halved from 29 percent in 2017 and has been flat since 2022 (Synergy Research Group 2025). The CMA's market investigation measured the behavioral print of that structure, fewer than 1 percent of customers switching per year and the two leaders earning sustained returns above their cost of capital (CMA 2025). Core Internet & Web Protocols/Standards (block 3) is a public-goods block: standards are nonrival rule-books whose value is adoption, and the web platform that implements them is supplied as zero-price products funded by adjacent rents (search defaults, advertising). There is no market size by construction; the measurable structure is usage, Chrome at 62.7 percent of European page views and no EU-headquartered browser-engine vendor anywhere in the table (row b3-browser-share, derived from StatCounter's June 2026 shares).

Enabling Capabilities (4 to 8). Digital Identity, Authentication & Trust (block 4) is a mandatory complementary input: every application above it requires authentication, marginal cost is near zero, and switching costs are the highest this programme has measured anywhere (mean F_exit 0.77 on the crosswalked services, Section 6.2), which makes identity a textbook lock-in locus; the only published Europe-specific market size is USD 6.1bn for 2025 from a low-confidence analyst source, no EU-provider share is published at all, and the public-option response, the EU Digital Identity Wallet, carries a Member-State availability deadline of end-2026 (rows b4-iam-market, b4-eu-share, b4-eudi). Cybersecurity & Resilience (block 5) is a risk market with information asymmetries: buyers cannot observe product quality, certification substitutes for observation, and demand is partly created by regulation (NIS2, DORA); the EU market is about EUR 30bn (2024) on the ECSO estimate, and 60 to 70 percent of the European civilian market is owned by non-European players (ECSO 2024). Data Management, Processing & Analytics including AI/ML foundations (block 6) runs on scale economies in data and compute: the EU27 data market was EUR 90bn in 2024 inside a EUR 579bn data economy (4.4 percent of GDP), European AI spending is forecast to approach USD 290bn by 2029, and while no EU-provider share of the model layer is published, more than 85 percent of GenAI value-chain funding goes to US companies (European Data Market Study 2025; rows b6-ai-spend, b6-genai-funding). Infrastructure & Service Management/Operations (block 7) and Software Development Tools, Frameworks & Platforms (block 8) are the commons blocks: their substrate is largely open-source software, a shared, unpriced input that is nonrival, nonexcludable, and maintained by a concentrated set of maintainers, the textbook under-provision configuration. Their market invisibility is itself a finding: no defensible European market size is published for either block (rows b7-market, b8-market), because the substrate's value never passes through a price. The proprietary margin and its rents carry a price: VMware held about 45 percent of the virtualisation market when Broadcom acquired it, and European cloud providers and buyers have since reported licence costs rising by a factor of 10 or more at renewal (CISPE and ECCO complaint documents, flagged as interested-party evidence, row b7-vmware-share). The development-tool forge layer concentrates on one US-controlled platform with 180 million-plus accounts (GitHub Octoverse, row b8-github).

Interaction & Application (9 to 11, plus 15). Communication & Collaboration Services (block 9) is a network-effects market locked through formats and identity coupling: Microsoft holds 73 to 80 percent of EU public-sector productivity software at EU level on the one published measurement (Compass Lexecon for the Open Cloud Coalition, a study commissioned by Microsoft's competitors, flagged as such), and the Commission's Teams decision records a preliminary finding of dominance in SaaS productivity applications worldwide (European Commission 2025). Digital Market Platforms & Services (block 10) is the two-sided-platform block: app stores, marketplaces, and payment rails price the passage between producers and users, the intermediation-rent configuration. The DMA designations put six gatekeepers and 22 core platform services on the list, none EU-headquartered at designation (European Commission 2023b), international card schemes process about 61 percent of euro-area card payments with 13 euro-area countries entirely reliant on them (ECB 2025), and the European B2C e-commerce turnover this block intermediates was EUR 842bn in 2024 (row b10-ecommerce). Domain-Specific Applications & Public Services (block 11) represents the meeting of digital value and heterogeneous sector demand (verticals, public services); no EU-wide market size is published, and the scale anchor is Europe's largest software vendor, SAP, at EUR 36.8bn total revenue in FY2025 (row b11-sap). End-User Computing & Edge (block 15, the revision's addition, Section 4.4) is the default-setting layer: client operating systems set the defaults for everything above them, and the measured structure is a US duopoly at 99.97 percent of European mobile OS usage, with Windows at 60.7 percent of desktop and community Linux, at 4.7 percent, the only non-US-vendor entry; European IoT spending, the block's edge half, was about USD 260bn in 2024 with no published EU-provider share (rows b10-15-mobile-os, b15-desktop-os, b15-iot-spend).

Cross-cutting concerns (12 to 14). The three lenses (Emerging Technologies; Standards & Interoperability; Strategic Dependency & Autonomy Risk) are methods for reading the substantive blocks and carry no market of their own; they are never scored (Section 3.4), and lens 14 is, in substance, the apparatus this paper applies.

2.2 The market table

#	TBB	European market size (anchor)	EU-provider share, or nearest published proxy	conf.
1	Network Infrastructure & Connectivity	telecom services EUR 407bn (2024, Connect Europe)	services EU-supplied by structure; equipment: Europe top-3 Ericsson, Nokia, Huawei; Chinese-vendor 5G RAN reach 41% (2022)	high / medium

#	TBB	European market size (anchor)	EU-provider share, or nearest published proxy	conf.
2a	Compute & Storage: silicon chain	chip sales into Europe USD 54.1bn (2025, WSTS/ESIA)	EU about 10% of global production; about 80% of suppliers non-EU (Commission factpage)	high
2b	Compute & Storage: cloud and datacenter ops	cloud infrastructure services EUR 61bn (2024, Synergy)	EU providers 15%; AWS+Microsoft+Google 70% (Synergy)	medium
3	Core Internet & Web Protocols/Standards	no market by construction (standards; zero-price browsers)	Chrome 62.7% of EU page views; EU browser-engine share about 0 (derived, StatCounter)	high
4	Digital Identity, Authentication & Trust	Europe IAM USD 6.1bn (2025, analyst source)	no published figure found	low / n/a
5	Cybersecurity & Resilience	EU EUR 30bn (2024, ECSO; definitions range to EUR 67bn)	non-European players 60-70% (ECSO)	medium
6	Data Mgmt, Processing & Analytics (incl. AI/ML)	EU27 data market EUR 90bn (2024, EC study); AI spend to USD 290bn by 2029 (IDC)	no share figure; >85% of GenAI funding to US firms (CEPS/EESC, proxy)	high / medium
7	Infrastructure & Service Mgmt/Operations	no published figure found	VMware about 45% of virtualisation (2023, CISPE, flagged)	medium
8	Software Development Tools & Platforms	no published figure found	GitHub 180m+ accounts (usage proxy, US-controlled forge)	high (proxy)
9	Communication & Collaboration Services	no published figure found	Microsoft 73-80% of EU public-sector productivity (Compass Lexecon/OCC, flagged)	medium
10	Digital Market Platforms & Services	B2C e-commerce context EUR 842bn (2024)	intl. card schemes about 61% of euro-area payments (ECB); app-store/mobile-OS duopoly 99.97% (StatCounter)	high
11	Domain-Specific Applications & Public Services	no published figure found	no share figure; SAP EUR 36.8bn FY2025 global (EU scale anchor)	high (anchor)
15	End-User Computing & Edge	IoT spend USD 260bn (2024, IDC via trade press)	OS layer: US vendors effectively 100% of mobile, Windows 60.7% of desktop; IoT share: no published figure	high / medium

Reading discipline, stated once and binding throughout: scopes differ by row (WSTS Europe, Synergy Europe, EU27, euro area, StatCounter Europe) and are never harmonized, so cross-block ratios computed from this table would mix scopes; usage shares (page views, accounts, installed base) stand in for revenue shares only where no revenue split is published, and are labeled; commissioned and interested-party sources (OCC, CISPE, Asterès) are flagged in their rows and weighted accordingly; forecasts carry their horizon and are never mixed with current-year sizes. Full provenance per cell, verbatim source wording included, is in `data/`

`market-anchors-v1.csv` (45 rows, schema validated by the suite, check M1); the no-figure cells are recorded there as explicit gap rows with the searches that failed. Coverage is partial and stated: eight of thirteen scored blocks carry a published market size, ten of thirteen a share or proxy, and the gaps concentrate exactly where the substrate is unpriced (blocks 7, 8) or the taxonomy is finer than published statistics (blocks 9, 11).

3. The instruments

The paper computes nothing new at the base layer; it composes four published instruments.

3.1 The measured graph
The applied companion (Fermigier 2026a) documents a dependency graph of the European digital stack: 52 categories in a computable core with 151 dependency and distribution edges, 37 typed jurisdictional-reach edges (which jurisdiction's legal instruments reach which category), 48 funding and influence edges, and 7 adversarial-threat edges, with per-edge sources and adversarial verification. The core carries six European demand categories, three capital categories, and jurisdiction, standards, talent, channel, and influence layers alongside the supply chain from raw materials to SaaS. All block scoring below runs on this graph read-only; the verification suite reproduces the graph's published totals before computing anything else.

3.2 The cut calculus and the certification theorem
The theory companion (Fermigier 2026c) defines, for a buyer over an AND/OR production structure with controller-labeled options: the option number N, the maximum number of controller-disjoint feasible realizations, and the coercion number κ, the size of the smallest coalition of outside controllers whose joint refusal makes every realization infeasible. Weak duality gives $N \leq \kappa$. The certification theorem is the negative result: a criterion that inspects one layer at a time certifies per-class breadth, and for every $m \geq 2$ there are structures with m independent suppliers at every layer and $N = 1$. Applied to the measured graph, the calculus is exact (the controller universe is small), and the published results include $\kappa = 1$ for every European demand category with a five-bloc solo-cutter list. This paper reuses the calculus unmodified, applying it to a virtual buyer per block (Section 5).

3.3 Service-level exit friction
The cloud-services matrix (Fermigier 2026b) scores 35 AWS services on four judgment axes (architectural coupling, data gravity, protocol proprietary-ness, migration friction) aggregated into an exit tax F_exit in [0, 1], under a two-coder protocol with a robustness suite (2000 seeded draws). Where a TBB is service-shaped, its crosswalked CSM rows supply a friction profile the graph's category resolution cannot see. The scores are read from the frozen published results and recomputed from their components as a cross-check.

3.4 The taxonomy
The TBB position paper (Fermigier 2025) names fourteen blocks: eleven substantive layers under the published group headings "Foundation Layers" (1 to 3), "Enabling Capabilities" (4 to 8), and "Interaction & Application" (9 to 11), and three "Cross-cutting Concerns" (12 to 14), with eight policy recommendations (the document's section 6, "Recommendations for Action"; title, grouping, and count verified against the published document, accessed 2026-07-17). The names of blocks 1 to 11 are used verbatim throughout. The published names of 12 to 14 carry the parentheticals "(Lens)", "(Principle & Practice)", and "(Analysis & Mitigation)" respectively; this paper calls all three lenses for short. The lenses are methods; they are carried as unscored rows, and lens 14 (Strategic Dependency & Autonomy Risk) is, in substance, the apparatus this paper applies.

4. The crosswalk

4.1 The mapping

Each block maps to the compute-core categories that implement it. The mapping was fixed in a working note before scoring and reconciled to the current 52-category core (the core grew from 49 categories by three splits, all in the demand, capital, and influence layers, none affecting block membership).

#	TBB	Compute-core categories
1	Network Infrastructure & Connectivity	eu-connectivity, us-network-hw-trust-infra, gnss-space, datacenter-physical
2	Compute & Storage Infrastructure	us-hyperscalers, eu-cloud, sovereign-wrappers, specialist-foreign-clouds, datacenter-physical; upstream: fab-leading-edge, litho-euv-eu, us-ai-silicon, eu-silicon, memory-hbm, cpu-isa, us-semi-toolchain, raw-materials-china, asia-materials-packaging
3	Core Internet & Web Protocols/Standards	standards-bodies, browsers-web
4	Digital Identity, Authentication & Trust	identity-iam
5	Cybersecurity & Resilience	cybersecurity (adversary carried as threat context, outside the calculus)
6	Data Management, Processing & Analytics (incl. AI/ML foundations)	us-ai-models, eu-ai-models, ai-input-substrate
7	Infrastructure & Service Management/Operations	us-firmware-os-virt, oss-commons-substrate
8	Software Development Tools, Frameworks & Platforms	software-supply-chain, oss-commons-substrate
9	Communication & Collaboration Services	us-saas, eu-saas-oss
10	Digital Market Platforms & Services	us-consumer-distribution, payments-rail
11	Domain-Specific Applications & Public Services	us-saas, eu-saas-oss (vertical segments); public-sector-demand as the reference buyer
12–14	Emerging Technologies; Standards & Interoperability; Strategic Dependency & Autonomy Risk (lenses)	method rows, never scored
15	End-User Computing & Edge (added by revision; Section 4.4)	client-os, iot-edge

4.2 Scoring defaults

Three defaults were adopted before scoring, each then pending a ruling by the taxonomy's authors; the rulings arrived at v0.3. Blocks 10 and 11 are scored at current graph resolution (confirmed as-is, 2026-07-17, revisable when the graph grows marketplace or vertical nodes): the graph resolves no marketplaces or verticals, so a thin verdict there is a finding about the graph-taxonomy mismatch. Block 2 is scored at sub-block resolution as the primary presentation (2a, the silicon chain of nine categories; 2b, cloud and datacenter operations of five), with the block total kept as a summary row (ruled 2026-07-17), because block-level averaging would smear exactly the chokepoints the analysis exists to find. Section 2.1′s economics states the same distinction in market terms. The graph's demand, capital, jurisdiction, knowledge, and influence layers (17 context categories after the Block 15 revision, plus the reference buyer) are reported as out-of-taxonomy context rows: the taxonomy is a supply-side list, and the scoring says so instead of forcing them in. The formerly unplaced client-os and iot-edge categories now form Block 15 (Section 4.4).

4.3 Open questions for the taxonomy's authors, and their status
The crosswalk surfaced four disagreements between taxonomy and graph; the taxonomy's authors own the resolution, and Section 6 attaches the numbers. Status as of v0.3 (rulings of 2026-07-17):

1. Blocks 10 and 11 are thinner in the graph than in the taxonomy; either the graph grows nodes (marketplaces, verticals) or the blocks merge for scoring purposes. Ruled: the as-is default stands, thin is the finding; revisable when the graph grows the missing nodes.
2. The graph's demand, capital, and jurisdiction layers have no TBB home; the present treatment (context rows) preserves them without classification. Standing treatment, unchanged.
3. Block 2 spans the silicon-to-cloud chain; sub-block scoring preserves the chokepoint structure. Ruled: sub-block scoring adopted as the primary presentation, block total kept as a summary row.
4. Two supply-side categories (client-os, iot-edge) had no TBB home; the v0.1 candidate placements (block 7 or 2 for client-os, block 11 or 1 for iot-edge) would each have imported the category's jurisdictional reach into the receiving block. Resolved by revision: both categories now form Block 15 (Section 4.4).

4.4 The taxonomy revised: Block 15
The taxonomy's author, exercising the revision mechanism this section exists for, added a fifteenth substantive block on the crosswalk's evidence: **End-User Computing & Edge**, housing client-os (desktop and mobile operating systems) and iot-edge (IoT devices, industrial IoT, edge compute, embedded systems), the two supply-side categories the fourteen-block taxonomy could not place. The numbers that motivated the ruling are the ones v0.1 attached to the open question: both categories carry jurisdictional reach edges (client-os US-reached, iot-edge China-reached, the only China reach edge outside the semiconductor floor), so any placement inside an existing block would have imported that reach and changed the receiving block's classification. A dedicated block scores the exposure where it lives. The published position paper (Fermigier 2025) keeps its fourteen blocks and is cited as published throughout; Block 15 is this revision's addition, flagged for the position paper's own next version. The frozen rule, the variants, and the perturbation design apply to Block 15 unchanged, and the suite pins every pre-revision result and reproduces it exactly (Annex B).

5. The classification rule, frozen before computation

The thresholds below were written down in the crosswalk note before the scoring script first ran, and were never adjusted afterward; the one variant reading was also pre-stated. The freeze discipline is in-repo attestation: the rule note, the script, and the first results entered the repository history together, so pre-statement is documented by dated note sections and the commit record, which is the witness for every later variant going forward. Per block, the suite computes: the set of cut-calculus categories; foreign reach exposure (count and share of the block's categories reached by United States or China jurisdiction edges; EU regulatory reach is reported as context and never counted as foreign); the option number N_b and coercion number κ_b of a virtual buyer whose sole requirement is the block, with the block's categories grouped into the calculus's functional classes (same class substitutable, distinct classes jointly required) and full upstream subtrees included, on the ownership channel under both the strict and lenient controller codings and on the legal channel; the forkability share (the flagship's binary substrate coding over the block's substrate-coded categories); F_exit summaries over the crosswalked CSM rows; and concentration and capture-risk coding summaries.

The rule, first match wins:

1. **absent**: the block has no cut-calculus category, or its function is infeasible at the full controller budget (the $N = 0$ case). Rationale: a function no one can supply is missing whatever the codings say.
2. **captured**: options exist, and foreign reach covers every category, or concentration sits at its worst level on every category, or capture risk sits at its worst level on every category. Rationale: options that all hang under one foreign legal lever, or all carry worst-level codings, are options in name only.
3. **thin**: $N_b = 1$, or foreign reach covers all but at most one category, or the forkability share is zero with at least one coded category. Rationale: a single controller-disjoint way through the block, a single unreached category, or the absence of any forkable substrate leaves no redundancy to lose.
4. **adequate**: otherwise.

Two companion specifications were frozen with the rule. The sensitivity variant V1 shrinks the reached set before the reach terms are evaluated: a category whose coded substrate is entirely forkable is removed from the reached set, with the block's category count (the denominator) unchanged, because forking severs the reach edge while the category itself remains in the block (an open substrate admits self-hosting, so that reach is severable by forking). The denominator convention carries the finding: removing such categories from the block itself would return blocks 7 and 8 to reach share 1 and leave them captured, inverting Section 6.3, and the suite computes both conventions and asserts the difference. V1 bounds how much of any captured verdict rides on reach into open-source stewardship. The perturbation design (Gate 3): 500 draws, seed 20260716, every ordinal coding the rule reads (concentration, capture risk, substrate) independently re-leveled one step up or down with probability 1/4 each way, graph structure and reach edges fixed; a block whose frozen classification holds in fewer than 80% of draws is flagged fragile.

A third specification, the above-floor variant V2, was pre-stated in the same working note on 2026-07-17, before it was first computed: the frozen thresholds applied unchanged to above-floor inputs, meaning the block's non-floor categories and the option calculus at the source model's above-floor scope (the nine semiconductor-floor layers waived throughout, exactly as the flagship's published above-floor configurations do), with the identical 500-draw stability stream re-used and the frozen v0.1 results pinned by a regression check in the suite. V2 is reported next to the frozen classification and never replaces it (Section 6.5).

Two further specifications were pre-stated at v0.4 (2026-07-17, crosswalk note section 11), after an independent referee round observed that the Gate-3 design perturbs only the ordinal codings and holds the binary reach edges and the block membership fixed, so the captured class faced no perturbation able to move it. The reach-edge perturbation V3 drops each of the graph's 37 jurisdictional-reach edges independently, at drop rates 0.10 and 0.25 (the latter equal to Gate 3′s per-direction re-level probability), 500 draws per rate on a dedicated seeded stream (2026071603) that leaves every existing result unchanged by construction; the frozen rule is then applied with every other input at its frozen value. The leave-one-category-out pass L1 removes each category of every multi-category block in turn and recomputes the full pipeline on the remainder; single-category blocks are reported n/a. Both are reported next to the frozen classification (Section 6.4) and never replace it, and the suite pins the complete v0.3 output as a byte-level regression.

Two additions were pre-stated at v0.5 (2026-07-17, crosswalk note section 12), both descriptive: the block-graph contraction B1 (Section 7), whose rule, membership multi-map, and cross-checks were fixed before computing and which no classification reads; and the market-anchor table (Section 2.2), which the suite validates for provenance schema and which no computed quantity reads. The complete v0.4 output is pinned byte-for-byte as a further regression (check R3), with the matrix CSV and figure byte-identical.

One correction and one pin were pre-stated at v0.6 (2026-07-17, crosswalk note section 13), after a second referee round found the reported cost-anchor "median" to be the upper median: the statistic is recomputed as the conventional median with a suite check (K4), the upper-median values stay computed in the results file under exact names, and the complete v0.5 output is pinned as a further regression (check R4), with the matrix CSV and every figure byte-identical.

6. Results: the classification

6.1 The matrix

TBB scored matrix v1: options, reach, forkability, exit friction (52-node compute core)

TBB block	N	kappa	foreign reach	fork	F_exit (n)	stability	classification
1 Network Infrastructure & Connectivity	1	1	1/4 = 0.25	0.00	n.c.	99%	thin
2a Compute & Storage: silicon chain	1	1	7/9 = 0.78	0.00	n.c.	99%	thin
2b Compute & Storage: cloud and datacenter operations	1	1	3/5 = 0.60	0.00	0.54 (35)	99%	thin
2 Compute & Storage Infrastructure (block total)	1	1	10/14 = 0.71	0.00	0.54 (35)	100%	thin
3 Core Internet & Web Protocols/Standards	1	1	1/2 = 0.50	n.c.	n.c.	81%	thin
4 Digital Identity, Authentication & Trust	1	1	1/1 = 1.00	n.c.	0.77 (2)	100%	captured
5 Cybersecurity & Resilience	1	1	1/1 = 1.00	n.c.	0.67 (1)	100%	captured
6 Data Management, Processing & Analytics (incl. AI/ML foundations)	1	1	1/3 = 0.33	0.00	0.76 (4)	96%	thin
7 Infrastructure & Service Management/Operations	1	1	2/2 = 1.00	0.75	0.63 (3)	100%	captured
8 Software Development Tools, Frameworks & Platforms	1	1	2/2 = 1.00	1.00	n.c.	100%	captured
9 Communication & Collaboration Services	1	1	1/2 = 0.50	0.00	n.c.	81%	thin
10 Digital Market Platforms & Services	1	1	2/2 = 1.00	0.00	n.c.	100%	captured
11 Domain-Specific Applications & Public Services	1	1	1/2 = 0.50	0.00	n.c.	81%	thin
15 End-User Computing & Edge	1	1	2/2 = 1.00	0.00	n.c.	100%	captured

N, kappa: ownership channel, strict coding, full scope (virtual block root, full upstream subtrees). foreign reach: nodes reached by US or China jurisdiction edges.
forkable: share of substrate-coded categories coded forkable; n.c. = no coded category. F_exit: mean over the CSM rows mapped by the crosswalk (n rows). stability: share of 500 seeded perturbation draws (seed 20260716) keeping the frozen class; * = fragile (< 80%). Blocks 12-14 are lenses and are never scored. Rule frozen in crosswalk-v1.md sec. 5.
Sub-blocks 2a/2b are the primary presentation; the block-2 total is a summary row. Block 15 was added by the taxonomy's author (v0.3 revision, crosswalk-v1.md sec. 10); the published position paper remains fourteen blocks.

#	TBB	N	κ	foreign reach	fork	F_exit (n)	stability	class
1	Network Infrastructure & Connectivity	1	1	1/4 = 0.25	0.00	n.c.	99%	thin
2a	Compute & Storage Infrastructure: silicon chain	1	1	7/9 = 0.78	0.00	n.c.	99%	thin
2b	Compute & Storage Infrastructure: cloud and datacenter operations	1	1	3/5 = 0.60	0.00	0.54 (35)	99%	thin
2	*(block total: summary row)*	1	1	10/14 = 0.71	0.00	0.54 (35)	100%	thin
3	Core Internet & Web Protocols/Standards	1	1	1/2 = 0.50	n.c.	n.c.	81%	thin
4	Digital Identity, Authentication & Trust	1	1	1/1 = 1.00	n.c.	0.77 (2)	100%	captured
5	Cybersecurity & Resilience	1	1	1/1 = 1.00	n.c.	0.67 (1)	100%	captured
6	Data Management, Processing & Analytics (incl. AI/ML foundations)	1	1	1/3 = 0.33	0.00	0.76 (4)	96%	thin
7	Infrastructure & Service Management/Operations	1	1	2/2 = 1.00	0.75	0.63 (3)	100%	captured
8	Software Development Tools, Frameworks & Platforms	1	1	2/2 = 1.00	1.00	n.c.	100%	captured
9	Communication & Collaboration Services	1	1	1/2 = 0.50	0.00	n.c.	81%	thin
10	Digital Market Platforms & Services	1	1	2/2 = 1.00	0.00	n.c.	100%	captured

#	TBB	N	$\varkappa$	foreign reach	fork	F_exit (n)	stability	class
11	Domain-Specific Applications & Public Services	1	1	$1/2 = 0.50$	0.00	n.c.	81%	thin
15	End-User Computing & Edge	1	1	$2/2 = 1.00$	0.00	n.c.	100%	captured

N and $\varkappa$ on the ownership channel, strict coding, full scope; foreign reach counts US and China jurisdiction edges at category level; fork is the share of substrate-coded categories coded forkable (n.c.: no coded category maps); F_exit is the mean over the crosswalked CSM rows (row count in parentheses); stability is the share of perturbation draws keeping the frozen class. Sub-blocks 2a/2b are the primary presentation and the block-2 total is a summary row (Section 4.3, ruling 3); Block 15 is the revision's addition (Section 4.4). Tally at the primary presentation: 6 captured, 7 thin (the block-2 summary row is also thin), 0 adequate, 0 absent, 3 lenses unscored.

6.2 Reading the classes, with the market behind them

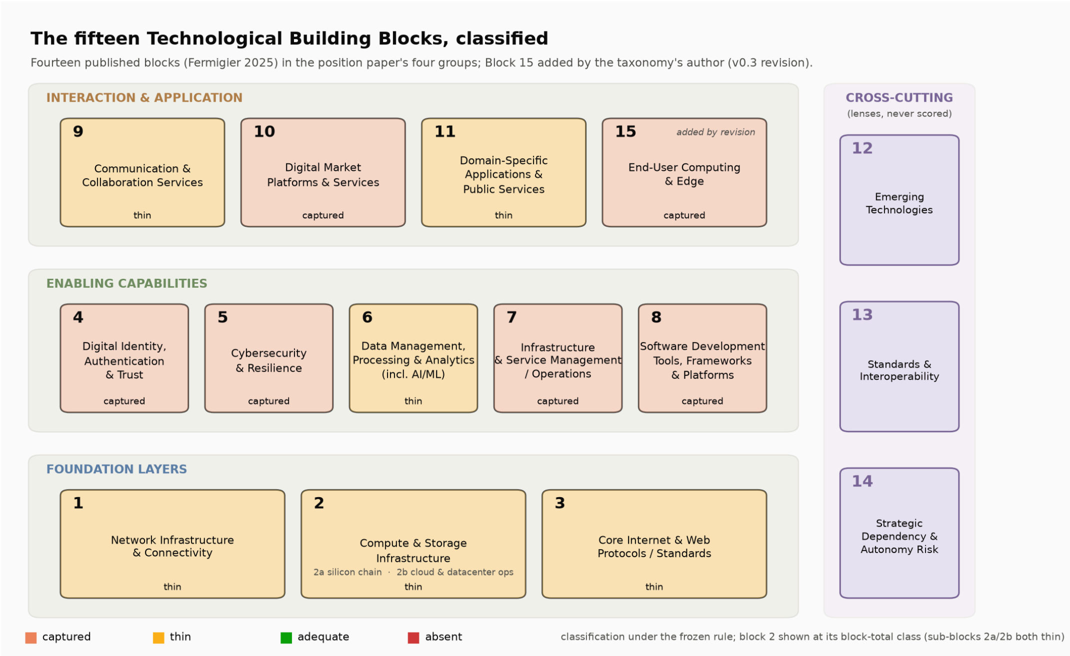

Captured without a forkable option: blocks 4, 5, 10, 15. Identity (block 4) is the compounding worst case: a single category, fully US-reached, concentration at its worst level, and the highest service-level exit friction of any crosswalked block (IAM 0.72, Cognito 0.81, mean 0.77); the market table adds that this mandatory input to every application above it has no published EU-provider share at all. Cybersecurity supply (block 5) is one US-reached category; its verdict is partly an artifact of single-category resolution, and its codings sit at medium. The published market structure matches the class, with non-European players at 60 to 70 percent of a EUR 30bn EU market (ECSO 2024). Market platforms (block 10) is captured on two independent triggers, full reach and worst-level concentration on both categories (app-store distribution and payment rails). The sourced receipts give the same picture from three directions, the DMA's designated gatekeepers (none EU-headquartered at designation), the ECB's 61 percent international-scheme share with 13 fully reliant countries, and the 99.97 percent mobile-OS duopoly that carries the app-store layer. Two categories stand in for the taxonomy's whole platform economy, which is disagreement 1 in numbers. End-user computing and edge (block 15, the revision's addition) is captured at reach 2/2 with a split exposure: the client-OS category is US-reached with concentration at its worst level and a proprietary substrate coding (and a measured usage structure to match, Section 2.1), while the IoT/edge category is China-reached, the only China reach edge outside the semiconductor floor, with European industrial options recorded inside the category under the lenient coding.

Captured through the commons: blocks 7 and 8. Operations tooling and development tooling classify captured because both of their categories, including the open-source commons substrate, carry US jurisdictional reach edges. These are the two verdicts the sensitivity variant flips (Section 6.3), and their forkability shares (0.75 and 1.00) mark them as a different policy object from blocks 4, 5, 10, and 15: the substrate exists and is forkable; the exposure runs through stewardship and distribution. The market table's blank cells for these blocks are the economic complement of the same fact: the substrate is unpriced, so market statistics cannot see it, while the priced proprietary margin next to it is currently demonstrating what repricing a locked base looks like (the VMware renewals, Section 2.1).

Thin: blocks 1, 2a, 2b (and the block-2 total), 3, 6, 9, 11. Each has exactly one unreached path element too many to be captured and no second controller-disjoint realization. The silicon sub-block 2a concentrates the damage: reach 7/9, two of the graph's three China reach edges (the third lands on Block 15′s iot-edge), seven of nine categories at worst-level concentration, and an EU-owned lithography category that is nonetheless US-law-reached. The cloud sub-block 2b (reach 3/5) keeps its unreached categories (eu-cloud, datacenter-physical) as the visible policy surface, and its market row prices the stakes: EUR 61bn growing near 24 percent a year, with the EU-provider share flat at 15 percent since 2022. The block-2 total, the summary row, carries the highest absolute exposure of any row (10 of 14 categories reached, 8 by the US and 2 by China; F_exit mean 0.54 over all 35 services, range 0.10 to 0.90) and shows exactly what the sub-block split exists to prevent: the average smears the chokepoint structure the two sub-blocks display separately. Blocks 9 and 11 are numerically identical (the graph resolves no verticals), which is the rest of disagreement 1. The market evidence differentiates them where the graph cannot, block 9 carrying the published 73-80 percent incumbent share and the dominance finding, block 11 carrying no published market size at all. Block 6 sits at reach 1/3 with heavy service friction where the CSM maps (DynamoDB 0.90, SageMaker 0.77, Bedrock 0.74, Redshift 0.64) and an EU model category coded at worst-level capture risk. Its market rows say the block's stake is large and growing (EUR 90bn data market, AI spend approaching USD 290bn by 2029) while the funding proxy points the model layer's gravity at the US. Block 3′s protocols exist in the graph mostly as edge types; its two categories split into one US-reached (browsers) and one open-participation standards category, and every realization still meets US control through the client-OS and supply-chain subtree.

6.3 The sensitivity finding

Under the pre-stated variant V1 (forkability-discounted reach), blocks 7 and 8 flip from captured to thin and nothing else changes. Two of the six captured verdicts therefore rest entirely on US jurisdictional reach into the open-source commons substrate. Whether reach into forkable stewardship (foundations, forges, registries) constitutes capture in the rule's sense is a modeling question, recorded for the taxonomy's authors; the frozen rule's answer (it does) stands as the headline classification, and both readings are reported. Blocks 4, 5, 10, and 15 are captured under both readings; no forkable category maps to any of them (block 15′s one coded category is proprietary).

6.4 Stability

No block's classification is fragile under the pre-stated threshold: the perturbation leaves every verdict in place in at least 80% of the 500 draws. Three blocks sit at the edge, 81% (3, 9, 11; block 3 at 80.8%), and the mechanism is the same in all three: a two-category block with one category at worst-level concentration flips to captured on any draw that lifts the other category one level. Every observed flip lands on captured; a flip to adequate would require N or reach to move, and the design holds the graph fixed. The six captured blocks and the block-2 total are at 100%: a reach share of one cannot be perturbed away by re-leveling codings.

Those 100% rows are partly a property of the design: the Gate-3 draws re-level codings and hold the reach edges fixed, so a captured verdict resting on reach share 1 faces no draw that could move it. The pre-stated reach-edge perturbation V3 (Section 5) supplies that draw, dropping each of the 37 reach edges independently with every other input frozen. Reach counts can only fall under edge drops, and every block keeps $N = 1$ whatever the edges do (the classification channel reads no reach edge; the suite asserts this), so all seven thin verdicts hold on every draw at both rates and the movement is confined to the captured class. At a 10% drop rate: identity (4) and market platforms (10) hold at 100%, because their worst-level concentration codings classify them captured with or without reach; cybersecurity supply (5) holds at 90.2%, the two tooling blocks (7, 8) at 80.4% and 81.4%, and end-user computing (15) at 81.8%. At a 25% drop rate the same four verdicts fall to 76.4% (5), 55.8% (7), 57.8% (8), and 54.2% (15), all below the 80% line the ordinal perturbation never crossed; stability tracks edge survival, near $1 - p$ for the one-edge block and near $(1 - p)^2$ for the two-edge blocks. Four of the six captured verdicts therefore rest entirely on one or two single-coder binary reach edges that no independent coder has re-coded, and a single miscoded edge would move any of them to thin. Blocks 4 and 10 are anchored by an independent worst-level coding and keep their class under any reach error.

The deterministic leave-one-category-out pass L1 (Section 5) covers the other input both perturbations hold fixed, the crosswalk membership. All six captured verdicts keep their class under every single-category removal, as do the thin verdicts of the multi-category blocks 1, 2a, 2b, 2, and 6. Three thin blocks flip under exactly one removal each: dropping the unreached standards category from block 3, or the unreached European SaaS category from block 9 or 11, leaves a single fully reached category behind, and the remainder classifies captured. Every membership flip lands on captured, mirroring the Gate-3 mechanism, while reach-edge error moves verdicts the other way (captured to thin), so the two error classes bracket the frozen matrix from opposite sides. Single-category blocks (4, 5) admit no removal and are reported n/a.

6.5 The uniform floor, and what remains above it

The floor, in one paragraph. Every scored block has $N = 1$ and $\varkappa = 1$, and the same five blocs (the United States, China, Japan, Korea, Taiwan) are each solo cutters: any one of the five, refusing alone, makes every realization of every block infeasible, because the unique minimal trace contains all five, every feasible way of obtaining every block's function crossing the semiconductor floor. No coordination is required, the adequate class is empty at this resolution, and this is the certification theorem operating on measured data: block-level adequacy is unattainable in isolation, whatever any per-block checklist reports, so the criterion's informative boundary moves down one level, to captured against thin, where the market structure, reach, forkability, and the codings separate the blocks (Section 6.2). Under the lenient controller coding the solo-cutter set is the same minus Japan (which enters only through the strict coding of the instruction-set category); on the legal channel, US law and Chinese law are each solo cutters for every block. The result is true, uniform, and by itself undiscriminating, which is why it stands here as a caveat; its policy content fits in one clause: no block-level instrument in Section 9 closes the floor, and floor policy conditions everything.

Variant V2 (Section 5) prices what remains above the floor: the frozen thresholds, unchanged, on the block's non-floor categories, with the option calculus run above the floor. The result, side by side with the frozen rule:

#	frozen class	V2 class	V2 N	V2 $\varkappa$	V2 solo cutter	V2 reach	V2 stability
1	thin	thin	1	1	US	1/4	99%

#	frozen class	V2 class	V2 N	V2 $\varkappa$	V2 solo cutter	V2 reach	V2 stability
2a	thin	n/a (floor)	·	·	·	·	·
2b	thin	thin	1	1	US	3/5	99%
2	thin	thin	1	1	US	3/5	99%
3	thin	thin	1	1	US	1/2	81%
4	captured	captured	1	1	US	1/1	100%
5	captured	captured	1	1	US	1/1	100%
6	thin	thin	1	1	US	1/3	96%
7	captured	captured	1	1	US	2/2	100%
8	captured	captured	1	1	US	2/2	100%
9	thin	thin	1	1	US	1/2	81%
10	captured	captured	1	1	US	2/2	100%
11	thin	thin	1	1	US	1/2	81%
15	captured	captured	1	1	China, US	2/2	100%

No classification moves, and no block gains an option. Above the floor, every applicable block still has $N = 1$ and $\varkappa = 1$. For every block but one, the solo-cutter set collapses from the five blocs to the United States alone, on all three channels (ownership strict and lenient, and US law on the legal channel). The exception is Block 15: its above-floor trace on the strict coding is {China, US}, each a solo cutter, because the IoT/edge platform layer keeps its contested US/China coding with the floor gone; under the lenient coding (European industrial options and embedded Linux inside the category) its cut is {US} like the rest, and on the legal channel US law and Chinese law each cut it. The silicon sub-block 2a is the floor itself and carries the pre-stated n/a; the block-2 summary row's above-floor category set coincides with sub-block 2b's by construction. Two readings follow. The five-bloc structure is almost entirely the floor's contribution: Japan, Korea, and Taiwan enter every trace through the semiconductor layers and nowhere else, and China's only footholds outside the floor's materials nodes are Block 15′s IoT/edge layer (a strict-coding controller and a reach edge) and nothing else. And conversely, no verdict was an artifact of the floor: the reach shares, the single-option structures, and every thin and captured classification are unchanged by the floor's removal, so the block-specific gaps are their own problem, addressable by instruments faster than capacity investment, while even complete floor autonomy would leave every block one US-severable path deep at this resolution.

6.6 Supplier resolution, where the pilot permits it

The flagship's supplier-resolution pilot (28 adversarially verified options over three layers) intersects three blocks, and there the category-level $N = 1$ can be checked against supplier-level structure; the suite recomputes the pilot's numbers from its data and cross-checks them against the flagship's frozen results. Beside the category resolution: sub-block 2b carries 12 supplier options with controller-independent breadth 9 under both codings (EU- or self-controlled providers counting one apiece, plus pairwise-disjoint foreign controllers); blocks 9/11 carry 7 supplier options with breadth 6, and the composed productivity workload (suite plus hosting) has unbounded $\varkappa$ at supplier resolution because a fully self-controlled realization exists (an EU suite on EU-controlled hosting); the AI-accelerator class of sub-block 2a carries 7 exercisable suppliers with breadth 1 and $\varkappa = 1$ (United States) at any resolution. The direction of the aggregation error is now measured per block: category resolution understates the option structure of 2b and 9/11 (supplier plurality, including one self-controlled composition, exists below the category grain) and is exact for the accelerator class (seven suppliers, one controller). Scope caveats carried from the source: control channel, above the silicon floor, measured classes only. The remaining blocks (1, 3, 4, 5, 6, 10) have no supplier-level cells in the pilot, so nothing is computable there at supplier resolution today; that finding stands in place of a number.

7. The dependency structure at block level

The category graph carries the stack's dependency structure at a grain the taxonomy cannot see; contracting it onto the blocks makes the block-to-block structure a computed object. The contraction rule was pre-stated in the crosswalk note (section 12a) before computing: the primary-presentation blocks with their cut-calculus category sets; membership as a multi-map (four categories sit in two blocks each: datacenter-physical in 1 and 2b, oss-commons-substrate in 7 and 8, us-saas and eu-saas-oss in 9 and 11); an edge contributes weight 1 to the contracted edge A → B for every ordered pair (A, B) in its block image with A distinct from B; relation types carried verbatim; edges with a context endpoint (demand, capital, jurisdiction, knowledge, influence layers) ledgered separately. The suite cross-checks the contraction against the category graph: the four endpoint ledgers partition the 243 raw edges exactly, the 151 spine edges decompose into 26 intra-block, 67 inter-block, and 58 context-touching, the contracted reach ledger reproduces every block's frozen reach counts, and an independent recount reproduces every degree sum (checks B1a to B1d).

Block-to-block dependencies: the 52-category graph contracted onto the blocks

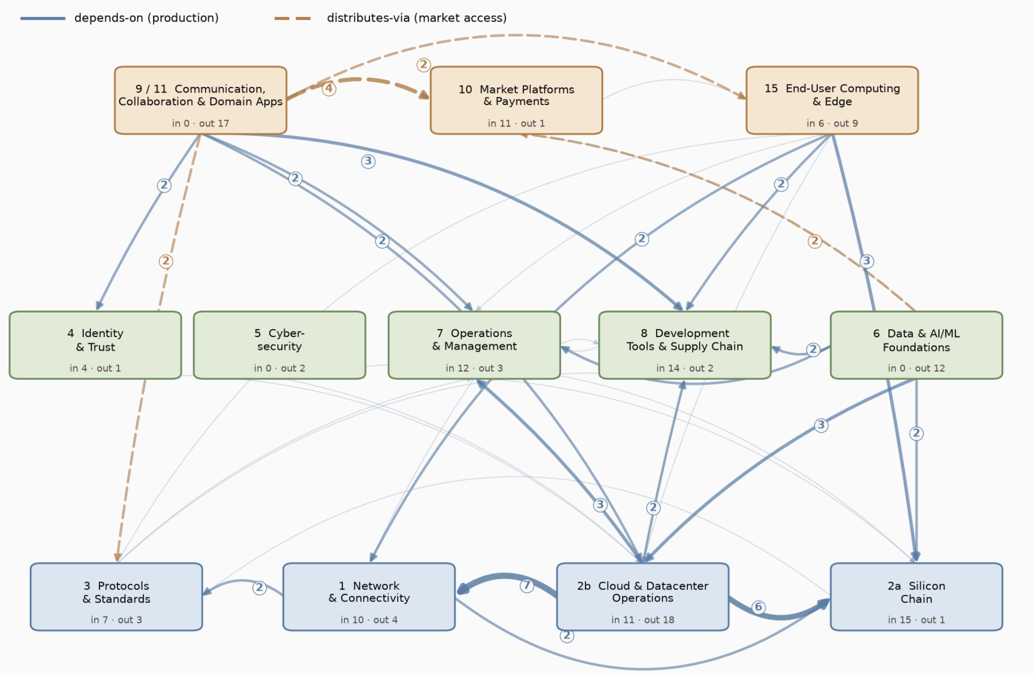

Arrow A→B: a category in block A depends on (or distributes via) a category in block B; width and circled label = category-edge count (labels on weight ≥ 2). Contraction rule pre-stated in crosswalk-v1.md 12a. Blocks 9 and 11 map to the same two categories and are drawn once; weights are per block. 'in' counts inter-block dependency and distribution edges into the block (the block supplies others); 'out' counts edges out of it (the block consumes). Intra-block edges (the silicon chain 2a keeps 19 of its own) and the context layers (demand, capital, jurisdiction, influence) are reported in the results JSON.

block	supplier in	consumer out	net	reading
2a silicon chain	15	1	+14	deepest net supplier; also keeps 19 intra-block edges (the chain itself)
8 development tools	14	2	+12	toolchain supplier to every layer, foundation included
7 operations	12	3	+9	operations substrate beneath cloud, applications, AI, edge

block	supplier in	consumer out	net	reading
10 market platforms	11	1	+10	supplier of market access: 10 of 11 inbound incidences are distribution edges
2b cloud and datacenter ops	11	18	−7	the hinge: supplies applications and AI, consumes every foundation block
1 network and connectivity	10	4	+6	physical substrate; 7 of 10 inbound incidences from cloud
3 protocols and standards	7	3	+4	rule-book supplier; also the target of the influence overlay (below)
15 end-user computing and edge	6	9	−3	mid-position: distribution target above, silicon consumer below
4 identity	4	1	+3	mandatory input: both SaaS categories require it
5 cybersecurity supply	0	2	−2	no block depends on the security category as coded
6 data and AI foundations	0	12	−12	pure consumer of silicon, cloud, tooling, and data inputs
9 and 11 applications	0	17 each	−17	pure consumers; identical by construction (same two categories)

Weighted degrees over inter-block dependency and distribution edges; blocks 9 and 11 carry identical rows because they map to the same two categories (asserted by the suite). Full edge lists, intra-block weights, and the context ledgers are in `verify/results/tbb_matrix.json`.

Four economic readings, each read directly off the contracted graph. First, seven blocks carry positive net degree, and dependencies concentrate on four of them: the silicon chain, the development toolchain, operations tooling, and market platforms (in-degrees 15, 14, 12, 11 against out-degrees of 1 to 3); the remaining net suppliers, network infrastructure, standards, and identity, sit at +6, +4, and +3. Three of the four leaders classify captured, and the fourth (2a) is the floor: the stack's upstream is exactly where the classification found the exposure. The heaviest single flows run through cloud: 2b draws 7 incidences from network infrastructure and 6 from the silicon chain, and everything above draws on 2b, which makes cloud the hinge (in 11, out 18) and explains in structure what the market row prices (a 70 percent big-three share at the hinge taxes the whole cone above it). Second, the direction of value and the direction of dependency are opposite: the application blocks (9, 11) and the AI block (6) have in-degree zero, nothing in the stack depends on them, while everything they do crosses the enabling and foundation blocks. They are where the demand layer attaches (below) and where the stack's output is sold, so the blocks that earn the revenue are the blocks with no structural leverage, and the blocks with structural leverage (8, 7, 2a, 10) earn theirs as inputs or tolls. Third, the typed edges separate two kinds of upstream position, and the data types them. Block 10′s supplier position is market access: 10 of its 11 inbound incidences are distributes-via edges (applications and AI reaching users through app stores, marketplaces, and payment rails), a rent collected at the door rather than an input supplied, which is the graph-structural face of the intermediation economics of Section 2.1 and the DMA and ECB receipts behind it. The influence overlay adds the socio-technical channel at block grain: the cloud category influences the open-source commons substrate of blocks 7 and 8, and the cloud, SaaS, and distribution categories influence the standards block (six inter-block influence incidences in all), so two of the upstream supplier blocks

(7, 8) and the rule-book block (3) are themselves influence targets of the hinge they supply. The funding overlay stays context-to-block (EU public programmes fund 1, 2a, 2b, 6, 7, 8; European private capital funds 2b, 6, 9, 11; US capital funds 1, 2b, 6). Within depends-on, the graph does not distinguish hardware supply from software licensing from operated service, the contracted weights are edge counts with no economic intensities behind them, and no euro flows ride the edges; the typed reading above goes exactly as far as the flagship's relation types go. Fourth, the floor's location is visible in the contraction itself: the silicon chain absorbs inbound dependency from six blocks and returns one edge (its instruction-set category depending on standards), and its 19 intra-block edges are the chain the option calculus walks when every trace ends at the same five blocs (Section 6.5). The contraction and the calculus agree on where the bottom is.

8. What rests on the stack

The stack matters for the value built on top of it, and the graph's demand layer plus sourced macro receipts let this be stated with measurements. Discipline first: every figure in this section prices exposure or value at stake; none measures an intervention moving anything, and nothing here claims an instrument delivers (Section 9′s opening discipline applies in full).

The demand layer attaches at the top of the cone. The six measured European demand categories (public sector, sovereign-critical, critical national infrastructure, enterprises, consumers, startups) concentrate their dependency edges on cloud and applications: 2b receives edges from five of the six demand categories, blocks 9 and 11 from all six, end-user computing and network infrastructure from five each, and cybersecurity supply from four. Blocks 6, 7, and 8 receive zero direct demand edges, and block 4 exactly one: AI foundations, operations tooling, and development toolchains reach the demand layer by composition, through the 17 inter-block incidences that blocks 9 and 11 each place on blocks 4, 7, 8, 2b, and below (Section 7). Identity alone touches it directly. What procurement sees is therefore the top of the cone, and what the criterion prices is the cone itself: adoption statistics say the visible layer is now the norm (52.7 percent of EU enterprises bought paid cloud services in 2025, and 72 percent of EU businesses reached basic digital intensity; Eurostat 2026), so the invisible blocks beneath it underpin the dependencies of most of the European enterprise population.

The value at stake is macroeconomic, and the receipts are sourced with their weights stated. On the priced side, the Asterès studies for Cigref (commissioned by a buyer association, flagged as such) estimate that 80 percent of European business spending on cloud software and services went to US vendors, EUR 265bn in the 2025 study, and, on the wider 2024 base of the 2026 study, EUR 330bn of EUR 400bn (83 percent), with observed price rises of 8.7 percent over the last three years on the locked base draining an estimated EUR 93bn per year from the continent out of an average EUR 140bn annual inflation cost (Asterès 2025, 2026). Those flows exit through the blocks this paper measures: the hinge and the application layer, where the published shares sit at 70 percent (cloud, big three) and 73 to 80 percent (public-sector productivity), and the market-access block, where the ECB documents 61 percent of euro-area card payments on international schemes and where the DMA's designated app-store gatekeepers price distribution. Per block, the leak concentrates where high non-EU share meets high supplier in-degree (Section 7): the toolchain and operations blocks (8, 7) tax by input, the platform block (10) taxes by toll, the identity block (4) locks by switching cost, and the hinge (2b) prices the pass-through. The VMware repricing episode (renewal costs up by a factor of 10 or more on complaint evidence) illustrates the exercise of such a position when ownership of a locked input changes hands. On the unpriced side, the Commission's own study estimates that around EUR 1bn of EU company investment in open-source software (2018) generated EUR 65 to 95bn of positive impact on the European economy, a benefit-cost ratio above 1:4, and that a 10 percent increase in contributions would add 0.4 to 0.6 percent of GDP annually (Blind et al. 2021): the estimate covers open-source software across the whole EU economy, and the commons substrate of blocks 7 and 8, invisible in the market table because unpriced and invisible to the demand layer because never procured, is the channel through which every layer of the stack draws on that base (Section 7). An economy-wide weight of this order riding on an unpriced input is the economic case for treating its stewardship as an instrument family in its own right (Section 9) without claiming any given engagement delivers it. The AI and data block's stake is of the same order (a EUR 579bn EU27 data economy at 4.4 percent of GDP, AI spend approaching USD 290bn by 2029), which prices what the thin classification of block 6 is protecting the option on.

The blocks where European value concentrates (applications, data, the demand layer's own operations) are pure consumers of the blocks where the classification finds capture and the contraction finds leverage. A per-block checklist would fund the visible layer; the measured cone says the visible layer's exposure is set upstream, in blocks the demand side never sees on an invoice.

9. From classification to instrument

The matrix's use is the assignment of instrument families to measured gaps. The discipline throughout: the structure indicates which instrument class addresses which gap; no instrument on this list delivers sovereignty by purchase or decree, and a certification that a block passed a checklist certifies nothing about paths (Section 3.2).

Captured without a forkable option (4, 5, 10, 15): build-or-buy. The binding gap is reach share 1 with no open substrate computed anywhere in the block. Only ownership or jurisdiction change moves that number: building a European option, or acquiring one, at the identity, security-supply, platform, and now end-user layers. Exit-friction scores say the identity layer compounds the problem (mean F_exit 0.77): the single option is also the hardest to leave, so option creation there buys the most friction relief per block. Block 15 splits internally: the client-OS category carries the proprietary coding and the worst-level concentration, while the IoT/edge category's lenient coding records European industrial options already inside the category, giving a demand-side instrument something to prefer there before any build. One stewardship anchor also touches this block: the Sovereign Tech Agency's EUR 1m GNOME engagement (`data/cost-anchors-v1.csv`, row sta-gnome) prices client-OS stewardship three orders of magnitude below the family's build-or-buy anchor, so the layer whose coded substrate is proprietary already carries a stewardship-priced instrument at its open margin. The build-or-buy assignment prices replacing the incumbent, and the anchor prices keeping the open alternative alive.

Captured through the commons (7, 8): substrate stewardship. The gap is jurisdictional reach into a substrate that is already forkable (shares 0.75 and 1.00). The indicated instrument is stewardship of the commons: European hosting of forges and distribution, foundation-level governance participation, and maintained fork capability, which severs the reach edge the V1 variant isolates. The macro receipts of Section 8 say what the substrate is worth to the EU economy; the contraction says every layer of the stack draws on these two blocks; and the cost anchors below say this is the cheapest anchored gap class per intervention. The institutional design for organizing diffuse beneficiaries into a provisioning club, with its free-riding and governance analysis, is the buyers'-club companion (Fermigier 2026d); this matrix supplies its target list.

Thin (1, 2a, 2b, 3, 6, 9, 11): demand aggregation and maintenance. The gap is $N = 1$: one controller-disjoint realization, with unreached categories present (eu-cloud, datacenter-physical, eu-saas-oss, standards participation, EU models and AI inputs) that a second-sourcing or aggregation instrument can build on. Procurement that preserves and exercises the existing European option, and maintenance funding that keeps it exercisable, addresses the measured gap directly; the option-value framing of why an unexercised option still needs paying for is the theory companion's Section 6 argument. The market rows add the scale of what the option protects: the thin blocks carry Europe's largest measured positions (telecom services EUR 407bn; cloud EUR 61bn with the EU share flat at 15 percent; the EUR 90bn data market), so a percentage point of share preserved or regained at these blocks is measured in billions per year, a statement about stakes rather than about any instrument's effect. The Commission's 2026 capacities study reports the supply side of two of those unreached surfaces as weak: on cloud, "fragmentation within the European cloud ecosystem limits the availability of scalable, interoperable alternatives capable of supporting cross-border services", and on standards, a SWOT weakness cell reading "Limited private incentive to adopt EU open standards" (European Commission 2026e). An unreached category that no private buyer has an incentive to adopt supplies an option on paper, which is the gap the aggregation half of this family addresses.

The floor condition (2a). All of the above is conditional. Every block's trace crosses the silicon chain, where all five blocs cut, two of the graph's three China reach edges live, and capacity investment is the slowest lever. Any block-level investment unaccompanied by compute-floor policy buys resilience theater: the block-level number improves while every realization still crosses the same five-bloc floor. The above-floor variant prices the converse too: complete floor autonomy would still leave every block at $N = 1$ with a US-only cut (Section 6.5), so floor policy and block-level instruments are complements, and neither substitutes for the other.

The ranking: family-level anchors from public records; the block level still waits. The intended criterion is option value restored per euro: for each candidate intervention, the change it induces in the block's measured gap (an N increment, a reach-share reduction, an F_exit reduction), divided by its cost, with leverage points ranked first. The numerators are computable from this paper's suite today. For the denominators, public records now anchor each instrument family's per-intervention price (`data/cost-anchors-v1.csv`, 21 rows, each carrying source URL, access date, and confidence flag, all accessed 2026-07-17; schema and ranges validated by the suite): substrate stewardship at EUR 99,060 to 1,056,673 per component engagement (16 Sovereign Tech Agency projects, sampled from a published portfolio of 108 supported technologies and EUR 37.3m invested since October 2022; Sovereign Tech Agency 2026); demand aggregation and maintenance at EUR 9m to 180m (Schleswig-Holstein's one-time migration budget for 2026, and the Commission's sovereign-cloud tender ceiling over six years; Land Schleswig-Holstein 2025; European Commission 2026a); build-or-buy at EUR 1.2bn (the IPCEI-CIS public envelope, the family's only public anchor, sitting at the cloud-edge layer rather than at the identity, security-supply, or platform layers where this section's build-or-buy family points; European Commission 2023a); capacity investment at EUR 32bn (thirteen Chips Act State-aid decisions, public and private combined; European Commission 2026b). The four ranges are pairwise disjoint in that order, separated by factors of 8.5, 6.7, and 27 at adjacent extremes; the stewardship median (EUR 289,715) sits a factor of about 31 below the demand-side family's floor and about 620 below its ceiling (31.1 and 621.3, computed and cross-checked by the suite, check K4). The factor of 27 also compares inconsistent bases: the IPCEI-CIS envelope is public aid only while the Chips Act total mixes public and private money, and the ordering holds unless the public share of the Chips total falls below about 4% (the suite's break-even, 0.0375). The coarse ordering this supports, stated with its condition: wherever the per-intervention gap values across families differ by less than the cost separation, the per-euro ordering inherits the cost ordering, with the commons-reach gap of blocks 7 and 8 the cheapest anchored gap class per intervention and floor capacity the dearest. A further juxtaposition, added at v0.5, motivates without computing anything: the market anchors of Section 2.2 put the annual markets riding on the stewardship-gap and thin blocks at EUR 10^10 to 10^11 while the stewardship family's anchored price is 10^5 to 10^6 per engagement; the two magnitudes live in different scopes and units (a market per year against an instrument per component) and no ratio between them is computed. The size of the disparity is the reason to build the per-euro criterion. The anchors are envelopes and ceilings rather than outturns; the numerators are not commensurable across families; stewardship prices one component while a block's substrate contains many (the stewardship-against-demand-side comparison holds up to roughly eight components per block gap at the range boundary, the stewardship-against-build-or-buy comparison up to roughly a thousand); every figure prices an instrument aimed at a gap, never the gap moving, so delivery remains a separate, unpriced event; and no family substitutes for another, the floor condition standing. Per-row provenance and the full failure-region list are in the cost-anchors note. The block-level ranking (which block first, within a family) remains unpublished: public records do not price per-block interventions, and no such ranking is published until they do.

10. What the criterion does not do

Category resolution. The graph carries one category per function family; the flagship's supplier-resolution pilot (28 options over three layers) shows aggregation can both overstate and understate coercion. Section 6.6 integrates the pilot where it intersects the blocks (2b, 9/11, and the accelerator class of 2a); for every other block the $N = 1$ results stay aggregation-sensitive in exactly the way the flagship documents, with no supplier-level data yet to check them against.

Inherited codings. Substitutability, concentration, capture-risk, and substrate codings come from the flagship's single-coder aggregation; the Gate-3 run covers their ordinal uncertainty, the v0.4 reach-edge perturbation prices the binary reach edges the ordinal design holds fixed (Section 6.4), and an independent re-code has yet to happen. The re-code is the revision's outstanding item: V3 measures how much a miscoded edge would move, and only a second coder can say whether any edge is miscoded.

Market anchors are anchors, and their scopes differ. The table of Section 2.2 mixes WSTS Europe, Synergy Europe, EU27, the euro area, and StatCounter Europe; usage shares stand in for revenue shares where no split is published; several rows are commissioned or complaint evidence and flagged; forecasts carry horizons. The table prices stakes per block and supports no cross-block arithmetic; no computed quantity in the suite reads it, by design.

The contraction counts edges, and euros do not ride them. Section 7's weights are category-edge counts under a pre-stated multi-map, with pair incidences exceeding raw edges where membership overlaps (both totals reported); they are not economic intensities, the depends-on type is not subdivided (hardware supply, software licensing, and operated service look alike), and the influence overlay is a sampled lower bound inherited from the flagship. The contraction locates concentration; it does not weigh it in money.

Forkability coverage. No substrate-coded category maps to blocks 3, 4, and 5, so their forkability is uncomputable; block 8′s share of 1.00 rests on the commons substrate's three coded categories, and the supply-chain category itself (forges, registries, tooling) carries no coding; block 15′s share of 0.00 rests on the single coded client-OS category, and its four IoT/edge categories carry no coding.

Exit friction coverage. F_exit exists only where the CSM scored AWS services; blocks 1, 2a, 3, 8, 9, 10, 11 have none, and one ops row the crosswalk names (a Systems-Manager-class service) has no scored CSM row.

Vertical and protocol resolution. Blocks 9 and 11 are indistinguishable in the graph, block 10 is two categories wide, and block 3′s protocols live mostly as edge types; the taxonomy sees structure the graph cannot yet score, and for 9 against 11 the market evidence currently discriminates where the graph cannot (Section 6.2).

A scope note on the reference buyer. The public-sector buyer's measured dependency cone does not reach the AI-model categories, so "public-sector demand through block 6" has no measurable path; block 6′s numbers are the block root's own.

Diagnosis, still. The matrix diagnoses; it does not deliver. The above-floor variant is computed (Section 6.5), the dependency structure is computed (Section 7), and the per-euro criterion carries family-level cost anchors (Section 9); the block-level ranking still waits for per-block cost data, and nothing in this paper measures an intervention actually moving a gap.

The taxonomy is revisable, and has been revised. Of the four open questions of Section 4.3, one produced a taxonomy revision (Block 15), two were ruled in favor of the recorded defaults, and one is a standing treatment; the numbers attached informed each ruling without preempting it, which is the mechanism working as designed. The revised block inherits the resolution limits above: two categories stand in for the whole end-user and edge economy.

11. Related lists and methods

The TBB position paper (Fermigier 2025) is this paper's own object: fourteen blocks, three lenses, eight recommendations. It enumerates and motivates; it does not rank, and its strategic-dependency lens is stated as a perspective without a measurement apparatus. This paper is that lens made computable, and its first output fed back into the taxonomy (the Block 15 revision of Section 4.4, flagged for the position paper's next version).

The EuroStack Initiative Foundation publishes a layered model of European digital infrastructure and policy proposals built on it. As a list it shares the TBB paper's granularity and its limitation: layers are named and advocated for, and no path-level quantity distinguishes a layer whose options are real from one whose options all cross the same controller.

The CNCF landscape is the most complete node enumeration in the cloud-native domain, at product granularity with categories and subcategories. It is a supply catalogue: it records that alternatives exist within a category and carries no controller, jurisdiction, or path structure, so it cannot see that a category's alternatives share a fab, a distribution chokepoint, or a legal reach.

The Sovereign Tech Agency portfolio is a funding list of critical open-source components, selected by criticality assessments. It operates at the component level inside what this paper treats as the commons substrate (blocks 7 and 8) and is the closest existing instrument to the stewardship family of Section 9; its published figures (108 technologies supported since October 2022, EUR 37.3m invested, per-project amounts on public project pages, accessed 2026-07-17) supply that family's cost anchors. It does not rank blocks against each other, which is the decision this criterion addresses.

The Commission's 2026 study of the EU's critical digital capacities beyond 2027 is the largest recent public evidence base over this terrain, and it holds both sides of the forkability question without joining them. Its main report records open source as the strategy: "Open source is seen as a strength and a potential strategy to close gaps where Europe lags behind global competitors, particularly the US and China. In areas like software, open source could provide a more agile and collaborative route to innovation" (European Commission 2026d). Its dependency table then lists among the EU's own dependencies "Opensource standards (Kubernetes, Spark)", "Orchestration Software (e.g. Kubernetes, OpenStack) for Cloud", "AI Frameworks (TensorFlow, PyTorch)", and "Key Blockchain Protocols (Ethereum)", and its protocol section records that the core protocols "are governed outside the EU, limiting influence over updates and standards". The same artifacts are the remedy in one chapter and the exposure in another, and the report nowhere states the condition under which an open licence yields control. A forkability coding adjudicates that condition, and the pre-stated variant V1 prices it: blocks 7 and 8 are captured on jurisdictional reach into a substrate whose licence is already open, and forking severs that reach only where the fork is maintained (Sections 6.3 and 9). The study's interoperability annex leaves the matching gap on the specification side, claiming an "Active role in international standards bodies for e-government" in a SWOT strength cell while naming no standards body anywhere in its text (European Commission 2026e).

The method itself has formal kin outside the programme, named in the theory companion and summarized here so the composition's novelty claim is scoped. Reliability engineering computes minimal cut sets on AND/OR fault trees and prices, as common-cause failure, the correlation that defeats per-component redundancy counts (Barlow and Proschan 1975); the five-singleton cut structure of Section 6.5 is that phenomenon with controllers in place of failure events. Network interdiction optimizes an attacker's arc removals under a budget (Wood 1993), the coercion number's adversarial reading. Critical-infrastructure analysis studies the interdependencies that make per-sector assessments miss system-level exposure (Rinaldi, Peerenboom and Kelly 2001), and component-level criticality assessment is the selection practice behind the Sovereign Tech Agency portfolio above. The delta, as the theory companion states it: controllers with identity and coalitions replace anonymous failure events, N supplies the dual quantity of interest, and the frozen four-class rule turns the calculus into a taxonomy-level instrument. On the market side, the anchors of Section 2.2 compose published market measurement (Commission studies, authority reports, analyst releases) with the structural criterion; the composition claims no new measurement, only provenance discipline.

12. Conclusion

A list of fourteen blocks became an economically grounded criterion under a rule frozen in advance, and the criterion discriminates. Six blocks are captured and seven thin, and the sourced market anchors line the classes up with market structure: where a share is published for a captured block, non-EU providers hold from 60 percent to effectively all of the measured European market, while the thin blocks hold Europe's largest measured positions (telecom services EUR 407bn, cloud EUR 61bn at a 15 percent EU share, chip demand USD 54bn against 10 percent EU production). The contraction of the measured graph onto the blocks gives the stack its economic geometry: dependencies concentrate on the silicon chain, the development toolchain, operations tooling, and market platforms, three of the four captured; cloud is the hinge; the platform block's upstream position is a toll on market access rather than an input; and the demand layer's edges never touch the toolchain, operations, or AI-foundation blocks and touch identity exactly once, so procurement sees the top of a cone whose exposure arrives by composition. Two of the six captured verdicts are stewardship problems misread as build problems until forkability is priced, at anchored costs three orders of magnitude below build-or-buy. The commons substrate they concern, the stack's channel onto an open-source base with a Commission-estimated EUR 65 to 95bn of economy-wide EU GDP impact, remains invisible to both market statistics and procurement. The verdicts hold under perturbation of every ordinal coding they read, the floor's removal, and membership deletion, with the residual weight located on single-coder reach edges by the pre-stated reach-edge perturbation. Beneath the discrimination, one paragraph of floor: every block has $N = 1$ and $\kappa = 1$ with five blocs each a solo cutter through the semiconductor floor, so no per-block checklist certifies the stack, and above the floor every block still hangs on a single US-severable path; block-level instruments and floor policy are complements. What the criterion can now say about the per-euro order is coarse and family-level, on public cost records in pairwise disjoint ranges from EUR 10^5 stewardship engagements to over 10^10 floor capacity; the block-level order still requires per-block cost data, and the criterion that will consume it (option value restored per euro) is specified. The revision mechanism has run once on the record: the taxonomy's author added Block 15 on the crosswalk's evidence, and the market table now prices what that block carries.

References

Asterès (2025). Technological dependence on American software and cloud services: an assessment of the economic consequences in Europe. Study for Cigref, May 2025. https://www.cigref.fr/technological-dependence-on-american-software-and-cloud-services-an-assessment-of-the-economic-consequences-in-europe (accessed 2026-07-17).

Asterès (2026). From technological dependency to economic capture: the cost of cloud and software services inflation to Europe. Study for Cigref, May 2026. https://www.cigref.fr/from-technological-dependency-to-economic-capture-the-cost-of-cloud-software-services-inflation-to-europe (accessed 2026-07-17).

Barlow, R. E., and Proschan, F. (1975). Statistical Theory of Reliability and Life Testing. Holt, Rinehart and Winston.

Blind, K., Böhm, M., Grzegorzewska, P., Katz, A., Muto, S., Pätsch, S., and Schubert, T. (2021). Study about the impact of open source software and hardware on technological independence, competitiveness and innovation in the EU economy. European Commission, DG CNECT. https://digital-strategy.ec.europa.eu/en/news/commission-publishes-study-impact-open-source-european-economy (accessed 2026-07-17).

CMA (2025). Cloud services market investigation: summary of the final decision. Competition and Markets Authority, 31 July 2025. https://assets.publishing.service.gov.uk/media/688b20e6ff8c05468cb7b120/summary_of_final_decision.pdf (accessed 2026-07-17).

CNCF (2026). CNCF cloud native landscape. https://landscape.cncf.io (accessed 2026-07-17).

Connect Europe (2026). State of Digital Communications 2026. February 2026. https://connecteurope.org/sites/default/files/2026-02/connectEU-Rapport-web-2026-onepage-v2.pdf (accessed 2026-07-17).

ECB (2025). Report on card schemes and processors. European Central Bank, February 2025. `https://www.ecb.europa.eu/pub/pdf/other/ecb.reportcardschemes202502~1614226b0a.en.pdf` (accessed 2026-07-17).

ECSO (2024). Cybersecurity market analysis and recommendations, v1.1. European Cyber Security Organisation, December 2024. https://ecs-org.eu/ecso-uploads/2024/12/ECSO-Cybersecurity-Market-Analysis-and-Recommendations-v1.1.pdf (accessed 2026-07-17).

ESIA (2025). Press release on the WSTS Autumn 2025 semiconductor market forecast. European Semiconductor Industry Association, 2 December 2025. https://www.eusemiconductors.eu/sites/default/files/ESIA_WSTS_PR_AuFc2025.pdf (accessed 2026-07-17).

European Commission (2023a). Commission approves up to €1.2 billion of State aid by seven Member States for an Important Project of Common European Interest in cloud and edge computing technologies. Press release IP/23/6246, 5 December 2023. https://ec.europa.eu/commission/presscorner/detail/en/ip_23_6246 (accessed 2026-07-17).

European Commission (2023b). Digital Markets Act: Commission designates six gatekeepers. Press release IP/23/4328, 6 September 2023. https://ec.europa.eu/commission/presscorner/detail/en/ip_23_4328 (accessed 2026-07-17).

European Commission (2025). Commission accepts commitments by Microsoft on Teams. Press release IP/25/2048, 12 September 2025. https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2048 (accessed 2026-07-17).

European Commission (2026a). Commission advances cloud sovereignty through strategic procurement. News, 17 April 2026. https://commission.europa.eu/news-and-media/news/commission-advances-cloud-sovereignty-through-strategic-procurement-2026-04-17_en (accessed 2026-07-17).

European Commission (2026b). European Chips Act (policy page). https://digital-strategy.ec.europa.eu/en/policies/european-chips-act (accessed 2026-07-17).

European Commission (2026c). European Chips Act (factpage). https://digital-strategy.ec.europa.eu/en/factpages/chips-act (accessed 2026-07-17).

European Commission, Directorate-General for Communications Networks, Content and Technology (2026d). Study on the EU's critical digital capacities deployment beyond 2027: final report. Prepared by Technopolis Group; K. Izsak, P. Simmonds, M. Ploeg, S. Viscido, A. Lotito, C. Moreno and R. Alasghar (eds.). Luxembourg: Publications Office of the European Union. ISBN 978-92-68-40079-1, doi:10.2759/0673677.

European Commission, Directorate-General for Communications Networks, Content and Technology (2026e). Study on the EU's critical digital capacities deployment beyond 2027, Annex 15: Interoperability and GovTech. Prepared by A. Murray, A. Adewusi and T. Marshall, Technopolis Group. Luxembourg: Publications Office of the European Union. ISBN 978-92-68-40089-0, doi:10.2759/5257213.

European Data Market Study (2025). European Data Market Study 2024-2026: first report on facts and figures. IDC, the Lisbon Council, and Carsa for DG CONNECT, March 2025. https://ec.europa.eu/newsroom/repository/document/2025-13/EDM_2024__2026__First_Report_on_Facts_and_Figures_9cU0iIjcuSZhjmwI9TchwqenidQ_114043.pdf (accessed 2026-07-17).

Eurostat (2026). Cloud computing: statistics on the use by enterprises (Statistics Explained); Digitalisation in Europe, 2026 edition. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Cloud_computing_-_statistics_on_the_use_by_enterprises and https://ec.europa.eu/eurostat/web/interactive-publications/digitalisation-2026 (accessed 2026-07-17).

EuroStack Initiative Foundation e.V. (2026). EuroStack. https://eurostack.eu (accessed 2026-07-17).

Fermigier, S. (2026a). Path sovereignty: dependence, control, and options in the European digital stack. Working paper, `../../eurostack-path-sovereignty/papers/power-network-wp.md`.

Fermigier, S. (2026b). A two-dimensional taxonomy of cloud service integrations: modeling architectural coupling and portability friction in managed infrastructures. Working paper, `../../cloud-services-matrix/papers/cloud-services-matrix.md`.

Fermigier, S. (2026c). Cuts and options: a theory of strategic autonomy in production networks. Working paper, `../../eurostack-path-sovereignty/papers/theory-wp.md`.

Fermigier, S. (2026d). Buyers' clubs for the open-source flow: collective provision in the FLOSS value network. Working paper, `../../oss-adoption/papers/collective-provision-wp.md`.

Fermigier, S. (2025). Position Paper: A Structured Framework for Europe's Digital Sovereignty: Technological Building Blocks for the Public Digital Infrastructure and Internet Commons. v1.1, May 2025; drafted within the NGI Commons initiative discussions. https://euro-stack.com/blog/2025/5/tbb-for-open-internet-stack (author-hosted; accessed 2026-07-17; title and recommendation count verified against this URL).

Land Schleswig-Holstein (2025). Press release of the state chancellery, 4 December 2025 (open-source migration status: LibreOffice share, mailbox migration, one-time budget, licensing savings). https://www.schleswig-holstein.de/DE/landesregierung/ministerien-behoerden/I/Presse/PI/2025/cds/251204_cds_open-source (accessed 2026-07-17).

Open Cloud Coalition (2025). Microsoft dominates 80% of public sector productivity software market in EU (Compass Lexecon study announcement), 24 July 2025. https://opencloudcoalition.com/microsoft-dominates-80-of-public-sector-productivity-software-market-in-eu-raising-competition-concerns-according-to-new-report/ (accessed 2026-07-17).

Rinaldi, S. M., Peerenboom, J. P., and Kelly, T. K. (2001). Identifying, understanding, and analyzing critical infrastructure interdependencies. IEEE Control Systems Magazine 21(6), 11–25.

Sovereign Tech Agency (2026). Technologies and projects (portfolio). https://www.sovereign.tech/tech (accessed 2026-07-17).

StatCounter (2026). Global Stats: browser, desktop OS, and mobile OS market shares, Europe, June 2026. https://gs.statcounter.com (accessed 2026-07-17).

Synergy Research Group (2025). European cloud providers' local market share now holds steady at 15%. 24 July 2025. https://www.srgresearch.com/articles/european-cloud-providers-local-market-share-now-holds-steady-at-15 (accessed 2026-07-17).

Wood, R. K. (1993). Deterministic network interdiction. Mathematical and Computer Modelling 17(2), 1–18.

Annex A: Working notes (temporary annex)

This annex lists internal working documents and will be removed from the final version.

- `notes/crosswalk-v0.md`: the initial 14-block mapping and metric plan (against the 49-category core).
- `notes/crosswalk-v1.md`: the reconciled mapping, scoring defaults, and the classification rule as frozen (section 5), with the sensitivity variant and perturbation design (5a/5b), the v0.2 pre-statements (section 9: above-floor variant V2, supplier-resolution reporting rule), the v0.3 taxonomy revision and presentation ruling (section 10), the v0.4 pre-statements (section 11: reach-edge perturbation V3, leave-one-category-out pass L1, convention and ratio checks), and the v0.5 pre-statements (section 12: block-graph contraction B1, figures, market anchors), and the v0.6 pre-statements (section 13: conventional median and translation map, check K4; full v0.5 snapshot pin, check R4).
- `notes/scored-matrix-v1.md`: the findings note behind Section 6, with per-block readings at full length, the v0.2/v0.3 additions (sections 9 and 10, including Block 15′s full row), the v0.4 results (section 11: V3 and L1), and the v0.5 results (section 12: the contraction in numbers, market-anchor coverage, figures), and the v0.6 correction record (section 13: the conventional median and its ratios).
- `notes/cost-anchors-v1.md`: the cost-anchor method, provenance discipline, and the full failure-region list behind Section 9′s ranking paragraph.
- `notes/market-anchors-v1.md`: the market-anchor method, coverage accounting, and reading discipline behind Section 2.2.
- `../notes/programme-extensions.md` (§3): the programme plan section that specified this paper.

Annex B: Reproduction: scripts, data, and figures

- `verify/tbb_matrix.py` (matplotlib only; seed 20260716): computes every number and every figure in this paper, Block 15 included; applies the frozen rule and the pre-stated variants V1 and V2; runs the 500-draw perturbation (V2 stability from the identical draw stream, which Block 15′s addition cannot shift because draws are generated over all nodes before any block loop); runs the v0.4 reach-edge perturbation V3 (drop rates 0.10 and 0.25, 500 draws per rate, dedicated stream seed 2026071603) and the deterministic leave-one-category-out pass L1, plus the V1 denominator-convention and derived-ratio checks; computes the v0.5 block-graph contraction B1 under the pre-stated rule with its partition, decomposition, reach-reproduction, and recount checks (B1a to B1d) and the 9/11-identity assertion; validates the market-anchor table's provenance schema (M1); pins every pre-v0.3 result (classification, V1 variant, stability, V2 classification, V2 stability), the complete v0.3 output byte-for-byte against `verify/results/tbb_matrix_v03.json` (check R2), the complete v0.4 output against `verify/results/tbb_matrix_v04.json` (check R3), and the complete v0.5 output against `verify/results/tbb_matrix_v05.json` (check R4, under the documented median translation map), with the matrix CSV and figure byte-identical; validates the cost-anchor table, recomputes its family ranges, and recomputes the stewardship median as the conventional median (check K4; the upper median reported through v0.5 stays in the results file as `median_upper`); recomputes the pilot supplier-resolution numbers by importing the flagship's `pilot_discount.py`; ends by asserting its 75-entry cross-check ledger and printing ALL CHECKS PASSED. Outputs: `verify/results/tbb_matrix.json`, `data/tbb-matrix-v1.csv`, `papers/figures/tbb-matrix-v1.png`, `papers/figures/tbb-architecture.png`, `papers/figures/tbb-architecture-classes.png`, `papers/figures/tbb-block-graph.png`.

- `data/cost-anchors-v1.csv`: the 21-row cost-anchor table behind Section 9's ranking paragraph; every row carries its source URL, access date (2026-07-17), and confidence flag.
- `data/market-anchors-v1.csv`: the 45-row market-anchor table behind Section 2.2; every row carries metric kind, value, unit, period, scope, source, source URL, access date (2026-07-17), and confidence flag, with verbatim source wording in the row notes and explicit no-figure rows where no defensible figure exists.
- Frozen upstream inputs, read-only and cross-checked in the ledger: the flagship graph CSVs (`../../eurostack-path-sovereignty/data/edges/compute_nodes.csv`, `compute_edges.csv`, `category_nodes.csv`) against the published totals (52 categories; 151 dependency and distribution, 37 jurisdictional-reach, 48 funding and influence, 7 threat edges); the cut-calculus results (`../../eurostack-path-sovereignty/verify/results/cuts_options_v1.json`, all 36 published configurations reproduced exactly by the imported model) and the compute-core results (`.../compute_core_v1.json`, jurisdiction reach counts 21/13/3); the CSM scores (`../../cloud-services-matrix/notes/services.csv`, all 35 F_exit values recomputed from components at w = 0.30) and robustness results (`../../cloud-services-matrix/verify/results/robustness.json`, seed 20260714, 2000 draws, all 35 baseline quadrants reproduced); the pilot supplier-resolution data (`../../eurostack-path-sovereignty/data/pilot/options.csv`) with its frozen results (`.../verify/results/pilot_discount_v1.json`, per-layer breadths and workload kappas recomputed by the imported module and matched exactly).
- The cut calculus itself is imported unmodified from `../../eurostack-path-sovereignty/verify/cuts_options_model.py`; the virtual-block-root construction and its exact semantics are documented in `notes/crosswalk-v1.md` section 4, and the contraction rule in section 12a.