

Sovereignty by Design

Engineering Provable Control in the Era of EU Data Sovereignty,
Vendor Neutrality, and Cloud Exit Readiness

A living advisory framework for financial services, payments, healthcare, telecommunications, manufacturing, and public sector organizations navigating DORA, the EU Data Act, the EU AI Act, Schrems II, and the emerging Cloud and AI Development Act (CADA).

Authors

Krishna Battula, Sunny Sharma, Mayank Gupta

Contents

Executive Summary

1. Opening Thesis
2. Why Sovereignty Matters Now
3. The Sovereignty Problem: Six Dimensions
4. Where Hyperscaler Sovereign Offerings Fall Short
5. The EU Is Not One Jurisdiction: Sub-Regional Regimes
6. From Cloud-First to Sovereignty-First: The Mosaic Architecture
7. Vendor Neutrality: Lock-In Taxonomy and Anti-Patterns
8. Infrastructure Decision Framework: Bare Metal to OpenShift
9. Datacenter and Cloud Exit Readiness
10. Cryptographic Sovereignty: The Customer-Held-Key Differentiator
11. Case in Practice: A Global Payments Network
12. Frequently Asked Questions: Making Informed Decisions Under Mandate
13. AI-Accelerated Cybersecurity and Infrastructure Risk
14. Modern Data and Metadata Governance
15. Governance Operating Model
16. The Mphasis Sovereignty Manifesto
17. Conclusion

Key Takeaways



Sovereignty is an engineering discipline, not a checkbox - True sovereignty means proving control over data residency, jurisdiction, operations, cryptography, supply chain, and exit readiness under regulatory or operational stress, not just passing a procurement audit.



No single cloud solves everything - Hyperscaler sovereign tiers solve data residency and operational control well, but they cannot deliver full jurisdictional immunity due to their US corporate parentage and exposure to the US CLOUD Act



Adopt a Mosaic Architecture - Winning enterprise architectures blend multiple "lanes"-using hyperscalers for elastic AI/analytics, EU-native providers for jurisdictional immunity, and private/hybrid infrastructure anchored by customer-held cryptographic keys for crown-jewel workloads.



Customer-held keys are the ultimate control - Externalizing cryptographic layers to independently validated Hardware Security Modules (HSMs) outside both the hyperscaler's and the organization's primary perimeters is the most defensible way to answer Schrems II, DORA, and CADA mandates.




Exit readiness must be tested, not documented - Under DORA and the EU Data Act, cloud exit is a design requirement; organizations must run real exit rehearsals, testing data portability, identity mapping, and key revocation, before a forced migration occurs.

Executive Summary

Sovereignty is no longer a procurement checkbox. It is an engineering discipline. Across financial services, healthcare, telecommunications, manufacturing, and public sector, regulators have converged on a single demand: the organization must prove control of its data, its keys, its operations, its AI, and its ability to leave. This must hold under stress, not just on paper.

This paper sets out Mphasis' position on that demand. It uses the Sovereignty by Design thesis as its foundation, extends it with field evidence from a live global payments-sector engagement, and generalizes both into a repeatable advisory framework, the Mosaic Architecture, that any regulated enterprise can apply.



Our central finding, validated against public regulatory sources and a live cross-industry case: no single cloud, vendor, or platform satisfies every sovereignty tier. Sovereign hyperscaler tiers solve data residency and day-to-day operations well. EU-native, nationally-qualified providers solve jurisdictional immunity. Private and hybrid infrastructure, anchored by customer-held cryptographic keys, solves the dimensions that hyperscale architecture cannot solve by construction. The winning enterprise architecture blends all three, governed by shared controls for identity, cryptography, metadata, and exit.

“Sovereignty is the ability to prove control under stress, during a regulator review, a cyber incident, a provider outage, a geopolitical disruption, or a forced migration.”


1. Opening Thesis

The next era of enterprise technology will not be defined by cloud adoption. It will be defined by an organization's ability to prove control over data, AI, infrastructure, cryptographic keys, suppliers, and operations, under regulatory, cyber, geopolitical, and operational stress.

DORA applies from 17 January 2025 and covers Information and Communication Technology (ICT) risk management, third-party risk, resilience testing, incident reporting, and oversight of critical ICT providers [1]. The EU Data Act applies from 12 September 2025 and mandates cloud and edge switching, interoperability, and safeguards against unlawful third-country government access, with switching charges, including data egress, eliminated from 12 January 2027 [2].

The EU AI Act (Regulation (EU) 2024/1689) entered into force in August 2024 and is moving AI governance into enforceable operating obligations, with the timeline for high-risk systems recently adjusted [3][22]. Schrems II requires exporters to assess third-country law and adopt supplementary safeguards, or suspend transfers, where equivalent protection cannot be ensured [4]. And the proposed Cloud and AI Development Act (CADA), published 3 June 2026, introduces a four-level EU cloud and AI sovereignty assurance framework tied to public-sector procurement [5].

Mphasis' position: sovereignty is not a product; it is an architecture and an operating model. It is proven when an enterprise can show where data and metadata live, who can access them, who controls the keys, how supplier dependencies are monitored, how AI lineage is evidenced, and how a critical service can move or be substituted without unacceptable disruption. This is a reusable industry thesis for financial services and payments, health-care, public sector, telecommunications, and manufacturing alike.



2. Why Sovereignty Matters Now

Cloud has become a regulated control surface. DORA recognizes that financial entities increasingly rely on ICT third parties, and that unmanaged ICT risk can disrupt financial services across borders [1]. The EU Data Act moves beyond data access into cloud-market structure, requiring cloud and edge providers to meet minimum switching and interoperability requirements so that switching becomes “free, fast and fluid” [2].

For a global payment network, the board-level question is no longer “Is EU data stored in Europe?”

It is: Can the organization prove legal control, operational control, cryptographic control, AI/data lineage, supplier substitutability, and exit readiness for its real-time payment and fraud workloads?

The same logic extends to patient records in healthcare, subscriber metadata in telecom, citizen services in public sector, and industrial telemetry in manufacturing. The Data Act explicitly reaches connected products, industrial data, and third-country government access to non-personal data [2].

2.1 EU Regulatory Timeline and Convergence

Date	Milestone	What changed
25 May 2018	GDPR applies	EU-wide personal data rights, accountability, and data portability.
16 Jul 2020	Schrems II ruling	Third-country legal access becomes a transfer-risk and architecture issue, not just a contractual one.
17 Jan 2025	DORA applies	ICT resilience, third-party risk, resilience testing, incident reporting, and critical-provider oversight become operational obligations.
12 Sep 2025	EU Data Act applies	Cloud switching, interoperability, and safeguards against unlawful third-country access become enforceable; switching charges phase out by 12 Jan 2027.
1 Aug 2024 → phased	EU AI Act phases in	Prohibited practices (Feb 2025) and GPAI transparency (Aug 2025) are already in force; the Annex III high-risk deadline, originally 2 Aug 2026, was deferred to 2 Dec 2027 under the Digital Omnibus (Regulation (EU) 2026/1744, adopted 27 Jul 2026).
3 Jun 2026	CADA proposed	A four-level EU cloud and AI sovereignty assurance framework enters the legislative process, tied to public-sector procurement.

Sources: EUR-Lex Regulation (EU) 2022/2554, 2023/2854, 2024/1689; EDPB Schrems II Recommendations 01/2020; European Commission CADA proposal, 3 June 2026. Full citations in References.

EU Sovereign Cloud & Financial Services Regulatory Timeline (2020 - 2027)




Figure 1. EU Sovereign Cloud and Financial Services Regulatory Timeline (2020–2027) — data-protection, resilience, switching, and sovereignty-framework milestones arrive on separate tracks but converge on the same architecture decision

A note on accuracy, current as of this paper's publication: the EU AI Act's Annex III high-risk deadline was pushed from 2 August 2026 to 2 December 2027 under the EU's Digital Omnibus simplification package, formally adopted as Regulation (EU) 2026/1744 on 27 July 2026, six days before the original deadline [22][23]. Transparency obligations for deployers, and the general-purpose AI model rules that took effect on 2 August 2025, remain on schedule and unaffected [23][24]. This is a live illustration of the thesis itself: regulatory timelines move, and an architecture built for provable control absorbs that movement without a redesign; an architecture built to a single compliance date does not.

3. The Sovereignty Problem: Six Dimensions

Sovereignty is not a single control. It is a composite of six dimensions, and hyperscaler, EU-native, and private/hybrid infrastructure each solve a different subset by construction - not by configuration.

Dimension	The question	Practical implication
Residency	Where are bytes stored and processed?	Necessary but not sufficient. Region choice alone says nothing about who can access data.
Jurisdiction	Whose law can compel access?	Schrems II requires assessment of third-country law and practice, not contracts alone.
Operations	Who can administer, support, or break glass?	Sovereign-cloud claims must be tested against support access paths, admin routes, and audit logs.
Cryptography	Who can decrypt?	The EDPB recognizes encryption, with keys held under exporter/EEA-equivalent control, as an effective supplementary measure in defined scenarios.
Supply chain	Who builds and updates the stack?	Regulators and NIST guidance both flag malicious functionality, counterfeit components, and reduced visibility as systemic risks.
Exit	Can the organization move or substitute the service?	DORA and the EU Data Act turn exit and switching into design requirements, not contractual afterthoughts.

Framework adapted and extended by Mphasis from public EU regulatory sources; see References [1]–[4], [9].

Classify sovereignty requirements before choosing a platform. A low-risk innovation sandbox may tolerate managed-cloud dependencies. A real-time payment authorization path, a clinical AI pipeline, a citizen registry, or a network-control workload may require jurisdictional, operational, cryptographic, and exit controls that no single hyperscaler region can deliver by construction.

4. Where Hyperscaler Sovereign Offerings Fall Short

The three major hyperscalers each have built genuinely useful sovereign features and none of them, by construction, delivers full EU jurisdictional sovereignty. Each solves residency and operations well; each remains structurally reachable under the US CLOUD Act and FISA §702 because its ultimate parent is US-incorporated.

4.1 AWS European Sovereign Cloud

Launched 15 January 2026, AWS ESC is a physically and logically separate environment with its first region in Brandenburg, Germany, an independent control plane, IAM, DNS, billing, and monitoring stack, and EU-resident-only day-to-day operations [6][7]. AWS has published an independently validated Sovereign Reference Framework aligning controls to governance independence, operational control, data residency, and technical isolation [8]. The ultimate parent, however, remains Amazon.com, Inc. (US), which keeps AWS ESC reachable under the CLOUD Act regardless of regional configuration.

4.2 Microsoft Sovereign Cloud

Microsoft's Sovereign Public Cloud adds Data Guardian requiring regional approval and tamper-evident logging for Microsoft personnel remote access - External Key Management for customer-controlled encryption, and Regulated Environment Management for policy-as-code guardrails, layered on the existing EU Data Boundary [9][10][11]. National partner clouds such as Bleu (France, with Orange and Capgemini) and Delos (Germany, with SAP) run Azure-derived stacks under separate EU legal entities. The gap remains the same as AWS's: the underlying technology and, in Microsoft's own public testimony to the French Senate in June 2025, the ability to

4.3 Google Cloud Sovereign Controls

Google Cloud's Sovereign Controls by Partners model routes operational and cryptographic control through EU domiciled partners - S3NS (majority Thales-owned) in France, T-Systems in Germany, and others, so that decrypting customer data requires a call outside Google's infrastructure to an externally managed key [12][13]. This partner-operated model is the strongest jurisdictional posture of the three hyperscalers precisely because it moves the operating entity outside the US corporate structure for the sovereign tier, however, the underlying software supply chain remains hyperscaler-originated.

4.4 Sovereign Features vs. Jurisdictional Sovereignty

Capability	All three Hyperscalers	EU-native / partner-operated tier
Data residency in the EU	Yes, by design	Yes
EU-resident operations staff	Yes, in the sovereign tier	Yes
Customer-held encryption keys	Available via BYOK/HYOK/EKM	Native to the model
Immunity from CLOUD Act / FISA §702	Structural gap - US parent entity	Achievable when the operating entity is genuinely EU-owned
Highest EU sovereignty-tier eligibility	Constrained by ownership	The realistic path to the top assurance tiers under emerging frameworks such as CADA

The main takeaway: Hyperscalers have built genuinely useful sovereign features. They have not and by corporate structure cannot, deliver full jurisdictional sovereignty on their own. That gap is not fixable with more engineering inside the hyperscaler. It is a company-ownership gap, and it is why emerging EU frameworks gate their highest assurance tiers on who owns the company and what law governs its staff.

5. The EU Is Not One Jurisdiction: Sub-Regional Regimes

DORA, the EU Data Act, GDPR, and the emerging CADA framework set a common EU-wide floor. Individual member states, plus Switzerland, apply materially stricter regimes on top. For any organization operating across borders, this changes workload placement decisions county by county, not just region by region.

5.1 Germany - BSI C5

The Cloud Computing Compliance Criteria Catalogue (C5), issued by Germany's Federal Office for Information Security (BSI), is the de facto standard for cloud security in Germany and is mandated in practice by German banking supervisors and public-sector procurement [14]. Its most recent iteration requires independent Type 2 audits over a sustained operational period and adds explicit physical-resilience, patch-SLA, and supply-chain controls that hyperscalers have historically found demanding to meet.

5.2 France - ANSSI SecNumCloud

SecNumCloud, issued by ANSSI, is the only EU national certification scheme that explicitly requires immunity to extraterritorial law, the reference document requires qualified providers to be able to resist an injunction issued under extraterritorial law that would compromise the availability, confidentiality, or integrity of the service [15]. This is the closest existing national precedent for CADA's proposed highest assurance tiers, and no direct hyperscaler offering currently holds the qualification.

5.3 Switzerland - FINMA, revFADP, and Swiss Banking Guidelines

Though outside the EU, Switzerland is relevant to any organization with Swiss financial-sector counterparties. FINMA Circular 2018/3 requires the outsourcing institution to retain full audit rights and a tested exit strategy before outsourcing begins; the revised Federal Act on Data Protection restricts the cloud provider to processing data only as instructed by the exporting entity. In practice, customer-held encryption keys on Swiss soil are the most common way Swiss institutions resolve this.

Jurisdiction	Effective floor	Effective ceiling
EU-wide (GDPR / DORA / Data Act)	Sovereign hyperscaler tier	Baseline EU sovereignty assurance
Germany (BSI C5)	Sovereign hyperscaler tier + C5 attestation	Mid-tier assurance via nationally partnered providers
France (SecNumCloud)	SecNumCloud-qualified provider	The strongest currently-qualified national tier, achievable by EU-native providers
Switzerland (FINMA / revFADP)	Customer-held keys or Swiss-domiciled provider	Private/hybrid infrastructure with Swiss sovereign providers

Sources: BSI Cloud Computing Compliance Criteria Catalogue (C5); ANSSI SecNumCloud qualification referential, cyber.gouv.fr; FINMA Circular 2018/3. Full citations in References [14],

6. From Cloud-First to Sovereignty-First: The Mosaic Architecture

The “cloud-first” era optimized for speed and elasticity. The “multi-cloud” era optimized for optionality, often without true portability. The emerging “sovereignty-first” era optimizes for provable control, cryptography, metadata, and exit.

Mphasis' recommendation is not to select a single sovereign provider. It is to design a Mosaic Architecture: place each workload in the lane that matches its sovereignty tier and govern every lane with the same shared controls for identity, cryptography, metadata, and exit readiness.

The Four Lanes of the Mosaic

Sovereignty is no longer a procurement checkbox. It is an engineering discipline. Across financial services, healthcare, telecommunications, manufacturing, and public sector, regulators have converged on a single demand: the organization must prove control of its data, its keys, its operations, its AI, and its ability to leave. This must hold under stress, not just on paper.

Lane	Best fit	What it delivers	What it cannot deliver alone
Hyperscaler innovation lane	AI/ML training, elastic analytics, non-sensitive workloads	Scale, GPU-backed inference, fastest time-to-market	Full jurisdictional and cryptographic immunity
Sovereign cloud lane	Regulated workloads needing cloud speed	Residency, EU-resident operations, assurance frameworks	Escape from parent-company jurisdiction
EU-native / partner operated lane	Workloads requiring extraterritorial-law immunity	SecNumCloud / C5-qualified, genuinely EU-owned operating entities	Hyperscaler-grade AI/ML service breadth at present
Private / hybrid / bare metal lane	Cryptographic core, PCI HSM key custody, low latency systems	Maximum physical and legal control	Elasticity and fast time-to scale




Figure 2. Mphasis Mosaic Architecture - workload classification and lane selection sit above shared controls for identity, policy-as-code, customer-held keys/HSM, metadata and lineage, and observability and evidence, which govern all four lanes uniformly.

An enterprise does not choose a cloud - it designs across clouds.

7. Vendor Neutrality: Lock-In Taxonomy and Anti-Patterns

Vendor lock-in is rarely a single event. It accumulates across five layers, and each layer has its own anti-pattern and its own mitigation

Lock-in type	What it looks like	Anti-pattern	Mphasis mitigation
Infrastructure	Cloud networking, IAM, block storage, load balancers	“Containers make exit easy.”	Provider adapters, portable infrastructure-as-code, storage abstraction, tested network mapping
Platform	Managed databases, queues, serverless, proprietary monitoring	Critical path embedded in provider-native PaaS	Use managed services by risk tier; require an exit map for every critical workload
Data	Proprietary formats, closed catalogs, egress economics	Data exported without metadata, lineage, or meaning	Open table formats, exportable catalogs, data contracts, lineage standards
AI	Model APIs, vector stores, feature stores, closed evaluation logs	Fraud, clinical, or citizen service AI tied to an opaque provider stack	Model registry, feature portability, prompt/inference logs, lineage evidence
Operational	Skills, runbooks, tooling, support model	People who can only operate one provider	Cross-platform SRE playbooks, exit drills, supplier substitution training

Kubernetes conformance helps because CNCF-certified Kubernetes supports the same required APIs across installations, enabling interoperability [16]. Conformance does not remove provider-specific identity, storage, networking, key-management, or observability dependencies.

Kubernetes conformance helps because CNCF-certified Kubernetes supports the same required APIs across installations, enabling interoperability [16]. Conformance does not remove provider-specific identity, storage, networking, key-management, or observability dependencies.

8. Infrastructure Decision Framework: Bare Metal to OpenShift

Portability is engineered at the infrastructure layer, not conjured by the orchestration layer. A recurring question in sovereignty engagements concerns managed Kubernetes: if an organization already runs a managed control plane such as AWS EKS or Azure AKS, why does the underlying infrastructure still matter?

Mphasis View

Managed Kubernetes simplifies operations; it does not change where, or under whose control, workloads run. Even with a managed control plane, the platform still depends on the provider's compute, networking, storage, and identity primitives. Real architectural control and real exit feasibility, remains at the infrastructure layer.

Option	Best fit	Strength	Caution
Bare metal	Crown-jewel data, HSM roots, low-latency systems	Maximum physical control, 100% performance baseline	Slow elasticity, highest operations burden
Virtual machines	Legacy modernization, private cloud	Mature isolation and tooling, ~85–98% of baseline performance	Hypervisor and licensing dependency
Kubernetes on bare metal	High-control cloud-native platforms	Open K8s APIs, ~90–98% of baseline performance, high anti-lock-in	Requires strong platform-engineering / SRE skill
Kubernetes on VMs	Enterprise hybrid platforms; the common cloud pattern	K8s portability plus VM flexibility; managed services (EKS/AKS) offload maintenance	Dependency shifts to the cloud's native constructs (VPC, IAM, storage)
OpenShift on bare metal	Regulated enterprise Kubernetes	Consistent hybrid foundation, built-in CI/CD, security, and compliance tooling	Platform cost and lifecycle governance
OpenShift on VMs	Multi-cloud standardization	“Write once, run anywhere” across on-prem and any cloud IaaS	Layered cost and complexity

Figure 3. Performance across all six models clusters between 78% and 100% of a bare-metal baseline once tuned — the real differentiators are portability, sovereignty, cost structure, and operational complexity, not raw throughput.

8.1 Microsoft Sovereign Cloud

Mphasis deliberately does not reduce this decision to a single ranked score. A universal winner would misrepresent the choice: the right infrastructure depends on the sovereignty tier, cost model, and operating maturity of each specific workload. The framework below rates each option qualitatively across six dimensions - including, distinctly, its CapEx profile and its OpEx profile, since capital and operating cost behave very differently across these six models and a single blended “cost” figure hides that difference.

Dimension	Bare Metal	VMs	K8s on Bare Metal	K8s on VMs	OpenShift Bare Metal	OpenShift on VMs
Sovereignty alignment	Maximum	High (if EU-hosted)	High	High	High	High
Portability / anti-lock-in	High, but physical only	Moderate (hypervisor dependency)	High (open K8s APIs)	High	High (minor Red Hat dependency)	High (runs on any IaaS)
CapEx profile	High - upfront hardware and colocation build out	Medium - hardware plus hypervisor licensing	High - hardware, no managed service offset	Low - consumption based, no hardware purchase	High - hardware plus platform subscription	Low - cloud consumption, no hardware
OpEx profile	Low - power, cooling, and maintenance only	Medium - licensing renewals and admin staff	Low-to-medium - lean once automated	Medium-to-high - pay-as-you-go compute plus management fees	Medium - platform support offsets operations headcount	High - subscription plus cloud spend layered together
Operational complexity	Highest	Moderate	High - requires SRE-level skill	Medium	Medium-to-high	High
Elasticity / speed-to-scale	Low	Medium	Medium	High	Medium	High

Ratings are qualitative and directional, drawn from a live cross-industry engagement; they are intended as a starting framework for an organization's own architecture review board to re-weight against its specific workloads, contracts, and regulatory tier — not as a definitive score.

8.2 Mphasis Advisory Recommendation

Rather than naming a single preferred platform, Mphasis advises that the organization's architecture review board apply the framework above per workload, guided by the following patterns observed across engagements:

- **1** **Cryptographic core and HSM-anchored workloads**
bare metal or Kubernetes on bare metal is generally justified despite the higher CapEx and operational burden, because sovereignty alignment and anti-lock-in dominate the decision for this workload class.
- **2** **Real-time, regulated transactional workloads**
OpenShift on bare metal offers a defensible balance: strong sovereignty alignment and portability, with built-in compliance tooling that reduces the operational burden bare metal alone would carry.
- **3** **Elastic analytics, model training, and non-sensitive data workloads**
Kubernetes on VMs or OpenShift on VMs is typically the more efficient choice, trading a higher OpEx run-rate for materially faster time-to-scale and lower upfront capital commitment.
- **4** **Legacy modernization and private-cloud consolidation**
virtual machines remain a pragmatic intermediate step, particularly where mature hypervisor tooling and existing licensing already exist within the estate.
- **5** **Multi-business-unit standardization**
OpenShift's consistent hybrid operating model, across both bare metal and VMs, is advisable where the priority is a single operating pattern across many workloads rather than the lowest cost for any one workload.

The advisory principle underneath every one of these patterns is the same: infrastructure selection should follow the sovereignty tier and lifecycle economics of the workload, and should be re-examined at each contract renewal, each DORA resilience-testing cycle, and each CADA assurance-level milestone - not fixed once and left unreviewed.

9. Datacenter and Cloud Exit Readiness

Exit readiness is an engineering discipline, not a contract clause. DORA requires that the ability to exit an ICT third party be demonstrated, not documented [1]. The EU Data Act requires cloud and edge providers to export data in commonly used, machine-readable formats, support interoperability, and eliminate switching charges - including egress - from 12 January 2027 [2].

Exit capability	Level 1 - Documented	Level 3 - Tested	Level 5 - Industrialized
Dependency inventory	Static spreadsheet	CMDB + service maps	Live dependency graph tied to telemetry and contracts
Data portability	Manual export	Restore tested	Automated export/import with validation
Configuration portability	App-specific scripts	Shared infrastructure-as code	Provider-neutral deployment factory
Identity portability	Manual mapping	Federated identity	Tested identity substitution
Key control	Provider-managed KMS	Bring-your-own-key (BYOK)	External HSM / hold-your own-key (HYOK) with revocation drills
Evidence	Tabletop record	Technical rehearsal	Regulator-ready evidence pack

Figure 4. An exit-ready hybrid architecture routes primary-provider workloads through a governed "exit factory" - container images, IaC, data and metadata export, identity mapping, key migration, runbooks, and an evidence pack - into either an alternate sovereign landing zone or a private/on-prem repatriation lane.

Mphasis' recommendation: treat every critical-function workload as requiring an exit rehearsal, not just an exit clause. A completed technical rehearsal - restored data, moved configuration, mapped identity, migrated keys, and a documented evidence trail - is the strongest exit control an organization can hold, and it is precisely the evidence a DORA or CADA-linked audit will ask for.

10. Cryptographic Sovereignty: The Customer-Held-Key Differentiator

Across every dimension of sovereignty, one control is repeatedly the most defensible and the most portable: customer-held cryptographic keys. The EDPB's Schrems II guidance recognizes encryption, with keys retained under exporter or EEA-equivalent control, as an effective supplementary measure in defined storage and backup scenarios [4].

	Provider-held keys (default cloud KMS)	Customer-held keys (external KMS + HSM)
Who controls the HSM	Cloud provider	Customer
Compelled disclosure risk	Provider can be compelled to assist	Provider physically cannot decrypt
Standard used	Standard cloud KMS only	Externally validated HSM standards (e.g. PCI PTS HSM, PCI PIN Security, for payments)

	Provider-held keys (default cloud KMS)	Customer-held keys (external KMS + HSM)
Exit implication	Exit is theoretical while keys remain with the provider	Revoking keys is a real, tested exit
Regulatory exposure	Schrems II / CLOUD Act exposure remains open	Immunity is architectural, not contractual

For any regulated data category - payment credentials, patient records, citizen identifiers, and industrial IP cloud native key management alone should not be treated as sufficient. Externalizing the cryptographic layer to a customer owned, independently validated HSM cluster is the single most reusable control across every vertical this paper addresses.

10.1 Where the Vault Itself Should Sit

Externalizing the key is necessary but not, on its own, sufficient - where the vault physically and administratively sits matters just as much as who legally owns it. Mphasis' advisory position on HSM placement is deliberately stricter than “bring your own key” alone.

The KMS/HSM vault should be an outside, independently governed solution under the organization's own control - never provisioned inside the Hyperscaler's perimeter, and not simply relocated inside the organization's own primary data-center perimeter either.


The first half of that principle is the one every sovereignty program already accepts: a key vault provisioned inside a hyperscaler's own perimeter - even a sovereign-tier perimeter - is still administratively reachable through that provider's operational stack, which is precisely the exposure Section 4 describes as structural rather than fixable. The second half is the one programs miss. Moving the HSM cluster into the organization's own primary data center does not, by itself, externalize the key -it only moves the single point of failure from the provider's perimeter to the organization's own. If that data center is the same facility running the workload, sharing the same network segmentation, the same physical access list, and the same disaster-recovery dependency, then a compromise, seizure, or outage affecting that facility compromises the compute and the cryptographic control together, at the same time, through the same failure path. The vault has changed custodianship without changing blast radius.

The practical implication Mphasis carries into every HSM reference architecture it scopes: the vault should be deployed as a distinct, independently network-segmented facility. This can be a dedicated HSM-as-a-service colocation, a separate sovereign-certified site, or a multi-site HSM cluster with its own access governance. It should be architecturally decoupled from both the hyperscaler's operational perimeter and the organization's own primary compute footprint..

Revocation, rotation, and audit of the keys should be exercises the organization can run without needing physical or administrative access to either perimeter. Only then does “customer-held” mean what a regulator, or an incident, will test it against.

11. Case in Practice: A Global Payments Network

The framework above is not theoretical. Mphasis applied it directly in a live advisory engagement with a global card payment network operating across the EU, Germany, France, and Switzerland. The engagement asked two linked questions that recur, in different words, across every regulated vertical: where do hyperscaler sovereign offerings fall short for the organization's workloads, and is the EU actually one jurisdiction or several?



11.1 What the Engagement Found

1	None of the three hyperscaler sovereign tiers, by construction, removed exposure to extraterritorial law. This confirmed the Section 4 analysis directly against the organization's real workload map.
2	Real-time authorization and clearing, and cardholder PII/KYC data, both required either customer-held PCI-HSM key custody or an EU-native/private placement, not a sovereign hyperscaler tier alone
3	Analytics, model training, and non-PII data-lake workloads were well served by the sovereign hyperscaler tier, preserving innovation velocity where sovereignty risk was genuinely low.
4	An existing private data-center footprint in France was identified as a strategic Lane-C asset for the cryptographic core, rather than a legacy liability to be decommissioned.
5	The infrastructure decision framework, the same structure shown in Section 8.1, gave the organization's architecture review board a repeatable way to place each new workload without re-litigating the cloud-vs sovereign debate every time.

11.2 The Resulting Recommendation

The recommended architecture followed the Mosaic model exactly. Sovereign hyperscaler tiers were retained for AI-native and analytics workloads. EU-native, nationally-qualified providers carried workloads where extraterritorial-law immunity was a hard requirement. Private/hybrid infrastructure, anchored by customer-held PCI-HSM keys, carried the cryptographic core and payment-authorization flows.

Customer-held PCI-HSM key ownership was identified as the single most defensible control, answering Schrems II, the CLOUD Act, DORA/NIS2 resilience obligations, and the emerging CADA framework in one architecture decision.

11.3 Why this Generalizes Beyond Payments

The specific regulation names change by vertical, but the six-dimension test in Section 3 does not. Substitute HIPAA equivalent clinical-data rules for PCI-DSS, a national health data authority for a banking supervisor, or a telecom regulator for a payments network, and the same lane-selection logic applies. This is deliberate: the case is presented here as a template, not a one-off.

Vertical	Analogous “crown jewel” workload	Analogous regulator lens	Mosaic lane most likely required
Financial services & payments	Real-time authorization, settlement, tokenization	DORA, PCI-DSS, national payments supervisors	Private/hybrid with customer held HSM keys
Healthcare	Patient records, clinical AI pipelines	National health-data authorities, GDPR special category data rules	EU-native or private, with AI Act-aligned model lineage

Vertical	Analogous “crown jewel” workload	Analogous regulator lens	Mosaic lane most likely required
Telecommunications	Subscriber metadata, network-control systems	NIS2, national telecom regulators	EU-native / sovereign hyperscaler hybrid, tested exit
Public sector	Citizen registries, classified workloads	CADA public-procurement assurance levels, national security rules	EU-native / private, highest CADA tier where applicable
Manufacturing	Industrial IoT telemetry, digital twins, proprietary IP	EU Data Act (connected products), trade-secret protection	Sovereign hyperscaler for scale; private for IP-sensitive telemetry

Figure 5. The same six-dimension sovereignty test, applied to a global payments network, generalizes directly to the “crown-jewel” workload in every other regulated vertical.

12. Frequently Asked Questions: Making Informed Decisions Under Mandate

The questions below are the ones Mphasis hears most often once a sovereignty mandate moves from a compliance memo into an architecture decision that a board is expected to sign off. This mandate might be DORA, NIS2, the EU Data Act, an emerging CADA assurance level, or a national scheme such as BSI C5 or SecNumCloud.

They are organized the way an implementation program actually experiences them: first the mandate itself, then the hyperscaler question everyone asks first, then the sub-regional detail that surprises most programs, then the architecture and cost trade-offs, and finally the governance and procurement questions that determine whether the decision survives contact with a real budget cycle.

Each answer reflects Mphasis' advisory position, tested against public regulation and against field experience advising a global payments network through exactly this sequence of questions.

12.1 Framing the Mandate: Where to Start

“We are under a DORA, NIS2, or CADA-linked mandate. Where should the organization actually start?”

Not with a vendor conversation. The most common mistake Mphasis observes is an organization opening a sovereignty program by asking a hyperscaler account team “what can you offer us,” before it has classified its own workloads. That sequencing hands the framing of the problem to the party with the least incentive to describe its own structural limits.

The correct starting point is an internal workload census, run against the six-dimension test set out earlier in this paper: data residency, operational sovereignty, jurisdictional sovereignty, supply-chain sovereignty, cryptographic sovereignty, and exit. Every workload gets scored against all six, independent of any current or planned provider. Only once that census exists does a placement conversation with any provider become meaningful, because the organization now knows what it is buying protection against, rather than accepting a provider's own description of what it protects against.

In practice this census takes contained effort. Mphasis typically scopes a first-pass sovereignty diagnostic in the order of four weeks for a mid-sized regulated estate. It should produce two artifacts a board can actually act on: a workload-by-workload sovereignty tier classification, and a jurisdiction-by-jurisdiction placement map showing which workloads are currently placed against tier requirements they do not meet. That gap map, not a vendor's roadmap, is what should drive the next twelve months of spend.

“Our hyperscaler already markets a ‘sovereign cloud’ tier. Doesn’t that already satisfy the mandate?”

It satisfies part of the mandate, usually the easiest part. Sovereign hyperscaler tiers are real engineering investments, not marketing labels: AWS European Sovereign Cloud runs an independent control plane, IAM stack, and EU-only Security Operations Centre out of Brandenburg, Germany [6][7][8]; Microsoft’s Sovereign Public Cloud adds Data Guardian’s tamper-evident, regionally-approved remote-access model [9][10]; Google’s Sovereign Controls by Partners routes decryption through an EU-domiciled partner such as S3NS or T-Systems [12][13].

What none of the three changes is the corporate parent. AWS ESC, Sovereign Public Cloud, and Google’s direct offering are all ultimately controlled by a US-incorporated entity. This keeps each one reachable under the US CLOUD Act and FISA §702 regardless of where the servers sit or who staffs the support desk.

CADA’s Level 3 and Level 4 assurance tiers, and France’s SecNumCloud qualification, are written specifically around jurisdictional immunity. For a mandate written around that requirement, a sovereign hyperscaler tier is a strong answer to the residency and operations questions and a structurally incomplete answer to the jurisdiction question.

The mandate is rarely satisfied by one tier alone. It is satisfied by placing each workload in the tier that matches what that specific workload actually needs immunity from.

12.2 Evaluating Hyperscaler Sovereign Offerings

“What, specifically, can AWS European Sovereign Cloud, Microsoft Sovereign Cloud, and Google Sovereign Controls not fix, no matter how much each vendor invests?”

This question deserves a precise answer rather than a general one, because the specifics differ by provider and matter for procurement:

1

AWS European Sovereign Cloud cannot remove the possibility of compelled disclosure through its US parent. At launch, it also carried real practical constraints: a narrower service catalogue with no GPU or AMD instance families, a reported 15–20% price premium over eu-central-1, and no cross-region peering back into global AWS. That last constraint matters most operationally. It means workloads must be re-provisioned into a wholly separate account structure rather than migrated in place[6][7]

2

Microsoft Sovereign Cloud carries the most publicly documented version of the same limit: Microsoft France’s own director of public and legal affairs testified under oath before the French Senate in June 2025 that Microsoft cannot guarantee French citizens’ data will not be transmitted to US authorities without French consent [25]. Its national partner clouds, Bleu (with Orange and Capgemini) and Delos (with SAP), run on Microsoft-licensed technology, which is exactly why EU trade bodies have flagged resale-through-aggregator structures as a route by which a foreign-controlled stack could claim a higher assurance tier without a genuine change of ownership.

3

Google Cloud's Sovereign Controls by Partners model, operated through EU entities such as S3NS and T Systems, is arguably the strongest jurisdictional posture of the three because operational and cryptographic control sits with a genuinely EU-domiciled partner. However the underlying software supply chain remains hyperscaler-originated. That gap is why, in Mphasis' assessment, even this model is expected to reach CADA's Level 3 at best, not Level 4, without deeper re-engineering.

Each vendor has solved the parts of sovereignty that engineering and regional operations can fix. None has solved, or can unilaterally solve, the part that comes from corporate ownership and the law attached to it.

“Is this really a technology gap, or something else?”

It is a company-ownership gap, not a technology gap - and this distinction should change how an organization spends its budget. If the gap were technical, more engineering investment by the hyperscaler would eventually close it, and the correct posture would be to wait. It is not technical. CADA's own Level 3 and Level 4 assurance criteria are explicitly written around who owns the company and what law governs its staff, not around encryption strength, uptime, or physical security controls - all of which the hyperscalers have already largely solved [5]. An organization waiting for a hyperscaler to "fix" jurisdictional immunity through a future product announcement is waiting for a change of ownership structure. That change is not on any vendor's roadmap, because it would require restructuring the parent company, not the cloud region. The organization's own architecture is the only lever it directly controls. The vendor's product cycle is not.

“Our fraud-detection and AI workloads need GPU-backed inference. Does moving to a sovereign tier constrain that?”

Potentially, yes, and this is a genuine trade-off rather than a solved problem, which is exactly why the mosaic model, rather than a single-provider migration, is the right response. Sovereign hyperscaler tiers have historically launched with a narrower instance catalogue than the parent region, specifically lacking GPU-backed instance families that AI/ML fraud-detection inference depends on.

Rather than accepting degraded model performance for the sake of a uniform sovereignty posture, or over-exposing cardholder PII to a tier that cannot fully protect it, the more defensible pattern separates the workload.

Train and run non-PII, non-sensitive inference in the sovereign or standard hyperscaler tier, where the GPU catalogue is richest and innovation velocity matters. Route the cardholder-PII-bearing inference path itself, along with the associated data lineage evidence the AI Act will require, through the EU-native or private lane.

This is a design decision at the data-flow level, not a single infrastructure choice. It is one of the most common places Mphasis sees organizations either over-restrict their AI roadmap unnecessarily, or under-protect PII unnecessarily, when they have not separated the two questions.

12.3 Navigating EU Sub-Regional Regimes

“We are already GDPR- and DORA-compliant. Why does a national scheme like Germany's BSI C5 or France's SecNumCloud matter on top of that?”

Because GDPR and DORA set a floor, not a ceiling, and national supervisors enforce above that floor as a matter of routine practice, not exception. BSI C5:2026 is not a voluntary best-practice document. It is the de facto standard German banking supervisors, healthcare authorities, and public-sector procurement teams check against. It has 168 controls across 17 domains and requirements that go well beyond GDPR's text: a 10-minute burglary-resistance standard on physical barriers, 48-hour self-sufficient operation during a power outage, and a 3-hour patch SLA for critical (CVSS 9.0–10.0) vulnerabilities [14]. An organization that is GDPR-compliant but lacks C5 attestation is not in a grey area in Germany, it is, in practice, excluded from significant procurement segments.

SecNumCloud goes a step further than C5 in one specific respect: it is the only EU national certification that explicitly requires a provider to be able to resist an injunction issued under extraterritorial law that would compromise the service's availability, confidentiality, or integrity [15]. That is precisely the CLOUD Act exposure that GDPR and DORA do not, and cannot, resolve on their own, since both instruments regulate how data is protected once accessed, not whether a foreign court can compel access in the first place. For any workload a French or German supervisor would classify as sensitive, national certification is the more relevant test, and it is a materially higher bar than the EU-wide floor

“We operate outside the EU, for example in Switzerland. Does any of this apply to us?”

Yes, and in some respects the bar is higher, not lower. This is precisely because Switzerland sits outside the EU's harmonized floor and regulates outsourcing on its own terms.

FINMA Circular 2018/3 requires the outsourcing financial institution, not the provider, to retain full audit rights and a tested exit strategy before outsourcing even begins. The revised Federal Act on Data Protection restricts a cloud provider to processing data only exactly as instructed by the institution that engaged it.

Hyperscaler Swiss regions exist, but a Swiss region alone does not resolve the client identifying-data question Swiss regulators actually ask. That question is closer to the SecNumCloud extraterritorial-law test than to a GDPR residency test..

In Mphasis' experience, organizations with Swiss financial-sector counterparties consistently resolve this the same way: customer-held encryption keys physically on Swiss soil, paired with a Swiss domiciled provider or private infrastructure for the client-identifying data itself, while less sensitive workloads remain on the standard hyperscaler estate.

“How do we know which of our workloads actually require the strictest tier, rather than defaulting everything to the most expensive option?”

This is the single most valuable question a program can ask early. The opposite failure mode, defaulting every workload to the highest sovereignty tier out of caution, is nearly as costly as under-protecting sensitive data. It is just less visible on an audit finding.

The organization should apply the six-dimension test from Section 3 workload by workload, not estate-wide, and expect the answers to vary sharply. Real-time authorization and cardholder PII will usually score high on jurisdiction and cryptography. An internal analytics dashboard built on aggregated, non-PII data usually will not.

A jurisdiction-by-jurisdiction placement map, which workload, in which country, against which regime, is what turns this from a subjective debate into a defensible, documented classification a regulator or an internal audit committee can actually review.

Organizations that skip this step tend to either over-spend on unnecessary sovereign placement for low-risk workloads, or discover a genuine gap only during a regulator's review, both outcomes this classification step is designed to prevent.

12.4 Making the Architecture and Cost Decision

“Should the organization just migrate everything to an EU-native or sovereign provider to be safe?”

No - and Mphasis actively advises against that instinct, even though it feels like the conservative choice. A wholesale migration trades one form of concentration risk for another. It exchanges dependence on a hyperscaler's US parent for dependence on a smaller EU-native provider's more limited service catalogue, AI/ML capability, and elasticity, without necessarily buying more sovereignty than the specific workload required.

The mosaic principle set out in Section 6 exists precisely to avoid this over-correction. Hyperscaler tiers remain the right home for elastic, non sensitive, AI-native workloads where innovation velocity has real business value. EU-native, nationally-qualified providers carry the workloads where extraterritorial-law immunity is a hard requirement. Private and hybrid infrastructure, anchored by customer-held cryptographic keys, carries the cryptographic core.

Migrating everything to one lane, in either direction, reintroduces the single-point-of-failure risk the mandate was written to eliminate.

“We already run managed Kubernetes, such as AWS EKS or Azure AKS. Doesn't that already give us the portability the mandate is asking for?”

Partially, and this is one of the most persistent misconceptions Mphasis encounters in architecture reviews. Managed Kubernetes gives an application real portability at the workload-definition level. The container images and manifests will run on any CNCF-conformant cluster.

It does not, by itself, give the organization jurisdictional, operational, or cryptographic portability, because the managed control plane still sits on the provider's own compute, networking, storage, and identity primitives.

An EKS-hosted workload can be redeployed to another Kubernetes cluster in an afternoon; extracting it from AWS's IAM model, VPC constructs, and KMS-backed storage encryption is a considerably larger undertaking. It is that second exercise, not the first, that DORA's exit-readiness requirement and the EU Data Act's switching regime are actually testing for.

The practical implication is not to abandon managed Kubernetes; it is to treat portability as a property that has to be deliberately engineered into the workload's data, identity, and key design, rather than assumed because the workload happens to run in containers.

“Is bring-your-own-key (BYOK) enough, or do we need full external HSM ownership?”

For most workloads, BYOK is a meaningful improvement over provider-default key management. It is not the same control as external HSM ownership, and regulated data categories should not treat the two as interchangeable.

BYOK typically still allows the key to be loaded into the provider's own key-management service. This means the provider's infrastructure, and by extension the jurisdiction attached to that infrastructure, retains a path to the key material - Hold-your-own-key (HYOK) architectures, where the key is generated, stored, and used inside an externally validated hardware security module the customer controls, close that path. The provider is physically unable to decrypt the data even under a compelling legal order, because the key never resides inside its infrastructure at all.

For payments specifically, this distinction has a further technical dimension. Standard cloud key management services are not, and cannot be, PCI PTS HSM-validated to the standard payment-card cryptography requires, regardless of how the key was originated.

Mphasis' recommendation is straightforward. BYOK is an acceptable interim control for workloads with moderate sensitivity. For cardholder data, patient records, citizen identifiers, or any data category a regulator has named explicitly, external HSM custody, validated to the relevant standard, should be the target state, not BYOK.

“How do we justify the cost of a private or hybrid lane to the board when the sovereign hyperscaler tier looks cheaper on paper?”

The honest answer is that on a pure unit-cost comparison, a hyperscaler sovereign tier will usually look cheaper. A board that only sees a unit-cost comparison will usually, reasonably, ask why the organization needs anything else.

The comparison the board should actually see is framed differently: not "cost per workload," but "cost of the exposure a given placement leaves open." A real-time authorization workload placed on a sovereign hyperscaler tier without customer-held keys is not merely somewhat less compliant. It carries open exposure under Schrems II, the CLOUD Act, and, once financial-sector procurement adopts the tiers, CADA Level 3 classification, all three simultaneously. A single architecture decision, externalizing the cryptographic layer to a customer-owned HSM cluster, resolves this in one move.

Presented this way, the private/hybrid lane's incremental CapEx is priced against the cost of a specific, named regulatory exposure rather than against an abstract sovereignty preference. That reframing is usually what moves the conversation from a cost objection to a risk-acceptance decision the board can actually own, with the exposure made explicit.

12.5 Governance, Procurement, and the Board

“How does the organization evaluate a prospective sovereign provider without being misled by marketing language?”

Ask two questions that vendor marketing decks are structurally unable to answer well, because the honest answer to both is frequently unfavorable to the vendor.

First: "Who holds the CADA recognition or national qualification directly, the provider itself, or an intermediary or reseller arrangement?"

Second: "Given your current ownership structure today, what is your actual path to the next assurance level up, and is that path an engineering roadmap or a request for a political derogation?"

An answer that leans on an associated-country derogation, an aggregator relationship, or a national-partner-cloud structure that still runs licensed technology from a foreign parent is describing a political dependency, not an architectural one. A political dependency can be withdrawn by a decision the organization does not control.

A provider that can answer both questions with a direct ownership structure and a genuine engineering roadmap is describing something closer to real jurisdictional sovereignty.

This is deliberately a due diligence checklist, not a scoring rubric. The goal is to surface the honest gap, not to produce a pass/fail grade a sales cycle can optimize around.

“How do we sequence this transformation without stalling the organization's innovation roadmap?”

Sequence by workload risk, not by infrastructure lane. Run the highest-value, lowest-risk moves first to build institutional confidence before tackling the harder cryptographic-core migrations.

A workable order, consistent with what Mphasis has seen succeed elsewhere: begin with the workload classification and jurisdiction placement map described above. This alone produces decision-useful output without moving a single workload.

Externalize the cryptographic layer for the highest-exposure workloads next. This is the single control that resolves the most regulatory exposure per unit of engineering effort, and it can typically be implemented without a full infrastructure migration.

Run a contained exit-capability pilot: one representative workload, moved end-to-end between two lanes with a documented, tested playbook. This converts the exit-readiness requirement from a paper exercise into a rehearsed capability before it is needed under real pressure.

Only then extend the pattern estate-wide. Each stage produces a concrete artifact, a classification, a key-custody model, a tested runbook, that a board can review independently, rather than asking the board to approve a multi-year migration on faith.

13. AI-Accelerated Cybersecurity and Infrastructure Risk

Vulnerability discovery, exploit generation, and adversary automation are all compressing under AI-accelerated tooling. ENISA's 2025 EU Threat Landscape analysis, covering incidents from 1 July 2024 to 30 June 2025, documents the scale of this shift across the EU threat environment.

NSA and CISA's jointly published Kubernetes Hardening Guidance identifies Kubernetes clusters as a frequent target for data theft, computational-resource theft, and denial of service. It recommends vulnerability scanning, least privilege execution, network separation, strong authentication, log auditing, and periodic security review as baseline mitigations [18][19][20].

The architectural implication for this paper's framework: **infrastructure choice is now a cybersecurity decision, not only a sovereignty decision**. Stronger sovereignty tiers require stronger patch velocity, signed artifacts and provenance, hardened Kubernetes control planes, identity segmentation, runtime detection, customer-held keys, and rehearsed recovery paths. A managed platform can reduce day-to-day operational toil, but it does not remove the enterprise's responsibility for supply-chain governance, identity blast-radius, or cryptographic separation.

14. Modern Data and Metadata Governance

Data sovereignty fails without metadata sovereignty. Poor metadata management, fragmented data silos, and limited semantic interoperability are recognized barriers to compliant data sharing under the EU Data Act [2].

Open standards make governance operational rather than documentary. Apache Iceberg supports in-place schema and partition evolution as metadata-only operations, without rewriting underlying data files [21]. This is valuable for long lived, regulated data products that must remain auditable for years.

OpenLineage provides an open framework and standard API for capturing lineage events across datasets, jobs, and pipeline runs. Metadata platforms built on it can serve as a context graph for lineage, data contracts, and access policy across a data estate.

The practical operating pattern Mphasis recommends: domain-owned data products, machine-readable data contracts, a shared business glossary, end-to-end lineage, explicit quality SLAs, defined retention and allowed processing-location rules. Model lineage should be treated as a first-class extension of data lineage, not a separate discipline bolted on for AI Act compliance.




Figure 6. Sovereign Data and AI Governance Architecture - customer-held keys/HSM and privacy controls anchor a data management layer and an AI and governance layer, both of which feed a single audit-evidence plane for GDPR, DORA, the EU Data Act, and the AI Act

15. Governance Operating Model

Sovereignty by design requires a named owner for each control, not a diffuse “everyone's responsibility.” Mphasis recommends the following minimum governance structure, adaptable to existing risk and architecture committees:

Governance body	Core accountability
Sovereignty office	Workload tiering, jurisdictional controls, regulatory evidence
Architecture review board	Platform patterns, portability rules, anti-lock-in exceptions
Cryptographic authority	Key custody, HSM policy, rotation, revocation drills
Data governance council	Data products, metadata, lineage, retention
AI governance board	Model risk, AI lineage, human oversight, monitoring
Supplier risk forum	Concentration risk, subcontractors, exit readiness
Platform / SRE guild	Resilience testing, observability, runbooks, recovery

A Quarterly Board Question Set

Can the organization prove control under stress, with data-flow maps, key-custody records, resilience tests, exit drills, and a current supplier dependency register?

Is the organization buying sovereignty, or engineering it? Provider features alone are not sufficient evidence.

Where does the organization's real attack surface sit today, across applications, containers, Kubernetes, identity, hypervisors, CI/CD, and the supply chain?

Can the organization prove where its data came from and where it went, with data products, metadata, lineage, and governed catalogs?

Which suppliers create concentration risk today: critical function, no substitute, proprietary format, key control, or provider-native operations?

16. The Mphasis Sovereignty Manifesto

Principle	In practice
Sovereignty by design	Classify every workload before placement, using the six-dimension test.
Portability by design	Avoid embedding proprietary services in the critical path of a workload.
Exit by design	Test the exit before it is needed - a rehearsal, not a clause.
Cryptography by design	Separate keys from the provider wherever risk demands it.
Open standards first	Use open APIs, table formats, lineage, and metadata standards where they reduce risk.
Metadata as the control plane	Governance without metadata is not provable governance.
AI governance by design	Model lineage is data lineage - govern it as one discipline.

Principle	In practice
Mosaic over monolith	No single provider satisfies every sovereignty tier; design across lanes, not within one.

17. Conclusion

Sovereignty is the ability to prove control under stress: during a regulator review, a cyber incident, a provider outage, a geopolitical disruption, a model audit, or a forced migration. The evidence across public EU regulation, hyperscaler sovereign-cloud commitments, and a live cross-industry engagement all point to the same conclusion.

The winning architecture is not one cloud, one vendor, or one platform. It is a governed mosaic of portable workloads, sovereign data, customer-controlled cryptography, open metadata, hardened infrastructure, tested exit, and evidence-ready operations.

This applies consistently whether the workload is a payment authorization flow, a clinical AI pipeline, a citizen registry, a telecom control system, or an industrial digital twin.

The question every regulated enterprise should be able to answer today: can it prove control, revoke access, continue operations, explain its AI and data lineage, and exit its providers?

It should be able to answer that within its own risk tolerance, before a regulator or an incident forces the question.

References

All sources below are publicly available. Superscript numbers in the text correspond to the numbered list here; each entry indicates what it supports in this paper.

1

European Commission / EUR-Lex. Regulation (EU) 2022/2554 of 14 December 2022 on digital operational resilience for the financial sector (DORA). <https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng> - Supports: ICT risk management, third-party risk, resilience testing, exit-readiness obligations (Sections 1, 2, 9).

2

European Commission / EUR-Lex. Regulation (EU) 2023/2854 of 13 December 2023 on harmonised rules on fair access to and use of data (the Data Act). <https://eur-lex.europa.eu/eli/reg/2023/2854/oj/eng> - Supports: cloud switching, interoperability, egress-fee removal, metadata barriers to data sharing (Sections 1, 2, 9, 13).

3

European Commission / EUR-Lex. Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (the AI Act). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng> - Supports: AI governance obligations and phased applicability (Sections 1, 2).

4

European Data Protection Board. EDPB adopts recommendations on supplementary measures following Schrems II (41st plenary session, 11 November 2020). https://www.edpb.europa.eu/news/news/2020/european-data-protection-board-41st-plenary-session-edpb-adopts-recommendations_en - Supports: third-country transfer risk assessment, encryption as a supplementary measure (Sections 3, 10).

5

European Commission - Shaping Europe's Digital Future. Cloud and AI Development Act (CADA) - proposal of 3 June 2026. <https://digital-strategy.ec.europa.eu/en/policies/-cloud-and-ai-development-act> - Supports: four-level EU cloud sovereignty assurance framework (Sections 1, 4, 6).

6

AWS (Amazon Web Services). AWS Launches AWS European Sovereign Cloud and Announces Expansion Across Europe, 15 January 2026. <https://press.aboutamazon.com/aws/2026/1/aws-launches-aws-european-sovereign-cloud-and-announces-expansion-across-europe> - Supports: AWS ESC launch details, EU-resident operations (Section 4.1).

- 7 Cloud Security Alliance. AWS ESC: What to Know and Do, 16 January 2026. <https://cloudsecurityalliance.org/blog/2026/01/16/aws-launch-es-european-sovereign-cloud-what-you-need-to-know-and-what-you-need-to-do> - Supports: independent third-party analysis of AWS ESC governance model (Section 4.1).
- 8 AWS Security Blog. Exploring the new AWS European Sovereign Cloud: Sovereign Reference Framework. <https://aws.amazon.com/blogs/security/exploring-the-new-aws-european-sovereign-cloud-sovereign-reference-framework/> - Supports: AWS's own sovereignty control mapping (Section 4.1).
- 9 Microsoft. Announcing comprehensive sovereign solutions empowering European organizations, Official Microsoft Blog, 16 June 2025. <https://blogs.microsoft.com/blog/2025/06/16/announcing-comprehensive-sovereign-solutions-empowering-european-organizations/> - Supports: Microsoft Sovereign Cloud, Data Guardian, External Key Management (Section 4.2).
- 10 Microsoft Learn. Data Guardian overview - Microsoft Sovereign Cloud. <https://learn.microsoft.com/en-us/industry/sovereign-cloud/sovereign-public-cloud/capabilities/data-guardian> - Supports: Data Guardian's tamper-evident, regionally-approved access model (Section 4.2).
- 11 Microsoft Learn. What is Microsoft Sovereign Cloud? <https://learn.microsoft.com/en-us/azure/azure-sovereign-clouds/microsoft-sovereign-cloud> - Supports: Sovereign Public Cloud, Sovereign Private Cloud, and National Partner Cloud models (Section 4.2).
- 12 Google Cloud. Sovereign Controls by Partners - product page. <https://cloud.google.com/security/products/sovereign-controls-by-partners> - Supports: partner-operated key management and access-justification model (Section 4.3).
- 13 Google Cloud Documentation. Overview of Sovereign Controls by Partners. <https://docs.cloud.google.com/sovereign-controls-by-partners/docs/overview> - Supports: list of live EU sovereign-controls partners including S3NS and T-Systems (Section 4.3).
- 14 Kiteworks. BSI C5 - Germany's Cloud Security Framework Requirements. <https://www.kiteworks.com/regulatory-compliance/bsi-c5-germanys-cloud-security-framework-requirements/> - Supports: BSI C5:2026 control catalogue and German market implications (Section 5.1).
- 15 ANSSI (Agence nationale de la sécurité des systèmes d'information). FAQ - Qualification SecNumCloud. <https://cyber.gouv.fr/enjeux-technologiques/cloud/faq-qualification-secnumcloud/> - Supports: SecNumCloud's explicit extraterritorial-law immunity requirement (Section 5.2).

- 16 Cloud Native Computing Foundation. Certified Kubernetes Conformance Program. <https://www.cncf.io/certification/software-conformance/> - Supports: Kubernetes API conformance and cross-installation interoperability claim (Section 7).
- 17 Red Hat. OpenShift - hybrid cloud application platform. <https://www.redhat.com/en/technologies/cloud-computing/openshift> - Supports: consistent hybrid-cloud foundation and CNCF-certified compatibility claim (Section 7).
- 18 Cybersecurity and Infrastructure Security Agency (CISA). CISA and NSA Release Kubernetes Hardening Guidance. <https://www.cisa.gov/news-events/alerts/2021/08/02/cisa-and-nsa-release-kubernetes-hardening-guidance> - Supports: Kubernetes-specific attack-surface and hardening recommendations (Section 13).
- 19 National Security Agency / CISA. Kubernetes Hardening Guide, Version 1.2, August 2022. https://media.defense.gov/2022/Aug/29/2003066362/-1/-1/0/CTR_KUBERNETES_HARDENING_GUIDANCE_1.2_20220829.PDF - Supports: detailed mitigation list - scanning, least privilege, network separation, auditing (Section 13).
- 20 European Union Agency for Cybersecurity (ENISA). ENISA Threat Landscape 2025. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025> - Supports: scale and complexity of the EU cyber threat environment (Section 13).
- 21 Apache Software Foundation. Apache Iceberg — documentation. <https://iceberg.apache.org/docs/latest/> - Supports: in place schema and partition evolution as a metadata sovereignty enabler (Section 14).
- 22 Gibson Dunn. EU AI Act Omnibus Agreement - Postponed High-Risk Deadlines and Other Key Changes, May 2026. <https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/> - Supports: deferral of the Annex III high-risk AI deadline (Section 2.1).
- 23 FelloAI. EU AI Act 2026: Enforcement, Deadlines and Fines. <https://felloai.com/eu-ai-act/> - Supports: confirmation of Regulation (EU) 2026/1744 (Digital Omnibus) adopted 27 July 2026 (Section 2.1).

24

Legiscope. EU AI Act Deadlines 2026–2027: Compliance Calendar and Fines. <https://www.legiscope.com/blog/eu-ai-act-timeline-deadlines.html> - Supports: GPAI and transparency obligations unaffected by the Annex III deferral (Section 2.1).

25

The Register. Microsoft admits it 'cannot guarantee' data sovereignty, 25 July 2025. https://www.theregister.com/2025/07/25/microsoft_admits_it_cannot_guarantee/ - Supports: Microsoft France's French Senate testimony on CLOUD Act exposure (Sections 4.2, 12.2).

Explore **Mphasis Next-Gen Data** Solutions



Scan to see how we help
regulated enterprises secure and modernize their data.

Ready to take control of your data?

Talk to our team about keeping control of your data and your cloud options.

Contact Us

For more information, contact: marketinginfo.m@mphasis.com

USA

Mphasis Corporation
41 Madison Avenue
35th Floor, New York
New York 10010, USA
Tel: +1 (212) 686 6655

UK

Mphasis UK Limited
1 Ropemaker Street, London
EC2Y 9HT, United Kingdom
T : +44 020 7153 1327

INDIA

Mphasis Limited
Bagmane World Technology Center
Marathahalli Ring Road
Doddanakundhi Village, Mahadevapura
Bangalore 560 048, India
Tel.: +91 80 3352 5000

Copyright © Mphasis Corporation. All rights reserved.



www.mphasis.com