EUROPEAN CLOUD SOVEREIGNTY

The human, economic, technological and industrial capabilities that shape Europe's digital autonomy



FRANCK MEVENGUE

European Cloud Sovereignty

The human, economic, technological and industrial capabilities that shape Europe's digital autonomy



Franck Mevengue

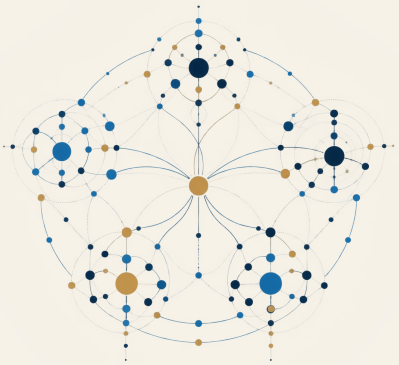
Note to the reader

This book distinguishes what sources establish, what remains partially documented, what is a plausible hypothesis, and what is still open. Empirical limits are stated when they shape the argument.

GRAPHIC LANGUAGE

The Nodes motif

Solid lines: documented relations. Dotted lines: hypotheses or trajectories. This motif carries no empirical data.



Nodes signature motif (conceptual illustration).

Preface

Europe is digitising its administrations, industries, and public services at a pace that leaves little room for illusion. A growing share of computing, storage, and applications now rests on infrastructures operated by a small number of global providers. This shift is not a technical detail. It redefines what an organisation can still control, substitute, or rebuild when an incident, a contract, or geopolitical tension reminds us how fragile a dependency can be.

Cloud computing, often reduced to a hosting question, has become an economic and strategic infrastructure. It concentrates issues of cost, skills, reversibility, industry, and institutions. An exclusively legal approach, contenting itself with locating data or signing clauses, leaves out what matters most: who can actually operate, exit, industrialise, and control. An exclusively commercial approach, comparing price catalogues without an equivalent scope, would confuse a tariff advantage with autonomy.

This book proposes a different thread. It treats European cloud sovereignty as a **degree of interdependent capabilities** (human, economic, technological, industrial, and institutional) in the face of dependencies judged critical. The thesis does not announce sovereignty already achieved. It invites the reader to examine, chapter by chapter, what the sources allow us to establish, what remains plausible, and what stays open.

The reader will find distinctions here rather than slogans: adoption and sovereignty, localisation and control, list price and total cost, legal obligation and tested migration, funding and industrialisation. These distinctions do not close the debate. They make it possible.

The reading asked for is critical. Where series are missing, where an experiment has not been run, or where causal evaluation remains out of reach, the text says so. The strength of a technological essay is measured partly by what it refuses to invent.

Author's Note

This book grew from a question simple in its wording and demanding in its consequences: what remains of control when digital activity relies on shared infrastructures, often located outside the direct perimeter of those who use them?

I chose to treat together schools and laboratories, prices and total cost of ownership, standards and reversibility, industry and public procurement. Not to conflate everything, but because cloud dependencies become critical precisely where these dimensions intersect. Training without real access to environments, a contract without a tested exit, funding without an industrialised offering: each link in isolation can give the illusion of progress.

The project was built with an apparatus of sources and claims, and with the discipline of distinguishing facts, estimates, interpretations, and hypotheses. Some propositions deliberately remain working hypotheses. That is not a hidden weakness: it is a rule of reading.

The reader may use this work as a map of distinctions, as a method for reading an indicator or a political announcement, or as a framework for setting research priorities. They will find neither a general ranking of suppliers, nor global scientific validation decreed by the mere composition of the book.

I thank in advance those who read with rigour, correct a formulation, point out a limit, or reject an overly broad conclusion. An essay on cloud sovereignty deserves to exist only if it remains refutable.

Contents

Note to the reader . . . iii
Preface . . . iv
Author's Note . . . v
Introduction . . . 1

I Cloud, sovereignty and dependencies . . . 5

1 The cloud as economic and strategic infrastructure . . . 6
2 What does cloud sovereignty really mean? . . . 14
3 Mapping critical dependencies . . . 19

II Human capital and skills . . . 25

4 Training from school: digital culture and computer science . . . 26
5 Universities, laboratories, and teaching infrastructure . . . 30
6 Training, certifying, and retaining operational skills . . . 33

III Economics of cloud choices . . . 36

7 Comparing storage, compute, and network at equivalent scope . . . 37
8 Understanding total cost of ownership . . . 43
9 Experimenting, migrating, and reducing exit costs . . . 48

IV Technology, industry and resilience . . . 53

10 Open source software, standards, and interoperability . . . 54
11 Datacentres, semiconductors, energy, and supply chains . . . 58
12 Cybersecurity, resilience, and operational mastery . . . 63

V Research, markets and competition . . . 67

13 Research, funding, startups, and industrialisation . . . 68
14 Public procurement and industrial policy . . . 72
15 Adoption, competition, cooperation, and network effects . . . 77

VI Measurement, scenarios and autonomy . . . 82

16 Building a dashboard of capabilities and dependencies . . . 83
17 European scenarios, trade-offs, and possible policies . . . 88
18 Toward progressive, measurable, realistic cloud autonomy . . . 92
Conclusion . . . 96
Epilogue . . . 98

Introduction

Imagine an organisation that announces it has "migrated to the cloud." On-premises servers are gone. Monthly invoices have replaced depreciation. Teams work faster. Yet no one yet knows what a full exit would cost, who would respond to a serious incident outside business hours, or whether the skills needed to switch to another provider actually exist on the team. This scenario is not empirical proof. It is a typical situation. It shows why cloud sovereignty cannot be reduced to a box checked on a hosting contract.

This book starts from a simple observation: the cloud has become an economic and strategic infrastructure, not merely an IT tool choice. When a growing share of computing, storage, and software rests on third-party operators, questions of dependency, cost, skills, and industrial capacity cease to be peripheral. They become political in the broad sense: they concern what a society, an administration, or a company can still control, substitute, or rebuild.

1 Why the term "cloud sovereignty" invites confusion

The word sovereignty evokes, first of all, the power of a state. Applied to the cloud, it has drifted toward floating vocabulary. Sometimes it means locating data in Europe. Sometimes the nationality of a provider. Sometimes a certification. Sometimes an industrial ambition. These usages are not equivalent. Conflating them turns debate into slogans.

This book does not use "cloud sovereignty" as an already established European legal standard, nor as a single score. It is an **analytical framework**: a degree of capabilities in the face of dependencies judged critical. This framework distinguishes, without merging them, localisation, ownership, operational control, industrial autonomy, technical reversibility, and institutional levers. An organisation can comply with a localisation requirement and still be unable to exit. It can buy from a European provider and depend on components, skills, or platforms outside Europe. The reverse is also possible: dependence on a non-European actor is not automatically critical if substitutability, timelines, and impacts remain manageable.

2 Thesis and limits

The central thesis of this work is as follows. European cloud sovereignty results from the progressive development of interdependent human, economic, technological, industrial, and institutional capabilities, whose combination makes it possible to reduce critical dependencies on non-European actors.

This proposition does not claim that Europe has already achieved full cloud sovereignty. It does not claim that everything must be relocalised, nor that every European provider is safer, cheaper, or more sovereign. It does not claim that open source automatically guarantees reversibility, that a certification proves transferable operational competence, that a legal obligation demonstrates a successful migration, or that increased funding proves effective industrialisation.

The thesis is a **structuring hypothesis** to be examined chapter by chapter. Some parts of the reasoning rest on bounded, traceable facts. Others remain plausible mechanisms, private estimates, or propositions still insufficiently documented. The reader must be able to see this difference.

3 Five families of capabilities

To read the cloud other than through a label, the book groups capabilities into five families, themselves detailed in finer dimensions in chapter 2.

Human capabilities concern initial training, higher education, laboratories, and continuing education: without a pool able to operate, audit, and migrate, other levers remain fragile. **Economic** capabilities concern prices, total cost of ownership, lock-in effects, and market dynamics. **Technological** capabilities concern software, standards, portability, and exit protocols. **Industrial** capabilities concern datacentres, semiconductors, energy, supply chains, and operational mastery. **Institutional** capabilities concern law, public procurement, regulation, and cooperation among public actors.

None of these families suffices alone. A local datacentre without operational control does not produce mastery. A public budget without an industrialised offering does not create an alternative. Abundant training without an operational outlet does not automatically reduce critical dependencies.

4 Documentary method

The book relies on an apparatus of sources and claims. Each important factual claim is presented with an explicit confidence level. Levels range from solidly established facts to propositions still insufficiently documented. Sources examined are distinguished from pages merely consulted. When a document could not be archived or read in full, the text says so.

This method has a cost in reading: it slows conclusions. It has a benefit: it prevents turning a market-share estimate into a measure of sovereignty, or a regulation into proof of a successful migration. Footnotes refer to sources. They do not replace reasoning.

Two quantities, often cited together, illustrate the discipline required. In 2025, 52.7% of enterprises in the Union covered by the Eurostat survey declared that they used paid cloud services.[1] According to private estimates from July 2025, the share of "European" cloud providers in the regional infrastructure market would have remained around 15%, while AWS, Microsoft, and Google together would represent roughly 70%.[2] Adoption measures declared usage. Market shares estimate concentration of supply. Neither measures sovereignty.

5 Organisation of the book

The book comprises six parts and eighteen chapters.

The **first part** presents the cloud as infrastructure, clarifies the vocabulary of sovereignty, and analyses the mechanisms that make a dependency critical. The **second** deals with human capital: schools, universities, laboratories, continuing education. The **third** opens the economics of choices: price comparison at equivalent scope, total cost of ownership, migration and portability, free software and standards. The **fourth** addresses technology and industry: the material base, cybersecurity and resilience. The **fifth** links research, funding, public procurement, and market dynamics. The **sixth** proposes a dashboard without a composite score, analytical scenarios, and a synthesis of progressive autonomy.

Each chapter poses a question, develops a line of reasoning, presents objections, marks its limits, then opens toward what follows. The most advanced chapters already rest on an examined corpus. Others remain enriched analytical frameworks: their gaps are visible, not hidden.

[1] Eurostat, "53% EU enterprises used paid cloud services in 2025", and associated dataset.

[2] Synergy Research Group, 24 July 2025: private estimates, not auditable from the article alone.

6 What the reader may conclude, and what the book does not promise

By the end, the reader should be able to distinguish localisation and control, list price and TCO, legal text and industrialisation, adoption and sovereignty. They should have a method for reading a contract, a price basket, an indicator, or a political announcement without merging them.

The book does not promise a ranking of providers. It does not promise proof that such-and-such European policy has already succeeded. It does not promise an orchestration migration protocol as executed if the trial has not taken place. It does not fabricate a European percentage of ultimate ownership of operators when no homogeneous series has been identified. These absences are part of the inquiry's result.

What follows begins with the cloud itself: before speaking of sovereignty, one must understand what cloud computing has changed in the organisation of digital production.

PART I

Cloud, sovereignty and dependencies



CHAPTER 1

The cloud as economic and strategic infrastructure

For a long time, enterprise IT rested mainly on machines and software installed on the organisation's premises. One bought servers, depreciated them, and sized them for a peak load that was often hypothetical. *Cloud computing* shifted this relationship. A growing share of computing, storage, networking, and software is now supplied remotely, on demand, by third-party operators.

This shift is not an architectural detail. It changes how organisations invest, deploy, and evolve their systems. It also changes, more discreetly, what they can control, leave, or replace. Before speaking of sovereignty, one must therefore understand what the cloud has become: not a simple IT service among others, but an **infrastructure**, a technical and economic base on which other activities rest.

Two readings, drawn from digital policy analysis, help situate the issue without deciding it. On one hand, the cloud appears as a critical technology for the European Union's competitiveness in the digital economy. On the other, "data sovereignty" initiatives are linked there to the will to reduce strategic dependencies on foreign cloud providers dominant in the European market.[1] These formulations belong to an argued framing. They do not measure sovereignty, and they do not replace adoption or market-share data, which come from other sources.

1.1 What is meant by cloud

For this book, the cloud denotes a mode of supplying IT resources, compute, storage, network, platforms, applications, accessible over a network, made available on demand, and generally billed by usage or subscription.

This definition is primarily pedagogical. It converges with the synthesis offered by the Organisation for Economic Co-operation and Development (OECD) in its report on competition in the cloud, drawing on characteristics associated with Mell and Grance (2011): on-demand network access to a **shared pool** of configurable resources, rapidly provisioned and released.[2] The National Institute of Standards and Technology (NIST) primary publication is not treated here as an independently

[1] Filippo Gualtiero Blancato, "The cloud sovereignty nexus: How the European Union seeks to reverse strategic dependencies in its digital ecosystem", Policy & Internet, 2024; DOI 10.1002/poi3.358. Access examined via UCL abstract; full PDF not archived locally, no pagination invented.

[2] OECD, "Competition in the provision of cloud computing services", OECD Roundtables on Competition Policy Papers No. 323, 20 May 2025, DOI 10.1787/595859c5-en.

examined source; this institutional relay is relied upon without inventing NIST pagination.

Two properties explain much of the model's economic appeal. **Resource pooling** allows the provider to aggregate resources for several clients and allocate them dynamically: the client often chooses a region or a level of localisation, but does not necessarily control the exact physical machine. **Elasticity** allows capacities to be increased or reduced according to demand, which reduces the need to buy permanently "for the peak." Neither, on its own, implies legal control, easy reversibility, or any form of sovereignty.

1.2 Three service models, three responsibility splits

Behind the single word "cloud" lie very different responsibility contracts. Three service models are usually distinguished.

In **infrastructure as a service** (IaaS), the provider mainly supplies base servers, storage, and network; the client keeps operating systems, applications, data, and a large share of configuration. In **platform as a service** (PaaS), the provider runs the platform and execution tools; the client focuses on applications, at the cost of less control over lower layers. In **software as a service** (SaaS), the provider runs the complete application; the client mainly manages accounts, limited settings, and business data.

Model	What the provider mainly carries	What the client mainly carries
IaaS	Base infrastructure	Systems, applications, data, operations
PaaS	Platform and execution tools	Applications and their settings
SaaS	Complete application	Accounts, limited configuration, business data

The higher one moves toward SaaS, the more one delegates operations, and the more implicit choices (identity, region, interfaces, billing) may be imposed. The closer one stays to IaaS, the more technical levers one keeps, at the cost of heavier operational load. Adopting a "cloud service" therefore does not, by itself, say who controls what.

Deployment modes are added to these models. **Public cloud** serves a broad clientele, operated by a third party. **Private cloud** denotes an environment dedicated to one organisation. **Hybrid cloud** combines the client's infrastructure and third-party services. Public, private, hybrid on one side; IaaS, PaaS, SaaS on the other: the two

axes cross. Confusing them later obscures analysis of dependencies and dimensions of sovereignty.

Four functions run through almost all architectures useful to this book: **identity and access** (who can connect, with what rights); **network** (how services communicate); **orchestration** (automation of deployment and lifecycle); **data** (where they are stored, how they flow, who can extract or delete them). This is not an inventory of vulnerabilities: it is a compass for reading the layers.

1.3 A stack, not a flat product

The cloud can be pictured as a stack. At the bottom: equipment and physical sites. Above: virtualised infrastructure, often supplied as IaaS. Then platforms, often as PaaS. Then applications and software services, often as SaaS. Finally data and uses.

Each layer rests on those below. This image, conceptual and not quantified, justifies treating the cloud first as systemic infrastructure, before evaluating it in terms of sovereignty or lock-in.[3]

1.4 Who gains what in the chain

Around this stack gravitate actors with distinct interests: equipment suppliers; site and datacentre operators; *cloud service providers* (CSPs), who operate and sell IaaS, PaaS, or SaaS; software publishers often hosted by a third party; integrators and service operators; finally clients, who pay for business value.

Service flows, economic flows, and possible dependencies do not always coincide. A client may depend on a SaaS publisher, who depends on a CSP, who depends on equipment and energy. The criticality of these links, substitutability, timelines, impacts, will be the subject of a later chapter. Here, it suffices to note that the cloud value chain is not flat: it stacks intermediaries.

[3] Conceptual figures for the chapter: multilayer architecture and value chain (assets/diagrams). No empirical values are encoded there.

1.5 What market shares say, and what they do not

Providers' market shares do not measure the same thing as enterprise adoption. According to a Synergy Research Group article dated 24 July 2025, the share of "European" cloud providers in the regional cloud infrastructure market would have fallen from 29% in 2017 to 15% in 2022, then remained around 15% "since then"; AWS, Microsoft, and Google together would represent roughly 70% of the market at the time of publication; the European market size is estimated at 61 billion euros for 2024. The product scope stated explicitly in the article covers IaaS, PaaS, and *hosted private cloud*.[4]

These figures are **reported private estimates**, not Eurostat statistics. The article does not formally specify the criterion for "European cloud provider" (headquarters, capital, brand) beyond examples. They establish neither sovereignty nor its opposite. They indicate a concentration of supply as this source describes it, nothing more, and nothing less.

1.6 Adoption: a bounded fact, not proof of control

In 2025, 52.7% of enterprises in the European Union covered by the Eurostat survey declared that they used paid cloud computing services. This proportion rose by 7.4 percentage points since 2023 and by 34.9 points since 2014.[5]

What this statistic establishes is wide and growing diffusion of declared use of **paid** cloud services in the survey population. What it does not measure, by construction: intensity of use; quality, cost, or security of services; providers' shares. Strong adoption can coexist with concentrated supply. The two quantities do not contradict each other; they do not substitute for each other either.

Three frequent confusions deserve to be set aside from the outset. First, adoption does not equal market shares. Second, adoption does not equal data localisation or operational control: hosting in a European region does not guarantee effective legal control or the ability to exit. Third, adoption does not equal sovereignty, nor automatic reduction of critical dependencies. These mechanisms will be analysed further on; they are not deduced from the Eurostat rate.

On-premises IT and *edge computing* remain relevant for certain workloads. Recognising this does not weaken the finding: for a growing share of organisations covered by the survey, the cloud has become a structuring option.

---

[4] Synergy Research Group, "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. Private estimates (SIA™ method); not auditable from the article alone.

[5] Eurostat, "53% EU enterprises used paid cloud services in 2025", 3 February 2026, https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20260203-1; dataset isoc_cicce_use (indicator E_CC, enterprises ≥ 10 employees, EU27_2020), extraction 2026-10-11. Raw values 2014 = 17.82; 2023 = 45.32; 2025 = 52.74.

1.7 What this chapter allows us to hold

The cloud, as presented here, is infrastructure: pooled and elastic resources; service models with variable responsibilities; distinct deployment modes; layered stack; chain of actors and flows. Adoption data describe bounded diffusion. Market-share estimates describe concentration of supply, with strong reservations. The strategic framing situates the European issue without concluding that sovereignty has been achieved.

A prior question remains, without which the rest of the book would float: what can "cloud sovereignty" mean when jurisdiction, operational control, technology, data, economics, and industry intertwine, and what must we refuse to put under this word?

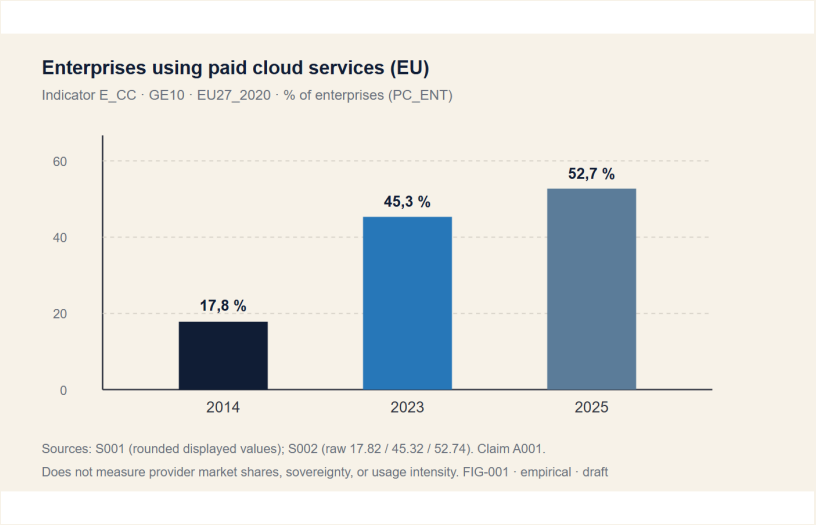


Figure 1.1: Declared adoption of paid cloud services (EU). Does not measure sovereignty.

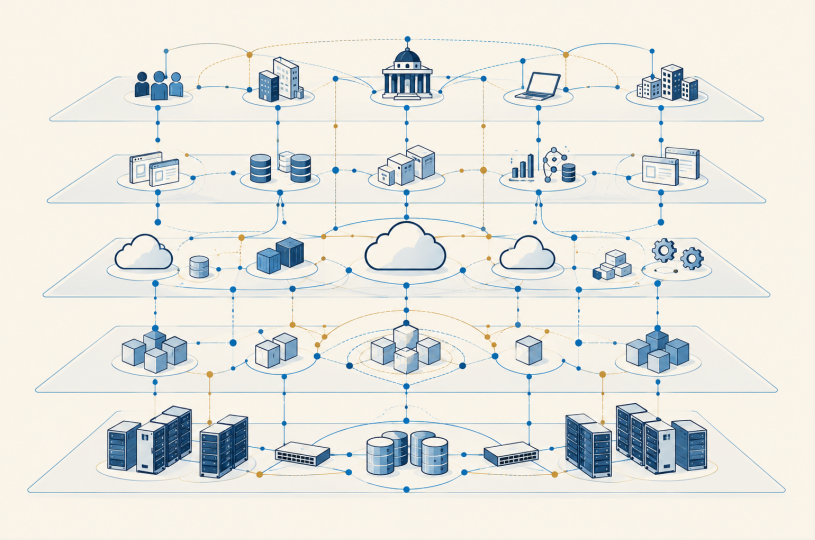


Figure 1.2: Multi-layer architecture - conceptual, not empirical

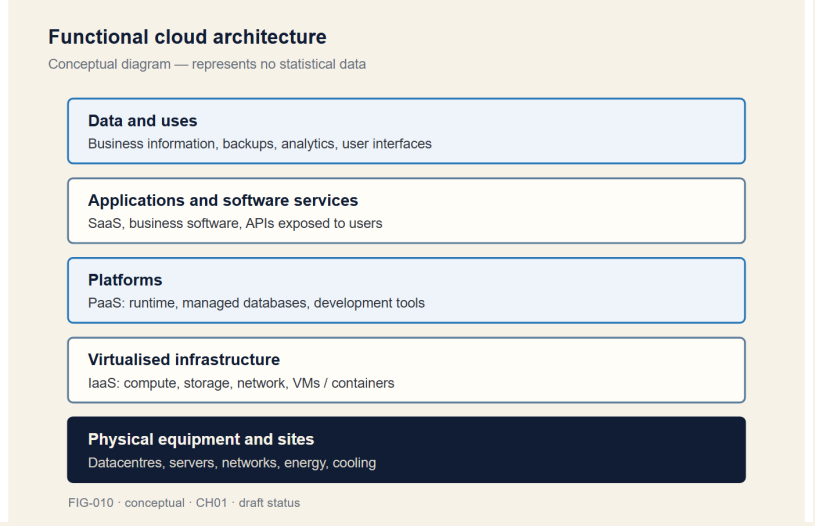


Figure 1.3: Five-layer architecture - pedagogical diagram, not empirical

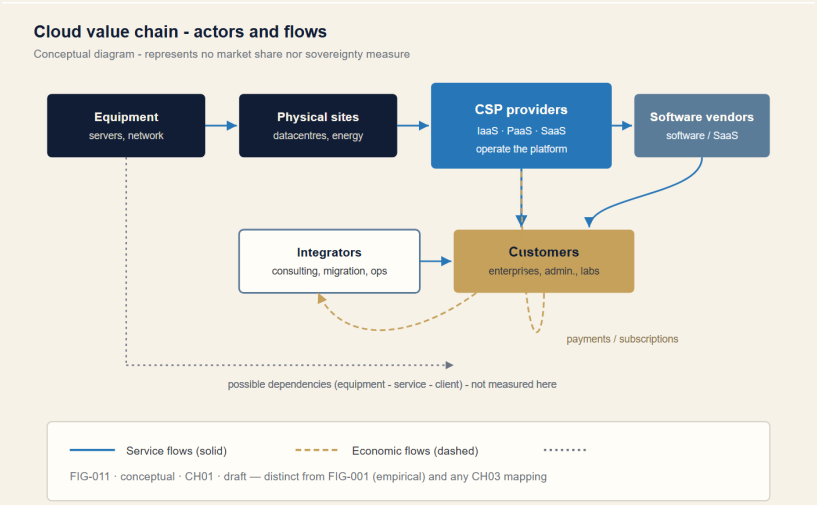


Figure 1.4: Value chain - pedagogical diagram, not empirical

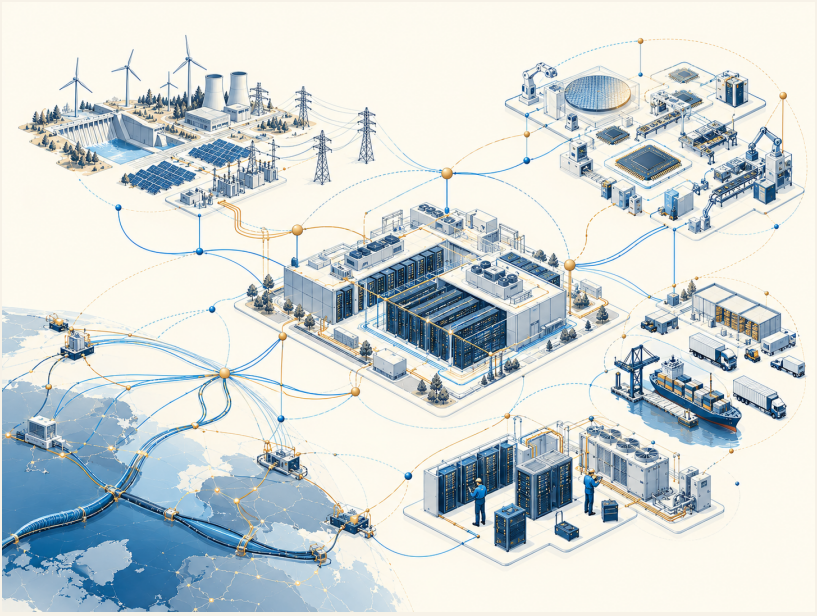


Figure 1.5: Physical infrastructure and network - conceptual illustration



CHAPTER 2

What does cloud sovereignty really mean?

The word "sovereignty" circulates today in public discourse, industrial strategies, and commercial catalogues. Applied to the cloud, it promises by turns control of data, technological autonomy, or simply locating servers on European territory. These usages are not interchangeable. As long as we do not specify *what* we are talking about, the thesis of this book, interdependent capabilities to reduce critical dependencies, remains undecidable.

The previous chapter showed that the cloud has become infrastructure, and that adopting a service is not equivalent to mastering it. Vocabulary must still be fixed. This chapter proposes a **reading framework**. It does not demonstrate that Europe has achieved cloud sovereignty, nor that it has missed it. It also refuses to turn this framework into a composite score.

2.1 What the political debate already covers

According to the framing proposed by Filippo Gualtiero Blancato in a 2024 article, cloud computing appears as a critical technology for the European Union's digital competitiveness, and "data sovereignty" initiatives there combine data protection and industrial policy instruments in the face of dominant foreign cloud providers.[1]

This framing, as it emerges from the public presentation of the article, without support from pagination of the full text, clarifies a political use of the vocabulary: linking protection and reduction of strategic dependencies. It does not establish the eight-dimension framework proposed below, nor an official norm, nor market shares, nor the finding that sovereignty would already be achieved. Figures on concentration of supply, when they appear, come from other sources, and from another chapter.

Three rhetorical slippages deserve to be named from the outset. First slippage: treating "data sovereignty" as synonymous with "cloud sovereignty." Mastery of data does not exhaust operational control, industrial base, or reversibility. Second slippage: treating European localisation as proof of control. Hosting is not deciding. Third slippage: treating a "European" provider as automatically more sovereign. A brand's nationality is not equivalent to substitutability of stacks, components, or skills.

[1] Filippo Gualtiero Blancato, "The cloud sovereignty nexus: How the European Union seeks to reverse strategic dependencies in its digital ecosystem", *Policy & Internet*, 2024, DOI 10.1002/poi3.358.

2.2 An institutional precedent: DigComp

Before introducing this book's own framework, a contrast is useful. The European DigComp framework, developed by the Commission and the Joint Research Centre, describes digital competences for daily life, learning, work, and social participation. DigComp 3.0 is the fifth edition; the framework presents itself as *technology-neutral* and adaptable.[2]

DigComp shows that an institutional multidimensional reference can exist without being a slogan. It does not define cloud sovereignty, however, and it does not measure skills for operating cloud infrastructures. Multiplying dimensions is therefore not, in itself, an arbitrary invention, provided the object measured is specified.

This contrast has pedagogical utility. It allows the book to propose several dimensions without imitating DigComp, and without suggesting that a framework for citizen competences already measures cloud autonomy. DigComp will be mobilised in the school chapter as a contrast; here it mainly serves to avoid confusing objects.

2.3 A degree of capabilities, not a label

In this book, **cloud sovereignty** denotes a **degree**, not a binary label, of capabilities making it possible to reduce **critical dependencies**. This analytical framework is specific to the work; **it is not an official European norm**.[3]

Eight dimensions organise the reading. Each is formulated as a question, not as a score.

Dimension	Operational question
Legal	Which rules and jurisdictions effectively apply to data, access, and contracts?
Operational	Who configures, supervises, responds to incidents, and controls identities?
Technological	Who masters stacks, interfaces, critical components, and their substitutability?
Data	Who can extract, port, delete, analyse, or restrict flows?
Economic	Are alternatives viable at relevant costs and timelines?
Industrial	Is there a base of datacentres, components, energy, and supply that can be mobilised?
Resilience	Does the system absorb shocks and failures without unacceptable rupture?

[2] European Commission / Joint Research Centre, *Digital Competence Framework* (DigComp) 3.0, JRC institutional page.

[3] Analytical framework of the work *La souveraineté cloud européenne*.

Dimension	Operational question
Reversibility	Has an exit been **tested**, and not merely declared?

These dimensions deliberately overlap. Data localisation is not applicable law; technical tools are not enough without operations; reversibility crosses both technique and economics. Figures associated with this chapter are conceptual: they illustrate these dimensions, families of capabilities, and the idea of substitution, without encoding any measure or score.

The framework refuses several frequent shortcuts: localisation alone; provider nationality; compliance alone; adoption rate; market shares; DigComp as an indicator of cloud sovereignty; any single composite score.

2.4 Useful tensions between dimensions

A buyer may require European localisation and discover, in an incident, that privileged access and supervision tools remain out of reach. An organisation may have a robust legal base and lack economically viable alternatives. An administration may fund datacentres and remain captive to non-substitutable PaaS interfaces. A team may succeed at exporting data and fail to rebuild the service.

These tensions are not invented case studies. They are analytical configurations. They justify keeping dimensions separate in the book's dashboard, rather than merging them into a single index.

2.5 Limits on the legal dimension

The legal dimension, to be consolidated beyond this reading framework, would call for Union law texts examined as primary sources, as well as doctrine and, where applicable, comparison with other capability frameworks. As long as these elements are not established in the same source apparatus, strong legal claims remain out of reach of an established fact. The *Data Act*, public procurement, and buyer frameworks will be treated further on as distinct sources, without being absorbed here into a general definition.

2.6 Objections

One objects that eight dimensions make the debate too complex. Complexity is the price of precision: a single word has already produced enough confusion. One objects that a score is needed to decide. A score without a validated method decides poorly; this book forbids it by default. One objects finally that political vocabulary suffices: political vocabulary clarifies intentions; it does not measure capabilities.

2.7 Vocabulary landmarks

The debate on "data sovereignty" links protection and strategic dependency, without measuring sovereignty by itself. DigComp reminds us that an institutional multidimensional framework can exist for another object, citizen digital competences. This book's vocabulary fixes a degree of capabilities in the face of critical dependencies, without an official norm and without a single score.

A distinct question remains, which the definition does not settle: *which* dependencies become critical, and *by what mechanisms* does lock-in operate once infrastructure is adopted?

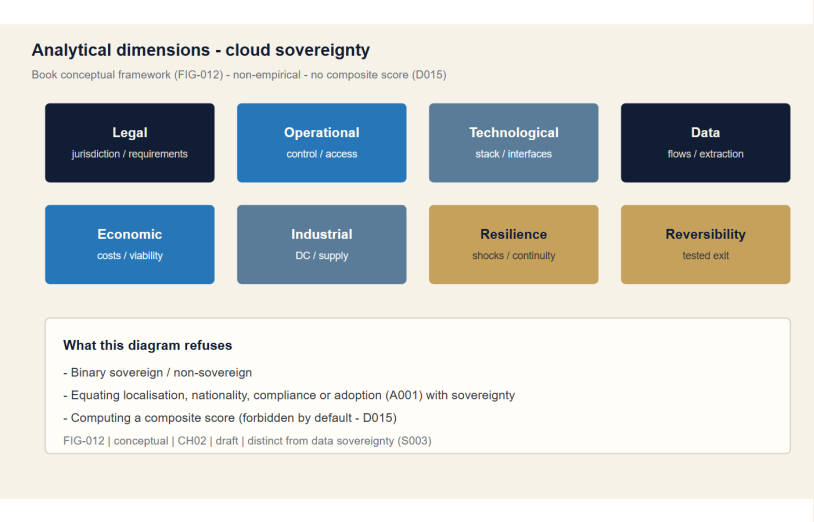


Figure 2.1: Eight analytical dimensions - not empirical



Figure 2.2: Reversibility and substitution - conceptual illustration



Figure 2.3: Five families of capabilities - conceptual illustration

CHAPTER 3

Mapping critical dependencies

Every organisation depends on other actors. The cloud does not create dependency: it redistributes its form. Once a service is adopted, the question is no longer only "who hosts?", but "can we substitute, in how much time, at what cost, with what impact?". When the answer is unfavourable, the dependency becomes **critical**.

The previous chapter fixed vocabulary: cloud sovereignty, in this book, denotes a degree of capabilities in the face of such dependencies. What remains is to understand *how* lock-in occurs. This chapter maps mechanisms. It does not quantify exit costs, nor validate a migration protocol: those questions, when treated, are elsewhere in the work.

3.1 Families of dependencies and criticality criteria

Several families can be distinguished, analytically: data; runtime and platform; software and SaaS; skills; legal; supply and physical. The figure associated with this typology is conceptual: it orders the reading; it measures no criticality.

Four criteria form here an **analytical grid specific to the book**, not an official method: limited substitutability; high exit timelines or costs; high impact in case of unavailability or rupture; information asymmetry between client and provider. Market concentration, difficulty of substitution, and strategic criticality must not be conflated. A systematic measure, workload by workload, at Union scale, is not yet available.

This grid has an immediate consequence. A dependency can be broad without being critical, if exit remains fast and inexpensive. Conversely, a narrow dependency, a single managed service, a single region, a single trained team, can become critical if the impact of a rupture is high. Volume of use does not measure criticality.

3.2 What the OECD documents on lock-in

The OECD report on competition in the provision of cloud services draws attention to several mechanisms that, in several jurisdictions, concern authorities.[1]

Concentration. Supply concentrates around a few large providers. This finding describes market structure; it does not prove, by itself, that *every* European organisation's dependency is critical. Private market-share estimates, "European" cloud providers around 15% after 2022, AWS–Microsoft–Google trio around 70% at the

[1] OECD, "Competition in the provision of cloud computing services", OECD Roundtables on Competition Policy Papers No. 323, 20 May 2025, DOI 10.1787/595859c5-en.

time of a July 2025 publication, provide **context**, not a map of non-substitutability.[2]

Complexity of switching. Barriers to changing provider (*switching*) include complex migration processes and technical, contractual, and financial obstacles, as well as limits to interoperability and standardisation. An ordinary dependency can become critical simply because exiting is long, costly, or risky. Reproducible, comparable European migration timelines and costs are still missing.

Data egress fees. Fees for moving data out of a cloud (*egress*) can make changing provider more expensive and contribute to *lock-in*. National studies sometimes report very high margins: these are secondary indices, not taken up here as established facts for the Union. A dated, homogeneous comparison would be necessary before any quantified generalisation. The chapter on prices will return to the need for a protocol before any ranking.

Credits and discounts. Free credits and volume discounts can influence initial choice and reinforce anchoring. The mechanism is commercial, distinct from technical competence. Its causal effect on European lock-in is not measured here. It suffices to note that a low entry price does not inform about exit cost.

Limited interoperability. Lack of standardisation between services complicates multi-cloud strategies. It is plausible, but remains a **hypothesis**, not demonstrated in this chapter, that dependence on differentiated managed interfaces and services, especially in PaaS and SaaS, often exceeds that of a more homogeneous IaaS. No map of critical interfaces at EU scale is produced here. The associated figure illustrates the interoperability issue; it does not measure its extent.

Multi-cloud coordination. Multi-cloud often appears as a response to dependence on a single provider. When implemented, it may, **as a hypothesis**, reduce certain risks while **increasing** integration costs and skill requirements. Documented European cases are missing here; no company case is invented to fill this gap. The associated figure contrasts a single hub and a multi-environment mesh, without claiming to describe empirical terrain.

3.3 Skills as a plausible mechanism

Even when technical portability is partial, lack of transferable operational skills can block an exit. An organisation may export objects and still be unable to rebuild the service for lack of teams trained on another stack. This point belongs to this chapter's typology; it is not demonstrated here. Training, certifications, and retention of skills will be the subject of the next part of the book.

[2] Synergy Research Group, "European Cloud Providers' Local Market Share Now Holds Steady at 15%", 24 July 2025.

3.4 What this chapter refuses to conclude

It refuses to say that every dependence on a non-European actor is critical. It refuses to say that concentration alone proves absence of sovereignty. It refuses to quantify an average European migration cost without a comparable series. It refuses to invent a workload-by-workload map. These refusals are not stylistic weaknesses: they protect the thesis against a demonstration too broad for the available evidence.

3.5 Useful objections

One sometimes objects that portability is "declared" in contracts or marketing pages; that *lock-in* is exaggerated as long as one stays on relatively standard IaaS; that multi-cloud "solves" the problem. These objections are serious. They are not settled without evidence of **tested** reversibility, duration, cost, functional gaps, rather than promised.

One also objects that egress fees will fall under regulatory pressure. Possible for certain items. That does not eliminate application debt, missing skills, or double running. Law can facilitate; it does not replace the trial.

3.6 What is documented, what remains open

Concentration, switching barriers, egress, commercial credits, and limits to interoperability are documented at institutional level by the OECD, with market context from private estimates. Their quantified European impact, workload by workload, remains open. Skills form a plausible mechanism, not yet proved on these pages.

The rest of the book leaves the register of technical and commercial mechanisms alone to examine a family of capabilities often neglected in sovereignty debates: human capital, beginning with schools, without claiming that schools alone produce cloud autonomy.

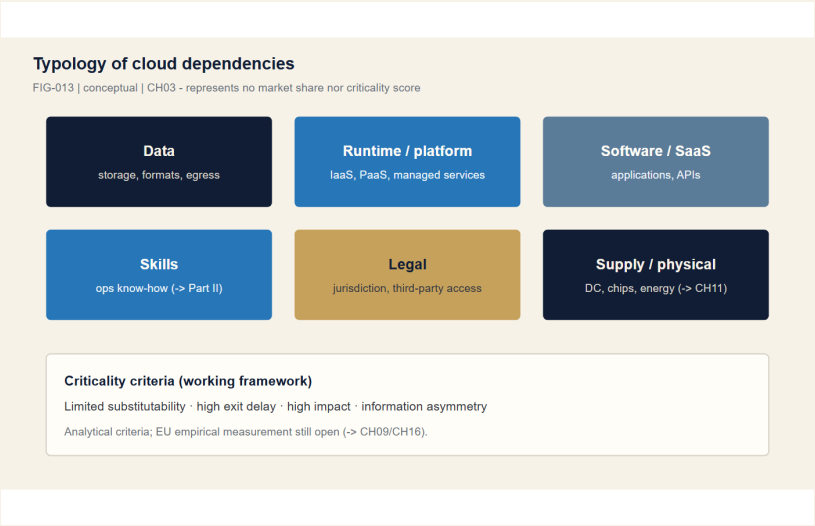


Figure 3.1: Six families of dependencies - conceptual

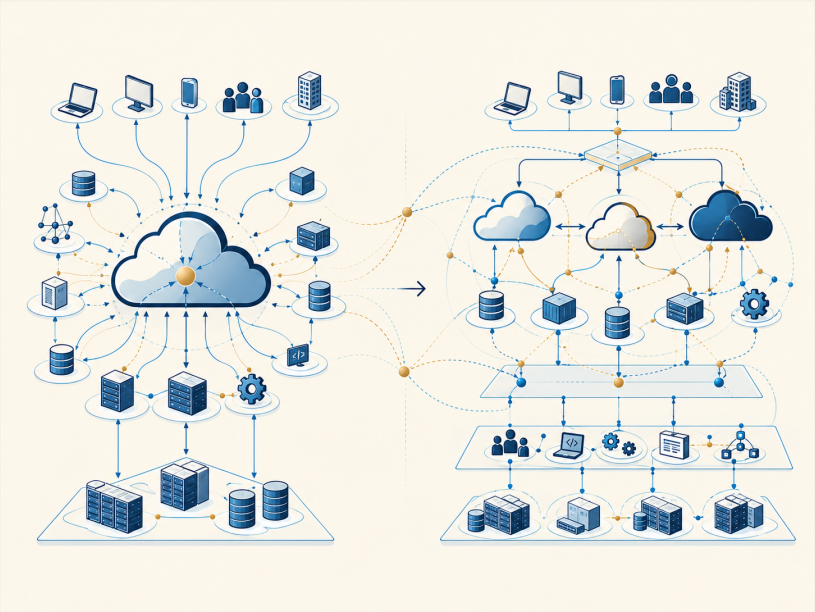


Figure 3.2: Concentration and multi-cloud - conceptual illustration

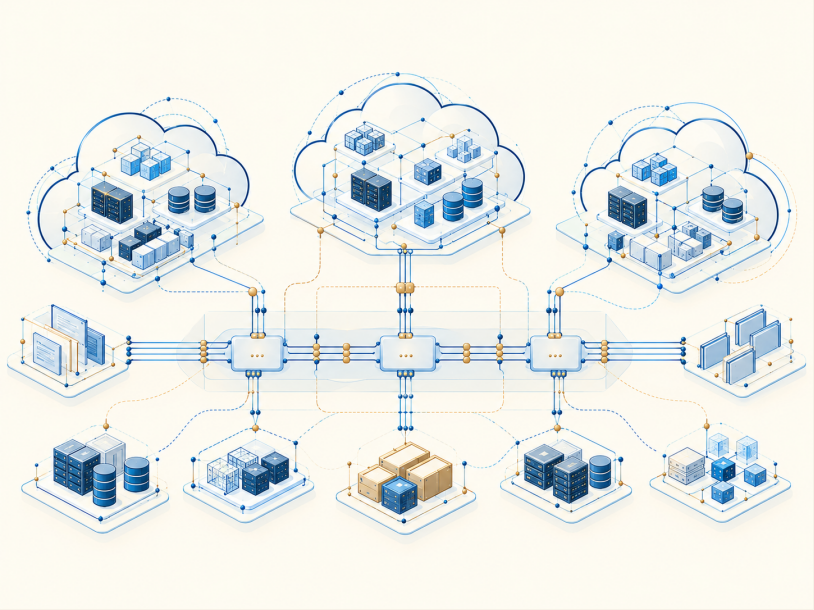


Figure 3.3: Multi-environment interoperability - conceptual illustration



PART II

Human capital and skills



CHAPTER 4

Training from school: digital culture and computer science

If the cloud has become infrastructure, where will those come from who can architect it, operate it, and, if need be, leave it? The most tempting answer is also the most misleading: "from school." Tempting, because a skills pool is not created in a single budget year. Misleading, because school mainly trains citizens and pupils, not platform operators.

Previous chapters fixed vocabulary, a degree of capabilities in the face of critical dependencies, and lock-in mechanisms. This chapter opens the **human capital** family through the school foundation. It examines a *potential* long-term role. It does not claim that school guarantees cloud sovereignty.

4.1 Three objects that must not be conflated

Before any promise, three notions must stay distinct.

Notion	Meaning in this chapter	What it is not
Digital skills	Uses for living, learning, working, participating (DigComp framework)	Cloud operations skills
Computer science education	Curricula, computational thinking, *informatics*	Proof of a pipeline toward cloud jobs
Operational cloud skills	Architecture, operations, security, cost control	What school measures today

Conflating the three lines produces excessive slogans: an introductory digital course is not operator training; a computer science module is not, by itself, a reversibility capability.

4.2 What European policy already pushes

According to the Council recommendation of 23 November 2023, Member States are invited in particular to improve the offer of digital skills and to support quality education in **computer science** or **computational thinking** at primary and secondary levels.[1]

This text orients. It is not binding. It does not define cloud administration

[1] Council of the European Union, Recommendation of 23 November 2023 on improving the provision of digital skills and competences in education and training, Official Journal of the EU.

skills. It does not measure national implementation. It does not prove an effect on sovereignty. The gaps it mentions, basic digital skills, information and communication technology specialists, international assessments, provide **context**, not demonstration of a school pipeline toward the cloud.

In the same vein, the Commission signals the development of **guidelines** for computer science teaching, intended to help schools strengthen digital competence in the classroom.[2] The existence of these guidelines, as it emerges from the institutional presentation, is neither an inventory of national curricula nor proof of homogeneous implementation. The pedagogical detail of the document is not taken up here as a line-by-line established fact.

4.3 DigComp: useful, but for another object

The DigComp framework describes citizen digital competences, *technology-neutral*.[3] It helps precisely to avoid confusing "digital skills" with skills for operating infrastructures. As the definitional chapter recalled, DigComp is not an indicator of cloud sovereignty.

Returning to DigComp here is not gratuitous repetition. In chapter 2 it served as a definitional contrast. Here it serves as a pedagogical safeguard: any progress measured in citizen digital skills cannot automatically be counted as progress in cloud capabilities.

4.4 A pipeline hypothesis, not a demonstration

Within this work's analytical framework, early school exposure to an **informatics** culture, and not only a digital one, is treated as a plausible long-term pipeline factor for European cloud skills.[4]

This hypothesis remains a hypothesis. No survey mobilised here measures a chain "school → cloud operations skills → reduction of critical dependencies." A conceptual figure may illustrate a training–practice–industry path; it does not establish causality. The evidential status of this proposition remains cautious.

The chain breaks in several places. That the Union encourages digital and computer science at school is documented. That States implement homogeneously remains unknown without a national inventory. That pupils acquire measurable informatics culture is partially observable elsewhere, but not treated here as independent proof. That this culture increases entry into operational cloud training is plausible,

[2]European Commission (DG Education and Culture / European Education Area), institutional page on computer science guidelines and EU support.

[3]European Commission / Joint Research Centre, *Digital Competence Framework* (DigComp) 3.0, JRC institutional page.

[4]Analytical framework of the work *La souveraineté cloud européenne* (school exposure factor).

not demonstrated. That these trainings then reduce critical dependencies belongs to the book's thesis, not to a result established in this chapter.

4.5 Inequalities: what could be examined

The policies themselves signal fractures, gender, basic skills, access. For the school foundation, candidate indicators, *not invented as results*, would for example be: presence of distinct computer science teaching; teacher training; territorial inequalities in equipment; results in informatics literacy. None of these indicators is produced here.

Naming these candidate indicators has a use: preventing a future school dashboard from inventing values to fill the void. An empty cell remains preferable to a fabricated European average.

4.6 Objections

One sometimes objects that the real lever is mainly continuing education for adults. The objection is serious: school's generational timelines are long, while technical stacks evolve fast. School remains, in this book, a possible **upstream condition**, not a unique lever.

One also objects that DigComp would suffice. Not for cloud operations jobs. Finally, the idea that school would "sovereignise" the cloud by itself is refused: an upstream human capability is not equivalent to mastery of critical infrastructure.

4.7 What school establishes, and what it does not prove

The Union pushes digital skills and computer science teaching at school; DigComp frames citizen digital life. A school pipeline toward cloud skills is **plausible**, not demonstrated. Homogeneous national implementation and the empirical link up to critical dependency remain open.

School alone does not produce cloud sovereignty. A distinct question remains, closer to jobs: does higher education train architects and operators, and under what conditions of real access to practice environments?



Figure 4.1: From learning to industry - conceptual illustration

CHAPTER 5

Universities, laboratories, and teaching infrastructure

A diploma may mention the cloud. A catalogue may promise laboratories. A harsher question remains: did the student actually deploy, observe, break, and repair something, or only read slides about distant services?

The previous chapter situated school as a possible foundation for long-term informatics culture, without conflating it with operations skills. This one changes scale. It examines higher education, universities, and access to means of practice. It does not claim that higher education alone guarantees cloud sovereignty.

5.1 What one risks counting too quickly

Three objects intersect here, easily merged.

Training in higher education, systems, networks, cybersecurity, cloud modules, can exist on paper without producing operational capability. **Real access** to laboratories or environments, credits, sandboxes, clusters, is another matter: a catalogue is not access, access is not use, use is not demonstrated competence. **Research** and transfer, finally, matter for the ecosystem, but counting publications as industrialisation would be a confusion; that thread will be taken up elsewhere in the work.

Continuing education for adults, vendor certifications, and transferability of operations skills belong to the next chapter, not this one.

5.2 A chain that must not be flattened

Between advertised offer and professional insertion lie several levels. The offer of courses, modules, or degrees is not inventoried here at European scale. Access to a usable environment, practice time, evaluation on real scenarios are not either. What can be documented, for now, is mainly a **context** of a workforce trained in information and communication technology (ICT), distinct from operational cloud skills.

This chain has an implication for the thesis. Strengthening only the nominal offer of cloud modules can produce diplomas without operators. Strengthening only research can produce articles without an industrialised offering. Higher education counts for cloud sovereignty mainly when it links training, real access, and demonstrable skills, something the current corpus does not yet measure pan-Europeanly.

5.3 A tertiary ICT pool, not a “cloud” counter

According to Eurostat’s statistical overview on ICT education, the number of people in the Union’s workforce with **ICT training** rose from more than 2.5 million in 2016 to 3.6 million in 2025; among them, 3.3 million were employed; in 2025, 74.8% of employed people with ICT training had **tertiary** education.[1]

These figures support a growing, strongly tertiary ICT training pool, a context relevant for reasoning about higher education. They do not establish the volume of “cloud” graduates, nor access to laboratories, nor operational quality, nor fine fit with large-scale infrastructure needs or sovereignty strategies. ICT training, in the survey sense, denotes qualifications at least at upper secondary level in specified computing fields; it is neither a cloud certification nor multi-cloud operations competence.

5.4 Real access as a hypothesis

Within this work’s analytical framework, **real** access to cloud laboratories and environments in higher education is a plausible condition for training competent architects and operators, **distinct** from the mere nominal existence of modules.[2]

This hypothesis is not measured here. No pan-European inventory of cloud laboratories is produced. As a reading tool, not an official method, three criteria help avoid confusing the poster with practice: an environment actually usable; operations scenarios (deployment, observability, incident); evaluation based on practice.

5.5 Pedagogical trade-offs, without inventory

Institutions that want students to practice the cloud encounter trade-offs that can be stated without inventing figures.

Option	Plausible advantage	Plausible risk
Credits or sandboxes from a major provider	Low entry cost, stacks close to the market	Anchoring in tools and skills with little transferability
Open software / local clouds	Transferability, control	Operations load, possible gap with the market
On-site institution laboratories	Local control	Investment, obsolescence, limited scale

No volume of educational credits or laboratories is invented. Primary sources for cloud education programmes remain to be examined before any mapping. The trade-off is not decided here: it is named to prevent a sovereignty discourse from

[1] Eurostat, “ICT education, a statistical overview”, Statistics Explained.

[2] Analytical framework of the work *La souveraineté cloud européenne* (access to practice environments).

ignoring the pedagogical cost of transferability.

5.6 Objections

One sometimes objects that the offer is already abundant. That may be true locally; without inventory, no pan-European verdict is grounded here. One also objects that the tertiary ICT pool proves that higher education "suffices." No: being employed with ICT training is neither cloud operations capability nor documented access to laboratories. Finally, research can enrich the ecosystem; it does not alone replace industrialisation of skills and infrastructures.

5.7 What higher education documents

A growing tertiary ICT pool is documented. Real access to practice environments as a condition of operational training remains a plausible **hypothesis**, not measured. Mapping of laboratories, actual use, and demonstrated cloud skills remain open. The associated claim keeps a cautious status.

Universities and laboratories do not suffice, in isolation, to guarantee cloud sovereignty. A link often decisive once diplomas are obtained remains: continuing education, certifications, and transferability of operations skills, where human lock-in can tighten or loosen.



CHAPTER 6

Training, certifying, and retaining operational skills

A badge may open an HR door. It does not say whether one can exit a platform, handle an incident on another stack, or retain skills when tools change. Between the pool of information and communication technology (ICT) specialists and **transferable** operational capability, the gap is often ignored, and politically comfortable.

Previous chapters situated school and higher education. This one treats the operational link of human capital: continuing education, certifications, practice, and retention of skills. It does not redefine cloud sovereignty. It does not claim that training suffices to guarantee it.

6.1 Four levels that must not be merged

Level	Object	What it does not prove
General digital skills	Frameworks such as DigComp	Ability to operate a critical cloud
ICT specialist employment	Statistical aggregates	Multi-vendor cloud skills
Vendor-specific skills	Certifications, consoles	Transferability or reversibility
Demonstrated operational capability	Incidents, runbooks, migrations, service objectives	Automatic sovereignty

An ICT specialist is not, by construction, a cloud operations specialist. A vendor badge is not, by construction, proof of exit.

These four levels also structure the book's dashboard. Mixing DigComp, ICT employment, and cloud badges in a single indicator of "sovereign human capital" would produce exactly the false precision that chapter 16 refuses.

6.2 A large ICT pool, absent ops measure

In 2025, 10.45 million people were employed as ICT specialists in the European Union, or 5.0% of total employment.[1]

This *employment* pool differs from the pool of people *trained* in ICT mentioned in the previous chapter: these are two quantities, not interchangeable. The Union

[1] Eurostat, "Number of ICT specialists in the EU continues to grow", 27 May 2026.

has a large ICT reservoir. These aggregates **do not measure** operational cloud skills, nor their transferability from one ecosystem to another.

6.3 Certification as signal, not proof of transfer

Official pages for foundational programmes from AWS, Azure, and Google Cloud describe exams centred on concepts and services of each ecosystem, Cloud Practitioner, Azure Fundamentals, Cloud Digital Leader.[2][3][4]

Within this work's analytical framework, a vendor certification is treated as a plausible **market signal**, not as automatic proof of transferable operational capability outside the vendor's ecosystem.[5]

This hypothesis remains a hypothesis. No empirical multi-cloud transfer survey is mobilised here. The existence of badges does not authorise concluding effective skill lock-in, only taking seriously *plausible* mechanisms: learning curve of consoles and programming interfaces; procedures tied to one provider; HR incentives around the dominant badge.

The point is not to accuse certifications. They can be useful for initial hiring or upskilling on a given stack. The point is to refuse counting them as proof of reversibility or sovereignty.

6.4 Continuing training, retention under penalty of depreciation

Continuing education aims to maintain operational employability as platforms evolve. It differs from school exposure, higher-education laboratories, and research transfer.

Without regular practice, incidents, recovery exercises, controlled migrations, skill capital **may** depreciate. Plausible claim, not measured on these pages. Depreciation is not a physical law demonstrated here; it is an operational risk that any serious dashboard must leave visible rather than fill by default.

[2] Amazon Web Services, "AWS Certified Cloud Practitioner" page.
[3] Microsoft, "Microsoft Certified: Azure Fundamentals" page.
[4] Google Cloud, "Cloud Digital Leader certification" page.
[5] Analytical framework of the work *La souveraineté cloud européenne* (certifications and transferability).

6.5 Objections

One objects that the market "validates" certifications. That holds for the local labour market, not for infrastructure reversibility. One objects that general skills suffice: faced with the specificity of managed stacks, the framing argument is serious, but it is not a new empirical demonstration here. One objects finally that training is not the main lever: even then, absence of operators with transferable skills can block an exit that is technically conceivable.

6.6 Badge, practice, transferability

This chapter fixes a rule of reading: distinguish badge, practice, and transferability. The ICT pool describes an order of magnitude of employment, not operational autonomy. The hypothesis of non-automatic transfer of certifications keeps a cautious status.

Another family of capabilities remains, often mobilised in public debate with as many slogans as proofs: **economic accessibility**. Before speaking of "cheaper", one needs a method to compare storage, compute, and network at a truly equivalent scope.



PART III

Economics of cloud choices



CHAPTER 7

Comparing storage, compute, and network at equivalent scope

"Cheaper" is a short sentence. An honest cloud comparison is a long one: which date, which region, which service class, which redundancy, which transfer volume, which currency, which taxes, which invisible discounts? Without these details, the ranking is only numbered rhetoric.

The previous chapter treated operational human capital. This one opens economic accessibility through **method**. It presents a partial catalogue basket scenario, limited, dated, illustrative study, reproducible. Economic accessibility, in this book's sense, begins here: not with a winner, but with the ability to read a price.

7.1 What does a price per GB, per vCPU, or per hour measure?

A public tariff is never a bare number. A "per GB" price assumes a unit. In cloud billing, the usual unit is often the *gibibyte* (GiB, 2^30 bytes), even when the marketing page writes "GB". The gap with the decimal gigabyte (10^9 bytes) is about 7%: negligible for conversation, not negligible for a basket of several tebibytes. A "per vCPU" price assumes a processor generation, a memory ratio, sometimes a network or I/O performance tier. A "per hour" price assumes a machine switched on, an operating system, sometimes a system disk not included in the compute line.

Comparing two catalogue lines without aligning these assumptions amounts to comparing different objects. Two "2 vCPU / 4 GiB" machines may belong to distinct families. Two "standard" storages may differ in redundancy, single zone versus multiple zones, and thus in operational risk, even if the price per GiB-month looks similar.

7.2 Why a protocol before any figure

Without an explicit protocol, date, region, configuration, volume, availability, transfer, taxes and currencies, assumptions, a price ranking is not sustainable as an empirical result.[1]

Lock-in mechanisms mentioned earlier, data egress fees, switching barriers, make this requirement clear: a low storage price can mask high transfer cost. Ap-

[1] Methodological framework of the work *La souveraineté cloud européenne* (comparison at equivalent scope).

parently cheap compute can become expensive once the workload massively moves data out or depends on a non-substitutable managed service. Competition analyses already stress that data extraction fees (*egress*) can make changing provider more expensive.[2]

Four kinds of amount are distinguished. **List price** is the public grid. **Discount**, commitment, volume, credits, is often opaque. **Observed cost** is the real invoice: it is **absent** from this chapter. **Estimate** multiplies an explicit quantity by a list price. Confusing estimate and invoice is the first error; confusing partial estimate and general ranking is the second.

7.3 Minimal dimensions of a comparison

Dimension	Examples	Risk if omitted
Quote date	Calendar day of price	Obsolete comparison
Pricing region	Central Europe, Western Europe	Intra-European gaps
Service and class	Standard storage vs infrequent access; machine family	Non-equivalence
Performance	Virtual processors, memory, I/O operations	“Cheaper” but slower
Capacity / duration / volume	GiB-month, hours, GiB egress	Incomparable scalars
Availability / redundancy	Multi-zone vs locally redundant storage	Risk not priced
Transfer	Internet egress, inter-region	Underestimation of lock-in
Taxes / currencies	Excl. tax / incl. tax; USD / EUR	Illusion of gap
Commercial terms	Catalogue vs discount	Catalogue ≠ invoice

These dimensions do not guarantee perfect equivalence. They make inequivalence **visible**. An organisation that publishes these bounds can be contradicted on facts; one that publishes a podium without bounds cannot even be audited.

7.4 Region, redundancy, storage, compute, transfer

Five technical levers often weigh more than the commercial slogan.

Region fixes a catalogue, sometimes latency, and sometimes a localisation constraint. Two “European” regions are not a single pricing market. The basket uses

[2] OECD, “Competition in the provision of cloud computing services”, OECD Roundtables on Competition Policy Papers No. 323, 20 May 2025, DOI 10.1787/595859c5-en.

`eu-central-1` for Amazon Web Services and `europe-west1` for Google Cloud: this choice is documented; it is not a claim of strict geographic equivalence.

Redundancy transforms the meaning of a storage price. Lowering price by lowering redundancy is not savings: it is risk transfer. **Storage** comes in classes and operation costs; the basket simplifies by retaining a "standard" class. **Compute** combines machine family, operating system, and uptime duration. **Transfer**, especially Internet egress, links the month's invoice to the difficulty of exiting later. The next chapter will integrate it into total cost; the reversibility chapter will treat it as a migration component. Here, it suffices to include it so it ceases to be invisible.

7.5 Free tiers, discounts, credits: why the catalogue is not the invoice

Free-tier mechanisms complicate comparisons. One provider may offer a global egress allowance; another a very limited one; a third external traffic advertised as free on a commercial page. As long as these allowances are not aligned, and as long as real account traffic is unknown, comparing "egress" lines at marginal rate is only an approximation.

Commitment discounts and entry credits often lower the apparent unit price. They can also reinforce commercial anchoring: low entry price does not inform about exit cost. Competition analyses already signal the plausible role of credits and discounts in initial choice; their exact causal effect on European lock-in is not measured here.[3] This chapter refuses to invent an "average European discount". It compares list prices without commitment, then warns that the real invoice may differ.

7.6 Mechanism: how "cheaper" becomes "more expensive"

Imagine, **illustrative hypothetical scenario**, not an observed case, an organisation that chooses cheaper catalogue storage but with local redundancy, while its workload requires multi-zone resilience. The initial price differential disappears once redundancy, backups, or incident absorption must be added. Add underestimated egress volume: the egress line can exceed savings on storage. Add finally a team trained only on the dominant provider's console: exit cost does not appear on month 1's invoice, but on month 18's migration project.

The causal mechanism reads in four steps. First, the visible tariff steers purchase choice. Second, omitted items, redundancy, egress, performance, correct operating cost. Third, skills and proprietary services fix future exit cost. Fourth, the sum of these items can reverse the apparent advantage of month 1.

[3] OECD, "Competition in the provision of cloud computing services", *op. cit.*

7.7 Illustrative study: basket SCE-BASKET-01

The scenario retains, for a conventional month of 730 hours, excl. tax and **without** commitment discounts: 1,024 GiB-month of storage, a Linux machine of roughly 2 vCPU and 4 GiB memory, and 100 GiB egress. This basket is a **limited illustrative study**. It is neither a quote nor a production architecture.

Working formulas: hourly price × 730 for compute; price per GiB-month × 1,024 for storage (or equivalent hourly × 730 × 1,024); billable quantity × marginal rate for transfer, after documented allowance. These multiplications are reproducible from the cited quotes. They do not fabricate missing values to invent a multi-vendor total.

7.7.1 Basket coverage: by line item, without consolidated total

Line item	AWS (`eu-central-1`)	Google Cloud (`europe-west1`)	Azure	OVHcloud
"Standard" object / blob storage	Complete (catalogue)	Complete (catalogue)	Partial (Hot LRS ≠ multi-zone)	Partial (multi-zone; €)
Small Linux compute	Complete (*t3.medium*)	Complete (*e2-medium*, not identical)	Absent	Absent
100 GiB egress	Complete (marginal rate)	Complete (Premium EU)	Absent	Advertised free on page consulted, not comparable
Currency / taxes	USD, excl. tax	USD, excl. tax	USD, excl. tax (fragment)	EUR, excl. tax (fragment)
Quote date	2026-10-11	2026-10-13	2026-10-11	2026-10-11

Read this matrix as a map of *documentary availability*, not as a ranking. AWS and Google Cloud machines are close in label, not in identical performance guarantee. Redundancies and allowances differ. Currencies and equivalences still oppose any single total.

7.7.2 Separate catalogue readings (estimates, not a podium)

Amazon Web Services, Price List API quote **11 October 2026**, region `eu-central-1`: standard object storage (0.0245 USD/GiB-month × 1,024) **25.088** dollars; *t3.medium* Linux compute (0.048 USD/hour × 730) **35.04** dollars; transfer at marginal rate (0.09 USD/GiB × 100) **9.00** dollars. Catalogue infrastructure subtotal for *these three lines only*: **69.128** dollars.[4] A 100 GB transfer *may* fall within a global

[4] Amazon Web Services, Price List API, object storage, compute and transfer, Europe (Frankfurt) region, extraction 2026-10-11.

allowance depending on the account; the basket publishes the marginal rate to keep the line visible.

Google Cloud, quote **13 October 2026**, region `europe-west1`: regional Standard *Cloud Storage*; `e2-medium` machine; 100 GiB *Premium Tier* toward Europe. Estimate: storage **20.480** dollars; compute **24.459** dollars; transfer **11.880** dollars. Subtotal for *these three lines only*: about **56.819** dollars.[5] This subtotal is not a ranking against AWS. Non-equivalences in the matrix above forbid it.

Azure, fragment: *Hot* storage in locally redundant storage (Western Europe), extension of about **20.07** dollars for 1,024 units, LRS ≠ multi-zone.[6] **OVHcloud**, fragment: standard multi-zone storage, order of magnitude about **12.80** euros excl. tax; external traffic advertised as free on page consulted, units and scope not comparable to AWS quote.[7]

A UK regulator documented, within its national scope, high egress fee gaps between certain actors.[8] This index contextualises the importance of egress; it does not calibrate the European basket.

7.8 What data to gather before concluding an advantage?

Before concluding that a service is economically advantageous, an organisation should at minimum gather: date and URL of the quote; region; service class and redundancy; exact compute configuration; storage and transfer volumes estimated on the real trajectory; tax regime and currency; existence or not of commitments; and, if possible, an invoice extract or quote. Without these elements, it has a catalogue estimate, useful, but insufficient for an architecture decision.

[5] Google Cloud, public pages Compute Engine, Cloud Storage pricing and VPC network pricing, consultation 2026-10-13.

[6] Microsoft Azure, Retail Prices API, Blob Hot LRS, West Europe, extraction 2026-10-11.

[7] OVHcloud, Public Cloud pricing page, consultation 2026-10-11.

[8] Ofcom, *Cloud services market study final report* (extract), 2023: United Kingdom, not a European Union statistic.

7.9 Frequent objections

One objects that organisations always obtain discounts and that the catalogue is therefore misleading. The objection is fair: that is why the chapter distinguishes catalogue and invoice. It does not justify inventing average discounts. One objects that a “European” provider would necessarily be cheaper. **That is not demonstrated** on these pages. One objects finally that the method is too heavy for a buyer in a hurry. Without it, the buyer mainly presses a premature conclusion.

Missing data remain open: strict multi-vendor equivalence, taxes, discounts, observed costs, Azure and OVHcloud machines aligned line by line. They will not be filled with invented values.

7.10 What this changes for sovereignty

Economic accessibility is a capability: it conditions adoption of alternatives. But an alternative only “cheaper” on one catalogue line, without equivalence or exit, does not automatically reduce a critical dependency. The comparison method is an instrument of sovereignty in this book’s sense, a capability to read prices, not an industrial ranking.

What remains is to aggregate these lines, and others often invisible on a public grid, in a transparent total cost of ownership, without double counting and without inventing missing parameters. That is the subject of the next chapter.



CHAPTER 8

Understanding total cost of ownership

One can "save" on the list price of a virtual machine and lose on data egress, on-call staff, or a poorly amortised migration. Total cost of ownership (TCO) is not the tariff of a single invoice line. It is the sum, made explicit, of what an organisation commits to run, secure, and, if need be, leave a system.

The previous chapter showed that price comparison requires a protocol. This one aggregates categories. As long as determining parameters remain uncalibrated, the model publishes partial subtotals and gaps, not a disguised total.

8.1 Infrastructure spend, operating cost, total cost

Three notions stack.

Infrastructure spend groups what the public grid often bills first: compute, storage, operating transfer. This is the territory of the previous chapter's basket. It is necessary; it is not sufficient.

Operating cost adds what keeps the service running day to day: monitoring, backups, patches, support, on-call, incident management. Part may appear on the cloud invoice; part in salaries, third-party tools, or internal time.

Total cost of ownership widens further: training, security, dedicated connectivity, initial migration project, double running during a transition, and exit costs. It looks at a trajectory, not only a month.

Reducing TCO to three catalogue lines is to confuse the month's receipt with the cost of the trajectory. Economic accessibility, in this book's sense, concerns the trajectory.

8.2 Categories in plain language

Before any formula, here is what each cost family denotes concretely.

Infrastructure: machines, storage, operating transfer in steady state, what the catalogue displays most clearly.

Operations: keeping the service running without major interruption, monitoring, backups, patches.

Personnel: salaries and on-call of people who operate, diagnose, and respond to incidents.

Security: identities, detection, audits, security incident response.

Connectivity: dedicated links, content delivery, peering, networks outside the base catalogue.

Training: certifications, laboratories, upskilling beyond a single console.

Initial migration: switchover project amortised over the analysis horizon.

Double running: period when old and new environments coexist.

Reversibility: export, rewrite, exit tests, and possible unavailability when leaving later.

A transparent model must keep these categories **disjoint**, separate what is observed, catalogue, estimated, hypothetical, or scenarised, and refuse to fill gaps with invented averages.[1]

It also specifies four bounds: **analysis duration** (twelve months does not equal thirty-six); **scope** (IaaS does not equal SaaS); **load** (stable or seasonal); **assumptions** labelled as such.

8.3 A single formula

Once categories are named, the working model SCE-TCO-01 reads:

C_TCO = C_infra + C_ops + C_personnel + C_securite + C_connectivite + C_formation + C_migration + C_double_run + C_reversibilite

This formula is a **calculation framework**, not a complete empirical estimate. As soon as a critical term remains hypothetical or unknown, **no TCO total** is announced: the model refuses to publish a disguised complete TCO.

Category	Examples	Status in working model
Infrastructure	Compute, storage, operating transfer	Partial estimate (AWS catalogue from basket)
Operations	Monitoring, backups, patches	Unknown / to separate if measured

[1] Analytical framework of the work *La souveraineté cloud européenne* (total cost of ownership model).

Category	Examples	Status in working model
Personnel	Operational salaries, on-call	Hypothetical, not calibrated
Security	Identities, detection, audits	Hypothetical, not calibrated
Connectivity	Links, CDN, peering	Hypothetical, not calibrated
Training	Certifications, laboratories	Hypothetical, not calibrated
Migration / startup	Initial project amortised	Hypothetical, not calibrated
Double running	Coexistence during switchover	Hypothetical, not calibrated
Reversibility	Exit transfer, exit tests	Hypothetical, not calibrated

Anti-double-counting rule: operating transfer in steady state stays in infrastructure; **one-off** exit transfer belongs to reversibility **or** amortised migration, never both at once.

8.4 Partial infrastructure subtotal

To anchor reasoning, the AWS catalogue infrastructure subtotal from the previous basket (transfer at marginal rate) is:

25.088 + 35.04 + 9.00 = **69.128 dollars**.[2]

This amount is a **partial catalogue infrastructure estimate**. It is neither a client invoice nor a provider ranking. Azure, OVHcloud, and Google Cloud are not aggregated here into a single total.

8.5 Why architecture choices shift costs

Personnel is not an accessory: it often decides whether a low catalogue stays advantageous. Monitoring and support turn an architecture into a maintained service. Security and availability shift amounts between infrastructure, operations, and, indirectly, reversibility. Under-investment can lower the visible invoice and increase the cost of an outage, rarely visible before the outage.

A highly managed architecture can reduce need for operations personnel while increasing dependence on proprietary services, and thus rewrite cost on exit. An architecture closer to infrastructure may require more on-call while facilitating substitution. These trade-offs are plausible mechanisms; they are not quantified here as European constants.

[2]Amazon Web Services, Price List API, catalogue basket lines Europe (Frankfurt) taken from previous chapter.

Automation follows the same logic. Infrastructure-as-code scripts can lower marginal operations cost and facilitate a later migration, *if* they are not tied to proprietary extensions. Otherwise they mainly accelerate deployment *within* an ecosystem.

8.6 Migration, double running, and reversibility

Initial migration amortises a project: design, tests, switchover. **Double running** denotes the period when both environments coexist: it protects continuity, at the price of temporary doubling of certain lines. **Reversibility** denotes what will be needed to exit later. It will be the heart of the next chapter; it already appears in the formula because a TCO without exit is a TCO of assumed, or forgotten, captivity.

8.7 Qualitative sensitivity: without invented simulation

Certain parameters, **once sourced**, can reverse an apparent catalogue advantage.

Parameter	Plausible direction	Status
Recurrent transfer volume	Increase favours low egress fees	To calibrate
Intensity of proprietary platforms	Increase raises rewrite cost	Hypothesis
Personnel / infrastructure ratio	Increase reduces advantage of low catalogue	Hypothesis
Multi-year commitment	Lowers unit price, raises exit cost	Qualitative
Analysis duration	Lengthening amortises or reveals exit costs	Method

This table is not a quantitative simulation. No numeric value is invented for uncalibrated parameters. When data are incomplete: list the gaps; estimate only what is measurable; identify lines that would reverse the trade-off; decide whether remaining uncertainty is acceptable.

8.8 Worksheet to complete: illustrative hypothesis

Suppose, **unobserved hypothesis**, an organisation has the catalogue infrastructure subtotal of 69.128 dollars for one month. It does not yet know its on-call, monitoring, or cost of an exit test. Rather than invent an “average TCO”, it can fill the following sheet.

This sheet is not a real case. It shows how to refuse false precision.

Field	Content to fill	Current status (example)
Observed data	Invoices, tickets, incident time	To complete
Catalogue data	Dated grid lines	Partial (infra 69.128 USD)
Missing data	On-call, monitoring, exit quote…	Open list
Explicit assumptions	Analysis duration, scope, load	To write
Validation owner	Name / internal role	To designate

8.9 Objections and pitfalls

FinOps, cloud spend management, is useful for the current invoice; it remains incomplete for reversibility and skills. Industry benchmarks enter here only if sourced, dated, and scoped. A TCO that is too broad can become unusable: an honest subtotal is better.

A subtler pitfall: by loading exit cost, one might make every "European" alternative artificially expensive, while staying with a hyperscaler would avoid this immediate cost. The objection is serious. It does not justify excluding exit: it justifies separating it, dating it, and comparing it to other risks, unavailability, change in commercial terms, regulatory obligation.

8.10 Bridge toward exit

An economic capability for sovereignty is not the lowest price of the month. It is the ability to see the lines that fix dependency: skills, double running, exit. TCO does not alone establish sovereignty. It makes visible economic conditions without which the sovereignty discourse remains undecidable.

A distinct question remains: **exit**, what European law requires, what competition analyses describe, and what must be *tested* to speak of real reversibility rather than promised portability.



CHAPTER 9

Experimenting, migrating, and reducing exit costs

A contract may promise portability. A marketing page may claim that data "exit". A legal obligation may even forbid certain provider-switching fees. The only question that decides remains: has one **tested** an exit, duration, cost, functional gaps, or settled for declaring it?

This chapter reports **no** migration experiment already run for this work. It proposes a proof grid, a reminder of European law, and a checklist to *prepare* a credible test. A containerised protocol was formalised on paper; **it was not executed**.

Previous chapters mapped lock-in mechanisms, a comparison method, and a total-cost grid. This one asks: *can we exit, in how much time, at what cost actually experienced?*

9.1 Four levels of reversibility

Level	Meaning	Example of proof
Declared	Marketing promise or generic clause	The word "portable" on a page
Documented	Procedure, formats, described export interfaces	Specification, requirements document
Verified interoperability	Same workload tested on at least two environments	Test log, criteria met
Successful migration	Measured switchover or double execution	Duration, cost, functional gaps

A legal obligation can strengthen documentation and reduce certain switching fees. It does **not**, by itself, constitute a tested migration. To speak of capability, this book requires at least the third level, and, for a critical workload, the fourth.

9.2 Why being able to leave is not being able to migrate

Contractual possibility to terminate opens a legal door. It does not rebuild an application.

Proprietary application programming interfaces (*APIs*) often fix calls with no exact equivalent elsewhere. Managed services, queues, databases with specific extensions, ecosystem-tied functions, speed initial development and weigh down rewrite. Incomplete export formats sometimes allow recovering files without recov-

ering schema, metadata, or rights. Skills may be tightly tied to one console: even with a right to exit, the team may lack know-how to rebuild elsewhere.

Data export cost, *egress*, is only one component. Rewrite, tests, coexistence of both environments, possible unavailability, and team time are added. Recurrent operating egress from the previous chapter is not the same object as one-off exit transfer.

9.3 What competition analyses document

According to the OECD report already used for dependency mechanisms, the sector draws attention in particular to switching barriers, migration, interoperability, and practices such as high fees for transferring data out of the cloud.[1]

These barriers do not prove that *every* European workload is captive. They show why reversibility deserves to be treated as a capability. A British index also links high exit fees to multi-cloud brakes; its scope is the United Kingdom, not a Union statistic.[2] No average European migration cost figure was found for this chapter; none is invented.

9.4 What the Data Act changes, and what it does not guarantee

Regulation (EU) 2023/2854, the *Data Act*, provides, in Chapter VI, obligations aimed at facilitating switching of data processing service providers.[3]

These provisions modify the contractual and pricing landscape of switching. Underestimating the law would be as misleading as overestimating it.

Critical limit: prohibition of *switching* fees **does not** automatically eliminate all migration costs, rewrite, double execution, personnel, tests, debt tied to managed platforms, unavailability. Documentary compliance is not tested technical reversibility.

9.5 Mechanism: an exit that fails “in law”

Consider a plausible chain, **illustrative scenario**, not an observed case. An organisation has a termination clause and a more favourable *switching* fee regime. It exports data. Files arrive. But the application depends on a proprietary queue, cloud-specific identity, and non-standard extensions. Rewrite takes months. Both environments coexist. Tests reveal gaps. Switchover is postponed. Cost far exceeds extraction fees.

---

[1] OECD, “Competition in the provision of cloud computing services”, OECD Roundtables on Competition Policy Papers No. 323, 20 May 2025, DOI 10.1787/595859c5-en.

[2] Ofcom, *Cloud services market study final report* (extract), 2023.

[3] Regulation (EU) 2023/2854 of the European Parliament and of the Council (Data Act), Official Journal of the European Union, CELEX 32023R2854.

Provision	Content retained
Article 25	Written contract; notice of initiation at most two months; transition in principle at most thirty calendar days (otherwise motivated alternative deadline, at most seven months after notification); data recovery period of at least thirty days after transition
Article 29	*Switching* fees (including extraction fees within the regulation's meaning) reduced from 11 January 2024 to 12 January 2027 (cap = direct costs), then **prohibited** from 12 January 2027
Article 31	Exceptions: highly customised services outside commercial catalogue; non-production test or evaluation versions for limited duration
Article 50	General application from 12 September 2025

The law opens the door; architecture, formats, and skills decide whether one can cross the threshold.

9.6 Measurable evaluation grid

Indicator	How to measure	Limit of interpretation
Migration duration	Time between exit decision and validated switchover	Depends on scope and workload
Direct cost	Project budget, transfer, tools, contractors	Often excludes opportunity cost
Downtime	Measured interruption window	May be masked by double run
Rewrite effort	Person-days on non-portable services	Subjective without work log
Non-substitutable dependencies	Inventory of services with no tested equivalent	Inventory often incomplete
Data export capability	Formats, completeness, extraction delays	Export ≠ service reconstruction
Test coverage	Critical scenarios replayed after migration	Coverage ≠ absence of defects
Skills required	Profiles actually mobilised	Rarely documented
Post-migration operation	Service objectives met over N days	Horizon too short misleading

No composite score is produced. Nine partial, dated measures are better than a cosmetic ranking.

9.7 Checklist: preparing a credible exit test

Before speaking of tested reversibility, an organisation can chain, concretely:

1. **Inventory** services, data, and skills actually used (not only those in the contract).
2. **Choose** a representative but bounded workload, critical enough to instruct, small enough to test.
3. **Document** export formats, non-substitutable dependencies, and success criteria *before* the test.
4. **Prepare** a minimal target environment and necessary access rights.
5. **Execute** export, redeployment, and test scenarios while logging duration, gaps, and incidents.
6. **Measure** possible unavailability and direct cost of the test (without double-counting operating egress).
7. **Decide** explicitly whether the level reached is "verified interoperability", "successful migration", or only "documented".

This checklist is not a trial result. It turns intention into procedure.

9.8 Containerised protocol: proposed, not executed

Kubernetes documentation describes a portable platform for containerised workloads.[4] A test protocol was formalised within this work (minimal workload, standard manifests, success criteria, timestamped log).

It was not executed. No cluster was deployed for this book. No trial log exists. Verified interoperability and successful migration are therefore **not achieved** here. Kubernetes may facilitate certain moves; it does not make all architectures portable. The next chapter situates this software opening in a broader chain, without rerunning the protocol.

[4] Kubernetes / CNCF, documentation, Concepts Overview, kubernetes.io.

9.9 What the reader can retain

Reversibility is a **testable capability**, not a slogan. The Data Act can facilitate certain aspects of provider switching; it does not replace the trial and does not remove all migration costs. Testing has a cost: it is measurable insurance. An organisation may choose not to pay it; it cannot then claim a demonstrated capability.

Without tested reversibility, one can accumulate legal and pricing capabilities without reducing dependency criticality. Being able to exit does not oblige one to exit. Staying without knowing how to exit, however, is a potential critical dependency.

Under what conditions do free software, open standards, and compatible interfaces actually reduce these frictions, and where does that contribution stop? That is the subject of the next chapter.



PART IV

Technology, industry and resilience



CHAPTER 10

Open source software, standards, and interoperability

Opening the code, publishing a standard, adopting a "portable" orchestrator: three gestures often presented as proof of sovereignty. They can reduce certain dependencies. They can also shift the problem, to an integrator, an underlying cloud, a single maintainer, or a specification never put to the test.

The previous chapter distinguished declared portability from tested migration, and formalised a containerised protocol **not executed**. This chapter returns to that only to situate software openness within the same discipline of proof. It examines what open-source software, open standards, and compatible interfaces **can** contribute, and what they do not prove. It refuses to treat "open source", "open", "European", and "sovereign" as synonyms.

10.1 Defining open source without idealising it

According to the Open Source Initiative's Open Source Definition, open source is not reducible to access to code: distribution must meet a set of criteria, redistribution, source code, derivative works, non-discrimination, technology-neutral licence, among others.[1]

This definition prevents counting as "open source" any code viewable under restrictive conditions. Distinguishing **accessible source code** from a **genuinely free licence** (in the sense of the definition) is the first analytical step. It does not, by itself, imply that an organisation can maintain, secure, and operate the software at the pace of incidents.

10.2 Six distinctions that avoid the slogan

Notion	Meaning	Frequent trap
Definition-compliant licence	Distribution and usage rights per OSI criteria	Confusing with free of charge
Code availability	Readable, recoverable code	Without maintenance capacity
Open standard	Documented, accessible specification	Divergent implementation

[1] Open Source Initiative, *The Open Source Definition*.

Notion	Meaning	Frequent trap
Theoretical compatibility	Interfaces described as compatible	Not tested
Portability of a component	One module moves	The full system remains captive
Technical openness	Accessible tools and formats	Without operational mastery

An open licence, code availability, or adoption of a standard **do not suffice**, on their own, to establish the absence of critical dependency or a degree of sovereignty.[2]

Reducing a dependency is not making it disappear. Theoretical compatibility is not tested compatibility.

10.3 Conditional benefits

Open standards, documented formats, stable interfaces, automation tools, and open-source software can facilitate inspection and audit, broaden the pool of skills beyond a proprietary console, reduce certain rewrite costs when the implementation is genuinely compatible, and enable reproducible interoperability tests, *if* the organisation resists proprietary extensions.

Each of these benefits depends on architecture, governance, skills, and tests actually run. A documented format that no one can rebuild elsewhere does not lighten exit.

10.4 Limits and costs often overlooked

Openness has a price. Integration, maintenance, vulnerability response, operator training, and implementation divergence can exceed the cost of a proprietary licence over a given horizon. This is not an argument against open source; it is an argument for treating it in TCO rather than as automatic savings.

Large-scale interoperability is hard to test: what works for a demo service may fail for a business system with state, compliance, and strict service objectives. Two implementations of the same standard can diverge on details that block real migration. Without a test log, the standard remains a potential resource.

[2] Analytical framework of *European Cloud Sovereignty* (limits of open source and standards).

10.5 Hypothetical example: open foundation, proprietary dependency

Imagine, **explicitly hypothetical scenario**, not an observed case, an application deployed on an orchestrator whose core is open source and documented. Base manifests are standard. But over sprints, the team adds: a load balancer tied to the cloud provider; a workload identity specific to that cloud; a storage class with no tested equivalent elsewhere; and a private image registry authenticated only in that ecosystem.

The foundation remains "open". The *system* becomes hard to move. The openness of the central component has not vanished; it has been bypassed by extensions. Reducing dependency on the runtime has not removed identity, network, and storage dependencies. This mechanism explains why licence audit and exit audit must not be conflated.

10.6 Institutional rhetoric, bounded evidence

An institutional page on the Union's open-source strategy presents open source as a tool for "technological sovereignty".[3] Promotional figures appear there **without displayed methodology**: they are not reproduced here as established facts. The 2020–2023 strategy document could not be downloaded when sources were examined; no invented pagination is drawn from it.

Political intention and industrial outcome remain separate. This book evaluates capabilities, not press releases.

10.7 Link with barriers to switching

Competition analyses already relate limits of interoperability and standardisation to barriers to switching providers.[4] It is plausible, but **not measured** here, that improving standards and interfaces would reduce certain exit costs. The net effect depends on implementation, skills, and reversibility tests described in the previous chapter.

[3] European Commission, factual page "EU Open Source Strategy", Shaping Europe's digital future.

[4] OECD, "Competition in the provision of cloud computing services", OECD Roundtables on Competition Policy Papers No. 323, 20 May 2025, DOI 10.1787/595859c5-en.

10.8 From openness to capability: in four steps

First, a standard or free licence creates a technical *possibility*. Second, a concrete implementation creates *availability*. Third, an interoperability test creates *partial proof*. Fourth, a successful migration, with skills and operations, creates *capability*.

Jumping from the first step to the last is the most frequent shortcut in debate on software sovereignty. For a containerised case, Chapter 9 already established that announced ownership does not equal demonstrated migration: the associated protocol remains not executed. Here, the issue is broader: openness is a lever, not a guarantee.

Kubernetes documentation illustrates the same discipline without re-detailing it: announced portability of a container orchestration platform, distinct from a monolithic application platform; no interoperability or migration measured in this book.[5]

10.9 Objections

"All open source" would reduce dependency: dependency can shift to integrators, the underlying cloud, or maintainers. The standard would suffice: without a test, the text does not carry the load. Open source would be low cost: integration and maintenance belong in TCO. On some horizons, a proprietary service may remain lower in TCO, which does not make it a "sovereign" solution.

10.10 Toward industrial substance

Open-source software and standards are **levers of technological capability**. They can support reduction of critical dependencies when combined with skills, tests, and an industrial base. Presenting them as guarantees of sovereignty turns a facilitating capability into a slogan. Ignoring them because they are not enough would be the symmetric error.

Industrial substance, datacentres, components, energy, then conditions what these software levers can actually support. Without physical and industrial capabilities, software openness remains a partial necessary condition, not a complete strategy. That is the subject of the next chapter.



[5] Kubernetes / CNCF, documentation, Concepts Overview, kubernetes.io.

CHAPTER 11

Datacentres, semiconductors, energy, and supply chains

A climate-controlled building, server racks, a high-capacity power line: the datacentre image reassures. It suggests that cloud sovereignty can be seen from the street. Yet megawatts can host services whose software stack, components, maintenance, and incident decisions remain out of reach. Conversely, claiming that no cloud policy makes sense without total autonomy on semiconductors turns a real bottleneck into an excuse for inaction.

Between these two extremes, this chapter poses an industrial question: **which material and energy capabilities condition**, without guaranteeing, a reduction of critical dependencies? Open-source software and standards, treated in the previous chapter, do not replace a physical base. Here, four linked subsets, hosted compute, components, energy and networks, supply chains, then a grid that separates what is located, owned, controlled, and truly mastered.

No homogeneous European industrial statistic of ultimate ownership is invented. The chapter's central industrial claim remains a working hypothesis.

11.1 Installed capacity, available capacity

Physical compute and hosting capacity forms a material base for cloud services. Without sites, without available electrical power, without connectivity, there is no cloud offer at scale. A study commissioned for preparation of the *Cloud and AI Development Act* estimates, for 2025, datacentre capacity on the order of 12.4 GW in EU-27, with strong concentration in a few Member States. Germany, France, the Netherlands, and Ireland would represent, according to this modelling, roughly 63% of the total.[1]

These figures deserve sober reading. They describe **capacity**, often modelled, sometimes cross-checked with surveys. They do not say who owns the sites, who operates the hypervisors, who controls privileged access, or whether hosted services could be rebuilt elsewhere. A gigawatt located in Europe can serve a non-European stack. More modest capacity, better mastered operationally, can reduce certain critical dependencies without matching hyperscaler size.

One must also distinguish **installed** from **effectively available** capacity. A full

[1] Cloud & AI Study, EU Publications Office, 2026, DOI 10.2759/9828163, modelled capacity / CATI survey; ≠ ultimate ownership of operators.

rack, a saturated power line, a permit not obtained, or a cooling constraint can make unusable part of what appears on a power map. Geographic concentration adds another nuance: if a significant share of capacity sits in a small number of Member States, national energy, planning, and tax policies weigh on European supply as much as pan-European industrial discourse.

11.2 Location, ownership, ultimate control, operations

An external BEREC study distinguishes five broad categories of ownership and business models: independent carrier-neutral colocation; sites owned by hyperscalers; public initiatives or public–private partnerships; centres held by telecommunications operators; hybrid or emerging models.[2]

This typology prevents treating "European datacentre" as a single category. A colocation site can host clients of very diverse nationalities and stacks. A hyperscaler site can sit in Europe while falling under a globalised decision chain. A public–private partnership can combine public funding and private operation. A telecom operator can own the building and subcontract part of the cloud layer.

Four questions then separate:

1. **Where** are the machines and data? (location)
2. **Who owns** the assets and contracts? (ownership)
3. **Who is the ultimate beneficiary** of economic control rights? (ultimate ownership)
4. **Who configures, supervises, and decides** in an incident? (operational mastery)

Location and ownership do not coincide. Neither do ownership and ultimate beneficiary: a European registered office can mask a control chain outside the Union, and the reverse. Legal control and operational mastery diverge further: a contract can assign rights without teams being able to exercise them under constraint. No homogeneous European series of ultimate ownership of operators was identified in the current corpus. No European percentage of "ownership" or "mastery" is therefore fabricated here. This absence is not a footnote detail: it is a structural gap for anyone wanting to turn the megawatt map into a sovereignty map. A homogeneous industrial statistic of ultimate ownership remains to be assembled; it was not invented.

[2] BEREC, External Study on Data Centres, BoR (26) 01, 2026: typology; no EU series of beneficial owners.

11.3 Semiconductors: upstream bottleneck of the cloud

Access to critical chips and components constrains what can be produced, evolved, or secured. Without components, no servers. Without credible manufacturing or supply capacity, announcements of cloud autonomy quickly meet the reality of lead times, technology nodes, and global chains.

The *Chips Act* sets a policy framework to strengthen the European semiconductor ecosystem.[3] JRC and European Court of Auditors reports document shares of fabrication or global capacity on the order of a few percent, with definitions that **do not overlap**. Production, installed capacity, Digital Decade objectives: these quantities must not be merged. Orders of magnitude often cited, around 7%, 9%, or 11.7% depending on definitions, must stay attached to their sources and methods.[4]

These figures contextualise industrial dependency. They do not measure cloud sovereignty. A Europe more capable in semiconductors could reduce certain supply risks for digital infrastructure. It would not automatically turn every cloud contract into operational mastery, nor every location into reversibility. Symmetrically, total autarky on the most advanced nodes does not forbid mapping bottlenecks and diversifying supply.

The link with cloud is mechanical: availability of servers, accelerators, and network equipment depends on chains that extend beyond the cloud provider. A cloud policy that ignores this bottleneck treats the building as if it were self-sufficient.

11.4 Energy, grid connection, and networks

Without energy and connectivity, no durable cloud capacity. Datacentres are major and growing electricity consumers. Eurostat publishes national series of datacentre electricity consumption; EU aggregation is not homogeneous.[5] A JRC estimate places EU-27 datacentre consumption between 45 and 65 TWh in 2022.[6] Report COM(2026) 500 cites, for 2024, a value of 68 TWh referring to the IEA. This is not a verified sum of Eurostat cells, and it is **not mechanically mergeable** with the JRC range.[7]

Energy is a condition, not the whole sovereignty problem. A constrained electrical region, connection queue, demand peak, or local acceptability can limit capacity expansion even if cloud demand grows. Abundant capacity powered without mastery of software stacks does not produce more sovereignty. Networks complete the

[3] Regulation (EU) 2023/1781 (Chips Act).

[4] JRC, Semiconductors in the EU (JRC133850), 2023; European Court of Auditors, Special report 12/2025. Do not merge 7%, 9%, and 11.7% without specifying definitions.

[5] Eurostat, isoc_env_ict_nrg.

[6] Kamiya & Bertoldi (JRC), 2024.

[7] European Commission, COM(2026) 500 final, IEA citation in footnote; do not merge with JRC.

picture: an isolated, poorly interconnected datacentre does not support the same services as a node near exchange points and international routes.

11.5 Supply chains: resilience without autarky

Geopolitical shocks, logistics breaks, delivery delays, dependence on network equipment or spare parts: resilience of a *supply chain* limits single points of failure. Resilience is not autarky. A Europe that diversifies suppliers and strategic stocks does not become independent of every global chain. It reduces certain rupture risks.

A unified quantitative map of the European cloud chain, from wafer to rack, firmware to maintenance, remains to be documented. It will not be replaced by unlabelled illustrative values. While that map is missing, the chapter can formulate mechanisms and distinctions; it cannot claim to measure the degree of industrial resilience of European cloud supply.

11.6 Grid: locating is not controlling

Dimension	Question	Frequent confusion
Location	Where are machines and data?	Local = sovereign
Ownership	Who owns assets and contracts?	Local ownership = technical independence
Ultimate control	Who holds effective economic power?	Registered office = ultimate beneficiary
Operations	Who configures, supervises, decides in incident?	Compliance = operational mastery
Supplier independence	Can critical components and services be substituted?	A second supplier always suffices

These questions are not answered by a single indicator. A located site can be owned by a non-European actor. Local ownership can coexist with non-substitutable software dependency. Documentary control can mask operational incapacity. Dual sourcing can remain insufficient if both sources share the same upstream bottleneck.

In this book's analytical framework, the European industrial base, datacentres, components, energy, *supply*, **constrains** the degree of cloud dependency reduction, without conflating it with location alone or technological autarky.[8] This proposition remains a **working hypothesis**. Contextual sources do not replace an inventory of capacity and ultimate ownership.

[8] Analytical framework of the book (industrial capabilities), working proposition, industrial evidence still partial.

11.7 Objections

One objects that local datacentres “suffice”. Only if they come with operational control, software substitutability, and skills, what location does not automatically provide. One objects that chip autonomy is impossible in the medium term: the objection may hold for total autarky; it does not forbid mapping bottlenecks. One objects that energy constraints make expansion undesirable: that reinforces the idea that cloud capacity is not a free good.

A subtler objection holds that ultimate ownership does not matter if contracts and certifications are robust. The next chapter shows why contracts and certifications alone do not measure capacity to act under incident. Here, the absence of a European inventory of beneficial owners prevents empirically settling this objection at Union scale.

11.8 Bridge to operations

Without a readable industrial base, cloud autonomy remains constrained. Location, ownership, ultimate control, and operational mastery do not coincide. The corpus provides energy, semiconductor, and typological context; inventory of ultimate ownership and measured control remains open.

Substance is not enough: without operational mastery and resilience to incidents, the legal framework and megawatts remain hollow. That is the subject of the next chapter.



CHAPTER 12

Cybersecurity, resilience, and operational mastery

A data-localisation clause can reassure a buyer. It does not say who holds the administration keys, who responds at three in the morning, or whether extraterritorial access can be demanded. Without operational control and resilience, the legal framework risks remaining an empty envelope, elegant on paper, inert under constraint.

The previous chapter recalled that substance, datacentres, chips, energy, conditions without guaranteeing. This chapter addresses **capacity to act**. Institutional sources contextualise the debate. They do not measure aggregated European operational mastery. No sensitive incident case is invented. No supplier security ranking is established.

12.1 Eight concrete questions

Before labels, the questions:

1. Who can **administer** systems, and with what rights?
2. Who holds **privileged access**, and how is it logged?
3. Who can **intervene** in an incident, in what timeframe, from where?
4. Who controls **logs** and supervision tools?
5. Who can **restore** data and services, and has that restoration been tested?
6. Who masters **crisis procedures**, isolations, failovers, communications?
7. What **dependencies** remain on providers and subcontractors?
8. How to **demonstrate** these capabilities, tests, logs, exercises, instead of merely declaring them?

These questions distinguish four layers often conflated: **regulatory compliance**, **contractual commitments**, **documented arrangements**, and **capabilities actually tested**. An organisation can excel on the first two and fail on the fourth.

12.2 What operational control denotes

Configure, supervise, manage identities, respond to incidents, decide on isolation or failover: that is the core of operational control. It differs from documentary compliance and from servers merely being on a territory. A successful audit can coexist with a team unable to rebuild a critical service. A certification can attest to a process without guaranteeing its mobilisation under stress.[1]

The **shared responsibility** model between customer and provider redistributes these tasks by service, infrastructure, platform, or software, without abolishing the question: *who can actually do what?* In IaaS, the customer retains a large share of systems and applications. In PaaS, the platform escapes their hand more. In SaaS, most operations are delegated. Moving up the stack does not eliminate the need for control; it changes its nature.

The associated figure illustrates these boundaries. It is a **pedagogical model**, not a universal contract, nor an incident statistic.

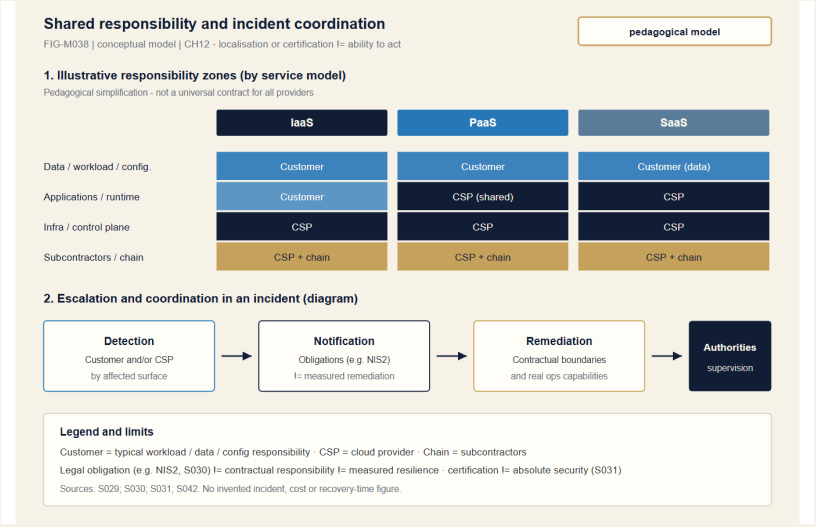


Figure 12.1: Shared responsibility and incident coordination (pedagogical model, not empirical).

The diagram does not prove operational independence. It shows why conflating contractual responsibility and capacity to act is an error.

[1] ENISA / CSA, limits of certification.

12.3 Location is not control

Locating data or machines can be necessary for certain legitimate requirements: proximity, jurisdiction, latency, sector policies. That does not guarantee mastery of privileged access, independence of supervision tools, or absence of non-substitutable software dependencies. The "local therefore safe" argument is a confusion, not a demonstration.

An institutional buyer framework can require declared levels of "sovereignty". That alone does not prove measured capacity to act.[2] This book's framework and buyer tools must not be merged. One proposes an analytical grid. The other structures a procurement procedure.

12.4 Resilience: absorbing without claiming autarky

Resilience here means the capacity to absorb shocks and failures without unacceptable rupture, major outage, supply break, security incident, prolonged unavailability. It does not require total absence of dependency. It requires knowing what breaks, how long one holds, and which alternatives actually exist.

Threat reports and notification obligations structure the regulatory landscape. For certain cloud entities and datacentres, the NIS2 Directive introduces cybersecurity and reporting obligations.[3] An obligation to notify is not a measure of observed resilience. Comparable continuity indicators across actors, recovery time, on-call coverage, substitutability of supervision tools, are still missing in the corpus. They will not be invented. No homogeneous ranking of providers' "operational mastery" is therefore published here.

Resilience also crosses the supply chain of the previous chapter. An incident is not only malicious software. It can be unavailability of a component, firmware, maintenance provider, or entire cloud region.

12.5 Extraterritoriality of access: open question

Regimes of access by foreign authorities to data or infrastructure located in Europe constitute an issue often cited. Older guidelines for Union institutions recall buyer responsibilities; their scope is narrow and dated.[4]

The question belongs to law, operational models, and privileged-access architectures. It is not settled by anecdote. It does not disappear because a site sits in the Union.

---

[2] Cloud Sovereignty Framework, Commission buyer tool; distinct from this book's measurement framework.

[3] ENISA Threat Landscape; NIS2 Directive, obligations ≠ measured resilience.

[4] EDPS, cloud guidelines for EU institutions, 2018: do not generalise to all administrations.

12.6 Working hypothesis

In this book's analytical framework, **without operational mastery and resilience**, a legal framework or location alone does not produce capacity to act on critical cloud dependencies.[5]

This proposition is a **working hypothesis**. Contextual sources do not replace incident comparison or measurement of operational mastery.

12.7 Objections

One objects that large global providers are "safer" than local alternatives. That may hold for certain perimeters, global redundancy, security budgets, process maturity. It is not proof that the associated dependency is acceptable, nor that exit remains possible. Relative security and dependency criticality are two distinct questions.

One objects that security is mainly about standards and audits. Useful, that view does not replace incident practice. One objects that sensitive incidents are poorly documented publicly: precisely why this chapter refuses invented examples.

An inverse objection holds that local operational mastery always suffices. Yet a local team can depend on proprietary tools, extraterritorial support, or rare skills. Declared mastery and mastery exercised under constraint must not be conflated.

12.8 From acting to supply

Cybersecurity and resilience modify the degree of cloud sovereignty as **capacities to act**, not as labels. Location ≠ control. Certification ≠ capacity to act. Notification obligation ≠ measured resilience. The institutional corpus contextualises; measured operational mastery remains open.

What remains is the passage from technical capability to **viable supply**: research, funding, startups, and industrialisation, without which alternatives stay laboratories. That is the subject of the next chapter.



[5] Analytical framework of the book, working proposition on operational mastery and resilience; aggregated evidence still absent.

PART V

Research, markets and competition



CHAPTER 13

Research, funding, startups, and industrialisation

A publication is not a product. A prototype is not a viable cloud alternative. Between the laboratory and a provider capable of lasting stretch bottlenecks, capital, industrialisation, demand, operations skills, often conflated in "sovereignty" discourse. Counting articles, demonstrators, or budget envelopes as if they already measured industrialised cloud supply is one of the costliest confusions in the debate.

Chapters on school and higher education opened research and transfer. This one changes scale: it looks at **supply**, how a technical capability does or does not become a credible economic alternative. No invented success story. No fabricated startup-to-scale conversion rate.

13.1 Seven steps, six possible ruptures

For a scientific or technical capability to become durable industrial capability, a minimal chain emerges:

1. **Research funding**, resources to explore and experiment.
2. **Knowledge production**, publications, patents, tacit know-how.
3. **Prototypes and demonstrators**, bounded proof of concept.
4. **Technology transfer**, passage to an organisation able to sell and operate.
5. **Company creation and growth**, structure, capital, teams.
6. **Customer adoption**, real use, payment, retention.
7. **Industrialisation and operation at scale**, availability commitments, support, security, sustainable unit costs.

Between these steps, ruptures are frequent. Discontinuous funding cuts research before the prototype. A brilliant prototype dies for lack of product owners or operations teams. A startup obtains capital without reaching compliance and support costs comparable to major providers. Initial adoption collapses if the service cannot bear load. Industrialisation hits non-substitutable components, skills, or third-party services, those of Chapters 11 and 12.

These ruptures are not anecdotes. They explain why a public programme can succeed on spend indicators and fail on cloud supply actually available.

13.2 What we risk counting too early

Level	Meaning	Trap
Scientific output	Publications, patents	Counting as industrialisation
Prototype / demonstrator	Proof of concept	Confusing with an operated service
Startup	Nascent organisation	Ignoring viability
Scale-up	Ability to grow	Underestimating capital and operations
Viable provider	Durable alternative	Using market share as sovereignty

These levels do not form an automatic ladder. An economy can excel in publications and lack managed offer. It can fund startups without their reaching operational scale needed to compete with full stacks.

13.3 Intention, act, outcome

The *Digital Europe* programme sets general objectives, including strategic autonomy, and budget envelopes by domain (high-performance computing, AI, cybersecurity, skills, deployment).[1] Reading this regulation as proof of cloud industrialisation would be a category error. A budget framework authorises actions. It does not demonstrate that durable European cloud supply exists.

The *Chips Act* targets the semiconductor ecosystem. It is distinct from managed cloud supply, even if it conditions part of the industrial base discussed in Chapter 11.[2] Commission policy pages announce instruments and alliances. An announcement is not an adopted act, nor a market outcome.[3]

Three boxes stay separate: **intention** (discourse, roadmap), **act** (regulation, decision, committed budget), **outcome** (industrialised capability, operated service, sustained adoption). Passing from the first to the third without intermediate proof produces narratives of sovereignty already achieved.

[1] Regulation (EU) 2021/694 (Digital Europe Programme), OJ text examined. Funding ≠ industrialised cloud capability.

[2] Regulation (EU) 2023/1781.

[3] Commission policy pages / communications; distinguish intention and outcome.

13.4 Capital, skills, and demand

Access to funding can enable industrialisation of a European cloud offer. It does not alone create *product-market fit*. Industrialising here means moving from a demonstrator to a service able to meet availability, support, security, and scale commitments. That requires capital, but also operations skills, those of Chapter 12, a sales chain, integration with existing ecosystems, and solvable demand.

Conversely, public or private demand without scalable supply leaves buyers facing the same dependencies. The two bottlenecks, capital and demand, must stay distinct until proven otherwise. The next chapter treats public procurement. This one refuses to solve the supply problem by referring only to purchasing.

Capital is not homogeneous. Funding a prototype is not funding a platform operated twenty-four seven. Funding a startup is not ensuring capacity to face compliance costs comparable to major providers.

13.5 Observable indicators: without invented values

While transfer, survival, and scale-up rates are missing in the corpus, one can at least **name** what should be observed, without inventing figures:

Candidate indicator	What it would make visible	Status in this book
Share of projects moving from demonstrator to operated service	Transfer rupture	Not measured here
Time and cost to reach contractual availability commitment	Ops industrialisation	Not measured here
Customer retention rate after 12 / 24 months	Real adoption	Not measured here
Dependence on non-substitutable components or third-party services	Industrial fragility	Qualitative only
Tested on-call and restoration capacity	Link with CH12	Not measured here

These indicators are a measurement agenda, not a results table.

13.6 Working hypothesis

In this book's analytical framework, industrialisation and funding condition the existence of **viable cloud alternatives**, distinct from scientific output alone or mere existence of university modules.[4]

This proposition remains a **hypothesis**. Examined funding frameworks contextualise. They do not measure completed industrialisation of a European cloud offer.

13.7 Objections

One objects that the bottleneck is demand, not capital. Possible: that is why the next chapter treats public procurement. One objects that the market "selects" the best offers: concentration documented elsewhere in the book reminds that selection does not imply substitutability. One objects that Europe "innovates enough": innovating is not industrialising at the scale of critical managed stacks.

A more radical objection holds that Europe need not produce hyperscalers, only regulate. That political position exists. It does not cancel the book's question: if dependencies are judged critical, which capabilities allow reducing them? Regulating without capabilities can frame; it does not substitute.

13.8 Toward public demand

Without industrialisation and provider viability, human and software capabilities remain upstream. Documented European programmes alone do not prove durable cloud supply. Intention ≠ act ≠ outcome. Publication ≠ product. Budget ≠ industrialisation.

Can public demand create a credible outlet without closing the market or promising automatic industrialisation? That is the next question.



[4] Analytical framework of the book, working proposition on industrialisation and funding; outcomes not measured.

CHAPTER 14

Public procurement and industrial policy

The State and local authorities are often first to speak of cloud sovereignty, and among the first to buy services. Public procurement can structure demand, impose clauses, create an outlet. It can also generate litigation, fragmentation, or local preference without real capabilities. This chapter reads these trade-offs. It is not legal advice. It is not a pan-European industrial impact assessment either: national volumes and effect studies are still missing.

The previous chapter situated supply and industrialisation. This one looks at a **demand** lever. Public procurement law, the *Data Act*, an institutional buyer framework, and a documented market provide a partial corpus. They alone do not demonstrate that public purchasing caused measured progress in sovereignty.

14.1 What public demand can do, and not do

Lever	Plausible effect	Limit
Purchase volume	Outlet for alternatives	Without viable offer, volume = dependency
Reversibility clauses	Exit requirement	Clause ≠ migration test
Localisation criteria	Territorial compliance	Location ≠ ops control
Security criteria	Minimum requirements	Audit ≠ capacity to act
Dependency transparency	Visibility of subcontractors and stacks	Inventory ≠ substitutability
Total cost of ownership	Trajectory reading (see Chapter 8)	Incomplete TCO if parameters absent
“Local” preference	Industrial support	Risk of inefficiency / litigation; ≠ capabilities

These levers are not interchangeable. Requiring testable reversibility is not the same as requiring a façade nationality. Imposing purchase volume on a non-existent offer does not magically create a viable provider. Supporting industry through purchasing can be tempting; it can also lock in mediocre solutions or expose the buyer to challenges.

14.2 Objectives, rules, criteria, clauses, execution, outcomes

Six layers are distinguished:

1. **Political objectives**, sovereignty discourse, roadmaps.
2. **Legal rules**, texts in force (directives, regulations).
3. **Award criteria**, what the specification concretely requires.
4. **Contractual clauses**, what the contract formalises.
5. **Execution**, what is actually delivered, tested, controlled.
6. **Observed outcomes**, effects on capabilities, competition, dependencies.

Conflating layer 1 with layer 6 produces undemonstrated success narratives. A political objective is not an industrial outcome.

14.3 Law in force, projects, and rhetoric

One must distinguish a text **in force**, a draft, and a political intention. Directive 2014/24 sets principles of non-discrimination and award, including the most economically advantageous tender (*MEAT*).[1] It does not forbid all consideration of quality, security, or performance linked to the subject of the contract. It alone does not authorise a generalised national preference presented as sovereignty.

The *Data Act* imposes, within its scope, *switching* obligations. It alone does not create industrial supply.[2] A legal obligation to facilitate provider change is not a successful migration, nor operational exit capability at the buyer.

A *Cloud Rulebook* and general procurement guidance are announced on policy pages. They were not published as acts at the time of examination.[3] Announcement ≠ applicable law.

A *Cloud Sovereignty Framework* and its implementation guidance let an institutional buyer structure criteria for a given procedure.[4] This framework is buyer practice. It is not this book's measurement framework. A documented market for European institutions, order of magnitude cited: roughly €180 million, illustrates a practice. It does not generalise to all Union administrations, and a contract amount is not a measured industrial effect.[5]

The figure distinguishes what law **imposes**, what it **authorises**, what is **observed** in documented practice, and whose **effect** is **not demonstrated**. It is not a procurement manual. It is a compass against confusions.

[1] Directive 2014/24/EU.

[2] Regulation (EU) 2023/2854 (Data Act), CELEX 32023R2854, detail also in reversibility chapter.

[3] Commission policy page "Cloud computing".

[4] Cloud Sovereignty Framework; Implementation guidance, buyer option / practice, not general EU obligation.

[5] Institutions market documentation / Cloud III.

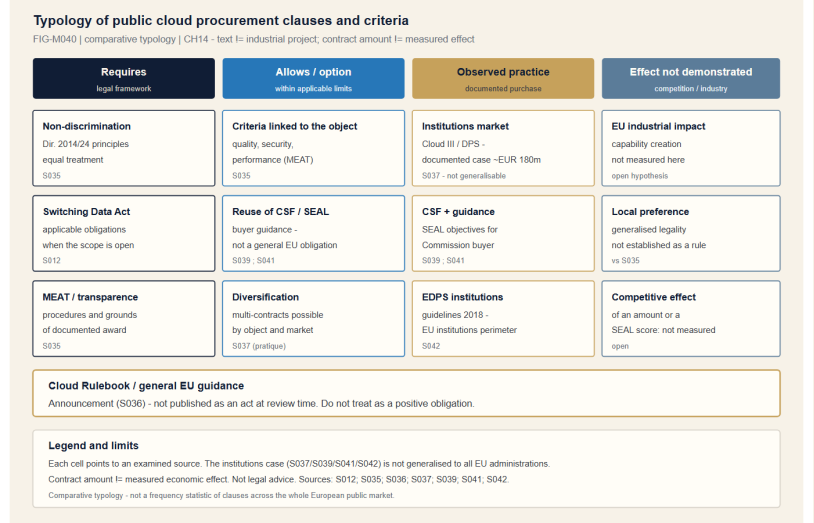


Figure 14.1: Typology of public cloud procurement clauses and criteria (analytical framework; institution practice partially documented).

14.4 Criteria that can influence capabilities, if verified

Among criteria useful in this book's sense, when actually controlled after award, are:

- **interoperability** and **reversibility** (formats, timelines, exit assistance, required tests);
- **security** and **continuity** (service objectives, restoration exercises);
- **dependency transparency** (subcontractors, critical components, regions);
- **total cost of ownership** on an explicit trajectory, not catalogue price alone;
- the provider's **real capacity to respond**, teams, operational references, not slides alone.

Without verification after award, these criteria remain contractual intentions. Public procurement can require; it does not, by specification alone, manufacture operations skills absent at the buyer, nor an alternative offer non-existent on the market.

Similarly, a localisation criterion can be relevant for certain workloads. It does not measure operational control. A buyer who localises without mastering privileged access buys territorial compliance, not capacity to act.

14.5 Effects on competition and small enterprises

High security, continuity, or reversibility requirements can exclude actors unable to meet them, including small European firms. Too-weak requirements can favour already dominant stacks. The net effect on competition is not demonstrated here as a European constant. It depends on the market, lot splitting, possibility of consortia, and proportionality of criteria to the subject.

The figure and table in this chapter therefore recall a limit: what is **legally authorised** is not automatically **industrially desirable**, and what is **desirable** is not automatically **measured**.

14.6 Working hypothesis

In this book's analytical framework, public procurement **can** structure cloud demand and reversibility requirements, but it does not automatically create industrial or operational capabilities; local preference does not equal industrial success.[6]

This proposition is a **working hypothesis**. Public cloud procurement volumes by Member State and dedicated impact studies are still missing. No claim that public procurement caused measured progress in European cloud sovereignty is established in these pages.

14.7 Objections

One objects that local preference is inefficient or legally risky. The objection is serious: it justifies preferring capability, testable reversibility, and security criteria to façade nationality. One objects that the single market forbids all industrial policy through purchasing: legal reality is nuanced and beyond the scope of advice in these pages. One objects that European cooperation suffices: cooperation will be treated with market dynamics, without conflating it with the act of purchase.

An inverse objection holds that without local preference no European offer will emerge. Possible as a market mechanism. Not demonstrated here as a general law, and not demonstrated that preference produces capabilities rather than rents.

14.8 Opening toward the market

Public demand is a conditional lever. It can require; it alone does not manufacture. Law in force ≠ draft ≠ rhetoric. Contract amount ≠ industrial effect. Exit clause ≠ tested migration. The legal and institutional corpus is partial; industrial effect remains undemonstrated.

Declared adoption, market shares, and European cooperation interact next, none

[6] Analytical framework of the book, working proposition on public procurement and industrial effect; impact not demonstrated.

of these quantities alone equals cloud sovereignty. That is the terrain of the next chapter.



CHAPTER 15

Adoption, competition, cooperation, and network effects

Three figures often circulate in the same paragraph: how many firms use the cloud, what market share major providers hold, what degree of "sovereignty" Europe has reached. These are three distinct objects. Merging them produces the most tenacious confusion in the debate, and one of the most useful to discourses that want to conclude too quickly.

Chapters on industry, operations, funding, and public procurement situated supply and a demand lever. This one looks at **market dynamics**: what attracts customers, what retains them, what strengthens or weakens credible competition, and what European cooperation can or cannot change. It recalls already documented quantities. It invents none. It refuses to turn them into a sovereignty verdict.

15.1 Three quantities not to merge

Quantity	What it measures	What it does not measure
Adoption	Declared use of paid cloud services	Provider shares; sovereignty
Market shares	Concentration of supply (estimates)	Operational mastery; reversibility
Cloud sovereignty (book framework)	Degree of capabilities facing critical dependencies	An adoption rate or a share

In 2025, 52.7% of enterprises in the Union covered by the Eurostat survey declared use of paid cloud computing services.[1] This proportion has risen since 2023 and since 2014. It describes adoption. It does not say who captures value, whether one can exit, or which layers, infrastructure, platform, or software, dominate among respondents.

According to private estimates of July 2025, the share of "European" cloud providers in the regional infrastructure market would have remained around 15% after 2022, while AWS, Microsoft, and Google together would represent roughly 70% at publication.[2] These are market estimates, not an official statistic, and not a

[1] Eurostat, "53% EU enterprises used paid cloud services in 2025", and associated dataset (paid cloud usage indicator).

[2] Synergy Research Group, "European Cloud Providers' Local Market Share Now Holds Steady at 15%",

measure of sovereignty. Reading them as a mastery barometer would be a category error.

Cloud sovereignty, in this book's framework, is neither the Eurostat rate nor the Synergy share. It is a degree of capabilities facing critical dependencies. High adoption can coexist with strong concentration. A higher European share could coexist with low reversibility. The three quantities can move in different directions.[3]

15.2 Network effects, switching costs, economies of scale

A cloud market is not only a confrontation of catalogues. The more an ecosystem is adopted, the more it attracts tools, skills, training, and integrators. These **network effects** make the ecosystem more useful for the next customer, and costlier to leave for the customer already committed. The mechanism is plausible and widely discussed. It does not imply concentration is "optimal" for reversibility. It does not imply forced fragmentation would automatically produce European capabilities.

Switching costs extend this reasoning. They are not reducible to the price of a technical migration. They include application rewrite, team retraining, temporary loss of functionality, coordination across business units, and sometimes data egress fees. Chapters on price and total cost already showed why the catalogue alone misleads. Here, the issue is dynamic: the higher and less measured these costs, the more initial adoption can become structural dependency, even if competitors exist on paper.

Economies of scale play the other way in the debate. Large operators spread fixed costs, regions, security, research, support, over a global base. That can explain part of their price and functional attractiveness. It does not erase the question of dependency criticality, nor that of exit possibility. Saying "it is cheaper because bigger" describes an economic mechanism. It is not proof that the associated dependency is acceptable for every critical workload.

24 July 2025: private estimates; partial "European provider" definition.

[3] Analytical framework of the book, distinction adoption / shares / sovereignty.

15.3 Trust, adoption, and multi-cloud

Declared adoption measures minimal operational trust: enough to pay for a service. It does not measure trust in capacity to exit, nor in mastery of privileged access, nor in continuity under constraint. An organisation can adopt massively while indefinitely postponing reversibility tests. The reversibility chapter already separated legal obligation, contractual clause, and migration actually tested. The market amplifies this confusion when it sells adoption as maturity.

Multi-cloud, often presented as remedy, is not proof of substitutability. Operating two providers can increase complexity and coordination costs without reducing criticality if critical workloads remain captive to one. Again, plausible mechanism ≠ outcome observed for each organisation. Without tests, multi-cloud may be only bill diversification.

15.4 Competition, interoperability, and resilience

Competition analyses highlight concentration, barriers to switching, and interoperability limits in several jurisdictions.[4] Competition is not judged only by the number of actors on a slide. It is also judged by real ability to switch, formats, timelines, skills, total costs.

Three requirements cross without conflating. **Competition** seeks credible alternatives. **Interoperability** seeks portability of workloads and data. **Resilience** seeks capacity to absorb shock without unacceptable rupture. Strengthening one can weaken another if care is not taken: standardisation too slow can slow supply; a race to adopt can sacrifice exit; very high resilience requirements can exclude smaller actors, including European ones, unable to respond immediately. These trade-offs are not resolved by slogan. They are read workload by workload, market by market.

15.5 European cooperation: scale without magic

Coordination among Member States and institutions can reduce fragmentation of requirements and increase project scale. It can also produce façade harmonisation without shared capabilities. Without programme evaluation recorded here as proof of industrial effect, one holds to this tension: cooperation ≠ forced uniformisation, fragmentation ≠ successful experimentation.

Cooperation can bear on standards, common markets, research infrastructure, security requirements, or industrial projects. Each object has its conditions of success. Conflating them under the sole word “European cooperation” produces undifferentiated optimism. The book prefers linking each form of cooperation to the capabil-

[4] OECD, “Competition in the provision of cloud computing services”, 2025.

ities it is supposed to strengthen, human, economic, technological, industrial, institutional, and to available proof. Public procurement, treated in the previous chapter, is a demand lever. It does not replace industrialised supply or operations skills.

15.6 What the market alone does not prove

Economic mechanisms recalled here, network effects, switching costs, economies of scale, trust, cooperation, illuminate why certain dependencies become critical. They alone do not demonstrate the degree of European cloud sovereignty. They do not replace the missing industrial inventory, nor measured operational mastery, nor exit tests still absent from the corpus.

In this book's analytical framework, adoption, competition, and cooperation **condition** the trajectory of cloud dependencies, without any of these quantities alone equalling sovereignty.[5]

This proposition articulates bounded facts (Eurostat adoption), private estimates (market shares), and documented competition mechanisms. It assigns no probabilities to future scenarios. It does not elevate gaps from industrial and operational chapters.

15.7 Objections

One objects that concentration reflects mainly efficiency. Possible for certain services. That does not eliminate the question of exit costs and criticality. One objects that adoption proves social acceptance of cloud: it proves declared use, not mastery. One objects that European cooperation is too slow: slowness does not authorise inventing results.

An inverse objection holds that only high European market shares would prove sovereignty. The book refuses this reduction. Higher shares could accompany low operational mastery or low reversibility. They would remain a useful quantity, among others, for reading supply dynamics.

[5] Analytical framework of the book, working proposition on adoption, competition, and cooperation.

15.8 Toward measurement

Adoption, market shares, and cloud sovereignty do not merge. Network effects and switching costs illuminate dependency criticality. European cooperation can increase scale; it does not magically create capabilities. Quantities recalled here remain bounded to their sources.

To make capabilities and dependencies visible without fabricating a single score, a disciplined dashboard is needed. That is the subject of the next chapter.



PART VI

Measurement, scenarios and autonomy



CHAPTER 16

Building a dashboard of capabilities and dependencies

A single "cloud sovereignty" figure would reassure political dashboards. It would also deceive. Arbitrary weightings, hidden missing values, merging incomparable objects: the composite score promises false precision. This chapter proposes something else, a **monitoring framework** of distinct indicators, with admitted gaps and separated families.

The previous chapter separated adoption, market shares, and sovereignty, and recalled market mechanisms. This one operationalises the book's measurement framework. It assigns **no composite score**. It invents no missing series. It refuses to automatically elevate still insufficiently documented claims.[1]

16.1 Four indicator families

Before dimensions, four analytical roles:

Family	Question	Example of meaning
Input	What is injected?	Training budgets, research envelopes, purchase volumes
Capability	What can one do?	Accessible labs, on-call teams, industrialised offer
Outcome	What is observed?	Tested migration delays, calibrated total costs, alternatives used
Resilience	What holds under constraint?	Tested restoration, substitutability, single points of failure

Conflating an *input* (spending) with an *outcome* (actually exiting) is one of the most frequent errors in cloud policy. Conflating a *declared capability* with *tested resilience* is another. The dashboard must make these confusions visible, not smooth them.

These families extend the multidimensional framework already set: cloud sovereignty there denotes a **degree** of capabilities allowing reduction of critical

[1] Measurement framework of the book, distinct indicators and profiles; default prohibition of composite score.

dependencies, not a binary label.[2]

16.2 Dimensions and indicator sheets

For each dimension, the two tables below propose a *candidate* indicator, not a populated value. The first fixes object and unit; the second fixes source, availability in this book, and manipulation risks. When data are absent or partial in the corpus, the cell says so. No interpolation is proposed.

Dimension	What is measured	Candidate operational indicator	Unit / population
Human capital	Pipeline and real access to environments	Share of establishments with accessible cloud labs; share of ops trained outside a single ecosystem	% establishments / % professionals (EU or State)
Economy	Price and total cost	Total cost gap on explicit trajectory (catalogue + exit + lock-in)	Currency / typical workload; bounded basket
Technology / reversibility	Exit and interoperability capability	Number and success of tested migrations on priority workloads	Count / delay / cost per workload
Industry	Material capacity and control	Available DC capacity; UBO / operator control inventory	GW / MW; documented share of operators
Operations / cyber	Operational mastery	Successful restoration exercises; documented on-call coverage	Success rate / delays; service perimeter
Governance / procurement	Requirements and verification	Share of contracts with exit clauses *and* post-award verification	% contracts / volume
Market	Adoption and concentration	Declared adoption; estimated shares, **in separate columns**	% enterprises; % market (explicit definition)

[2] Analytical framework of the book, multidimensional definition.

Dimension	Required source	Desirable frequency	Availability in this book	Limits / manipulation risks
Human capital	Establishment surveys; skills surveys	Annual	**Absent** (upstream hypotheses only)	Counting nominal modules as real access
Economy	Catalogues + calibrated exit parameters	Semiannual	**Partial** (method + basket; total not calibrated)	Displaying catalogue price alone
Technology / reversibility	Executed and logged protocols	Annual + after incident	**Absent** (protocol not executed here)	Turning legal clause into successful test
Industry	Capacity studies + company registers	Annual	**Partial** (capacity/energy/typology; UBO absent)	Conflating location and ultimate ownership
Operations / cyber	Exercise logs; ops audits	Annual + post-crisis	**Absent** (no homogeneous comparison)	Conflating certification and capacity to act
Governance / procurement	Public buyer data	Annual	**Partial** (law + bounded institution practice)	Counting political intentions as outcomes
Market	Official statistics + labelled market study	Annual	**Partial** (Eurostat + distinct private estimates)	Merging adoption, shares, and sovereignty

These tables form an **analytical framework**. They do not constitute a European empirical population. They authorise no general ranking of States or providers.

16.3 Missing-value rules

Four minimal rules avoid turning the dashboard into fiction.

First, a missing value stays missing: no implicit zero, no invented European average. Second, a private estimate carries its label and does not merge into an official statistic. Third, a non-executed protocol does not appear as a migration outcome. Fourth, a still hypothetical claim does not change status because it sits in a handsome table.

An institutional buyer tool can compute a score for a given procedure. That is not this book's measurement framework. The two must not be merged.[3]

16.4 Profiles rather than podium

A multidimensional profile, with explicit gaps, allows comparing situations without claiming a single grade. An administration can be strong on contractual localisation and weak on tested reversibility. A provider can be "European" in the sense of a market estimate and poorly documented on operational control. A Member State can have electrical capacity without a clear inventory of ultimate ownership.

The profile makes these tensions readable. The podium erases them. By default, **no global composite score** is calculated. Eventual weightings in the measurement framework remain provisional; they do not justify ranking until tested in sensitivity with suitable data.

16.5 State of data: without interpolation

Book family	Intention	Current documentary state
Human capital	School pipeline, higher education, ops	Hypotheses and partial context; certification ≠ ops capability
Economy	Equivalent prices, total cost, exit	Methods + partial basket; total not calibrated
Technology / reversibility	Interop, exit tests, open source	Protocols; orchestration trial **not executed**
Industry	Datacentres, chips, energy, supply	Capacity, energy, typologies documented; ultimate ownership absent
Operations	Ops control, resilience	Contextualised; measured mastery absent
Governance	Law, public procurement, cooperation	Partial texts and practices; industrial effect not demonstrated

[3] Distinction between the book's measurement framework and institutional buyer tools.

Book family	Intention	Current documentary state
Market	Adoption, concentration	Distinct quantities, partially populated

When a cell says “absent” or “not executed”, the dashboard must show a gap, not discreet interpolation.

16.6 Working proposition

The dashboard useful to the thesis is a set of **documented, separable indicators**, with missing-value rules, not a single sovereignty index.[4]

16.7 Objections

One objects that a score is needed to decide. An unjustified score decides badly. One objects there are too many indicators: better few populated indicators than an opaque mosaic. One objects finally that without complete series the chapter is empty: on the contrary, structured admission of gaps is a condition of seriousness.

A political objection holds that a dashboard without score does not convince. Perhaps. The book prefers not to convince through false precision.

16.8 Toward trajectories

Measuring, here, means making visible, including what is missing. No ranking. No composite score by default. Partial populations stay separate.

On this methodological base, one can formulate **analytical scenarios**, not predictions, to expose trade-offs and conditions of failure. That is the subject of the next chapter.



[4] Analytical framework of the book, dashboard of capabilities and dependencies.

CHAPTER 17

European scenarios, trade-offs, and possible policies

Scenarios are not forecasts. They expose plausible trajectories, trade-offs, and observable signs, without invented probabilities. This chapter proposes three, analytically. It **does not close the book's thesis**. It also does not turn documentary gaps from industrial and operational chapters into scenarised certainties.[1]

The previous chapter's dashboard supplies vocabulary: inputs, capabilities, outcomes, resilience, visible gaps. Here, these elements are combined, or failure to combine them is shown. Each scenario must be confrontable with observables. Otherwise it is only strategic fable.

No calibrated simulation is produced. No probability is displayed. Trajectories are neither "total success" nor "total failure": each admits sectoral progress and persistent dependencies.

17.1 Scenario A: Persistent dependency

Hypotheses. Adoption continues to grow. Switching costs remain high and poorly measured. Exit tests remain rare. The industrial base and transferable operations skills stay insufficiently articulated. Policies multiply labels and localisations without verifying capacity to act.

Mechanisms. Network effects and economies of scale strengthen already dominant ecosystems. Public procurement buys territorial compliance more than tested reversibility. Training stays centred on single stacks. Local industrialisation produces demonstrators without operated offer at scale.

Decisions or conditions leading there. Indefinitely postponing migration exercises; treating multi-cloud as proof; merging adoption and sovereignty in political dashboards; funding research without tracking passage to supply.

Critical dependencies. Weakly substitutable platform layers; proprietary supervision tools; non-transferable skills; components and regions hard to replace quickly.

Observable signals. Estimated shares of major providers stable or rising; migration fees and delays unpublished; reversibility protocols not executed; restoration exercises absent or purely documentary.

Risks and trade-offs. High service continuity at global actors can coexist with

[1] Analytical framework of the book, analytical scenarios, not predictive.

increased criticality. The scenario is not moral accusation: it is a trajectory coherent with strong adoption and weak exit.

What would weaken it. Migrations, tested and repeated, on priority workloads; exit costs actually measured and falling; industrialised alternative supply used outside niches.

17.2 Scenario B: Targeted autonomy

Hypotheses. Europe does not aim for autarky. It chooses perimeters, sensitive data, certain public processing, stacks where interoperability is testable, and reduces critical dependencies there. Elsewhere, global dependencies persist and are mapped.

Mechanisms. Public procurement requires exit proof and verifies after award. Continuing education targets transferability, not badges alone. Industry and energy remain constraints, but bottleneck inventory improves. Reversibility tests become practice, not decorative clause.

Necessary decisions or conditions. Prioritise workloads; fund operations as much as prototype; accept sometimes higher total costs for credible exit; refuse single scores that hide gaps.

Remaining critical dependencies. Advanced semiconductors; certain global software ecosystems; rare talent; international interconnections.

Observable signals. Tested migrations on an explicit portfolio of workloads; verified reversibility clauses; partial but filled industrial indicators; adoption and shares still published separately from mastery.

Risks and trade-offs. Targeting can leave dependencies outside perimeter that prove critical too late. It can also fragment requirements if each Member State invents its own perimeter without coordination.

What would weaken it. Labels without tests; local preference without real capability; abandoning exercises at first budget difficulty.

Scenario B is not lukewarm compromise. It demands more empirical discipline than total-autonomy discourse. It also describes an **intermediate** trajectory: possible sectoral progress, assumed persistent dependencies.

17.3 Scenario C: Integrated ecosystem

Hypotheses. Human, economic, technological, industrial, and institutional capabilities reinforce each other at significant scale. European cooperation reduces fragmentation of requirements without abolishing all experimentation. Viable alternatives exist for a growing share of critical workloads, without claiming to eliminate every global dependency.

Mechanisms. The operations pool becomes documented and mobile across stacks. Total costs and exits are calibrated on public baskets. Industrialised supply dialogues with demand that verifies. The material base, capacity, energy, chains, is read with control inventories, not only building maps.

Necessary decisions or conditions. Funding continuity beyond demonstrator; cooperation on standards actually implemented; refusal to conflate market shares and mastery; investment in tested resilience as much as in communication.

Plausible critical dependencies even here. Upstream components; certain firmware and equipment; global network and talent interdependencies. Denying them would caricature the scenario.

Observable signals. Documented transferable ops pool; repeated migrations and restorations; industrial base and supply in measurable dialogue with demand; populated dashboard without deceptive composite score.

Risks and trade-offs. Cost and time of such articulation are high. Poorly designed harmonisation can stifle useful experiments. Partial success on one segment can be wrongly presented as global sovereignty.

What would weaken it. Non-reproducible success stories; fictitious indicator population; forgetting persistent dependencies.

Scenario C is not a promise. It describes a demanding articulation of capabilities, still not demonstrated at European scale in the current corpus.

17.4 Trade-offs across the three trajectories

Trade-off	Tension
Adoption speed vs reversibility	Adopting fast can increase lock-in
Catalogue cost vs exit cost	Apparent savings / hidden debt
Local preference vs real capability	Industrial support / façade
Harmonisation vs experimentation	Scale / diversity
Open source vs industrialisation	Open licence / operated offer
Required resilience vs SME access	High requirement / entry barrier

These trade-offs do not designate a winner. They recall that each trajectory pays a price, sometimes in dependency, sometimes in cost, sometimes in slowness.

17.5 Levers and conditions of failure

Levers evoked in the book, school, higher education, operations, price, total cost, exit tests, open source, industry, cybersecurity, funding, public procurement, cooperation, do not act in isolation. A single lever can fail if others are missing: training without labs; labs without industrialisation; law without operational control; open source without maintenance; purchasing without supply.

Scenarios serve precisely to make these failures **readable** before they become surprise narratives.

17.6 Objections

One objects that industrial policy is ineffective. Possible: scenarios must then show observable conditions of failure, not slogans. One objects that laissez-faire suffices: concentration and documented barriers to switching elsewhere make this objection debatable, not settled here. One objects finally that probabilities are needed: without calibrated model, displaying them would be imposture.

17.7 Toward the conclusion

Three scenarios, trade-offs, observables, no verdict. Intellectual closure of the book belongs to the next chapter.



CHAPTER 18

Toward progressive, measurable, realistic cloud autonomy

This book's thesis was not a disguised conclusion. It proposed examining an idea: European cloud sovereignty results from progressive development of interdependent human, economic, technological, industrial, and institutional capabilities, whose combination allows reducing critical dependencies on non-European actors.

After the journey, infrastructure, definitions, dependencies, human capital, costs, reversibility, open source, industry, cybersecurity, supply, public demand, market, measurement, scenarios, what can one **hold analytically**, and what must one leave open? This chapter is **framework synthesis**: what is established, plausible, or undemonstrated; interdependencies; conditions of refutation. It does not declare the book scientifically validated. The final address to the reader belongs to the general conclusion.

18.1 What the journey established

Certain distinctions hold solidly. Cloud is infrastructure; adopting it is not mastering it. Cloud sovereignty, in this book, is a **degree** of capabilities, not a label. Dependencies become critical through mechanisms, concentration, switching costs, data egress, limited interoperability, not through slogan. Declared adoption, estimated market shares, and degree of sovereignty remain distinct quantities. The dashboard refuses the single score. Scenarios expose trajectories without predicting.

Bounded facts anchor reasoning: declared adoption measured by the European survey; private concentration estimates to handle with sobriety; legal texts in force on markets and provider change; price comparison and total cost protocols that refuse invented gaps. These elements do not "prove" the thesis. They prevent discussing it in a vacuum.

18.2 What the journey made plausible

Other propositions remain plausible without being demonstrated at European scale. School and higher education form upstream conditions of the talent pool. Certifications do not prove transferability. Open source and standards are conditional levers. The industrial base and operational mastery **condition** dependency reduction. Industrialisation and funding condition the existence of viable alternatives. Public procurement can structure demand without automatically creating capabilities. Adoption, competition, and cooperation orient trajectories without dictating them alone.

These articulations form the core of the thesis. They explain why sovereignty reduces neither to **location** of machines nor to **preference** for a provider. Locating without controlling access, without ability to restore, without ability to exit, is not mastering. Preferring a "local" actor without industrialised supply or operations skills is not reducing a critical dependency; it sometimes displaces it.

18.3 What the journey did not demonstrate

The thesis did not prove that Europe has achieved cloud sovereignty. It did not prove that complete autonomy is possible or desirable. It did not produce a pan-European inventory of ultimate ownership of operators, homogeneously measured operational mastery, nor orchestration migrations executed in this repository. The price basket and total cost model remain partial. Industrial effects of public procurement are not established as a European constant. No calibrated simulation of scenarios was produced.

These absences are not footnote details. They define the boundary between what the book can assert and what it must leave open.

18.4 Interdependencies: why a single lever fails

Skills without labs, labs without industrialisation, industrialisation without solvable demand, demand without tested reversibility, law without operational control, location without control inventory: each isolated link can give the illusion of progress. The thesis insists on **combination**. Reducing a critical dependency often requires, at once, human know-how, honest economic calculation, technical possibility of exit, a readable material base, capacity to act under incident, and institutions that verify what they require.

That is why scenarios of the previous chapter are not political menus. They are ways of seeing how these links assemble, or come apart.

18.5 Progressive autonomy

In this book's analytical framework, **progressive cloud autonomy** denotes measurable, partial, and revisable reduction of critical dependencies, distinct from autarky and a single score.[1]

Measurable: separate indicators, tests, admitted gaps. Partial: certain global dependencies may remain. Revisable: what was critical yesterday may not be tomorrow, and the reverse. This definition is not an empirical law. It is a compass against two extremes: promising total independence, or renouncing all reduction because autarky is impossible.

18.6 Conditions of refutation

The working thesis would be weakened or revised if, notably:

- interdependent capabilities developed without any observable reduction of critical dependencies on tested perimeters;
- dependency reduction were obtained without these capabilities, for example by simple provider displacement without mastery;
- measurement could never separate adoption, shares, and capabilities;
- autarky proved the only coherent reading, which the book refuses as framing, but which would then need arguing otherwise.

18.7 Research needs

The corpus calls for work this book does not replace: inventories of ultimate ownership and control of infrastructure; homogeneous comparisons of operational mastery and restoration; public execution of reversibility protocols; calibration of total costs on representative baskets; tracking of research → operated supply passages; impact evaluations of public procurement on capabilities, not only contract amounts. While this evidence is missing, caution is not weakness. It is method.

18.8 Final objections

One objects that autonomy is impossible: then one must say which dependencies are ineliminable and why. One objects it is already achieved: lock-in mechanisms and gaps in tested reversibility make that claim premature. One objects it is undesirable: that is a normative debate distinct from measuring capabilities.

[1] Analytical framework of the book, progressive autonomy and limits.

18.9 Proposition to examine

The thesis remains a **proposition to examine**, enriched by vocabulary, distinctions, and protocols, not an established empirical law. Progressive, measurable, realistic cloud autonomy denotes here a direction of work: combine capabilities, test what can be tested, refuse deceptive scores, and admit what is still missing.

This chapter closes **analytical synthesis**. The general conclusion takes up the central question for the reader and opens on priorities the available corpus allows, without imposing them.



Conclusion

The starting question was simple to formulate and hard to settle: under what conditions can European cloud sovereignty be understood as the result of interdependent capabilities, rather than as a label already achieved or a mere hosting location?

The book answers first with a **method**. It refuses to merge adoption, market shares, location, ownership, operational control, catalogue price, total cost, legal obligation, and industrialisation. It treats cloud sovereignty as a **degree of capabilities** facing critical dependencies. This answer is not a verdict of sovereignty achieved. It is a framework for reading what is established, estimated, plausible, and still open.

The previous chapter drew the analytical inventory: what the journey established, made plausible, or left undemonstrated, as well as conditions of refutation. Here, the task is to assume this framework facing decisions and research to come.

1 What the reader can take away

Several distinctions hold. Adoption is not sovereignty. Location is not control. A catalogue price is not a total cost. A legal text is not industrialisation. A certification is not a transferable skill. Funding is not operated supply. These reading rules do not replace missing proof. They prevent concluding too quickly.

The thesis insists on **combination** of human, economic, technological, industrial, and institutional capabilities. An isolated lever can give the illusion of progress. Global dependencies can remain even on a favourable trajectory. The book promises neither autarky nor a single score.

2 Assumed limits

The homogeneous European inventory of ultimate ownership of operators is missing. Aggregated operational mastery is not measured comparably. The industrial effect of public procurement is not demonstrated. The price basket and total cost model remain partial. A reversibility protocol can be formulated without having been executed. Several industrial and institutional claims remain working hypotheses.

These gaps are part of the inquiry's result. They indicate what should be observed to strengthen, nuance, or refute the thesis, without filling them by invention.

3 Priorities proportionate to evidence

Without proclaiming a political roadmap, the book allows prudent analytical priorities. Clarify quantities before merging them. Require price and total cost protocols before rankings. Test reversibility rather than relying on clauses alone. Distinguish intention, act, and outcome in public programmes. Treat location, ownership, and control as separate questions. Refuse composite sovereignty scores without validated method.

These priorities do not replace democratic debate. They discipline it.

4 Opening

The cloud will remain, for long yet, a space of partial dependencies, economic tradeoffs, and unevenly distributed capabilities. The question is not whether Europe can produce everything alone. It is which dependencies it judges critical, which capabilities it actually builds, and how it avoids taking labels for proof.

European cloud sovereignty, if the term is to keep meaning, is not a flag planted on a datacentre. It is a trajectory of capabilities, demanding, measurable in pieces, always incomplete, whose gears this book sought to make visible.

Epilogue

A digital infrastructure becomes manageable only when one can understand its components, train its operators, measure real costs, test exits, and distinguish what one owns from what one controls. Without these capabilities, the cloud remains a promise of efficiency and a risk of dependency.

Europe need not choose between technological naivety and the fantasy of autarky. It must build, piece by piece, verifiable capabilities. Some evidence is still missing: control inventories, migrations actually executed, calibrated total costs, measured industrial effects. Admitting that is not giving up. It is refusing to take a label for mastery.

Cloud sovereignty, if the word is to keep meaning, is not decreed on a cover. It is built in laboratories, contracts, factories, budgets, and skills, where one finally accepts to **measure** what one claims to master.



Visual continuity (conceptual illustration).

END OF THE BOOK

End of the volume

The interior pages end here. The empirical limits stated in the book remain open: they belong to the argument.

Back cover

Why speak of cloud sovereignty if hosting in Europe guarantees neither mastery, nor exit, nor industrialisation? This book offers a framework: interdependent human, economic, technological, industrial and institutional capabilities, whose combination may reduce critical dependencies, without proclaiming sovereignty already achieved.

Adoption, market shares, localisation and control: distinct quantities.



FRANCK MEVENGUE