RefleXio Intelligence
Strategic Dossier · Issue No. 01

Sovereign Cloud Migration in Europe

Stack reality, hidden costs and strategic trade-offs

A decision-support dossier on jurisdiction, localization and sovereignty — calibrated for European enterprises navigating GDPR, NIS2, DORA, the AI Act and the post-Schrems II regulatory perimeter

refleXio

STRATEGIC DOSSIER

Issue No. 01 · April/ May 2026 · v4.0

Sovereign Cloud Migration in Europe

Stack reality, hidden costs and strategic trade-offs

A decision-support dossier on jurisdiction, localization and sovereignty, calibrated for European enterprises navigating GDPR, NIS2, DORA, the AI Act and the post-Schrems II regulatory perimeter.

Publication date	April/ May 2026
Publisher	RefleXio Intelligence
Series	RefleXio Intelligence Strategic Dossiers, Issue No. 01
Version	4.0, expanded framework with CLOUD Act quantification, reference pricing, opportunity cost, AI Act operational implications, energy constraints and weak signals section
Prepared by	RefleXio Intelligence Editorial Desk

Independent strategic intelligence on sovereign cloud, AI compute and the European digital stack.

Contents

Key numbers at a glance

Nine figures that frame the rest of the dossier. Each is sourced and repeated, with context, in the relevant section.

70% **US hyperscaler share** *European cloud infra (2025)*	**€264B** **Annual EU→US outflow** *Cloud & software spend (≈1.5% of EU GDP)*	**15%** **European provider share** *European cloud infra (2025)*
$6.7B → $23.1B **European sovereign IaaS** *Triples 2025 → 2027*	**$80B** **Global sovereign cloud** *2026 spend, +35.6% YoY (Gartner)*	**20–30%** **AWS ESC premium** *Over standard AWS pricing*
<0.01% **CLOUD Act low-risk** *Annual probability for most enterprises*	**4–5%** **European AI compute** *Global capacity share*	**7–10 yrs** **FLAP-D grid waits** *Frankfurt, London, Amsterdam, Dublin*

Reading note. Market-share and spend figures come from Broadcom, Gartner via CIO Dive, Yahoo Finance/TechRadar, and Grand View Research. CLOUD Act probability is RefleXio analysis of Microsoft Government Requests Report. Grid-wait figures are from Reuters AWS interview and IEA European grid analysis. Full sources in Section 17.

Executive summary

Market context

€264 billion in annual cloud and software spending flows from Europe to US hyperscalers[1], equivalent to ≈1.5% of EU GDP. Three American companies hold 70% of the European cloud infrastructure market[2]; European providers hold 15%. Global sovereign cloud spend will reach $80 billion in 2026, up 35.6% year-on-year[3]. Europe's sovereign IaaS market alone will triple from $6.7 billion in 2025 to $23.1 billion by 2027[4]. The broader European sovereign cloud market (IaaS, PaaS, SaaS and managed services) stood at $32.8 billion in 2025 and is projected to reach $191.6 billion by 2033[5], growing at 25% CAGR. 60% of Western European CIOs want to increase use of local cloud providers[6]. The directional shift is structural and accelerating.

Core problem

This movement is not driven by technology preference. It is driven by three converging forces. A regulatory framework (GDPR, NIS2, DORA, AI Act, Data Act, the EU Cloud Sovereignty Framework) renders standard public cloud deployments legally non-compliant for broad categories of workloads. Geopolitical tension (the unresolved conflict between the US CLOUD Act and GDPR Article 48) has sharpened awareness of extraterritorial legal exposure. And a market concentration so extreme that meaningful diversification requires rebuilding architectural foundations, not simply switching vendors, completes the picture. AI is accelerating this problem by deepening dependency on hyperscaler-native services (model training, inference and data pipelines) at precisely the stack layers where European alternatives are weakest.

Why sovereign strategies fail

Most organizations pursuing "sovereign cloud" are solving the wrong problem. They conflate three distinct concepts (jurisdiction, which legal system governs the provider; localization, where data is physically stored; and sovereignty, who controls access and operations) and treat them as interchangeable when they are not. A German datacentre operated by a US company resolves localization but not jurisdiction. A European provider with US-sourced firmware resolves jurisdiction but not supply-chain sovereignty. This misdefinition produces predictable failures, namely migration cost overruns, architectural regressions and compliance gaps. McKinsey has estimated more than $100 billion in global cloud migration cost overruns[7]. The Bank of England's Oracle migration (budgeted at £7 million, ultimately exceeding £18.5 million[8]) is a reference case, not an outlier. The sovereign imperative is real; the path to it demands precision.

What this report delivers

A decision-support framework for C-level executives, technology leaders and operational managers, structured around seven analytical pillars: (1) regulatory drivers and their binding operational implications by sector; (2) a conceptual framing that separates jurisdiction, localization and sovereignty as independent design variables; (3) a layer-by-layer assessment of what European providers can and cannot deliver today, including where migration actually breaks; (4) a structured cost decomposition exposing the categories organizations systematically underestimate; (5) a risk

framework calibrated by industry vertical; (6) three deployment scenarios (full European migration, hybrid sovereignty, sovereignty overlay) with explicit trade-offs; and (7) a decision checklist that converts analysis into action. The report concludes with a forward-looking outlook on infrastructure, policy and market developments expected through 2028.

Notes. 1 Broadcom on €264B EU→US outflow · 2 Broadcom on 70% hyperscaler share · 3 CIO Dive / Gartner on $80B in 2026 (+35.6% YoY) · 4 Yahoo Finance / TechRadar on EU IaaS $6.7B→$23.1B · 5 Grand View Research on $191.6B by 2033 · 6 Broadcom on 60% of CIOs · 7 McKinsey via SoftwareSeni on $100B+ overruns · 8 SoftwareSeni on Bank of England £7M→£18.5M overrun

1. Introduction and scope

1.1 Defining sovereign cloud and AI infrastructure

Sovereign cloud is not a product category. It is a governance posture applied to cloud infrastructure, a set of technical, legal and contractual controls that ensure data, workloads and the AI systems operating on them remain subject to a defined legal jurisdiction and are protected from unauthorized extraterritorial access or control.

For operational purposes, sovereign cloud infrastructure encompasses six functional layers:

- **Compute** comprises general-purpose virtual machines, bare metal and high-performance computing nodes with GPU clusters for AI inference and training workloads.
- **Storage** comprises object, block and file storage with cryptographic controls on key custody.
- **Networking** comprises private connectivity, load balancing, content delivery and edge capabilities.
- **Data platforms** comprise managed databases, data warehouses, streaming and analytics pipelines.
- **AI/ML services** comprise model training platforms, inference endpoints, MLOps toolchains and foundation model APIs.
- **Governance layer** comprises IAM, KMS, audit logging, policy enforcement and the contractual perimeter.

Each layer carries its own sovereignty obligations, and each has its own implementation reality on European providers. The layers are not independent: an organization that resolves sovereignty at the compute layer but leaves AI/ML dependencies untouched has not resolved sovereignty at the stack level.

1.2 What this report is and is not

It is an analytical and strategic dossier for decision-makers who must price regulatory, jurisdictional and operational risk across the cloud stack. It synthesizes primary reporting, EU regulatory texts, provider documentation and independent market research into a decision framework that distinguishes the three sovereignty dimensions and maps them to deployment models.

It is not legal advice, a procurement recommendation, a detailed vendor evaluation or a trading or investment recommendation. Organizations should engage qualified legal, compliance and procurement professionals before committing architectural or contractual decisions.

1.3 Scope

This edition covers the European Union and EEA, with primary focus on France, Germany, the Netherlands, the Nordics, Italy, Spain and Poland. It addresses cloud infrastructure, platform services and AI compute, including the intersection with the European AI Act and the EU Cloud Sovereignty Framework v1.2.1. It excludes sector-specific compliance frameworks not directly

intersecting with cloud governance (e.g., MDR for medical devices), except where they materially alter provider selection (NEN 7510 in the Netherlands, HDS in France).

Notes. Section foundational; no new citations. Layer definitions synthesise European Commission AI Factories taxonomy and provider documentation.

2. Methodology and analytical approach

2.1 Sources

The analysis synthesizes five source families: (i) EU primary legal texts (GDPR, NIS2, DORA, AI Act, Data Act, EU Cloud Sovereignty Framework v1.2.1) and European Commission digital strategy publications; (ii) primary reporting from Reuters, CIO Dive, Lawfare, ERP Today, The Register and European trade press; (iii) analyst and market research from Broadcom, Gartner (via CIO Dive), Grand View Research, MarketsandMarkets, Yahoo Finance/TechRadar and BCG (via Civo); (iv) provider documentation from AWS, OVHcloud, Scaleway, STACKIT, IONOS, Exoscale, Hetzner, Civo, Kiteworks and Sovereign Cloud Stack; and (v) independent technical/policy analysis from CISPE, Euro-Stack, SoftwareSeni, CleanAim, Peerobyte, Raconteur and the EuroHPC JU. For new v4 material: US CLOUD Act legal analysis draws on Microsoft Transparency Report, Justia (US v. Microsoft Corp.), ComplianceHub and Lawfare. Grid-constraint analysis uses Reuters AWS interviews and The Register. Full APA 7 references appear in Section 17.

2.2 Analytical lens

Three analytical principles shape the argument:

- **Jurisdiction, localization and sovereignty are separable.** The dossier treats them as independent design variables and warns against the rhetorical move of using one word to mean any of the three.
- **Stack-aware migration economics.** Cost analysis is decomposed by workload layer rather than by provider. This framing exposes where refactoring cost concentrates and why infrastructure-line pricing calculators systematically understate total cost.
- **Scenario-based planning, not vendor ranking.** The deployment framework produces three scenarios (full migration, hybrid, overlay) and a decision tree; it does not rank providers, because optimal provider choice depends on which scenario the organization has selected.

2.3 Limitations

Market-share, spending and cost-percentage figures are rounded for readability and reflect third-party research available at the publication date. Forward-looking projections (sovereign cloud market trajectories, AI compute evolution, regulatory implementation timelines) are directional, not deterministic. The AWS European Sovereign Cloud launched in January 2026 and has only three months of operational track record at the time of writing; commercial performance data is limited and projections should be read as directional. The CLOUD Act risk stratification in Section 4.2 uses Microsoft transparency data (the most comprehensive public dataset) but other providers' transparency reports are less granular, so the base rate may be understated. RefleXio Intelligence maintains no commercial relationship with any provider analysed herein.

Notes. Source families summarised inline. Full APA 7 list in Section 17.

3. The sovereignty problem is misdefined

The most consequential error in European sovereign cloud strategy is linguistic, not technical. Three distinct concepts are systematically conflated:

Dimension	Core question	Typical indicator	Is it resolved by a German datacentre run by a US firm?
Jurisdiction	Which legal system governs the provider?	Parent-entity incorporation; CLOUD Act exposure	No (parent jurisdiction remains US)
Localization	Where is the data physically stored?	Datacentre location, replication zones	Yes (data is physically in Germany)
Sovereignty	Who controls access and operations?	Key custody, ops staff citizenship, audit rights	Partially (contractual but structurally unresolved)

The practical consequence is that organizations buy what the market labels "sovereign" without verifying which of the three dimensions is actually resolved. A European provider with US-sourced hardware firmware is jurisdictionally European but not supply-chain sovereign[18]. A hyperscaler sovereign partition resolves data residency and, contractually, operational access, while leaving parent-jurisdiction exposure structurally intact[19]. The EU Cloud Sovereignty Framework's SOV-5 (Supply Chain Sovereignty) objective[20] captures this dimension explicitly, weighted at 20% of the total SEAL score.

The strategic implication is that organizations must specify which dimension is their actual objective before selecting architecture. The three dimensions rarely carry equal weight for a given institution:

- **Public-sector buyers** typically prioritize jurisdiction, because constitutional protection of data of record outweighs operational convenience.
- **Regulated financial institutions** prioritize documented sovereignty controls, because DORA demands auditable operational posture, not necessarily a European parent entity.
- **Healthcare organizations** often prioritize localization, because national certification frameworks (NEN 7510, HDS) operate on geographic data residency.
- **Industrial firms** prioritize supply-chain sovereignty for OT data and IP, less about jurisdiction, more about whom one is dependent on along the stack.

An organization that selects a fully EU-sovereign provider but builds deeply on proprietary managed services without exit provisions has achieved provider sovereignty without operational sovereignty, which means that it has traded one form of lock-in for another.

The insight that should shape every architecture decision in this report is that **sovereignty is not a binary property of a provider. It is a multi-dimensional property of an architecture**. No single vendor selection resolves all dimensions. The correct framing is not "which provider is sovereign?"

but "which combination of architectural choices produces the sovereignty profile that matches the organization's regulatory obligations, risk tolerance and operational capacity?"

This framing (jurisdiction, localization and sovereignty as independent design variables) is the analytical foundation for the sections that follow.

Notes. 18 Exoscale on CLOUD Act applying despite localization · 19 SoftwareSeni (providers) on closed-source management planes · 20 European Commission on SOV-5 Supply Chain Sovereignty (weighted 20%)

4. The sovereign imperative

The shift toward sovereign cloud in Europe is not an expression of technology preference. It is the operational consequence of four convergent pressures, namely a maturing EU regulatory stack, an unresolved US–EU legal tension, sector-specific exposure asymmetries and a severe workforce constraint. This section documents each in turn.

4.1 Regulatory drivers

European organizations face a layered regulatory environment in which cloud architecture is no longer an IT decision made in isolation from compliance obligations. The following table summarizes the primary regulations affecting cloud infrastructure choices:

Regulation	Effective date	Primary scope	Direct cloud impact
GDPR	May 2018	Personal data processing, all sectors	Data residency, lawful transfers, controller-processor contracts; Article 48 conflicts with CLOUD Act
NIS2 Directive	Oct 2024 (national transposition)	18 critical sectors including energy, transport, health, finance, digital infra	Supply-chain security, incident reporting obligations for cloud providers
DORA	17 Jan 2025	Financial entities: banks, insurers, investment firms, payment processors	ICT third-party risk management, mandatory contract clauses specifying data location, audit rights
Data Act (Chapter VII)	Sep 2025	Cloud service providers, all sectors	Providers must implement measures to prevent unlawful non-EU government data access
AI Act	Aug 2024 (force); Aug 2026 (full)	AI system providers and deployers	Cloud delivery of GPAI models = placement on EU market; high-risk AI audit trails and data governance
EU Cloud Sovereignty Framework v1.2.1	Oct 2025	European Commission procurement	8 sovereignty objectives, SEAL scoring, €180M Cloud III DPS (procurement-grade sovereignty standards)

The regulatory trajectory is unambiguous. Each successive instrument expands the population of organizations with binding cloud compliance obligations and tightens the specific technical requirements those obligations impose. DORA alone brought the entirety of EU financial services under structured ICT third-party oversight with contractual teeth. The AI Act (entering full application in August 2026) will extend equivalent scrutiny to the AI services layer sitting atop cloud infrastructure.

Figure 1 - European Cloud Sovereignty Regulatory Timeline

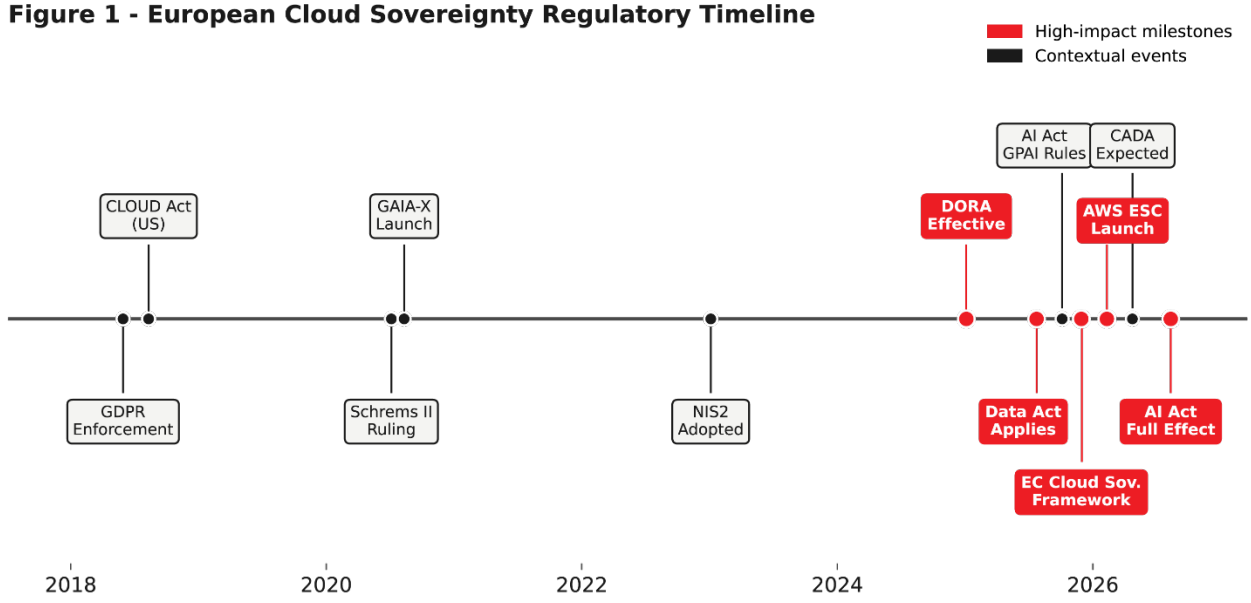

Source: RefleXio Intelligence synthesis of EU legal instruments and US CLOUD Act jurisprudence.

Figure 1. European Cloud Sovereignty Regulatory Timeline (2018–2026). High-impact milestones shown in red.

4.2 Geopolitical and legal tension

The fundamental jurisdictional conflict in European cloud strategy is not rhetorical; it is structural and legally unresolved. The US CLOUD Act (2018) grants American authorities the power to compel US-headquartered cloud providers to produce data stored anywhere in the world, regardless of the data's physical location. GDPR Article 48[25] prohibits transfers of personal data to third countries except pursuant to an international agreement. These two legal instruments impose directly opposing obligations on US cloud providers operating in the EU.

The distinction between jurisdiction and localization is critical and routinely conflated. Localization (storing data on servers physically located in the EU) does not resolve the CLOUD Act exposure because the exposure is jurisdictional, not geographic. A US provider's German datacentre remains subject to CLOUD Act demands because the parent entity is subject to US jurisdiction. Schrems II (2020)[26] sharpened this conflict by invalidating the Privacy Shield framework and establishing that US national security law creates surveillance access incompatible with EU fundamental rights standards.

The EU Data Act Chapter VII (effective September 2025)[27] directly addresses this by requiring cloud service providers to take technical and legal measures to prevent third-country government access that violates EU law. This shifts the burden explicitly onto providers (including US hyperscalers operating in Europe) to architect against the very access mechanisms their parent-country law may compel.

The architectural resolution that has emerged in practice is customer-controlled encryption with keys held entirely outside US-entity infrastructure. If a US provider cannot access the plaintext keys, a CLOUD Act demand for readable data cannot be technically fulfilled. This is the approach embedded in the AWS European Sovereign Cloud (launched January 2026)[28], Microsoft's EU Data Boundary and Oracle's EU Sovereign Cloud. Whether this architecture withstands a legal challenge in which the US government demands the keys themselves (not just the ciphertext) remains untested in court.

Quantifying the CLOUD Act risk

The rhetorical force of the CLOUD Act concern is well established; the empirical force is less often stated. Microsoft's transparency reporting (the single most comprehensive public dataset) shows that across all CLOUD Act content requests received from US law enforcement for non-US enterprise customers, Microsoft produced data for five non-US enterprise customers whose data was stored outside the US, and only one of those five was located in the EU, and it was a US Government contractor[29]. This is the only quantitative baseline available for enterprise EU data actually reached by CLOUD Act process.

The jurisprudence is equally illustrative. *United States v. Microsoft Corp.* (the Microsoft Ireland case) was rendered moot by the CLOUD Act itself after seven years of litigation[30]. Since then, the CLOUD Act's extraterritorial reach has not been tested at the Supreme Court level against an EU-specific shielding architecture (customer-held keys, EU-resident operations, contractual prohibitions). Providers retain statutory rights to challenge orders on comity grounds, and the UK-US agreement has absorbed the bulk of cross-border data traffic, with more than 20,000 orders transmitted under it since entry into force in October 2022[31], demonstrating that executive agreements meaningfully channel law enforcement demand.

A calibrated risk framework for organizations evaluating their exposure:

Risk profile	Organization type	Est. annual probability of CLOUD Act reaching EU data	Recommended posture
Low	Standard commercial enterprises, SMEs, consumer SaaS, non-sensitive industrial	< 0.01%	Sovereignty overlay or hybrid is proportionate
Medium	Financial institutions with US exposure, healthcare handling US patient data, multinationals with US subsidiaries	0.01%–0.1%	Architectural separation for regulated data; hybrid sovereignty
High	Defence, intelligence, government ministries, critical infrastructure with classified workloads, organizations under active US sanctions or investigation	> 0.1%, and reputationally catastrophic if realized	Full European migration; air-gapped or on-premise for classified workloads

This framework does not eliminate the sovereignty concern; it calibrates it. For most European enterprises, the CLOUD Act risk is empirically low but legally unresolved. For a narrow class of organizations, it is both legally acute and operationally unacceptable. Conflating these categories produces over-investment in the first case and under-investment in the second. The February 2025 Apple / UK decryption demand episode demonstrated the residual abuse potential, because the UK used its CLOUD Act agreement to compel Apple to disable Advanced Data Protection globally, triggering US legislative reform efforts[32]. The risk is not zero; it is stratified.

A diplomatic resolution is in progress but far from complete. The CLOUD Act authorises bilateral executive agreements that could, in principle, harmonize US and EU legal frameworks for cross-border data access. The US has concluded such agreements with the United Kingdom (2022) and Australia (2024)[33], providing a precedent. Negotiations with the European Union are underway but have proved lengthy[34], complicated by the question of whether the EU as an institution qualifies as a "Qualifying Foreign Government" under the Act or whether agreements must be concluded with individual Member States. A framework agreement (in which the EU sets parameters and Member States sign individually) has been proposed as a pragmatic path. Until such an agreement enters into force, the CLOUD Act / GDPR conflict remains structurally unresolved, and organizations should plan their architectures on the assumption that resolution is years away, not months.

4.3 Sector exposure

Sovereign urgency is not uniform across sectors. The table below maps primary regulatory drivers, urgency level and current hyperscaler dependence by sector.

Sector	Primary regulatory drivers	Sovereignty urgency	Current hyperscaler dependence
Public administration	Cloud Sovereignty Framework, GDPR, NIS2, procurement rules	Critical (public trust, data of record)	High; most agencies on at least one hyperscaler service
Banking / finance	DORA, GDPR, EBA guidelines, central-bank oversight	Critical (systemic risk, audit obligations)	Very high; core banking increasingly cloud-native
Healthcare	GDPR special category, NIS2, NEN 7510 (NL), HDS (FR)	High (patient data sensitivity)	Moderate to high; SaaS and analytics adoption growing
Telecoms	NIS2 (essential entity), GDPR, national security clearance	High (critical infrastructure)	High; NFV on hyperscaler infra
Manufacturing / industrial	NIS2, AI Act (AI in production), Catena-X	Moderate–High (OT/IT convergence, IP protection)	Moderate; accelerating with Industry 4.0

Public sector and financial services face the most immediate compliance deadlines and the highest penalty exposure. Healthcare organizations, particularly in the Netherlands (NEN 7510) and France (HDS certification), operate under national health data frameworks that constrain provider selection

even beyond GDPR. Industrial organizations face a different but emerging challenge, because as OT systems connect to cloud platforms, manufacturing IP and process data acquires sovereignty obligations not previously relevant to discrete IT environments.

4.4 The talent constraint

The sovereign cloud transition faces a workforce bottleneck that compounds every other challenge in this report. Europe's cloud engineering talent pool has been trained predominantly on hyperscaler platforms. AWS, Azure and GCP certifications dominate the European IT labour market; engineers with production experience on OVHcloud, Scaleway, IONOS or STACKIT are a fraction of the available workforce. The European Commission has estimated the EU faces a shortage of over 1 million ICT specialists[35], a shortage expected to widen as NIS2, DORA and the AI Act generate demand for compliance-specialized cloud roles[36].

The gap is strategic, not only operational. An organization that migrates to a European provider but relies on a hyperscaler-trained team to operate it will underperform, because operational patterns, debugging workflows and cost optimization techniques differ between platforms. Organizations planning sovereign migration should initiate training and hiring programs 6–12 months before the migration program begins, not during it. The alternative (paying premium rates for scarce sovereign-skilled consultants under deadline pressure) is among the hidden costs identified in Section 7.

4.5 GAIA-X: conceptual ambition vs operational reality

4.5.1 Original intent

GAIA-X was launched in 2020 as the EU's flagship initiative to create a federated, interoperable European cloud ecosystem, a framework that would allow European data to flow between providers under common trust and technical standards without dependency on US infrastructure. The ambition was genuine and the problem it addressed was real.

4.5.2 Conceptual strength

GAIA-X's contribution to the ecosystem is in its **conceptual architecture**: federation standards for data spaces, a hierarchical trust framework (Standard, Level 1, 2, 3 labels) and the development of sector-specific data exchange ecosystems. The automotive data space **Catena-X**[37] (built on GAIA-X federation principles and now encompassing 180+ data spaces) is the most operationally significant implementation and demonstrates that the underlying federation model has genuine value for industrial supply-chain data sharing.

4.5.3 Execution gap

The gap between GAIA-X's conceptual architecture and operational delivery is substantial and well-documented. **Scaleway**, one of France's largest cloud providers and a founding GAIA-X member, withdrew in November 2021[38] after concluding the initiative had become too process-heavy to produce tangible results. **Agdatahub** (a GAIA-X agricultural data space initiative) was liquidated despite receiving €4.8 million in funding and holding Day-1 membership[39]. Frank Karlitschek, founder of Nextcloud, publicly described the organization as a "paper monster"[40], that is, an institution that generates documents and working groups without producing deployed infrastructure. Perhaps most damaging to the initiative's credibility, GAIA-X admitted major US hyperscalers as members[41] (the very organizations the initiative was positioned to provide an alternative to), generating "Trojan

horse" criticism that the framework was being shaped by incumbents to preserve their market position.

4.5.4 Real-world impact

The market data during the GAIA-X era does not suggest the initiative meaningfully changed competitive dynamics. European providers' share of the European cloud infrastructure market declined during this period, with three US hyperscalers now holding 70% of the market[42] and European providers holding approximately 15%. The initiative's own CEO acknowledged it had been too ambitious[43] in scope relative to its operational capacity.

4.5.5 Strategic interpretation

GAIA-X should be understood as a **framework initiative, not an infrastructure initiative**. Organizations that expected it to produce a deployable cloud platform were always operating on a category error. Its lasting contribution is in establishing the conceptual vocabulary and data space standards that sector-specific implementations (Catena-X being the primary example) are now operationalizing. Technology leaders should engage with GAIA-X artefacts as reference architecture rather than as implementation-ready specifications, and should treat the federation certification labels as a useful but not comprehensive signal of sovereignty posture.

The strategic takeaway is precise. GAIA-X informs sovereignty discourse at the conceptual and policy layer, but it does not resolve implementation challenges at the infrastructure or application layer. The gap between GAIA-X's trust framework and a production-ready sovereign cloud deployment must be filled by actual providers, actual architecture and actual investment, which is the subject of the sections that follow.

Notes. 21 EC Digital Strategy on 18 critical sectors (NIS2) · 22 EIOPA on DORA effective date · 23 EC Digital Strategy on AI Act dates · 24 Sovereign Cloud Stack on Framework v1.2.1 (Oct 2025) · 25 Exoscale on GDPR Article 48 · 26 Exoscale on Schrems II (2020) · 27 Kiteworks on Data Act Chapter VII · 28 AWS on ESC launch (January 2026) · 29 Microsoft Transparency Report: Government Requests · 30 Justia Supreme Court, United States v. Microsoft Corp. · 31 BSA TechPost on 20,000+ UK-US executive agreement orders · 32 ComplianceHub on Apple/UK decryption demand · 33 Eurojust on UK (2022) and Australia (2024) · 34 Lawfare on EU negotiation complexity · 35 EC Digital Strategy on 1M+ ICT shortage · 36 SoftwareSeni (providers) on technical skill demand · 37 Catena-X federation · 38 Euro-Stack on Scaleway withdrawal · 39 Euro-Stack on Agdatahub liquidation · 40 Euro-Stack on "paper monster" · 41 Euro-Stack on hyperscaler GAIA-X membership · 42 Broadcom on 70% US hyperscaler share · 43 Euro-Stack on GAIA-X CEO acknowledgement

5. The stack reality problem

Where migration actually breaks.

Before examining the stack layer by layer, three service categories consistently produce the most severe technical and commercial friction in sovereign migration programs. These are not edge cases; they are the dominant cost drivers and timeline risks:

- **Serverless / Functions-as-a-Service** (AWS Lambda, Azure Functions, Google Cloud Run). Event-driven architectures built on serverless primitives cannot be replatformed; they must be refactored. The trigger integration model, execution environment and state management are entirely proprietary. This is consistently the most underestimated refactoring effort.
- **Proprietary databases** (AWS Aurora, Azure Cosmos DB, Google Spanner). These offer performance and scalability profiles that have no European equivalent. Organizations with production workloads on these services face the choice between accepting capability regression or investing in significant application-layer rewriting.
- **AI/ML platforms** (AWS SageMaker, Google Vertex AI, Azure Machine Learning). The most strategically consequential gap. These are not individual services; they are integrated ecosystems covering model training, experiment tracking, feature stores, model serving and MLOps pipelines. **No European provider offers an equivalent platform**[44]. AI is the forcing function that deepens hyperscaler dependency at the fastest rate, because every new AI investment creates new proprietary service bindings.

Organizations should assess their exposure to these three categories before committing to a migration archetype or timeline. If dependency is deep in all three, Lift-and-Shift is architecturally impossible and refactor timelines will be 12–18+ months.

5.1 The myth of portability

A foundational assumption in cloud strategy (that workloads can be moved between providers as commercial or compliance conditions change) does not reflect the operational reality of modern cloud architectures. Portability is real at the container and standard data format layer: a Docker container, a PostgreSQL dump, a Kafka event stream can in principle run on any provider. But the value delivered by modern cloud platforms has migrated upward into the managed service layer, where portability breaks down.

71% of businesses report that vendor lock-in risks deter further cloud adoption[45], and the concern is well-founded. Lock-in in contemporary cloud architectures operates across four distinct dimensions:

- **Technical lock-in arises from proprietary service APIs (AWS Lambda invocation model, Azure Cosmos DB wire protocol, Google BigQuery SQL extensions) that have no portable equivalents. Applications built against these APIs require code changes, often substantial,** to migrate.
- **Economic lock-in arises from egress fees of $0.08–0.12 per GB46, which create structural cost barriers to data movement. At that rate, a 10TB dataset costs approximately $900 simply to exit a provider,** before any application-layer migration cost.

- **Legal lock-in arises from enterprise agreements with volume commitments, reserved instance purchases and enterprise support contracts, which** create financial penalties for early exit that are rarely modeled in initial migration business cases.
- **Operational lock-in arises as teams** build institutional knowledge, tooling integrations and runbooks around a specific provider's operational model. The organizational cost of retraining and rebuilding that knowledge is rarely quantified but is consistently one of the largest line items when it materializes.

The strategic implication is that the optimal time to architect for sovereignty is **before** deep platform adoption, not after.

5.2 Layer-by-layer stack comparison

Figure 2 - European Provider Stack Maturity vs Hyperscaler Equivalence

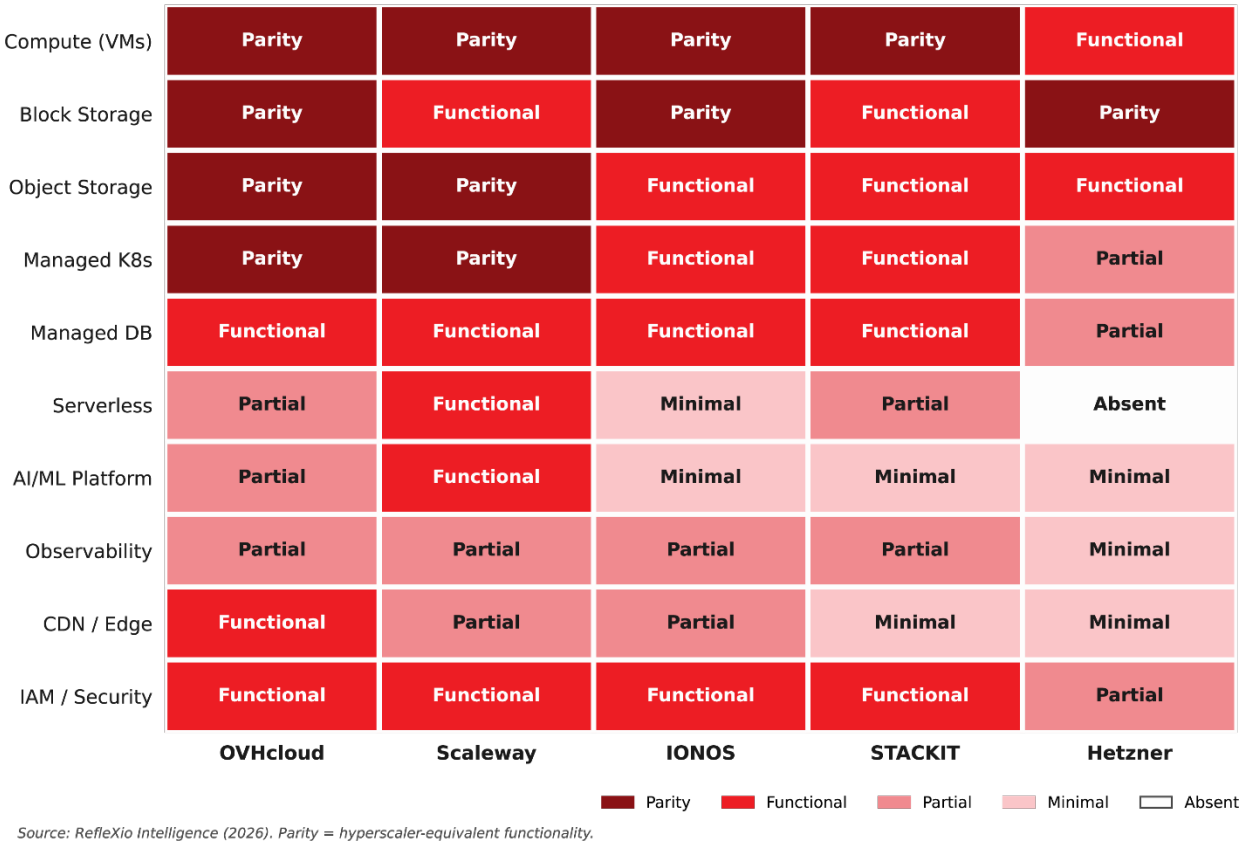

Figure 2. European provider stack maturity vs hyperscaler equivalence. Parity = hyperscaler-equivalent functionality.

The granular picture by layer:

- **Compute (VMs and bare metal) is functionally equivalent to hyperscalers. Hetzner in particular offers competitive bare metal at pricing that routinely undercuts hyperscaler equivalents. The gap is not compute hardware;** it is the surrounding operational services (auto-scaling, placement groups, dedicated host options, spot/preemptible instance markets).
- **Storage, meaning object storage (S3-compatible APIs), block storage and managed file systems, is** available at functional equivalence. OVHcloud's object storage and IONOS's S3-compatible offering cover the majority of enterprise use cases. Cold-tier archival pricing is less competitive than AWS Glacier or Azure Archive.

- **Managed Kubernetes is a** layer of genuine parity. OVHcloud Managed Kubernetes, Scaleway Kapsule, IONOS Managed Kubernetes and STACKIT Kubernetes Engine all provide managed control planes with functional equivalence to EKS/GKE/AKS for standard workloads. Differences emerge at scale (10,000+ node clusters) and in integrated tooling depth.
- **Managed databases, meaning PostgreSQL and MySQL managed services, exist across the European provider landscape. The gap is in proprietary database services (AWS Aurora, Azure Cosmos DB, Google Spanner),** which have no European equivalents.
- **For serverless and FaaS,** Scaleway Serverless Functions is the most mature European offering, but the ecosystem remains limited compared to AWS Lambda or Azure Functions in terms of trigger integration depth, runtime support breadth and ecosystem tooling.
- **AI/ML services present the** most significant capability gap. There is no European equivalent to AWS SageMaker, Google Vertex AI or Azure Machine Learning as integrated platforms. European AI infrastructure capability exists at the compute layer (EuroHPC's JUPITER exascale with 24,000 NVIDIA GH200 nodes, 19 AI Factories across Europe) but this is largely academic/research infrastructure, not enterprise-grade managed AI platform services. European LLMs (Token-7B (24 EU languages), Apertus, Mistral) are emerging, but Europe accounts for only 4–5% of global AI compute capacity.
- **Observability** remains largely self-managed on European providers. The practical resolution is deploying open-source stacks (Grafana, Prometheus, OpenTelemetry, Loki) on European infrastructure.
- **CDN and edge capability is limited** compared to CloudFront or Fastly. OVHcloud's CDN and UpCloud's network offer regional edge capabilities, but global content delivery at hyperscaler scale requires integration with a separate CDN provider or acceptance of performance trade-offs.
- **For IAM and security, industry**-standard IAM capabilities (RBAC, MFA, OAuth/OIDC, audit logging) are available, but the depth of fine-grained policy controls, conditional access and security tooling integration available in AWS IAM + Security Hub, Azure Entra + Defender or Google Cloud Identity + Security Command Center is not matched.

5.3 Equivalence analysis

The term "equivalence" requires precision. The following table distinguishes between three levels of equivalence assessment:

Dimension	Hyperscaler (AWS/Azure/GCP)	European provider (OVH/Scaleway/STACKIT)	Gap assessment
Compute (IaaS)	Full instance portfolio, global zones, auto-scaling	Competitive instance range, EU zones	Functional equivalence; operational maturity gap
Managed Kubernetes	EKS / GKE / AKS, deep integration	Managed K8s on all major providers	Functional equivalence for standard workloads
Object storage	S3, GCS, Azure Blob (mature, global)	S3-compatible, EU-based	Functional equivalence; global redundancy limited
Relational DB (PostgreSQL/MySQL)	RDS, Cloud SQL, Azure DB	OVHcloud, Exoscale, Scaleway DBaaS	Functional equivalence

Dimension	Hyperscaler (AWS/Azure/GCP)	European provider (OVH/Scaleway/STACKIT)	Gap assessment
Proprietary DB (Aurora, Cosmos, Spanner)	Native, mature	No equivalent	Critical gap; requires refactoring
Serverless / FaaS	Lambda, Azure Functions, Cloud Run	Scaleway Serverless; limited elsewhere	Significant gap; event-driven architectures affected
AI/ML platform	SageMaker, Vertex AI, Azure ML	No equivalent	Critical gap; EU research HPC not enterprise PaaS
Observability	CloudWatch, Azure Monitor, GC Ops	Self-managed (Grafana/Prometheus)	Operational gap; solvable with open-source investment
CDN / edge	CloudFront, Azure CDN, Cloud CDN	Limited; OVHcloud CDN	Material gap for performance-sensitive workloads
Security tooling	Security Hub, Defender, SCC	Basic; third-party supplementation needed	Significant for complex security architectures
Developer ecosystem	Extensive marketplace, integrations	Smaller ecosystem	Ecosystem maturity gap; slows onboarding

Functional equivalence (can the service do the job?) exists at the IaaS layers. **Operational equivalence** (does it behave with the same reliability, scalability and tooling depth in production?) is partial; most gaps emerge under demanding scale conditions or specialized requirements. **Ecosystem maturity** (third-party integrations, partner services, talent market) favours hyperscalers substantially.

5.4 Reference pricing comparison

Cost conversations at the architectural level require concrete numbers. The tables below consolidate published reference pricing for comparable workloads as of April 2026, split into hyperscalers and European providers for readability. All figures are on-demand list prices and should be treated as directional, because enterprise agreements, sustained-use discounts and reserved capacity routinely reduce these numbers by 30–60%.

Hyperscalers

Category	AWS	Azure	GCP
Compute (4 vCPU / 16 GB, monthly, on-demand)	$120–140 (t3.xl)	$140 (D4s_v5)	$98 (e2-std-4)
Object storage (per GB/month, hot tier)	$0.023	$0.018	$0.020
Egress (per GB, outbound internet)	$0.09 (first 100 GB/mo free)	$0.087 (first 5 GB free)	$0.08–0.12 (first 1 GB free)
Managed Kubernetes (3-node basic, monthly)	EKS: $73 control plane + nodes	AKS: free control plane + nodes	GKE Autopilot: $73 + nodes

European providers

Category	OVHcloud	Scaleway	STACKIT	IONOS	Hetzner
Compute (4 vCPU / 16 GB, monthly, on-demand)	€67–73 (b3-16)	€80–95	€75–100 (Azure-parity)	$35 (Cube L, 8 vCPU)	€30
Object storage (per GB/month, hot tier)	~€0.010–0.012	€0.012	Similar to Scaleway	$0.015	€0.005
Egress (per GB, outbound internet)	Bundled / unlimited on dedicated	€0.01	Comparable to Scaleway	€0.050	20 TB/month free
Managed Kubernetes (3-node basic, monthly)	Managed K8s included	Managed K8s	Managed K8s	Included	Not managed

Sources. GetDeploying AWS vs OVHcloud[58], VPSBenchmarks[59], TECHSY 2026 comparison[60], IONOS Cloud pricing[61], EU Cloud Cost[62], Finout storage comparison[63].

Three patterns emerge from the data:

- **Compute pricing** is roughly comparable, because European providers typically sit 20–40% below on-demand hyperscaler pricing, but hyperscaler reserved/Savings Plan pricing closes much of that gap. The compute layer is not where the sovereignty decision should primarily be economically justified.
- **Object storage** is cheaper on European providers, particularly Hetzner (€0.005/GB is 4–5× cheaper than AWS S3). For bulk storage workloads (archives, backups, media), the cost differential compounds materially at scale.
- **Egress** is the dramatic differentiator. Hyperscaler egress fees at $0.08–0.12/GB are a well-documented vendor lock-in mechanism. European providers either include egress in compute pricing (OVHcloud's dedicated servers have unlimited egress), provide generous free tiers (Hetzner 20 TB/month) or price an order of magnitude below hyperscalers (Scaleway €0.01/GB). For egress-heavy workloads, European providers can be 5–10× cheaper in total cost of ownership.

The pricing data does not support "European providers are always cheaper" narratives, nor the reverse. It supports a more nuanced conclusion. For compute-intensive, low-egress workloads with high utilisation, hyperscalers' enterprise discounts may deliver comparable or slightly better pricing. For storage-heavy, egress-intensive or bandwidth-sensitive workloads, European providers can deliver 40–70% cost reductions. Sovereign migration should therefore begin with workload profiling, that is, identifying which workloads are storage-heavy, which are egress-heavy and which are compute-bound, before any vendor selection is made.

Notes. 44 SoftwareSeni (providers) on lack of European AI/ML platform equivalent · 45 SoftwareSeni (migration) on 71% vendor lock-in deterrent · 46 SoftwareSeni (migration) on $0.08–0.12/GB egress · 47–57 SoftwareSeni (providers) on layer-by-layer parity assessments · 58 GetDeploying AWS vs OVHcloud · 59 VPSBenchmarks · 60 TECHSY 2026 comparison · 61 IONOS Cloud pricing · 62 EU Cloud Cost · 63 Finout storage comparison

6. Migration archetypes

Migration is not a single activity; it is a family of transformations with dramatically different cost, risk and timeline profiles. The archetype chosen determines whether the sovereign objective is achievable within the organization's capacity to absorb disruption.

Archetype	Description	Sovereign context	Risk level	Timeline
Lift-and-Shift	Rehost workloads with no architectural change; VMs to VMs	Works only for IaaS-native workloads; does not resolve proprietary service dependencies	Low–Medium	2–6 months
Replatform	Move to equivalent managed services (e.g., RDS → OVHcloud DBaaS) with minor changes	Resolves database lock-in; application code minimally changed	Medium	4–9 months
Refactor	Rearchitect applications to remove hyperscaler-native service dependencies	Necessary for Lambda/SageMaker/Cosmos DB; highest ROI long-term	High	9–18 months
Rebuild	Rewrite applications from scratch on cloud-native European stack	Appropriate for legacy systems already end-of-life	Very High	18–36 months

The critical strategic insight that most migration business cases miss is that **migration cost is transformation cost, not infrastructure cost**. The new provider's compute and storage bill is a small fraction of the total. The cost is in professional services consulting (≈30% of total migration budget), application refactoring (≈22%), parallel operation and testing (≈18%) and data transfer/egress fees (≈12%)[64]. Infrastructure setup and compliance certification together account for roughly 10% of total cost, the tail of the budget, not the head.

A typical five-phase sovereign migration (assessment, architecture, development/testing, staged migration, stabilization) spans 12–18 months for organizations of meaningful scale[65]. Organizations that compress this timeline to meet a compliance deadline without the parallel investment in refactoring and change management are accepting technical debt that resurfaces as operational incidents post-migration.

Archetype selection guide

The following matrix maps the two primary decision variables (proprietary service dependency depth and regulatory timeline pressure) to the recommended archetype:

Dependency depth ↓ Timeline →	> 18 months	6–18 months	< 6 months
Low (IaaS-native, containerized, OSS DB)	Lift-and-Shift to European provider; lowest cost, fastest path	Lift-and-Shift with parallel compliance certification	Lift-and-Shift; emergency migration feasible
Medium (managed DB, some PaaS, standard IAM)	Replatform; swap managed services to European equivalents	Replatform priority workloads; defer low-risk to phase 2	Sovereignty Overlay on hyperscaler; Replatform in background
High (serverless, proprietary DB, AI/ML)	Refactor; begin now, plan 12–18 month program	Refactor priority workloads + Overlay for immediate compliance	Sovereignty Overlay only; Refactor deferred (timeline does not permit)

Reading the matrix. Organizations in the bottom-right cell face the most constrained situation (deep proprietary dependency and an imminent compliance deadline). The only realistic short-term response is a Sovereignty Overlay (Section 10.3), with a parallel refactor program to achieve structural sovereignty over 12–24 months. Organizations in the top-left cell have the most favourable position and can achieve full sovereignty through direct migration.

Archetype selection should also be informed by a third factor, namely the organization's available transformation capacity. Refactor without a strong internal engineering team or a qualified systems integrator is a delivery risk, not just a cost risk. Organizations that lack this capacity should select a less ambitious archetype and invest in capability building before attempting higher-complexity migration patterns.

Notes. 64 SoftwareSeni (migration) on cost-category distribution · 65 SoftwareSeni (migration) on 12–18 month timeline

7. Hidden costs of sovereign migration

The most consistent failure mode in sovereign migration programs is an underspecified budget that models infrastructure costs but not transformation costs.

Why organizations systematically underestimate costs

The pattern is not accidental. Three structural biases produce cost underestimation in virtually every sovereign migration business case:

Infrastructure bias. Cloud migration is perceived as an infrastructure change, so the business case is built around infrastructure pricing (compute, storage, networking). But sovereign migration is primarily a transformation project. Professional services, refactoring, parallel operation and change management account for approximately 90% of total migration cost[66]; infrastructure setup accounts for roughly 6%. Organizations that model migration cost using cloud provider pricing calculators are modelling the tail of the budget, not the head.

Political pressure. Sovereignty migration programs are frequently initiated under regulatory deadline pressure or geopolitical urgency. The business case is constructed to secure approval, not to model reality. Costs that threaten timeline approval (refactoring complexity, parallel operation duration, training investment) are deferred to "phase 2" budgets that often never materialise. The result is a migration program that is funded to begin but not funded to complete.

Vendor narratives. Both hyperscaler sovereign partitions and European providers have commercial incentives to present migration as simpler than it is. Hyperscalers emphasise the continuity of their sovereign offerings ("same services, same APIs, sovereign controls") while understating the 15–30% ongoing sovereignty premium[67]. European providers emphasise infrastructure cost competitiveness while understating the ecosystem maturity gap and the operational investment required to compensate for missing managed services. Neither narrative is dishonest; both are incomplete.

Cost decomposition

The decomposition below reflects the actual cost distribution that emerges in post-migration analyses.

Figure 3 - Migration cost breakdown by category (typical distribution)

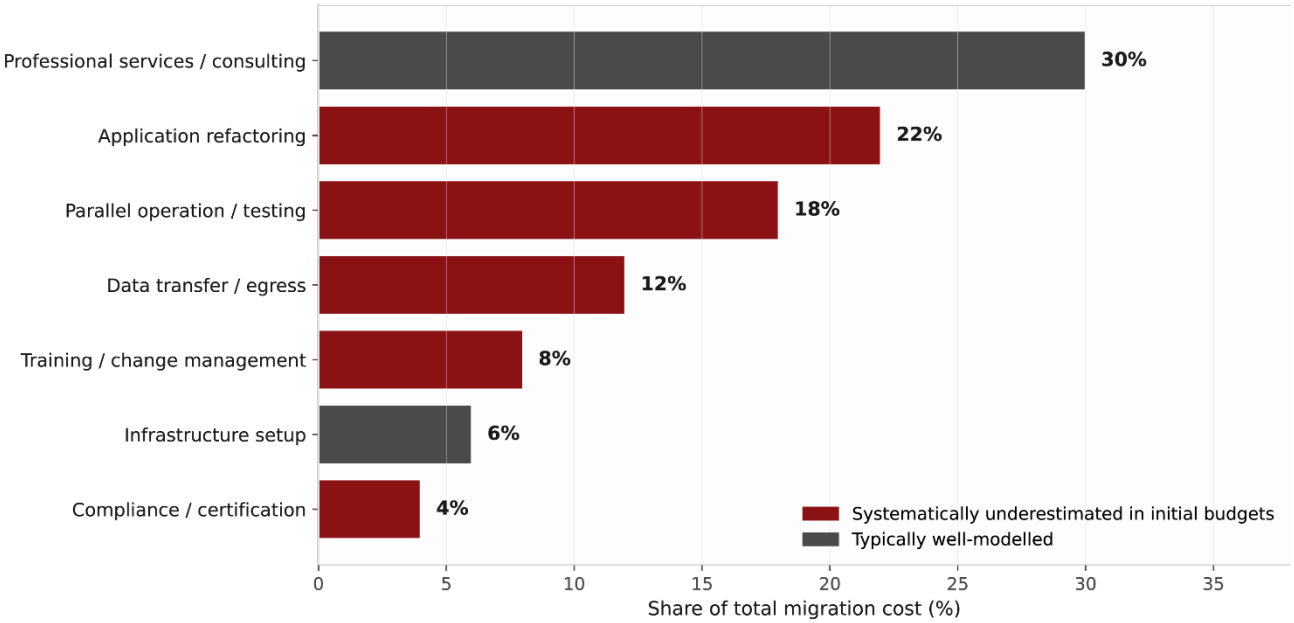

Source: SoftwareSeni migration analysis; BCG via Civo; RefleXio synthesis of post-migration budget reviews.

Figure 3. Migration cost breakdown by category. Red bars mark the five categories that are systematically underestimated in initial business cases.

Cost category	Description	Typical range (% of total)	Often overlooked?
Professional services / consulting	Architecture, program management, integration services, compliance advisory	≈30%	Partially (often underscoped)
Application refactoring	Code changes to remove proprietary service dependencies	≈22%	Yes (frequently not modeled)
Parallel operation / testing	Running source and target environments simultaneously during migration	≈18%	Yes (treated as temporary but extended)
Data transfer / egress	Exit fees from source cloud at $0.08–0.12/GB	≈12%	Yes (often discovered mid-program)
Training / change management	Developer reskilling, operational runbook rebuilding, organizational change	≈8%	Yes (almost always underfunded)
Infrastructure setup	New environment build, networking, security baseline configuration	≈6%	No (typically well-modeled)
Compliance / certification	Achieving required certifications on new provider (SecNumCloud, ISO 27001, BSI C5)	≈4%	Yes (timelines underestimated)

Five of the seven cost categories are systematically underestimated or not modeled at all in initial business cases. The Bank of England's migration to Oracle Cloud (budgeted at £7 million and

RefleXio
www.reflexio.es
press@reflexio.es

RefleXio Intelligence

Independent research on compute, infrastructure and sovereign AI